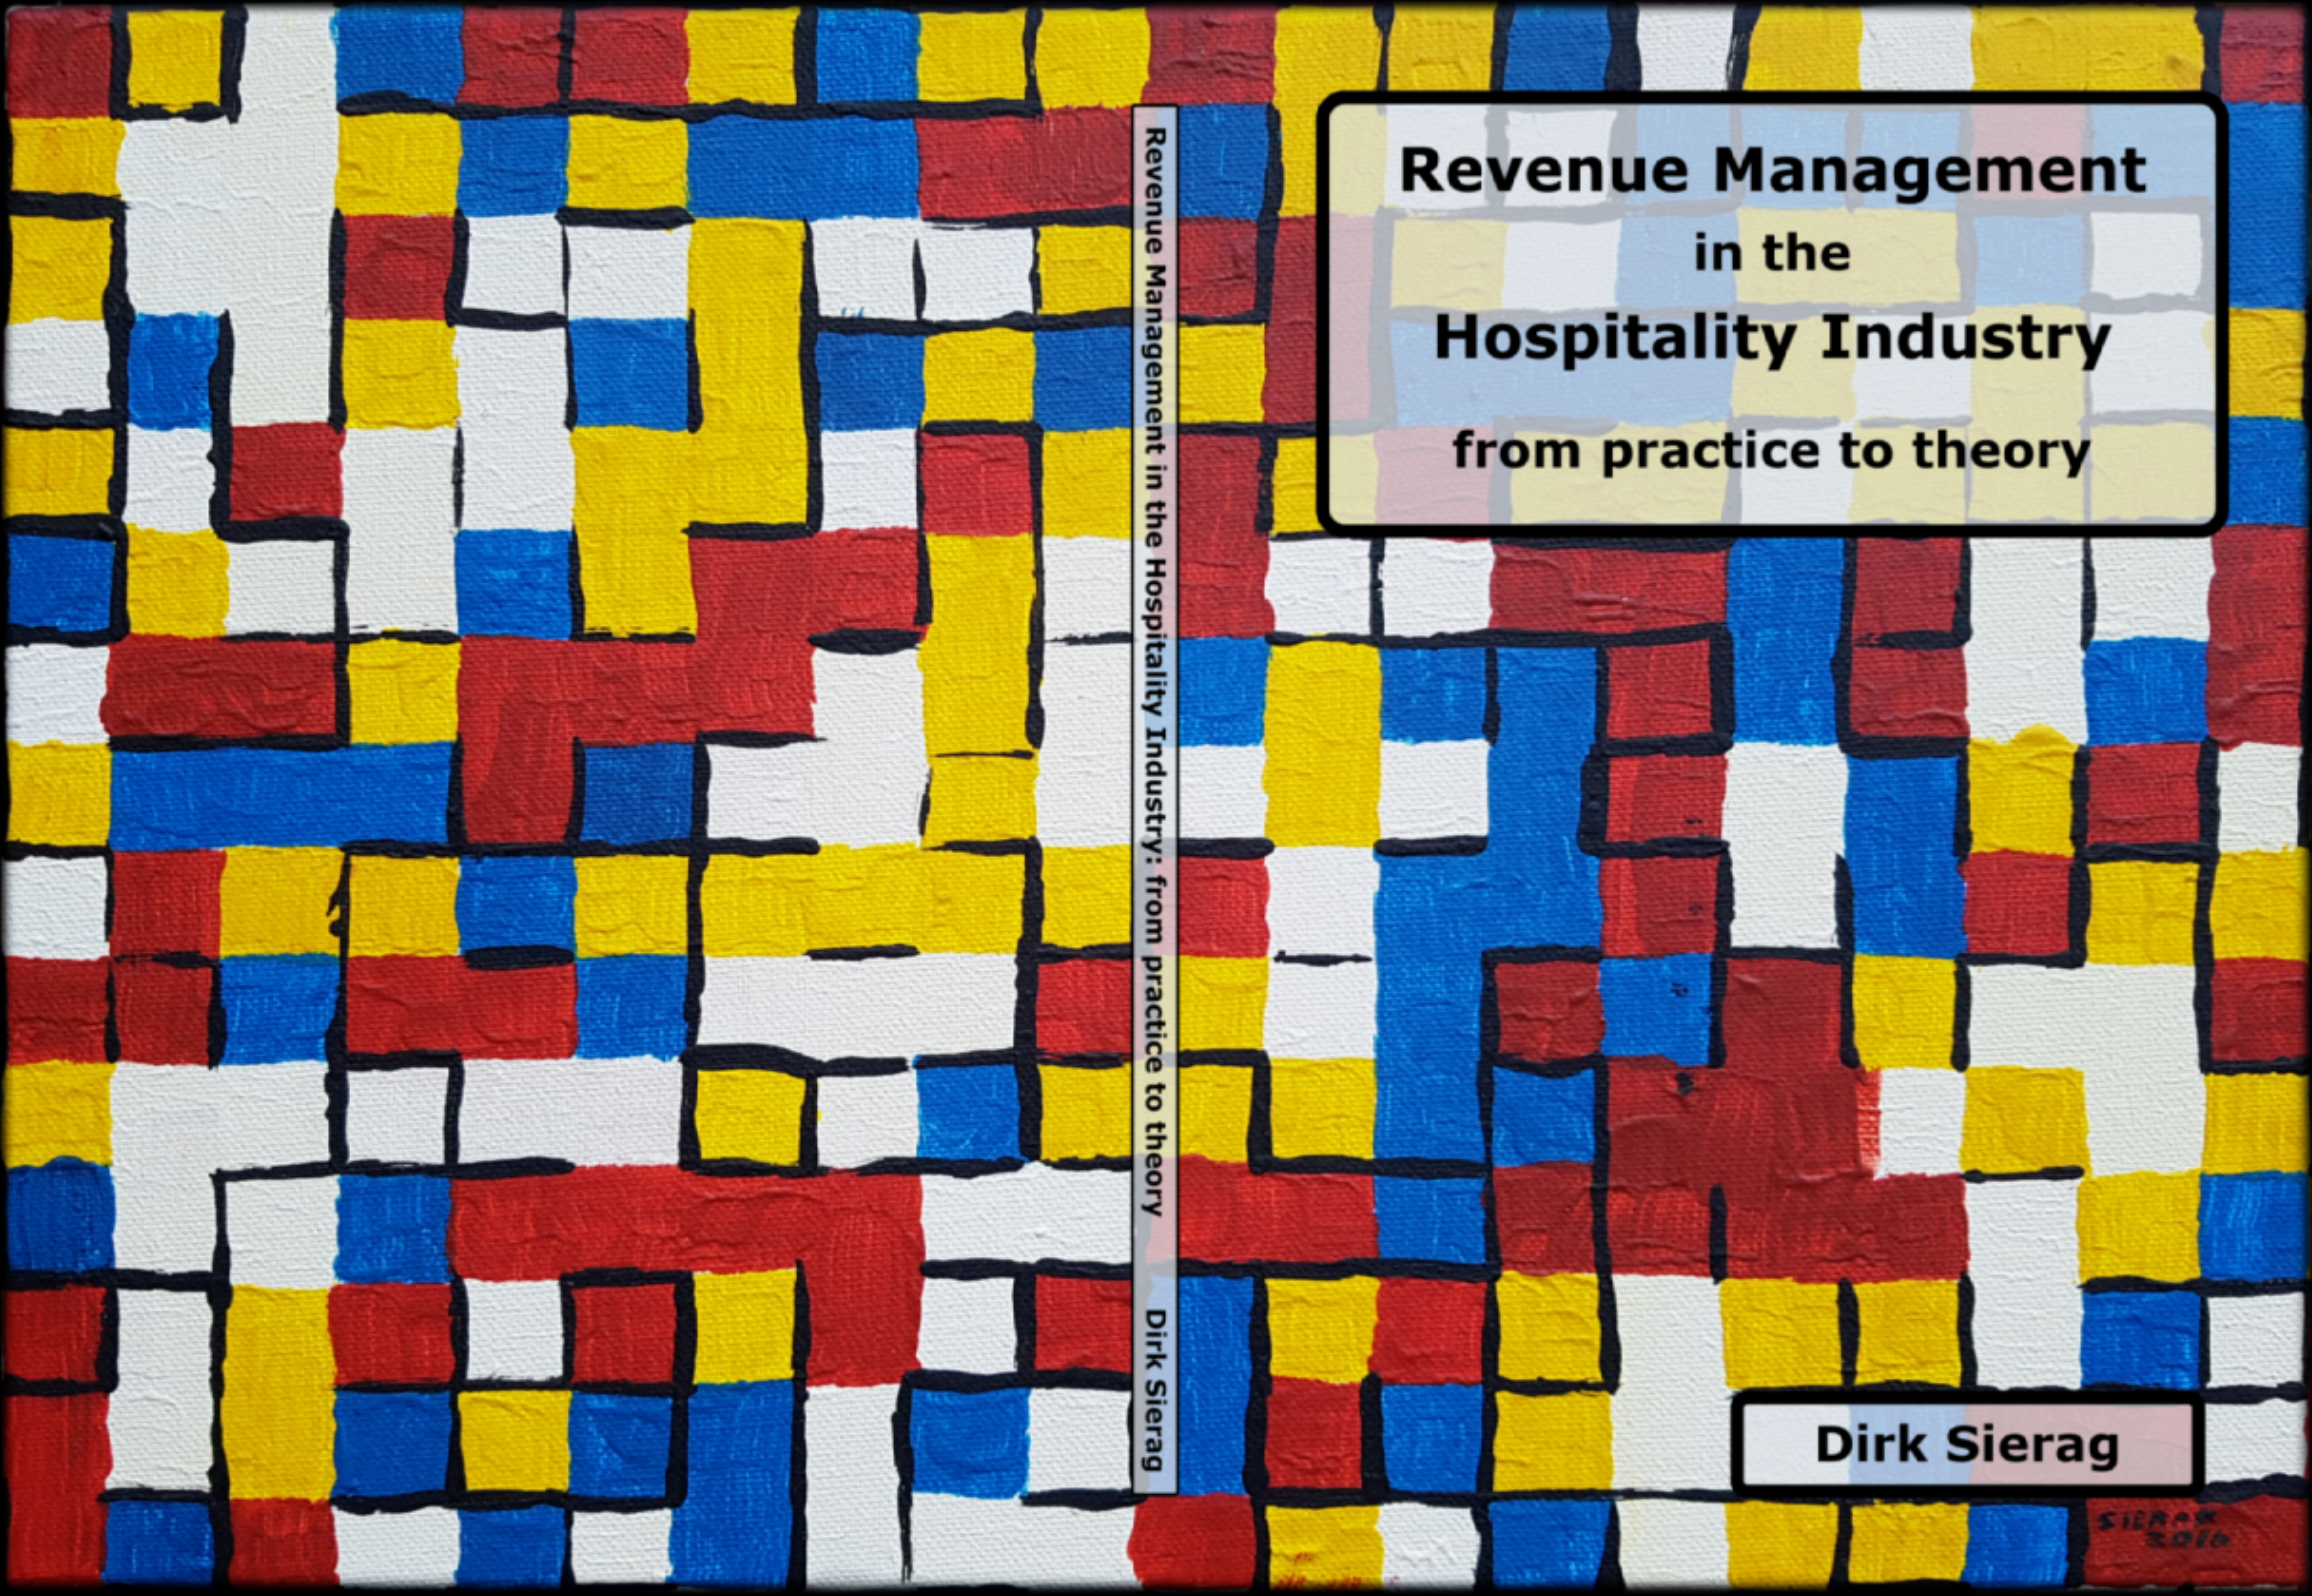

Revenue Management in the Hospitality Industry: from practice to theory   Dirk Sierag

Revenue Management
in the
Hospitality Industry
from practice to theory

Dirk Sierag

Revenue Management in the Hotel Industry

from practice to theory

Sierag, Dirk Daniël
Revenue Management in the Hotel Industry – from practice to theory

ISBN: 978-94-028-0551-2

The research in this thesis was supported by the Netherlands Organization for Scientific Research (NWO), as part of the research cluster 'Stochastics: Theoretical and Applied Research' (STAR); and by the Ministry of Education, Culture and Science (OCW) via the Regional Attention and Action for Knowledge circulation (RAAK), Foundation Innovation Alliance (SIA).

Published by Dirk Sierag, Leiden, The Netherlands.
Cover design by Dirk Sierag, Leiden, The Netherlands.
Printed by Ipskamp Printing, Enschede, The Netherlands.

VRIJE UNIVERSITEIT

Revenue Management in the Hotel Industry

from practice to theory

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad Doctor aan
de Vrije Universiteit Amsterdam,
op gezag van de rector magnificus
prof.dr. V. Subramaniam,
in het openbaar te verdedigen
ten overstaan van de promotiecommissie
van de Faculteit der Exacte Wetenschappen,
op donderdag 30 maart 2017 om 11.45 uur
in het auditorium van de universiteit,
De Boelelaan 1105

door

Dirk Daniël Sierag

geboren te Leiden

Promotoren:	prof.dr. R.D. van der Mei
	prof.dr. A.P. Zwart
Copromotoren:	prof.dr. G.M. Koole
	prof.dr. J.I. van der Rest

Acknowledgement

"Nowadays people know the price of everything, and the value of nothing."

The Picture of Dorian Gray
Oscar Wilde

With this dissertation I conclude three and a half years of research at Centrum Wiskunde Informatica, Vrije Universiteit Amsterdam, Hotelschool The Hague, and Leiden University. During this period I had the opportunity to interact with inspiring persons from all over the world, which made me realise how privileged I was. I want to thank the people who supported me, each in their own way, along the way and to the next step.

First and foremost, I want to thank my supervisors Ger Koole, Rob van der Mei, Jean-Pierre van der Rest, and Bert Zwart, who gave me the freedom I needed to develop my own ideas. Their different perspectives broadened my horizon and helped me solve the challenges I faced. I would also like to thank my collaborators from industry, in particular the cooperating hotel managers and Niels Schenk and Yvo Keuters, for sharing their experience and constructive input. I would like to thank the external committee members Sandjai Bhulai, Vivek Farias, Peter Grünwald, Jeroen Mulder, and Claudius Steinhardt for studying and evaluating my work.

I am grateful for all the great memories I have with my colleagues: playing pingpong, chilling at the awesome PhD weekends, chitchatting at *het theekransje*, or going out after conferences. However, without a doubt, M335 takes the cake. Finally, I am thankful for all the love and support that I received the last years, especially from my parents, Wim and Suzanne, and my sisters, Esther and Rebecca.

Dirk Sierag
Leiden, January 2017

Contents

CHAPTER 1

Introduction

In this dissertation revenue management (RM) and pricing methodologies are studied and developed, motivated by challenges and opportunities faced by the hospitality industry. RM is a discipline applying science and analytics in the relentless pursuit of increasing long-term profitability through better business decisions. In a capitalistic world with a highly competitive market it is essential for a company's survival and growth to apply RM. Classic examples can be found in the aviation industry and the hotel industry, which attempt to keep or increase their market share and revenue by dynamically deciding when to offer which seats or rooms, for what price, at what time, and to which customers according to information on amongst others demand forecasts, price sensitivity, and capacity constraints (Vinod, 2004; Clarke, 2004). However, RM can be applied to a wide range of fields, including but not limited to the car rental industry (Geraghty & Johnson, 1997), theatres (Langeveld, 2006), golf courses (Kimes, 2000), restaurants (Guerriero et alii, 2014), cruise ships (Maddag et alii, 2010), and TV and digital advertising (Gallego & Phillips, 2004). All application areas have in common that a limited number of perishable goods is sold, whether it is a hotel room, a seat on an aircraft, a rental car, a seat in the theatre, a timeslot on a golf course, a cabin on a ship, or an impression on TV or a website. RM models exploit this fact by dynamically adjusting the price and availability of a product according to the demand and remaining capacity of the products.

RM is a complex task that involves knowledge on numerous factors that impact price sensitivity and demand, i.e., factors that drive clients to purchase a company's product. For example, for a hotel the physical location, online rating and (recent) reviews, and available facilities, like a spa or fitness centre, have a great impact on demand volume and customers' willingness to pay. A company can pinpoint factors and the effects on demand by analysing available data (e.g., historical sales and online reviews) or by means of surveys (e.g., to explore the positive effect on demand of building a fitness centre). The company then implements these factors in an optimisation model with the objective of maximising revenue under certain conditions. This line of thought is followed here: the research contributions of this dissertation build directly on findings from practice, in particular the hotel industry. The source of the findings range from our own thorough data analysis, field studies reported in literature, and intense collaborations with practitioners. Below three important issues are given that hotel managers face, and which chapters address these issues in detail.

1. Hotel managers report that a large part of reservations are cancelled. This is sup-

ported by the results from the hotel data analysis, presented in Chapter 2, which shows that 21.7% of the reservations are cancelled between the years 2008 and 2012. Chapters 3, 4 and 6 exploit this by incorporating *cancellations* into the revenue optimisation process. Numerical experiments show that this leads to a revenue increase of up to 20%.

2. According to hotel managers, demand is very uncertain and varies a lot from day to day, even considering seasonality effects. Our data analysis supports this claim, by providing statistical evidence that demand data follows a nonhomogeneous Poisson distribution, a random process indeed. To this extend, all of our models assume that demand follows a (nonhomogeneous) Poisson distribution. Moreover, Chapters 4 and 7 provide solution methods to deal with uncertainty in demand by means of *robust optimisation* techniques. The numerical studies in these chapters show increased revenues of up to 3.25

3. Hotel managers claim that the wide available *online reviews* and *ratings* on booking websites like Booking.com, Expedia, and Tripadvisor have a big impact on demand. Studies in literature support this claim with empirical evidence (Pan & Zhang, 2011; Park & Lee, 2009; Yoo & Gretzel, 2011). On the other hand, the price/quality perception of the guest affects her motivation to write a positive or negative review about her experience (Zhou et alii, 2014). Chapters 6 and 7 model both effects of reviews explicitly in the optimisation process. It is shown that solely optimising revenue, and not considering ratings or reviews, reduces the rating and quality of reviews, which results in lower demand and lower revenues in the long run. Numerical studies in these chapters show an increase in revenue of up to 11% and 5.2%, respectively, when the effects of reviews are accounted for.

Next, a concise history of RM is given, with focus on the relation between practice and theory. This is then followed by an outline and contribution of this dissertation.

1.1 History of Revenue Management

The concepts and goals of RM date back to prehistoric times, when the first trade economies emerged around the world. Merchants at that time faced similar decision problems to optimise revenue as nowadays (Talluri & van Ryzin, 2004b). For instance, merchants had to decide what price to offer, on what market to sell, and whether it would be better to postpone offering the products until a future time, when one might get a better price. People have been very innovative and creative to come up with pricing RM strategies, improving economies to this very day. Hence, the *nature* of the questions that RM attempts to answer is not new, but rather the *methodology* to answer those questions is new: a data-driven, mathematical and technologically advanced approach.

The discipline of RM originated in the airline industry, in the 1970's, when sophisticated computer reservations systems provided detailed sales history data. Littlewood (1972) described the case of British Overseas Airways Corporation (BOAC, currently known as British Airways), who started offering a limited number of seats on flights at a discounted price for clients who purchased at least 21 days in advance. Through mathematical analysis BOAC was able to forecast how many seats it should reserve for the full fare in a way that no seats would remain empty and that full fare products would not be cannibalised, i.e., that full fare sales would not diminish.

The airline deregulation act of 1978 in the USA is often considered as the starting point for RM (Talluri & van Ryzin, 2004b). Between 1938 and 1978 domestic interstate flights of the USA were considered public utilities, just like electricity, postal services, and other forms of public transport, and routes, schedules, and prices were set by the federal Civil Aeronautics Board (CAB). After the deregulation, airlines were able to set their own routes, schedules, and prices. As a consequence, budget airlines emerged throughout the USA. A famous example is People Express Airlines (PEOPLExpress), which was founded in 1981 and offered flights at 50-70% lower prices than the establishes airlines at that time. PEOPLExpress was a great success, it increased market share and profits rapidly at the expense of established airlines, in particular American Airlines. To prevent losing more market share, American started implementing a large-scale RM system called *Dynamic Inventory Allocation and Maintenance Optimizer (DINAMO)*. With DINAMO American matched or undercut prices from PEOPLExpress by offering a limited number of ultimate super-saver fares which were non-refundable and had to be bought at least 30 days in advance. This way American offered competitive prices that could not be matched by PEOPLExpress, which caused American's revenues and profits to increase by 14.5% and 47.8%, respectively, and caused the abrupt bankruptcy of PEOPLExpress (Cross, 1997).

A pioneering example of hotel RM, following the airline industry, is the case of Marriott International, which camped with similar features as American did, such as perishable capacity constraints, budget competition, and advance reservations. Marriott created a demand forecasting system, taking into account the additional complexity of hotels caused by length of stays (LOS's), which helped increase revenues of the hotel chain by $150-200 million annually (Cross, 1997). Another prominent example and application area is found in the car rental industry, where the North American company National Car Rental, faced by liquidation in 1993, initiated a comprehensive RM program of analytical models to manage capacity, pricing, and reservations. Not only did National Car Rental manage to return to profitability, but it even managed to increase annual revenues by $56 million (Geraghty & Johnson, 1997).

Parallel to the development of RM systems in practice the research area of RM was established as a branch of operations research. Motivated by practical situations, mathematical problems were formulated and solution methods were studied. The pioneering work of Littlewood (1972) introduced a capacity control model for two fare-classes on a single flight leg. The natural extension of this model is to include multiple fare-classes, which can be solved exactly by using dynamic programming. However, these solution methods developed after practice was already using heuristical approaches, such as the popular *Expected Marginal Seat Revenue (EMSR)* solution methods by Belobaba (1987b,a, 1989), that are still used in practice. Extensions to these basic models include buy up/buy down effects and optimisation of *network* problems with multiple resources, e.g., optimisation of origin-destination journeys consisting of multiple flights, and multiple night stays in hotels; see McGill & van Ryzin (1999) for an overview.

One characteristic that is well studied in literature, but difficult to implement in practice, is the problem of modelling purchasing behaviour of clients by means of *choice models*. Earlier RM models assumed that clients were not comparing different products when making a purchase; rather, these RM models assumed that demand for products is independent from the *offer set*, i.e., the set of products that is offered. The seminal work by Talluri & van Ryzin (2004a) provides a detailed analysis of a single-leg RM model that takes customer choice behaviour into account. Following their work, Gallego et alii (2004) and Liu & van Ryzin (2008) provide analyses of choice-based network models, along with heuristics to overcome the curse of dimensionality caused by the state space, which becomes prohibitively large in practical examples. Despite the fact that literature has shown the benefits and potential revenue improvements of choice models, practice has been reluctant to apply such sophisticated models. One reason is that the choice models cannot simply be implemented as an add-on to old RM systems, because measuring choice behaviour requires investing in a new system that keeps track of more records than only sales data. InterContinental Hotel Groups (IHG) is a pioneer in implementing a choice-based RM system, and with great success: by optimally determining room rates based on occupancy, price elasticity, and competitor prices, IHG was able to increase revenues annually by 2.7% (Koushik et alii, 2012).

Over the years RM systems have become an integral part of many companies in a variety of industries. Many major companies have their own RM department, where new strategies and RM tools are developed and implemented. Moreover, specialised RM firms, such as PROS and IDeaS, have emerged, offering state-of-the-art RM tools for small and medium enterprises (SME) as well as assisting major companies with their RM system. Recent technological advancements sound the beginning of a new era in RM: a data-driven approach, where vast amounts of data ('big data') on client behaviour and sales are collected and analysed. Recent start-up companies, such as Uber and Airbnb, thrive on these data-driven RM methodologies. By collecting a unique data set of individual customers through their apps and websites, they can exploit unconstrained demand information to segment customers and to set prices accurately for maximal revenue.

1.2 Contribution of this Dissertation

This dissertation is the combined work of four years of research, which resulted in six journal papers, each represented here as a single chapter. Each chapter can be read independently, including references to relevant literature, while links between different chapters are marked where necessary. The following summarises the research contribution and gives an outline of the dissertation.

Using five years of data collected from a small and independent hotel in The Netherlands, the case study in Chapter 2 explores RM system data as a means to seek new insights into occupancy forecasting. The study provides an insight into the random nature of *group cancellations*, an important but neglected aspect in hotel RM modelling. The empirical study also shows that in a local market context demand differs significantly per point of time during the day, in addition to a seasonal monthly and weekly demand pattern. Moreover, the study presents evidence that demand follows an *nonhomogeneous Poisson distribution*, a crucial characteristic for forecasting modelling that is generally assumed but not reported in the hotel forecasting literature. This implies that demand is more uncertain for smaller than for larger hotels. By reporting the results of an in-depth case study, Chapter 2 seeks to draw attention to the critical and often overlooked role of exploratory data analysis in hotel RM forecasting.

In many application areas such as airlines and hotels a large number of bookings are typically cancelled. Explicitly taking into account cancellations creates an opportunity for increasing revenue. Motivated by this, Chapter 3 proposes a RM model is that takes *cancellations* into account in addition to customer choice behaviour. Moreover, overbooking limits are considered, as these are influenced by cancellations. The problem is modelled as a Markov decision process and three dynamic programming formulations are proposed to solve the problem, each appropriate in a different setting. The study proves that in certain settings the problem can be solved exactly by using a tractable solution method. For other settings, where the problem is intractable due to the curse of dimensionality, tractable heuristics are proposed. Numerical results show that the heuristics perform almost as good as the exact solution. However, the model without cancellations can lead to a revenue loss of up to 20%. Also, a parameter estimation method is provided that is fast and provides good parameter estimates. The combination of the model, the tractable and well-performing solution methods, and the parameter estimation method ensures that the model can efficiently be applied in practice.

A popular trend in RM captures the behaviour of customers that choose between different available products. The provided solution methods assume that there is no uncertainty in the parameters of the model. However, in practice the parameters may be uncertain, e.g., because of estimation errors. A relatively recent field of optimisation that takes into account uncertainty in the optimisation procedure is *robust optimisation*. Robust optimisation methods provide solutions where the worst-case scenario is optimised, taking into account uncertainty in parameters. Chapter 4 studies a robust optimisation approach to single-leg choice-based RM based on Talluri & van Ryzin (2004a) and Sierag et alii (2015). The problem is modelled as a Markov decision process and solved using dynamic programming. This chapter uses ϕ-divergence uncertainty sets to model the probability vectors of general choice-models. Novel robust optimisation techniques are applied to the dynamic program, taking into account uncertainty in the parameters. An important yet surprising insight from the numerical results is that the robust solution method performs better for smaller inventory than for larger inventory. Moreover, the robust solution method shows great performance when knowledge on cancellation behaviour is lacking: on average the expected reward then improves by 2.5-3.25%.

Chapter 5 proposes and analyses a pricing-based RM model that allows *flexible products* on a network, with a non-trivial extension to group reservations. Under stochastic demand the problem can be solved using dynamic programming, though it suffers from the curse of dimensionality. The solution under deterministic demand gives an upper bound on the stochastic problem, and serves as a basis for two heuristics, which are asymptotically optimal in capacity and demand. The numerical study, which is based on a problem instance from practice, shows that the heuristics perform well, even under uncertainty in demand. Moreover, neglecting flexible products can lead to substantial revenue loss.

Chapter 6 proposes a RM model that integrates *reviews* and *ratings*. The dependency between reviews and revenue is two-fold: on the one hand reviews impact demand, and on the other hand customers write reviews based on their price/quality perception. A complicating factor in this model is that the effects of reviews are delayed. For instance, by sacrificing revenue now in order to get better reviews, long-term revenue can be increased. Because the full planning problem of finding an optimal strategy for the proposed model is intractable, a novel solution methodology is proposed to solve the problem approximately by restricting the space of possible solutions to *equilibrium strategies*. The study shows that equilibrium strategies for the full problem can be found by viewing the full problem as a series of multi-objective Markov decision processes subproblems, while aiming to keep the balance between positive and negative reviews constant to a *target review ratio*. Numerical studies show that taking reviews into account in this manner can lead to an increase in revenue of up to 11% compared to the case where the sole objective is revenue.

Chapter 7 proposes a choice-based network RM model that integrates the effect of reviews. Faced by the complexity of the model, two heuristics are proposed, one of which uses robust optimisation techniques. Numerical results show a 3.5-5.2% improvement when reviews are taken into account. Moreover, the impact of reviews is larger under low demand intensity than under high demand intensity.

CHAPTER 2

Exploratory Data Analysis in Revenue Management Forecasting: a Case Study of a Small and Independent Hotel in The Netherlands

In this chapter a detailed analysis of data from a small and independent Dutch hotel is performed, stressing the importance of regularly analysing data in revenue management (RM). The study gives insight in the random nature of *cancellations* of (group) reservations and the effect on demand and revenue. The study also shows evidence that demand follows a *nonhomogeneous Poisson distribution*, an important characteristic that is often assumed in demand forecasting models without any evidence. A surprising result that follows from the analysis is that demand from business and leisure clients differs significantly per point of time during the day.

2.1 Introduction

Since the early 1990s hotel RM practice has evolved gradually (Ferguson & Smith, 2014) setting off large investments in sophisticated RM systems (RMS). Whilst varying in structure these RMS essentially calculate and update booking limits within a reservation system, extracting and processing information from various other systems (Phillips, 2005). One of these systems, lying at the heart of each RMS, is forecasting (Lemke et alii, 2013). As Talluri & van Ryzin (2004a, p. 407) observe: 'a RM system requires forecasts of quantities such as demand, price sensitivity, and cancellation probabilities, and its performance depends critically on the quality of these forecasts.' While there is ample research on forecasting, a major weakness of work in hotel RM is its focus on the model selection aspect of hotel forecasting, with notable exceptions such as Schwartz & Hiemstra (1997), Kimes (1999), Schwartz & Cohen (2004b), Schwartz & Cohen (2004a), Bendoly (2013), Koupriouchina et alii (2014), and Van der Rest et alii (2016). Forecasting comprises multiple facets including (a) problem definition, (b) information gathering, (c) preliminary (exploratory) data analysis, (d) choosing and fitting models, and (e) evaluating and adjusting the model (Makridakis et alii, 1998). In hotel RM all steps are critical and overlooking any of these steps can lead to under performance of the RMS. Moreover,

This chapter is based on Sierag, van der Rest, Koole, van der Mei, & Zwart (2016).

even after an initial round of model selection and evaluation, new data analysis will be required: hotels operate in a changing environment affecting the nature of the data, and thus adjustments to the model analysis may be required. Yet, most research focuses on defining a forecasting problem, developing or selecting a forecasting model, and testing the model. The crucial steps of information gathering and (preliminary) data analysis are often overlooked.

Motivated by this, this chapter aims to draw attention to the importance of regular data analysis by demonstrating how a real-life hotel can gain new forecasting insights by exploring and analysing data from its RMS. To this purpose, key factors of hotel demand, price sensitivity, and cancellations are identified, by analysing data from a small and independent hotel in The Netherlands. In particular, group cancellation behaviour, the effects of uncertainty in demand, and different dimensions of seasonality are studied. The remainder of the chapter is organized as follows. First the background of the research problem is explained in the remainder of this section. In section 2.2 the data set and hotel case study are described. Sections 2.3, 2.4 and 2.5, respectively, provide the case study findings, in particular insights into (1) different levels of seasonality; (2) group cancellation behaviour; and (3) uncertainty in demand and cancellations. Finally, the chapter discusses the findings, the limitations of the research and provides directions for future research.

2.1.1 Background

Forecasting is an area in operations research which over the years has grown into a whole discipline of its own with specialist research attention from a wide range of disciplines and sectors (Fildes et alii, 2008). For example, forecasting has received continuous research attention in tourism with work as early as Fritz et alii (1984), and with advanced contributions such as Li, Song, & Witt (2006), Li, Wong, et alii (2006), and Song et alii (2013). As Li et alii (2005) and Song & Li (2008) identified in two comprehensive literature reviews, 451 studies on tourism demand modelling and forecasting were published during the period 1960-2008. The hospitality literature has traditionally paid little attention to forecasting in hotel RM with the exception of Law (1998), Weatherford et alii (2001), Cranage (2003), Law (2004), Lim et alii (2009), Farouk El Gayar et alii (2011), Yang et alii (2014), and Koupriouchina et alii (2014), and Koupriouchina et alii (2017). In the operations research literature a stream of forecasting applications in hotels can be observed with work from Rajopadhye et alii (2001), Baker et alii (2002), Brännäs et alii (2002), Weatherford & Kimes (2003), Aghazadeh (2007), Chen & Kachani (2007), Yüksel (2007), Bermúdez et alii (2009), Guadix et alii (2010), Haensel & Koole (2011a), and Zakhary et alii (2011).

Hotel RMS traditionally assume that demand for each rate class is distinct and independent of the alternative options hotel guests have when booking a room. To challenge this common assumption and to incorporate other important buying behaviour aspects, customer choice models have been proposed in the RM literature (see Chapter 3 and Talluri & van Ryzin, 2004a; Meissner & Strauss, 2012; Liu & van Ryzin, 2008; Aydin et alii, 2012; Erdelyi & Topaloglu, 2010). When customer choice behaviour is incorporated, data analysis research will be especially important as in order to apply customer choice modelling to hotel RM practice successfully the appropriate choice (and estimation) of model parameters is crucial (e.g., van Ryzin, 2005; van Ryzin & Vulcano, 2013; Newman et alii, 2014). Bodea et alii (2009, p. 356) criticize the literature as 'the measurement of revenue benefits associated with choice-based RM has been based primarily on simulated data.' They argue that there is a need to test these models on real data sets to see if the customer choice concept really works as: 'choice-based systems are not simply an incremental improvement or 'add-on' to existing product-based systems, but are fundamentally different. Consequently, successful implementation of these systems will require a company to invest significant resources in developing new data collection procedures, RM algorithms, and user support systems' (Bodea et alii, 2009, p. 357).

Bodea et alii (2009) describe the laborious process of data collection and validation in order to provide a data set that could be used to benchmark the choice-based models proposed in the RM literature. They developed a data set based on five hotel properties and discuss its potential uses including 'proofing of concepts' and 'benchmarking'. Their study shows how crucial data collection is especially as a precursor to demand and forecasting model development. Studies focusing exclusively on data analysis, such as graphing data (for visual inspection), computing statistics (for relationships), decomposition analysis (for trends, unusual or extreme data points), however, are virtually non-existent. This is an important omission as exploratory analysis is key to the selection of the class of quantitative models (Makridakis et alii, 1998). Moreover, the academic literature on forecasting in hotel RM with an inclination for modelling makes many assumptions about the properties and nature of data, but which often are not supported by preliminary empirical research.

2.2 Case Description

Five years of data (2008-2012) was collected from a small and independent hotel. The utilization of such data is of theoretical and practical importance as little is known about RM in this type of hotel, which makes up the majority of all hotel properties in Europe (Luciani, 1999; Holverson & Revaz, 2006). Moreover, small and independent hotels generally do not employ a revenue manager who interacts with the RMS (Lee-Ross & Johns, 1997). This is an important criterion as the data was thus not limited by endogenous system effects.

The hotel is located in the countryside in The Netherlands and attracts business as well as leisure clients. As Table 2.1 illustrates, the hotel has 34 rooms which are partitioned into six room types each with a typical price.

Room type	Abbreviation	# Rooms	Typical price
Standard	STD	8	119
Garden view	GV	8	127
Large garden view	LGV	6	134
Old	STO	6	103
Private garden	PG	5	140
Bridal Suite	BRD	1	140

Table 2.1: Overview of room types and prices.

All rooms have a maximal capacity of two persons. The hotel has other facilities such as conference rooms and a restaurant. The restaurant not only serves hotel guests but also locals and tourists from the area.

Collecting the data was a lengthy process. Interaction with the hotel owner, the property system vendor, and two RM experts were needed to ensure data integrity. The data set had the following structure. Each data entry was a reservation for one hotel room. As a result, group bookings were recorded as separate reservations and further examination was required to identify which reservations were part of group bookings. Within each reservation several characteristics were recorded. First of all, the arrival date and the departure date of the booking was recorded, along with the check-in time and check-out time once the guest had stayed in the hotel. Also, the day and time of the booking were recorded. This characteristic proved to be essential for the data-analysis. If the reservation was cancelled, the cancellation date was recorded. The room type for which the reservation was made was present. The hotel regularly upgraded guests for free if a better room was available, but this was not recorded. The price that was paid for the reservation (room only) was recorded. The number of occupants of the room was also recorded, and it was even specified how many adults and children the room was booked for. Finally, the travel purpose (business or leisure), the name of the guest and if applicable the company name were present. A sample of the data set is presented in Table 2.2.

Booking		Arrival		Departure		Segment	Cancel.	Room		Occupancy	
Date	Time	Date	Time	Date	Time		Date	Type	Price	adult	child
2007-04-04	03:13	2008-02-14	14:00	2008-02-15	11:00	Business	2008-02-06	PG	140	1	1

Table 2.2: Overview of data set properties.

The following statistics were computed per room type (STD, GV, LGV, STO, PG, BRD) and for all data (TOTAL): (a) total number of reservations, (b) average occupancy, (c) average number of reservations, (d) average price that was paid for a room for one night, (e) percentage of nights that the hotel or room type was fully occupied, (f) percentage of rooms that was sold to groups, (g) average number of days between the reservation and the arrival day, (h) average length of stay, (i) percentage of guests that stayed more than one night, (j) total revenue, and (k) the percentage of bookings that were cancelled.

		Total March 2008
(a)	Total # reservations	443
(b)	Average occupation	1.38
(c)	Average # reservations	14.29
(d)	Average price	€121.74
(e)	% maximal occupancy	3.23%
(f)	% groups	65.99%
(g)	# days before arrival	46.03
(h)	LOS	1.48
(i)	% LOS > 1	28.91%
(j)	Total Revenue	€53,152
(k)	% cancellations	29.67%

Table 2.3: Example of key statistics.

Table 2.3 provides an example of these statistics (TOTAL, one month). These five-year statistics were determined for each month of the year, to capture average changes during the year, and per year, to capture changes from year to year. The hotel suffered from the recent economic crisis. Whereas the total number of reservations was 6,747 in 2008, in 2012 this was reduced by 26.6% to $4,952$. Total revenue reduced from €840,858.57 to €644,919.20.

2.3 Insight into Seasonality

An important aspect of demand is seasonality, i.e., a recurring pattern of demand across the year, week, or even during the day. In this section seasonality is analysed on these three different levels, with promising results.

2.3.1 Annual Seasonality

Changes in demand were first explored at the annual level. To compensate for seasonality within a week, the demand of different weekdays was aggregated in a week. Figures 2.1, 2.2, and 2.3 show the average annual demand pattern for all guests, for the leisure guests, and for business guests respectively.

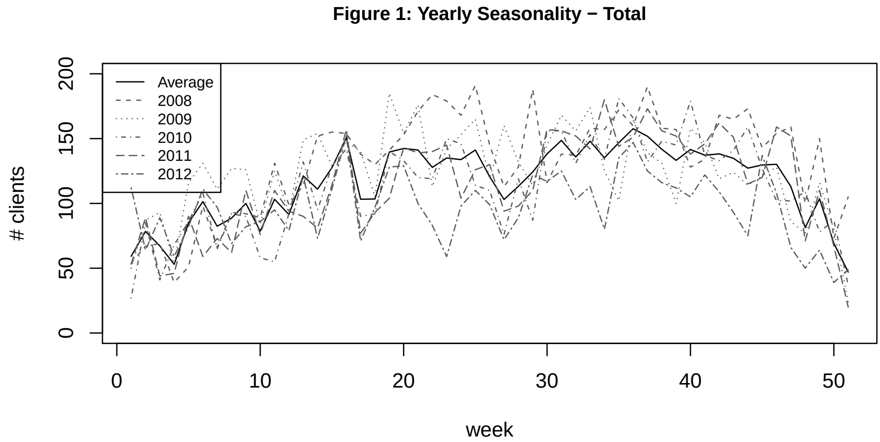

Figure 2.1: Yearly Seasonality – total.

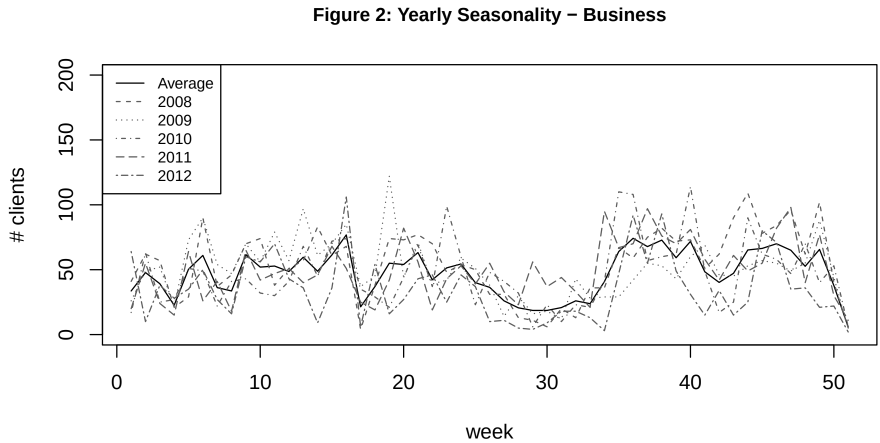

Figure 2.2: Yearly seasonality – business.

As Figure 2.1 illustrates, from January to September total demand increased and from September to January total demand decreased. Figure 2.2 shows that the demand from business guests was quite stable during the whole year, except for a gap in July and August. Figure 2.3 shows that the demand from leisure guests was low in winter, and steadily rose until a peak in July and August, in line with the Dutch summer holiday season. To examine whether leisure and business demand significantly differed a two-sided Kolmogorov-Smirnov test was applied to the time series of 2008-2012. The null hypothesis, stating that leisure and business demand were drawn from the same probability distribution, was rejected ($D = 0.2449$, $p < 0.001$). The annual seasonality of leisure and business thus differed significantly.

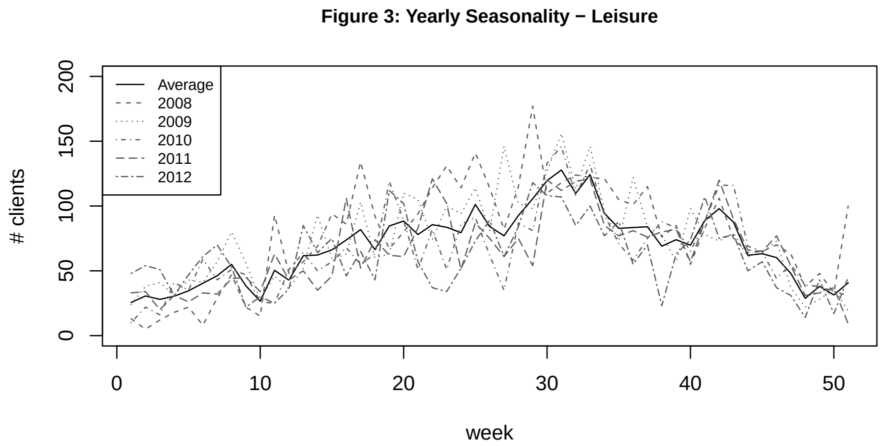

Figure 2.3: Yearly seasonality – leisure.

2.3.2 Weekly Seasonality

The hotel manager claimed to observe demand similarities at the week level. This phenomenon is not uncommon in hotel RM practice. Using seasonal-trend decomposition analysis the presence of weekly seasonality was verified. Decompositions were calculated with a frequency varying between 1 (no seasonality) and 366 (a whole year). Then the corresponding mean squared errors (MSE) were compared. The results are presented in Figure 2.4.

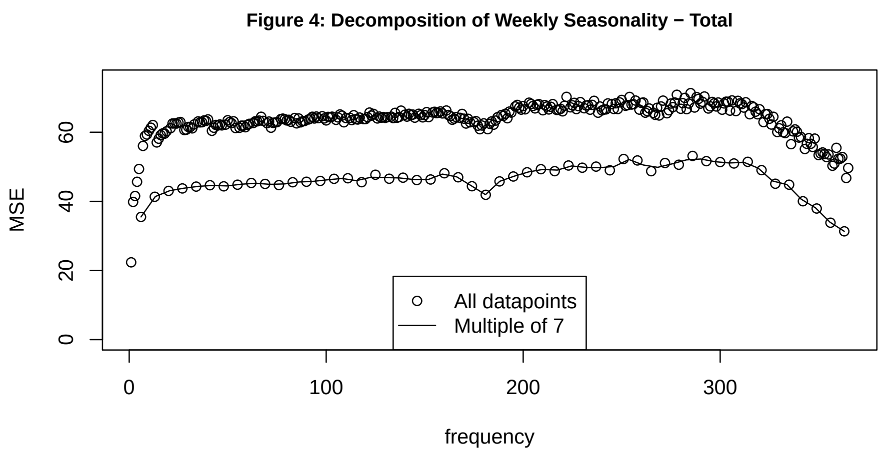

Figure 2.4: Decomposition of weekly seasonality – total.

It can be observed that the MSEs for decompositions with frequency equal to a multiple of seven are lower. This suggests that the observation of a weekly seasonality indeed is valid. Note that the MSE for values lower than seven are also low, but since multiples of these frequencies have high MSEs they are not true seasonality frequencies.

Figure 2.5 presents the average number of hotel guests per day of the week (DOW). To explore behavioural differences between business and leisure guests a distinction was made at the total, business and leisure level. A pattern was observed where leisure guests more

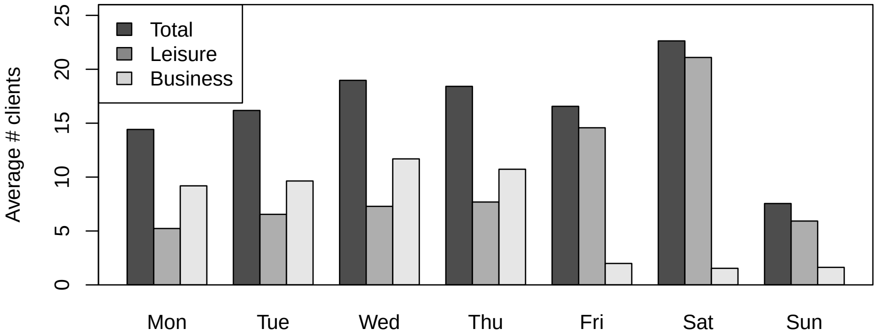

Figure 2.5: Average number of guests per week – total, leisure, business.

frequently booked for Friday and Saturday, and business guests for Monday, Tuesday, Wednesday and Thursday, with a low occupancy on Sunday. Another observation was that the occupancy in the weekend was higher than on weekdays. This did not imply, however, that the hotel served more leisure than business guests.

2.3.3 Daily Seasonality

A crucial observation was made about the booking behaviour at daily level. As Figures 2.6, 2.7, 2.8 and 2.9 illustrate, customer booking behaviour depended on both the weekday and the time of the day on which an advanced booking was made.

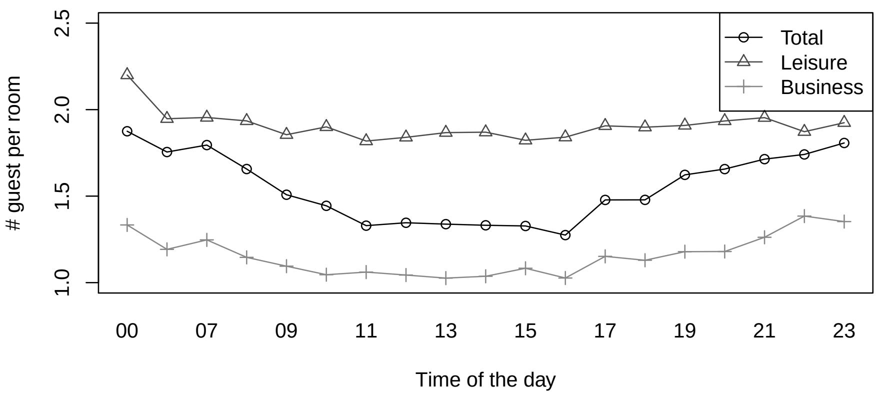

Figure 2.6: Hourly occupancy demand per weekday – total, leisure, business.

To examine whether the booking behaviour of business and leisure guests significantly differed a Mann-Whitney U-test was performed. Reservations that were made before 5pm on weekdays were more likely to have a higher price (Mdn = 120.50, M = 132.97) than reservations made in the weekend and on weekdays after 5pm (Mdn = 109.90, M = 120.01), U = 57031222, $z = -21.505$, $p < .000$, $r = -.13$. A chi-squared test confirmed a significant association between reservation moment and occupancy, $\chi^2(5, N = 25704) = 2497.756$,

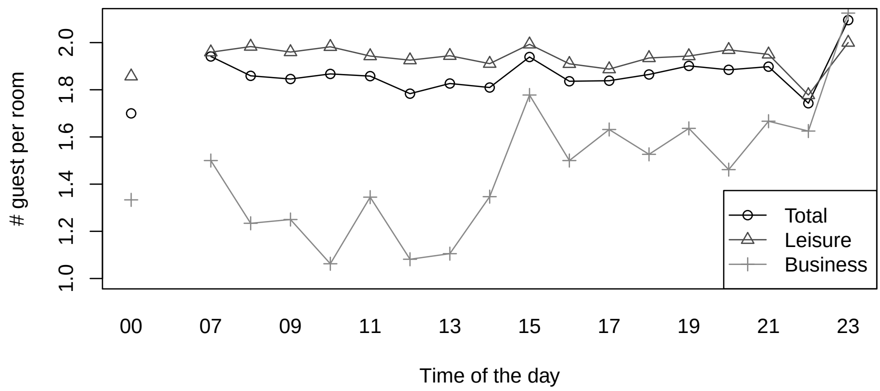

Figure 2.7: Hourly occupancy demand per weekend day – total, leisure, business.

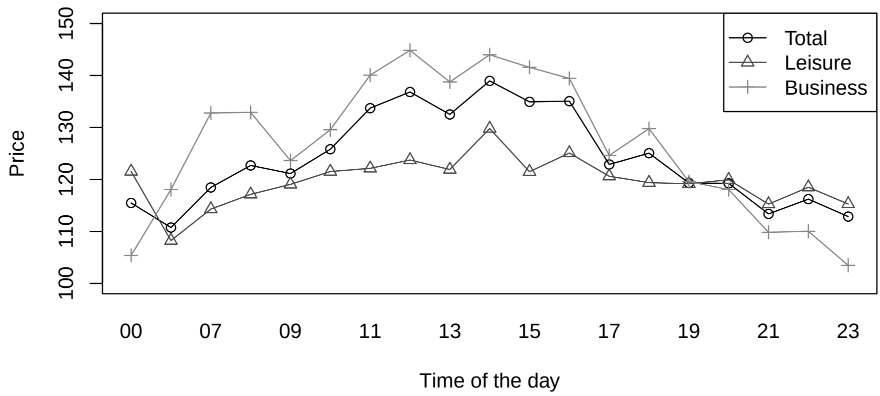

Figure 2.8: Hourly price per weekday – total, leisure, business.

$p < .000$. This seems to represent the fact that before 5pm on weekdays rooms were 4.12 times more likely (based on the odds ratio) to be occupied by one person; in the weekend and on weekdays after 5pm it was more likely to be two persons or more. A second chi-squared test confirmed a significant association between reservation moment and segment (business/leisure), $\chi^2(1, N = 25704) = 1948.420$, $p < .000$. Before 5pm on weekdays rooms were 3.49 times more likely (based on the odds ratio) to be occupied by a business guest; in the weekend and on weekdays after 5pm it was more likely to be a leisure guest. The findings thus indicated that business guests, who tended to make purchases during working hours, were willing to pay a higher price than leisure travellers, commonly with an occupancy of more than one person per room, who tended to make purchases outside working hours.

Figures 2.10 and 2.11 present business and leisure demand on an hourly basis. Purchases made in the weekend were dominated by leisure guests. Purchases made during the week consisted of a mix of business and leisure.

Table 2.4 shows that more than one third of all reservations were made in the weekend or

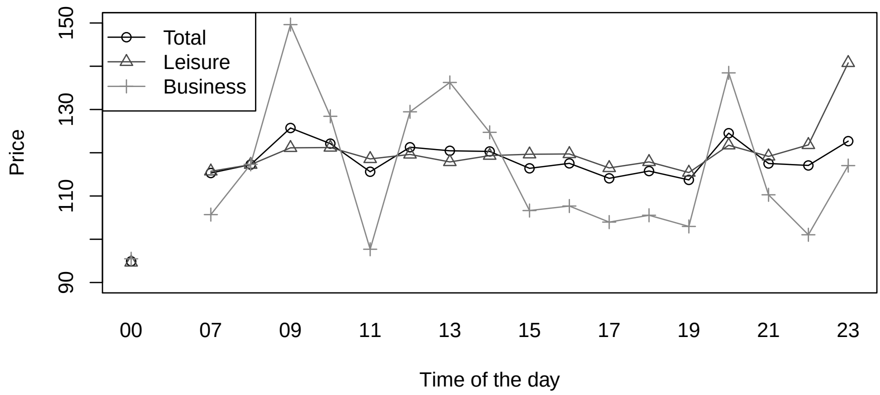

Figure 2.9: Hourly price per weekend day – total, leisure, business.

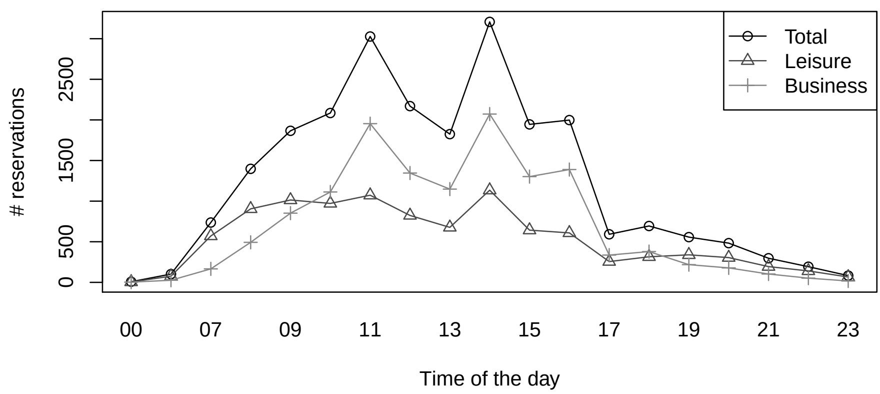

Figure 2.10: Hourly demand per weekend day – total, leisure, business.

after 5pm. on weekdays. The average price of those reservations was lower (M = 120.01) and the average number of guests per room was higher (M = 1.71) than the reservations that were made during weekdays before 5pm. (M = 132.97 and M = 1.36, respectively). This suggests that the hotel can take advantage of the two discrete segments by dynamically changing prices both in the weekends and during the day, instead of maintaining the one-price policy per room regardless of day and time of the day, a practice that is commonly observed in small and independent hotels.

The RM forecasting literature does not take into account that demand can vary at certain hours of the day. Whereas it is complicated to develop tractable solution methods and accurate parameter estimation methods that perform well on computation time, taking a relevant model extension (based on exploratory data analysis) into account can have substantial impact on revenue, as is reported in literature. For example, the seminal model by Talluri & van Ryzin (2004a) increased revenue up to 12% compared to 1-2% differences in other literature at the time, by incorporating customer choice-behaviour; and in Chapter 3 it is shown that by incorporating cancellations into the Talluri & van

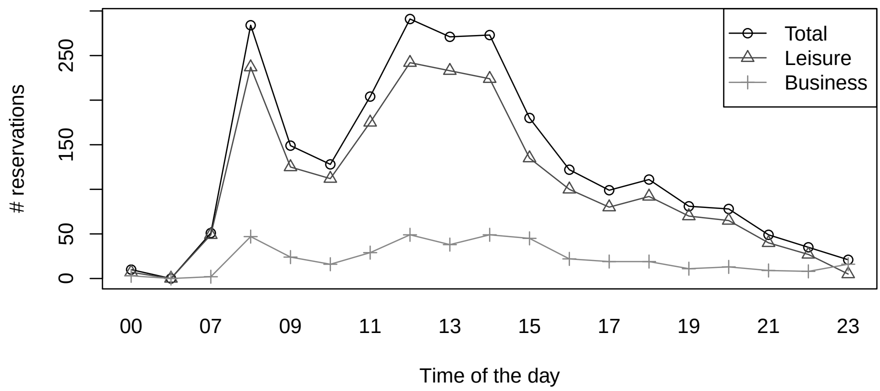

Figure 2.11: Hourly demand per weekend day – total, leisure, business.

	Volume	Leisure	Business	Guests	Price
Total	25704	47.22%	52.78%	1.46	€129.15
Weekend or from 5pm	7581	68.47%	31.53%	1.71	€120.01
Weekdays before 5pm	18123	38.33%	61.67%	1.36	€132.97

Table 2.4: Example of key statistics.

Ryzin (2004a) model, a substantial (additional) impact on revenue, up till 20%, could be achieved. The following claim is, therefore, formulated:

> **Claim.** There is room for optimization by bringing the RM strategy in line with demand for a specific month and day, as well as the observed booking behaviour at the point of time during the day.

2.4 Insight into Group Cancellation Behaviour

Hotels are vulnerable to demand and cancellation uncertainty (Chen et alii, 2011). This can lead to sudden increases and decreases in pickup which is why hotel revenue managers within the context of regular booking patterns tend to closely monitor the booking pace at both total and segmented level. This section analyses uncertainty in demand and cancellations.

2.4.1 Booking Pace

Figure 2.12 below shows the relationship between demand and time until arrival for a standard room at total and weekday level. It was found that demand increased as the time until arrival decreased. Moreover, the majority of guests tended to plan not too far ahead. To test if demand increased exponentially as the time until arrival decreased, as the visual inspection suggested, an ordinary least squares regression was performed on the log of the mean number of bookings against the weeks before arrival. The results of the regression showed that the number of weeks before arrival significantly predicted mean number of bookings, $\beta = -.908$, $t(38) = 13.32$, $p < .001$. The number of weeks before arrival also explained a significant proportion of variance in mean number of bookings, $R^2 = .824$, $F(1, 38) = 177.42$, $p < .001$. Moreover, as price behaviour was captured on the secondary axis, it was observed that during the last three months of the booking horizon the average price decreased as the day of arrival came closer. The hotel thus dropped prices as the booking window shortened. To identify whether the increase in the number of bookings was also affected by a drop in price, with respect to the three month booking window, a multiple regression (with time and price as the predictors) was performed. The results show that the number of days before arrival ($\beta = -.724$, $t(88) = -10.94$, $p < .000$) and price ($\beta = -.226$, $t(88) = -3.42$, $p < .001$) both significantly predicted the mean number of bookings. Days before arrival and price also explained a significant proportion of variance in the mean number of bookings, $R^2 = .804$, $F(2, 88) = 180.1$, $p < .000$. Therefore, in addition to the shortening booking window, the drop in price affected demand.

The booking pace was different per weekday, suggesting that forecasting and pricing models should take this behaviour into account. Arrivals on Monday through Thursday showed similar behaviour patterns as well as the arrivals on Friday and Saturday. The arrivals on Sunday behaved differently. The curve for Monday through Thursday was more flat compared to Friday and Saturday. This implied that these reservations were made earlier in the booking horizon (M = 37.06). On the other hand, reservations with arrival on Friday and Saturday tended to book closer to the day of arrival (M = 29.10). For Sunday this was even closer (M = 27.90).

2.4.2 Cancellations

Over five years of data on average about 21.71% of all reservations were cancelled. The number of cancellations varied per year and per month.

Year	Month												Total
	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	
2008	42.4%	26.7%	29.7%	34.0%	26.0%	23.1%	8.6%	13.6%	19.8%	32.9%	32.5%	28.5%	26.9%
2009	14.6%	18.2%	30.8%	32.5%	16.1%	38.4%	17.5%	23.0%	34.6%	18.9%	27.0%	13.9%	25.3%
2010	32.6%	16.8%	17.3%	15.3%	16.4%	22.3%	15.2%	18.9%	30.8%	16.4%	19.1%	15.3%	19.9%
2011	24.6%	14.2%	18.7%	6.9%	18.9%	12.6%	5.3%	9.5%	20.7%	10.4%	11.0%	16.6%	14.0%
2012	20.9%	22.9%	9.6%	29.9%	19.2%	31.8%	7.8%	13.7%	19.9%	16.4%	18.5%	20.5%	20.0%
Total	27.3%	20.0%	22.4%	25.9%	19.7%	26.5%	11.3%	16.2%	25.3%	19.9%	22.4%	19.6%	21.7%

Table 2.5: Cancellations per year per month.

As Table 2.5 illustrates, a higher cancellation rate was observed in 2008 (26.94%) than in 2011 (13.98%). Cancellation rates also varied per month. For example, in January the cancellation rate varied from 14.61% in 2009 to 42.38% in 2008.

Cancellations	Days before arrival				
	0-2	3-9	20-29	30-84	85+
2008	9.52%	11.85%	15.86%	30.02%	50.61%
2009	11.48%	8.57%	16.31%	27.70%	55.92%
2010	4.73%	10.13%	16.42%	31.29%	40.92%
2011	4.26%	8.51%	18.40%	24.35%	15.96%
2012	5.15%	6.36%	19.51%	34.66%	37.94%
Total	7.01%	9.13%	17.21%	29.43%	44.51%

Table 2.6: Cancellations and the booking window.

Table 2.6 shows that overall, cancellation rates lowered when the booking window shrank. For example, 44.51% of all reservations that were made at least 85 days before arrival were cancelled eventually. Up to two days this was 7.01%. These rates varied per year. For example, 15.96% of the reservations were cancelled in 2011 whereas this

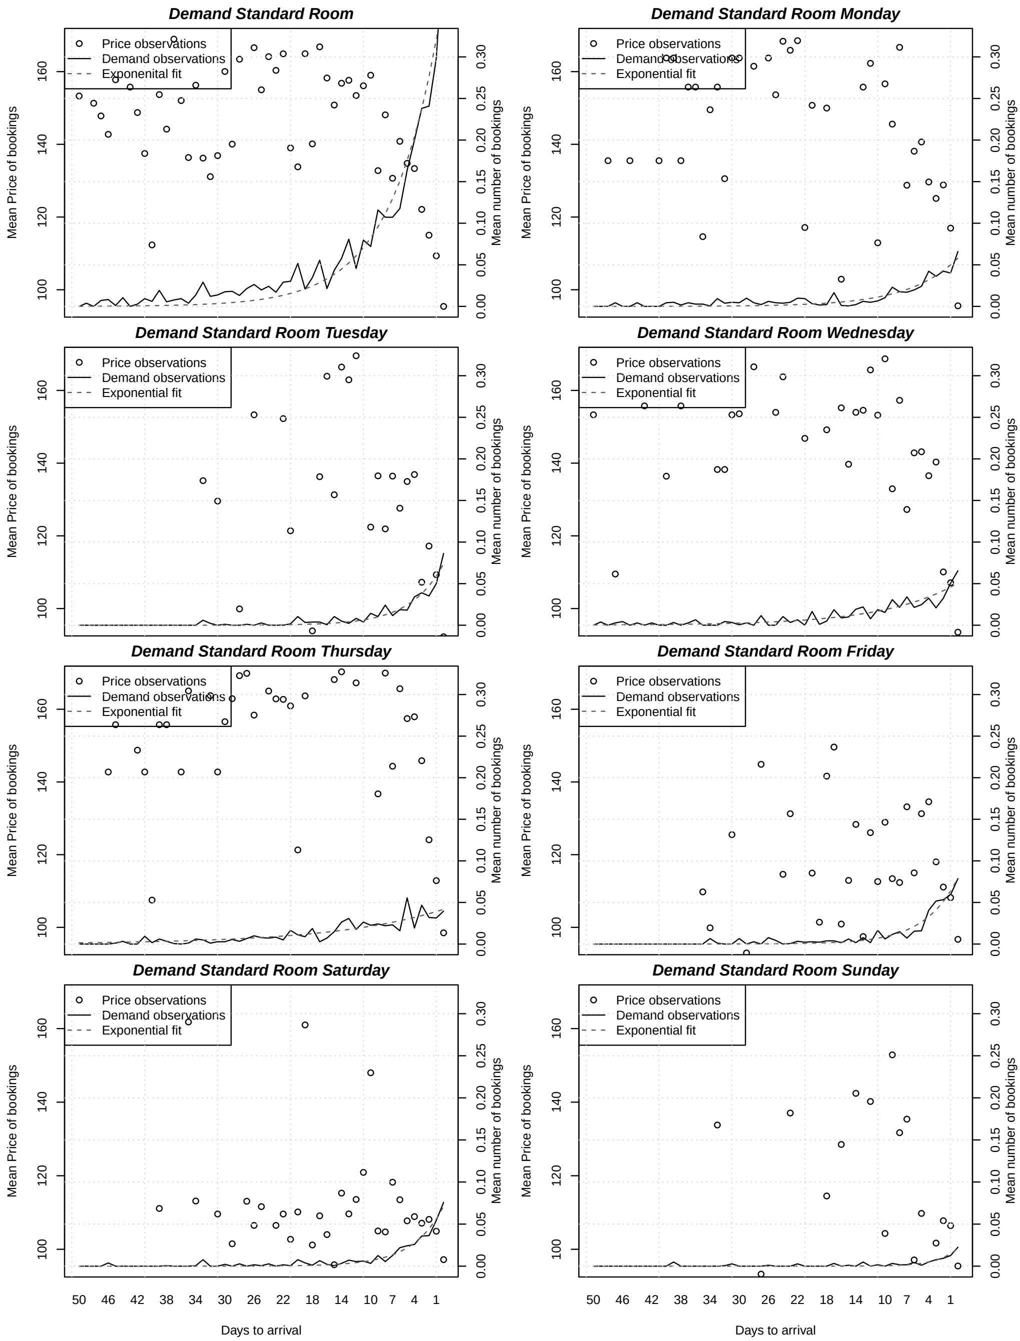

Figure 2.12: Advanced bookings and price of standard room – total.

was 55.92% in 2009. For individual days this variation was even higher as only monthly averages were considered. Also, it was observed that 55% of the hotel's occupancy came from group bookings. These bookings showed higher cancellation rates.

As Table 2.7 presents, on average 34.30% of the group bookings were cancelled, as compared to 21.71% of all bookings and 6.12% for transient bookings. A chi-squared

	Cancellations			Number of bookings		
	Transient	Group	Total	Transient	Group	Total
Business	8.26%	36.79%	30.78%	2857	10709	13566
Leisure	5.41%	26.72%	11.58%	8627	3511	12138
Total	6.12%	34.30%	21.71%	11484	14220	25704

Table 2.7: Cancellation behaviour per segment.

test confirmed a significant difference between group and transient, at the total, $\chi^2(1, N = 25704) = 2968.518$, $p < .000$, leisure, $\chi^2(1, N = 12138) = 1106.404$, $p < .000$, and business level, $\chi^2(1, N = 13566) = 861.630$, $p < .000$. Based on the odds ratio, it was found that groups were 8.00 times more likely to be cancelled than transient. For leisure and business groups this was respectively 6.37 times and 6.46 times more likely than transient. A similar pattern was observed for business versus leisure, with significant differences at total, $\chi^2(1, N = 25704) = 1390.407$, $p < .000$, group, $\chi^2(1, N = 14220) = 119.104$, $p < .000$, and transient level, $\chi^2(1, N = 11484) = 30.275$, $p < .000$. Using the odds ratio, it was found that business guests were 3.40 times more likely to cancel than leisure guests. Business groups were 1.60 times more likely to cancel than leisure groups. Business transient were 1.57 times more likely to cancel than leisure transient. Using Levene's test to identify differences in normalized variation between the segments ($p > .05$) it was found, with regard to group and transient cancellation behaviour, that group business represents a relatively large proportion of the uncertainty in demand and cancellations.

Cancellations have received wide research attention in the hotel RM (e.g., Chen & Xie, 2013). The recent customer choice models in hotel RM forecasting literature, however, do not take cancellation into account, with the exception of the work in Chapter 3. Also, group cancellations (and lost/ turn-down information) are not included in customer choice modelling. Group business, which – as one of four major areas – was identified 'as having the greatest growth potential in hotel RM' (Milla & Shoemaker, 2008, p. 110), has properties that make modelling very complex. In addition, 'transaction data, especially for the largest groups and smallest hotels, generally are sparse' (Holverson & Revaz, 2006, p. 49). Based on the exploratory data analysis the following claim is, therefore, formulated:

Claim. There is room for a model extension in the (e.g., customer-choice based) forecasting literature by bringing the RM strategy in line with the more variable and statistically uncertain nature of group cancellations.

2.5 Insight into Demand Uncertainty

The demand for hotel rooms varied a lot from day to day. To this extend the nature of the uncertainty of demand is studied.

2.5.1 Probability Distribution Function

One of the most crucial assumptions in any RM model is the probability distribution function that demand follows. As was found in the analysis, on average the closer to the day of arrival, the more clients booked, but this finding did not reveal the nature of the demand distribution.

Revenue management literature generally assumes a (nonhomogeneous) Poisson distribution (e.g., McGill & van Ryzin, 1999; Bitran & Caldentey, 2003; Talluri & van Ryzin, 2004a). That is, demand per time period is modelled as a homogeneous Poisson process. With the use of a likelihood ratio test as well as a chi-squared test, it was tested whether the data was Poisson distributed. All time periods (for the booking of a standard room) had p-values smaller than 0.001 so that the null hypothesis that the data was not Poisson distributed was rejected. Tests on the other six room types confirmed this finding. The finding that demand followed a nonhomogeneous Poisson process was in line with earlier work by Haensel & Koole (2011b) who found that airline data was Poisson distributed. Assuming a Poisson distribution in forecasting modelling has the advantage of containing the Markov (memory-less) property (i.e., future demand does not depend on the guests who booked a room for the same arrival day in the past). This is in accordance with reality, since it is reasonable to assume that hotel guests arrive independent from each other.

2.5.2 Implication: Logical Inferences about Hotel Size

When demand follows a Poisson process different consequences can be inferred for smaller and larger hotels. Suppose demand is Poisson distributed with parameter λ, i.e., the expected number of guests who book a room. The standard deviation is then equal to $\sqrt{\lambda}$ such that the 95% confidence interval of the actual demand D is given by:

$$D \in [\lambda - 2\sqrt{\lambda}, \lambda + 2\sqrt{\lambda}]. \tag{2.1}$$

The square root in formula 2.1 implies that the coefficient of variation decreases as increases, such that for smaller hotels the coefficient of variation in demand is higher than for large hotels.

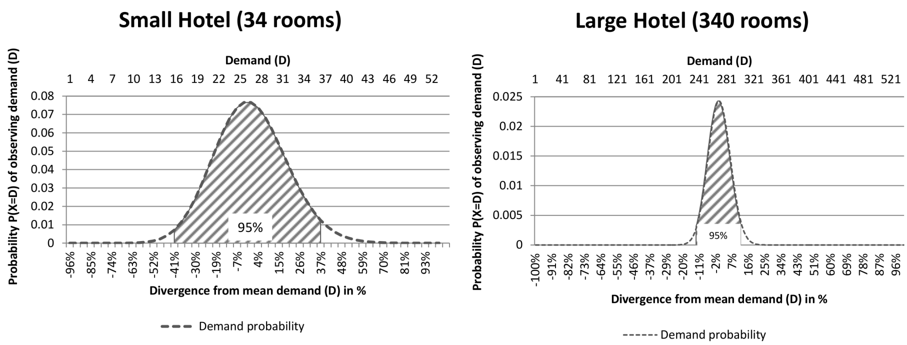

Figure 2.13: Different consequences for smaller and larger hotels when demand follows a Poisson distribution.

Figure 2.13 illustrates this size-implication inference for the case hotel (with 34 rooms) compared to a ten times larger hotel (with 340 rooms). For illustrative purposes, market demand is assumed at 297 rooms from with each hotel gets its fair market share (27 respectively 270 rooms). Then, in 95% of the cases the larger hotel would have a demand between 237 and 302 rooms while the small case hotel would have a demand between 16 and 37 in 95% of the cases. In the worst case, for the small hotel this leads to 38% less demand than the average case, while in the worst case for the large hotel this leads to 12% less demand than the average case. The finding that demand is Poisson distributed thus implies that a small hotel is more vulnerable to demand uncertainty than a large hotel.

As (simple) forecasting and pricing models only consider average behaviour, they provide an optimal strategy on average. However, when demand is volatile, as was found in this study, forecasts will be inaccurate, the errors being worse for smaller than larger hotels. To reduce forecasting error, forecasting and pricing models can take uncertainty into account, for example by considering the distribution of demand or by applying robust optimization techniques which take into account worst-case scenarios. This case study indicates that these methods are especially preferable for small hotels because demand variation is higher. Therefore, the following claim is formulated:

> **Claim.** A nonhomogeneous Poisson process is a good approximate model for hotel demand. As a consequence, demand is more uncertain for smaller than for larger hotels.

2.6 Discussion

This empirical study emphasizes the importance of preliminary and exploratory data analysis in hotel RM forecasting. Preliminary data analysis is key to the selection of the class of forecasting models, whereas exploratory data analysis is essential to evaluate whether a chosen model still is appropriate to capture changes that occur in the environment. As a whole, data analysis allows to determine whether a RM strategy is still optimal and to explore new opportunities for revenue optimization. In this context, this study found three overlooked or ill-researched aspects of data analysis in hotel RM forecasting, each with different theoretical implications for demand modelling, forecasting and revenue optimization.

First of all, it provides empirical evidence on a nonhomogeneous Poisson nature of the probability distribution function that demand follows. There is little evidence of this crucial and commonly assumed demand characteristic in the hotel forecasting literature, especially for small and independent hotels. This implies that especially for small hotels forecasting methods should be developed that take into account the uncertainty that comes with the Poisson distribution, for example by using robust optimization methods. Secondly, the study presents evidence on the random nature of group cancellations, an important but ill-researched segment in hotel RM. Optimization methods should take these cancellations into account. It is, however, unclear how such model would look like and therefore more research is needed. Thirdly, our results show that in a local market context business and leisure booking behaviour significantly differ per point of time during the day. As the study shows, forecasting models that take this behaviour into account can create a revenue increase. A further study that models this behaviour could reveal the extent of this potential.

As a whole the study finds support for the work of Koupriouchina et alii (2014) who argue that research in forecasting should take place at a more granular level. It supports three claims that answer to Bodea et alii (2009) who call for more work in forecasting based on real-world data. In this context, as hotel RM forecasting can be perceived as ‘a big data problem’, the study also supports Xiang et alii (2015, p. 120) who observe that ‘big data analytics approach in hospitality is yet to be well developed and established’, and who reveal the potential of big data analytics to generate new insights.

The study is, however, not without limitations. One concern relates to the data which was collected from a single hotel. While forecasting research tends to rely on simulated data, and in this respect this exploratory study is a positive exception, its contribution is case-based. In order to generalize the results and findings of this paper, similar studies with more hotels have to be performed. However, collecting data from a similar hotel is time consuming, since small and independent hotels generally collect and save data poorly and incompletely. Another limitation refers to the small and independent nature of the hotel and the specific local environment it operates in. The size and business mix in this study is specific to the context of the hotel. A comparative study, where hotels are grouped according to their location and client mix, could identify characteristics that can be generalized or are specific to certain hotels. Such study is quite involved, since it requires the cooperation of a lot of hotels and the collection and cleansing of their data. A final limitation is that the study did not include competitive data. The hotel industry is a highly competitive market, where hotels try to attract customer segments by means of positioning and offering a certain quality at a fair price. An important result of such study would be the effect of competitive prices on demand. However, again, the data collection process is quite involved. One would not only would need sales data of all involved hotels, but the whole pricing history, to identify price sensitivity when clients consider multiple hotels before making a purchase.

There are various implications for practice. Data analysis provides important insights in the booking and cancellation behaviour of hotel guests. When analysing at the segment level, data analysis can provide insights that are essential to maintain an optimal RM strategy and to explore new revenue opportunities. Data analysis also aids the process of evaluating the RM model as it tells how forecasting performs with respect to changes in demand. In this way, data analysis is vital for any hotel that seeks to stay competitive in a changing environment. In the case of the small and independent hotel cancellation was found to be much more severe than the hotel anticipated. Moreover, a daily booking pattern was identified. A rationale was thus provided for adjusting the RM strategy. Data analysis is however a laborious process (see also Bodea et alii, 2009). Especially for small and independent hotels, such investment in time and analytical skills is often perceived as not worthwhile.

The study suggests three directions for future research. First of all, the findings indicate that there is room for an extension to the customer choice modelling literature in forecasting. An existing attempt of such extension is given in Chapter 3, where it is shown that taking into account cancellation can impact revenue up till 20%. It would be interesting to examine whether their analysis holds for group bookings as well, and also whether their model can be extended to include differences in demand per point of time during the day. Secondly, the issue of variation in demand uncertainty as a result of differences in hotel size, and its implications for forecasting, can be further investigated. If demand and cancellations have a high variance, then conventional RM models are not appropriate. Empirical work could establish whether this variance indeed is higher for small hotels than for large hotels, as this exploratory study suggests. Finally, through systematic application of big data analytics techniques new sources of data could be analysed to learn which customer behaviour (such as day patterns) could be incorporated in forecasting modelling to further improve hotel RM performance.

CHAPTER 3

Revenue Management under Customer Choice Behaviour with Cancellations and Overbooking

Motivated by the data analysis in Chapter 2, this chapter proposes a revenue management (RM) model that takes into account cancellations and overbooking, in addition to customer choice behaviour. The study shows that the model can be solved exactly when cancellations rates are assumed to be linear and equal. For general cancellation rates the problem is intractable, and therefore effective heuristics are developed. Moreover, an efficient parameter estimation method is developed that provides good estimates. The numerics show that taking cancellations into account can lead to an increase in revenue of up to 20%.

3.1 Introduction

RM applications have in common that customers choose a product among different products. For example, a customer who is looking for a hotel room may compare rooms of different hotels before making a decision. The seminal paper by Talluri & van Ryzin (2004a) incorporates customer choice behaviour in RM models. Their idea laid the foundation for a new theory in RM: *customer choice models*. Since then the theory of customer choice models has been enriched with among others better solution methods (Strauss & Talluri, 2012; Meissner & Strauss, 2012), network models to take into account multiple night stays or multiple flight legs (Liu & van Ryzin, 2008), and outstanding methods to estimate the parameters of the model (Newman et alii, 2014). To the best of the authors' knowledge, however, an extension to a model including cancellations has not been made. Modelling cancellations by taking into account customer choice behaviour has been studied by Iliescu et alii (2008), but the cancellation process is not integrated with the decision process. Examples of studies on models that include cancellations but do not take into account customer choice behaviour are Lautenbacher et alii (1999), Aydin et alii (2012), and Bertsimas & Popescu (2003). Lautenbacher et alii (1999) provide an analysis of an RM model with cancellations and overbooking, but customer choice behaviour is not considered. Moreover, a satisfactory solution method for their intractable problem is not provided. Recently, Aydin et alii (2012) proposed an RM model with cancellations,

This chapter is based on Sierag, Koole, van der Mei, van der Rest, & Zwart (2015).

but they also do not take customer choice behaviour into account. Erdelyi & Topaloglu (2010) and, more recently, Kunnumkal et alii (2012) provide customer choice models that take into account overbooking but no cancellations.

The main contribution of this chapter is the inclusion of cancellations and overbooking in an RM model that also takes customer choice behaviour into account. Two other key contributions of this study are well-performing tractable solution methods and an accurate parameter estimation method with low computation time. The combination of these three contributions makes the customer choice cancellation model very suitable for practitioners.

Taking cancellations into account in the decision making process has a big impact on revenue. Numerical results (see Section 3.5) show that policies that do not take cancellations into account can lead to a substantial revenue loss of 20%. So far the customer choice models in existing literature do not take into account cancellations. However, in practice it is common that for example seats on an aircraft are cancelled, hotel rooms are cancelled, theatre tickets are cancelled, reservations for rental cars are cancelled, and reservations for spots on a golf course are cancelled. Therefore, we stress that effective RM systems should take cancellation into account in the decision making process. Moreover, overbooking makes more sense when taking cancellations into account.

The problem is modelled as a continuous-time Markov decision process and solved using dynamic programming. Some instances of the problem are too large to solve exactly because of the curse of dimensionality. To overcome this problem, three tractable methods are proposed, each appropriate under a different assumption. First it is shown that if the cancellation rates are equal and linear with respect to the number of reservations the problem can be reduced to a one-dimensional problem and thus it can be solved exactly. Second, an efficient heuristic is proposed which is appropriate in the case that the cancellation rates are linear with respect to the number of reservations, but not equal. Third, a heuristic is proposed that can be applied to general cancellation rates, under the only assumption that cancellations occur independent from each other. Also, a heuristic is provided which can be applied if it is not assumed that all the resources are identical. This occurs for example if the hotel has multiple room types or a theatre has multiple seat classes. Numerical results show that the heuristics perform well.

To apply this model in practice it is essential to estimate the parameters of the model accurately. The estimation of parameters of customer choice models can be challenging. A common attempt is to use variations of the expectation maximisation algorithm. The basic version of Talluri & van Ryzin (2004a), but also more sophisticated methods (e.g., van Ryzin & Vulcano, 2013) have the drawback of long computation time and bad parameter estimates. Recently Newman et alii (2014) proposed a different parameter estimation method that shows great potential. The parameter estimation method that is proposed for the customer choice cancellation model is based on this model. The parameters of the continuous-time Markov chain are estimated, in contrast to most literature (e.g., van Ryzin & Vulcano, 2013), which estimate the parameters of the discretised Markov chain. Our estimation method has the advantage that it has low computation time and it gives good estimates. The combination of the three effective solution methods and the efficient estimation method ensures that the customer choice cancellation model is well applicable in practice.

The remainder of this chapter is organised as follows. First the customer choice cancellation model is described in Section 3.2. In Section 3.3 solution methods are described: a dynamic programming formulation, which solves the problem exactly but is intractable; two heuristics, which are appropriate in different settings; and a heuristic to solve the problem for a hotel with different room types. Then Section 3.4 describes methods to estimate the parameters of the model. Numerical results of the model are presented in Section 3.5. Finally, in Section 3.6 and 3.7 we make some concluding remarks and propose topics for further research.

3.2 Model Description

In this section the *customer choice cancellation model* is introduced. It is an extension of the *customer choice base model* by Talluri & van Ryzin (2004a), which does not consider cancellations. First the model with no assumptions on the cancellation behaviour is provided. Second a first reformulation is provided based on the assumption that cancellations only depend on the current number of reservations. Third, a reformulation based on the additional assumption that cancellation rates are linear is provided. Finally, a reformulation based on the additional assumption that all cancellation rates are equal is provided. Section 3.3 provides effective solution methods for the models.

Consider a hotel with $C \in \mathbb{N}$ identical rooms. The rooms are sold for customer check-in at a fixed future time unit such that revenue is maximised. The customer check-in is in T time units. A common time unit is days, but some hotels sell rooms for a few hours to accommodate for short stays (e.g., at airport hotels, see Hopman, 2013). Overbooking is allowed up to $C_{\max}$ rooms. A *rate* or *fare product* j is a room in combination with a certain price r_j, also called the reward for product j, and certain conditions (e.g., cancellation conditions). Typical conditions in a hotel environment are refundable/non-refundable rooms, where a refundable room can be cancelled free of charge and a non-refundable not; or breakfast could be included or excluded from the fare product. Let N be the set of all fare products. Assume that there is a finite amount n of fare products, such that w.l.o.g. $N = \{1, \ldots, n\}$. Let $c_j(t)$ be the costs if product j gets cancelled in time period t. Note that $c_j(t)$ depends on the cancellation conditions.

Assume that cancellations only depend on the current number of reservations. This model can also be used with time of booking dependent cancellations by including more fare products. This problem is modelled as a finite-horizon continuous time Markov decision process over T time units. Define the state space by

$$X := \left\{ x \in \mathbb{N}^n \,\middle|\, x \geq 0, \sum_{j \in N} x_j \leq C_{\max} \right\}.$$

Each entry x_j of $x \in X$ corresponds to the pickup (number of reservations) for product j.

There is a penalty $q(x)$ involved if at the arrival day the state is x and there are more reservations than capacity, i.e.,

$$\sum_{j \in N} x_j > C.$$

Customers arrive according to a Poisson process with parameter λ. Continuously in time we decide which subset $S \subset N$ of products to offer. Set S is called an *offer set.* If set S is offered, the probability that an arriving customer buys product j equals $P_j(S)$. Furthermore, the probability that an arriving customer buys nothing under offer set S equals $P_0(S)$. Finally, for all products j the cancellations of reservations for product j are exponentially distributed with parameter $\gamma_j(x)$. In accordance with Markov decision process literature a feasible solution π to our problem is called a *policy.* A policy π that optimises total revenue is called an *optimal policy.* See Figure 3.1 for a visualisation of this model.

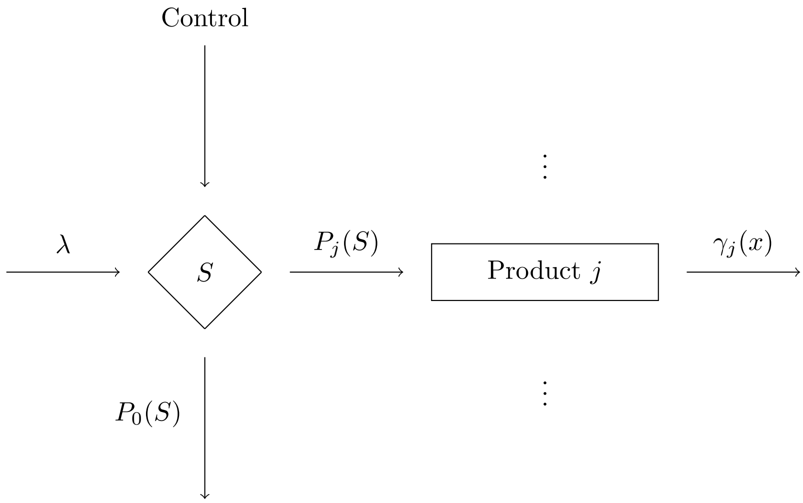

Figure 3.1: Visualisation of the customer choice cancellation model. The arrival process is Poisson distributed with parameter λ. There is control over the offer set S. Under this offer set an arriving customer buys product $j \in S$ with probability $P_j(S)$. With probability $P_0(S)$ the customer buys nothing. Finally, cancellations of product j follow an exponential distribution with parameter $\gamma_j(x)$.

Discretisation

To solve the customer choice cancellation model, consider the discretised Markov decision process. Divide time into T time periods, where the length of the intervals is such that the probability that more than one event occurs is very small. Therefore, assume that only one event occurs per time period, where an event is either an arrival, a cancellation, or neither arrival nor cancellation. Denote λ as the probability that a customer arrives in a time period; and $\gamma_j(x)$ as the probability that product j is cancelled in state x. The probability that no purchase occurs in a time period equals the sum of the probability that neither an arrival and nor a cancellation occurs, and the probability that an arrival occurs but the arriving customer makes no purchase. This is equal to

$$\left(1 - \lambda - \sum_{j \in N} \gamma_j(x)\right) + \lambda P_0(S) = 1 - \lambda \sum_{j \in S} P_j(S) - \sum_{j \in N} \gamma_j(x).$$

In each time period we decide which set S to offer. Note that time has to be scaled such that

$$\lambda + \max_{x \in X} \left\{ \sum_{j \in N} \gamma_j(x) \right\} \leq 1,$$

otherwise the probabilities are not well defined. This is possible as X is finite.

> **Remark.** It is possible to incorporate no-shows and walk-ins into this model. Lautenbacher et alii (1999) do this by incorporating it in the penalty $q(x)$. However, it is also possible to incorporate the walk-in probability in the arrival probability and the no-show probability in the cancellation probability in the last time step(s). We choose the latter method.

Reformulation under Assumptions on Cancellation Rates

It is reasonable to assume that cancellations occur independent of each other. Assumption 3.1 below states that the probability that product j is cancelled only depends on the current number of reservations x_j for product j. The assumption is equivalent to Assumption 1 in Lautenbacher et alii (1999).

Assumption 3.1. *Independence of Cancellations*

$\gamma_j(x) = \gamma_j(x_j)$ for all $x \in X$, for some function $\gamma_j : \mathbb{N} \to R^+$.

To emphasize the fact that we work under the independence of cancellations assumption, write $\gamma_j(x_j)$ instead of $\gamma_j(x)$. Note that x is a vector and x_j is a scalar.

Assumption 3.2 below is even stronger than Assumption 3.1.

Assumption 3.2. *Linear Cancellation Rates*

1. $\gamma_j(x_j) = \gamma_j x_j$, for $\gamma_j \in \mathbb{R}^+$, for all $j \in N$,
2. $q(x^1) = q(x^2)$ for all $x^1, x^2 \in X$ with $\sum_{j \in N} x_j^1 = \sum_{j \in N} x_j^2$.

The second part of Assumption 3.2 states that the penalty for overbooking only depends on the total number of products that are sold and is independent from the types of products that are sold. This assumption is realistic because we are dealing with identical rooms, so relocation to a similar hotel induces the same costs. This is also a very convenient step if we only want to keep track of the total number of reservations instead of the number of reservations per fare product. For $x \in X$, let y denote the total number of reservations, i.e.,

$$y = \sum_{j \in N} x_j.$$

Then under the linear cancellation rates assumption, write $q(y) = q(x)$ for all x with $\sum_{j \in N} x_j = y$.

Assumption 3.3 below states that the cancellation probability is independent from the product j, but it still depends on the number of reservations of product j.

Assumption 3.3. *Linear and Equal Cancellation Rates*

1. $\gamma_j(x) = \gamma x_j$, for $\gamma \in \mathbb{R}$ and for all $j \in N$,
2. $q(x) = q(y)$ with $y = \sum_{j \in N} x_j$.

Under the linear and equal cancellation rates assumption the state space of the problem is significantly reduced, as is shown in Theorem 3.1 in Section 3.3 below. However, this assumption is not realistic. Different fare products are likely to have different cancellation probabilities. For example, a refundable room is more likely to get cancelled than a non-refundable room.

Illustrative Example

An illustration of the model can be found in Example 3.1 below. This example is used throughout this chapter.

Example 3.1. [1]

Consider a hotel with $C \in \mathbb{N}$ rooms which offers three fare products. The reward for these products is given by

$$r = (r_1, r_2, r_3) = (160, 100, 90).$$

Both fare products 2 and 3 need to be purchased at least 21 days ahead. Furthermore, the cancellation conditions are as follows: if a product of type j is cancelled, the customer receives a refund of

$$c_j = \begin{cases} r_1 & \text{if } j = 1, \\ \frac{1}{2} r_2 & \text{if } j = 2, \\ 0 & \text{if } j = 3. \end{cases}$$

In other words, product 1 is refundable, product 2 only refunds half the price, and product 3 is non-refundable. Assume that overbooking is allowed but is bounded by 20% of the capacity, such that $C_{\max} = \lfloor C \frac{6}{5} \rfloor$. If we have $\sum_{j \in N} x_j - C > 0$ overbookings, then $\sum_{j \in N} x_j - C$ reservations are chosen at random and relocated to another hotel. The costs of relocating a customer is 170. Hence, the penalty $q(x)$ for overbooking at the arrival day in state x equals

$$q(x) = \begin{cases} 170 \left(\sum\limits_{j \in N} x_j - C \right) & \text{if } \sum_{j \in N} x_j - C > 0, \\ 0 & \text{otherwise.} \end{cases}$$

S	$P_1(S)$	$P_2(S)$	$P_3(S)$	$P_0(S)$
$\emptyset$	0	0	0	1
$\{1\}$	0.3	0	0	0.7
$\{2\}$	0	0.4	0	0.6
$\{3\}$	0	0	0.5	0.5
$\{1,2\}$	0.1	0.6	0	0.3
$\{1,3\}$	0.3	0	0.5	0.2
$\{2,3\}$	0	0.4	0.5	0.1
$\{1,2,3\}$	0.1	0.4	0.5	0

Table 3.1: Purchase probabilities under different offer sets.

The hotel wants to sell the rooms in a period of T days. The purchase probabilities $P_j(S)$ are given in Table 3.1. Assume that the purchase probabilities are the same for all time periods. The probability that a reservation for product j is cancelled is

$$\gamma_j(x) = \gamma_j x_j,$$

with $\gamma_j \in [0, 1]$ for all j.

[1] This example is based on Example 1 from Talluri & van Ryzin (2004a). The rewards are scaled to be representative values for hotels. Moreover, cancellation conditions were added, cancellation probabilities, and an overbooking policy.

3.3 Solution Methods

This section provides solution methods for the models derived in Section 3.2. First, an exact solution method is provided based on dynamic programming for the model with no restrictions on cancellation behaviour. Unfortunately, this problem faces the curse of dimensionality. To overcome this, heuristics for each model are provided. In particular a tractable solution method is provided that solves the model under the linear and equal cancellations assumption 3.3 exactly. Finally, a heuristic is provided that is applicable to a hotel with multiple room types.

3.3.1 General Solution Method: No Assumptions on Cancellation Behaviour

Let $V_t(x)$ be the maximal expected revenue from time t to the arrival day. To solve this problem, consider the following Bellman equation:

$$V_t(x) = \begin{cases} \max\limits_{S \subset N} \Bigg\{ \lambda \sum\limits_{j \in S} P_j(S)\big(r_j + V_{t-1}(x+e_j)\big) & \text{if } t > 0, \\ + \sum\limits_{j \in N} \gamma_j(x)\big(-c_j(t) + V_{t-1}(x-e_j)\big) & \\ + \left(1 - \lambda \sum\limits_{j \in S} P_j(S) - \sum\limits_{j \in N} \gamma_j(x)\right) V_{t-1}(x) \Bigg\} & \\ q(x) & \text{if } t = 0, \end{cases} \tag{3.1}$$

for all $x \in X$ with where $e_j \in \mathbb{R}^n$ is the j-th unit vector. With dynamic programming find an optimal strategy can be found. Unfortunately, this solution methods bears the burden of the curse of dimensionality. The size of the state space X is $O(C^n)$, so it increases exponentially if the number of fare products increases. For a small number of fare products the exact solution can be computed, like in Example 3.1. Obviously, in practice the number of fare products is larger than in our small example. Hence, this problem can not be solved for all practical purposes. To overcome this problem several options are proposed. First it is shown in Theorem 3.1 that under certain conditions the problem can be solved exactly. Second, several heuristics are proposed in case the conditions of Theorem 3.1 are not satisfied. It will turn out that each heuristic has its own advantages and disadvantages and is appropriate in a different setting.

Remark. In contrast to the dynamic programming formulation of Lautenbacher et alii (1999) the structure of an optimal policy is not clear beforehand. The main difference between the models is that customer choice behaviour is also considered. However, if the purchase probabilities follow an independent model, as described by Talluri & van Ryzin (2004a), we do get similar admission control policies. More precisely, if the purchase probabilities are given by

$$P_j(S) = \begin{cases} p_j & \text{if } j \in S, \\ 0 & \text{otherwise,} \end{cases}$$

then an optimal policy of Equation (3.1) will have the form

accept a fare product j request if and only if

$$r_j + V_{t-1}(x + e_j) > V_{t-1}(x).$$

In the remainder of this chapter we focus on dependent demand models.

Reformulation

The customer choice cancellation model can be reformulated such that the costs are incorporated in the rewards. We do this by adding the expected costs caused by cancellations to the value function. This reasoning is borrowed from Lautenbacher et alii (1999), who use it in their RM model which does not take customer choice behaviour into account. We show in Theorem 3.1 below that under certain conditions the customer choice cancellation model is equivalent to a one-dimensional problem.

Let $H(x,t)$ be the expected costs from cancellations from the reservations in state (x,t). The costs are fixed because we do not have control over the reservations. Then $H(x,t)$ is given by the recursive formula

$$H(x,t) = \begin{cases} \sum_{j\in N} \gamma_j(x)\big(c_j(t) + H(x-e_j, t-1)\big) + \left(1 - \sum_{j\in N} \gamma_j(x)\right) H(x,t-1) & \text{if } t > 0, \\ 0 & \text{if } t = 0. \end{cases}$$

Define the value function $\tilde{V}_t(x)$ by

$$\tilde{V}_t(x) := V_t(x) + H(x,t). \tag{3.2}$$

The value function $\tilde{V}_t(x)$ can be interpreted as the maximal expected revenue to go from reservations without costs from current cancellations. These costs are unavoidable and have no influence on future profits. Note that cancellations still occur, but it will turn out that this reformulation alters the cancellation term in a preferable way.

Define

$$\Delta H_j(x,t) := H(x,t-1) - H(x+e_j, t-1),$$

for all $x \in X$ and for all $t > 0$. Then we can rewrite $\tilde{V}_t(x)$ as follows:

$$\tilde{V}_t(x) = \begin{cases} \max_{S\subset N} \Big\{ \lambda \sum_{j\in S} P_j(S)\big(r_j + \Delta H_j(x,t) + \tilde{V}_{t-1}(x+e_j)\big) & \text{if } t > 0 \\ + \sum_{j\in N} \gamma_j(x)\big(\tilde{V}_{t-1}(x-e_j)\big) & \\ + \left(1 - \lambda \sum_{j\in S} P_j(S) - \sum_{j\in N} \gamma_j(x)\right) \tilde{V}_{t-1}(x) \Big\}, & \\ q(x) & \text{if } t = 0. \end{cases} \tag{3.3}$$

Equation (3.3) is equivalent to Equation (3.1) in the sense that it leads to the same policy. This follows from Equation (3.2) and the fact that the policy to go has no influence on $H(x,t)$.

Efficient sets

The customer choice base model has some elegant properties described in Talluri & van Ryzin (2004a) which reduce the action space and therefore the computation time. Unfortunately, these properties do not hold in the customer choice cancellation model. To show this, we introduce the concept of *efficient sets* in Definition 3.1 below (which is the same as Definition 1 in Talluri & van Ryzin, 2004a). We will show later on that some heuristics do satisfy some of the elegant properties described in Talluri & van Ryzin (2004a).

Definition 3.1. For each $S \subset N$, $x \in X$ and $t \in \{0, \dots, T\}$ define $Q(S)$ and $R(S)$ by

$$Q(S) := \sum_{j \in S} P_j(S), \qquad R(S,x,t) := \sum_{j \in S} P_j(S)(r_j + \Delta H_j(x,t)).$$

A set S^* is called *inefficient* at (x,t), with $x \in X$ and $t \in \{0, \dots, T\}$, if there exist $\alpha(S) > 0$ for all $S \subset N$ with $\sum_{S \subset N} \alpha(S) = 1$ such that both of the following inequalities hold:

$$Q(S^*) \geq \sum_{S \subset N} \alpha(S) Q(S),$$
$$R(S^*,x,t) < \sum_{S \subset N} \alpha(S) R(S,x,t).$$

Otherwise S^* is called *efficient* at (x,t).

The intuition behind efficient sets is as follows. An offer set $S^* \subset N$ is efficient if there is no randomisation of other offer sets $S \subset N$ such that the expected reward is strictly greater than $R(S^*,x,t)$ and the probability of a purchase is at most $Q(S^*)$. Efficient sets are on the efficient frontier of the trade-off between the expected revenue $R(S^*,x,t)$ and the probability of purchase $Q(S^*)$.

Efficient sets can be identified using the *largest marginal revenue procedure* described by Talluri & van Ryzin (2004a), which is the following: Initialise with the first efficient set $S_0 = \emptyset$. Then iteratively proceed as follows: Suppose S_i is the i-th efficient set found by the procedure. Then the $(i+1)$-th efficient set S_{i+1} is equal to

$$S_{i+1} = \underset{S:\ \substack{Q(S) \geq Q(S_i) \\ R(S,x,t) \geq R(S_i,x,t)}}{\arg\max} \frac{R(S,x,t) - R(S_i,x,t)}{Q(S) - Q(S_i)}.$$

The complexity of this procedure is $O(m2^n)$, where m is the number of efficient sets, which is at most 2^n. For large instances of n this procedure is computationally intractable, so then the use of heuristics is needed. One heuristic (proposed by Talluri & van Ryzin, 2004a) is to use only a subset of all efficient sets and then use the largest marginal revenue procedure. By looking at the specific instance certain offer sets can be ruled out. For example, if a hotel offers the same room under the same conditions at two different prices €80 and €120 at the same time, the price-sensitive customer would almost surely purchase the room for €80.

In the base model no cancellations exist, such that $\Delta H_j(x,t) = 0$ for all (x,t). Hence, the set of efficient sets is the same for each (x,t). In contrast, for the cancellation model the set of efficient sets is not necessarily the same for all (x,t), as is shown in Example 3.2.

Example 3.2. Consider Example 3.1. Let $\gamma = (0.05, 0.0025, 0.001)$, $C = 5$, $T = 10$. Consider the states (x,t_1) and (x,t_2), with $t_1 = 1$, $t_2 = 10$, and $x = (0,0,0)$. The efficient sets of state (x,t_1) and (x,t_1) are given in Table 3.2. Offer set $\{1,3\}$ is efficient in state (x,t_1) but not in state (x,t_2), and offer set $\{1,2\}$ is not efficient in state (x,t_1) but it is efficient in state (x,t_2).

S	$Q(S)$	$R(S,x,t_1)$	$R(S,x,t_1)$	Efficient set for (x,t_1)	Efficient set for (x,t_2)
$\emptyset$	0	0	0	Yes	Yes
$\{1\}$	0.3	48	30	Yes	Yes
$\{2\}$	0.4	40	40	No	No
$\{3\}$	0.5	45	45	No	No
$\{1,2\}$	0.7	76	69	No	Yes
$\{1,3\}$	0.8	93	75	Yes	No
$\{2,3\}$	0.9	85	85	No	No
$\{1,2,3\}$	1	101	95	Yes	Yes

Table 3.2: Efficient sets for Example 3.1.

An important property of the base model is that an inefficient set is never an optimal solution (see Proposition 2 in Talluri & van Ryzin, 2004a). Unfortunately, this property generally does not hold for the cancellation model. To this end consider Example 3.3.

Example 3.3. Consider Example 3.2. In state $x = (0,0,4)$ at time T offer set $\{1,2\}$ is an optimal offer set. However, offer set $\{1,2\}$ is not efficient in state x at time T, as was shown in Example 3.2.

3.3.2 Tractable Exact Solution: Linear and Equal Cancellations (LEC) Assumption

The model under the linear and equal cancellation Assumption 3.3 turns out to have the elegant property that it can tractably be solved exactly. In order to see this we use some derivations similar to Lautenbacher et alii (1999).

First consider the model under the independence of cancellations Assumption 3.1. The statement in Lemma 3.1 is taken from Lautenbacher et alii (1999), Lemma 2.

Lemma 3.1. Under Assumption 3.1 we have that

$$H(x,t) = \sum_{j \in N} H_j(x_j,t),$$

where the function $H_j(x_j,t)$, $1 \le j \le n$, satisfy the following recursive equations:

$$H_j(x_j,t) = \begin{cases} (1-\gamma_j(x_j))H_j(x_j,t-1) & \text{if } t > 0, \\ +\ \gamma_j(x_j)(c_j(t) + H_j(x_j-1,t-1)) & \\ & 0\text{if } t = 0. \end{cases}$$

for all $0 \le t \le T$.

Define $\Delta H_j(x_j,t)$ by

$$\Delta H_j(x_j,t) := H_j(x_j+1,t-1) - H_j(x_j,t-1).$$

Lemma 3.2 below states that under the independence of cancellations assumption Equation (3.3) can be simplified.

Lemma 3.2. Under Assumption 3.1 Equation (3.3) can be rewritten to

$$\tilde{V}_t(x) = \begin{cases} \max\limits_{S \subset N} \Big\{ \lambda \sum\limits_{j \in S} P_j(S)\big(r_j - \Delta H_j(x,t) + \tilde{V}_{t-1}(x+e_j)\big) & \text{if } t > 0, \\ \quad + \sum\limits_{j \in N} \gamma_j(x_j)\big(\tilde{V}_{t-1}(x-e_j)\big) & \\ \quad + \Big(1 - \lambda \sum\limits_{j \in S} P_j(S) - \sum\limits_{j \in N} \gamma_j(x_j)\Big) \tilde{V}_{t-1}(x) \Big\} & \\ q(x) & \text{if } t = 0, \end{cases}$$

where $\Delta H_j(x_j, t)$ satisfies the following recursive formula

$$\Delta H_j(x_j,t) = \begin{cases} (\gamma_j(x_j+1) - \gamma_j(x_j))c_j(t) & \text{if } t > 1, \\ +\ (1 - \gamma_j(x_j+1))\Delta H_j(x_j, t-1) & \\ +\ \gamma_j(x_j)\Delta H_j(x_j - 1, t-1) & \\ 0 & \text{if } t = 1. \end{cases}$$

Now consider the linear cancellations Assumption 3.2. Lemma 3.3 below simplifies the calculations of $\Delta H_j(x_j, t)$ even further. The statement is taken from Lautenbacher et alii (1999), Lemma 3.

Lemma 3.3. Under Assumption 3.2 $\Delta H_j(x_j, t)$ is independent of the number of reservations x_j such that we can write $\Delta H_j(t) = \Delta H_j(x_j, t)$ for all $j \in N$. Moreover, $\Delta H_j(t)$ satisfies

$$\Delta H_j(t) = \begin{cases} \gamma_j c_j(t) + (1 - \gamma_j)\Delta H_j(t-1) & \text{if } t > 1, \\ 0 & \text{if } t = 1, \end{cases}$$

for all $j \in N$.

Finally, consider the model under the linear and equal cancellations Assumption 3.3. Theorem 3.1 below shows that the states of the problem is significantly reduced. The statement is an extension of Theorem 1 in Lautenbacher et alii (1999) with customer choice behaviour.

Theorem 3.1.

Under Assumption 3.3 we have that $\tilde{V}_t(x) = \tilde{W}_t(y)$ for all (x,t), where $y = \sum_{j \in N} x_j$ and $\tilde{W}_t(y)$ is given by the following formula:

$$\tilde{W}_t(y) = \begin{cases} \max\limits_{S \subset N} \Big\{ \lambda \sum\limits_{j \in S} P_j(S)\big(r_j - \Delta H_j(t) + \tilde{W}_{t-1}(y+1)\big) & \text{if } t > 0, \\ \quad +\ \gamma y \tilde{W}_{t-1}(y-1) & \\ \quad + \Big(1 - \lambda \sum\limits_{j \in S} P_j(S) - \gamma y\Big) \tilde{W}_{t-1}(y) \Big\} & \\ q(y) & \text{if } t = 0. \end{cases} \tag{3.4}$$

Proof. We prove that $\tilde{V}_t(x) = \tilde{W}_t(y)$ by induction to t. For $t = 0$ we have that $\tilde{V}(x,0) = q(x)$ and $\tilde{W}(y,0) = q(y)$. By Assumption 3.3 we have that

$$\tilde{V}(x,0) = q(x) = q(y) = \tilde{W}(y,0),$$

so the statement holds for $t = 0$.

Suppose $t > 0$ and the statement holds for all values smaller than t. By Assumption 3.3, Lemma 3.2, and Lemma 3.3 we have that

$$\begin{aligned}\tilde{V}_t(x) = \max_{S\subset N}\Bigg\{ &\lambda \sum_{j\in S} P_j(S)\big(r_j - \Delta H_j(t) + \tilde{V}_{t-1}(x+e_j)\big) \\ &+ \gamma \sum_{j\in N} x_j \big(\tilde{V}_{t-1}(x-e_j)\big) \\ &+ \left(1 - \lambda \sum_{j\in S} P_j(S) - \gamma \sum_{j\in N} x_j\right)\tilde{V}_{t-1}(x)\Bigg\}.\end{aligned}$$

Using the induction hypothesis we can rewrite this to

$$\begin{aligned}\tilde{V}_t(x) = \max_{S\subset N}\Bigg\{ &\lambda \sum_{j\in S} P_j(S)\big(r_j - \Delta H_j(t) + \tilde{W}_{t-1}(y+1)\big) \\ &+ \gamma y \tilde{W}_{t-1}(y-1) + \left(1 - \lambda \sum_{j\in S} P_j(S) - \gamma y\right)\tilde{W}_{t-1}(y)\Bigg\},\end{aligned}$$

and hence $\tilde{V}_t(x) = \tilde{W}_t(y)$. □

Example 3.3 showed that in general inefficient sets can be optimal. However, under Assumption 3.3 inefficient sets are never an optimal solution, as Proposition 3.1 below states. The set of efficient sets still depends on (x,t), so unfortunately these efficient sets are not of much use to us.

Proposition 3.1. An inefficient set is never an optimal solution to Equation (3.4).

Proof. See Appendix 3.A. □

3.3.3 Independence of Cancellations Heuristic (IOC)

The *Independence of Cancellations* heuristic is appropriate for the model under the independence of cancellations Assumption 3.1. The goal of this heuristic is to find a solution under a model with general assumptions. In this approach we reduce the state space by only keeping track of the total number of reservations rather than the number of reservations per product. Again we make the assumption that the probability that product j is cancelled depends only on the number of reservations for product j, i.e., $\gamma_j^t(x) = \gamma_j^t(x_j)$ (see the independence of cancellations Assumption 3.1), but we make no further assumptions. The advantage of this heuristic is that it can be applied to problems with general cancellation probability functions.

Define the one-dimensional state space Y by

$$Y := \{y \in \mathbb{N}\,|0 \le y \le C_{\max}\}\,.$$

Each state $y \in Y$ corresponds to the total number of reservations. Let $\gamma(y,t)$ be the conditional probability that a reservation is cancelled in state (y,t). Let $c(y,t)$ be the expected costs for a cancelled reservation in state (y,t). The Bellman equation corresponding to this problem is

$$W_t(y) = \begin{cases} \max\limits_{S\subset N}\Bigg\{\lambda\sum\limits_{j\in S} P_j(S)\big(r_j + W_{t-1}(y+1)\big) & \text{if } t>0,\\ +\gamma(y,t)\big(-c(y,t)+W_{t-1}(y-1)\big) & \\ +\left(1-\lambda\sum\limits_{j\in S}P_j(S)-\gamma(y,t)\right)W_{t-1}(y)\Bigg\}, & \\ q(y) & \text{if } t=0. \end{cases} \tag{3.5}$$

The problem with using this one-dimensional state space is that we do not know the expected costs $c(y,t)$ and the conditional cancellation probability $\gamma(y,t)$. This follows from the fact that we do not know the mix of current reservations, since if we are in state $y \in Y$ we could be in any state $x \in X$ with

$$\sum_{j=1}^{n} x_j = y.$$

The current mix of reservations depends on the past strategy and therefore $c(y,t)$ and $\gamma(y,t)$ depend on the past strategy. To overcome this problem we propose the following procedure to estimate $c(y,t)$ and $\gamma(y,t)$.

Let π be an arbitrary strategy. Let $\bar{x}(y,t) \in \mathbb{R}^n$ be the expected mix of current reservations in state y at time t. Let $p(y,t)$ be the probability that we are in state y at time t. Let $p_t^y(x_j = k)$ be the probability that we have k reservations for product j in state (y,t). For strategy π we will explain a recursive procedure to calculate $\bar{x}(y,t)$, $p(y,t)$, $p_t^y(x_j = k)$, $\gamma(y,t)$, and $c(y,t)$.

Let y and t be arbitrary and suppose that we know $\bar{x}(y',t')$, $p(y',t')$, $\gamma(y',t')$, $p_{t'}^{y'}(x_j = k)$, and $c(y',t')$ for all $t' < t$, for all k and j, and for all y'. Then we can evaluate their values as stated in Lemmas 3.4 and 3.5 below.

Lemma 3.4. Define $\hat{p}_t^y(x_j = k)$ by

$$\begin{aligned}\hat{p}_t^y(x_j = k) = p(y+1, t+1) \Bigg[& p_{t+1}^{y+1}(x_j = k+1)\gamma_j^{t+1}(k+1) \\ & + p_{t+1}^{y+1}(x_j = k) \sum_{i \neq j} \sum_{k'=0}^{y+1-k} p_{t+1}^{y+1}(x_i = k')\gamma_i^{t+1}(k') \Bigg] \\ & + p(y, t+1) p_{t+1}^y(x_j = k) \Bigg[1 - \lambda \sum_{i \in S} P(S, i) \\ & - \sum_{i \neq j} \sum_{k'=0}^{y-k} p_{t+1}^y(x_i = k')\gamma_i^{t+1}(k') - \gamma_j^{t+1}(k) \Bigg] \\ & + p(y-1, t+1) \Bigg[p_{t+1}^{y-1}(x_j = k-1)\lambda P_j(S) \\ & + p_{t+1}^{y-1}(x_j = k)\lambda \sum_{i \in S \setminus \{j\}} P_j(S) \Bigg],\end{aligned}$$

for all k. Then $p_t^y(x_j = k)$ is given by

$$p_t^y(x_j = k) = \frac{\hat{p}_t^y(x_j = k)}{\sum_{k'=0}^{y} \hat{p}_t^y(x_j = k')}.$$

Proof. See Appendix 3.A. □

Lemma 3.5. $\bar{x}(y, t)$, $p(y, t)$, $\gamma(y, t)$, and $c(y, t)$ can be found via the following formulas:

$$\begin{aligned}\bar{x}_j(y, t) &= \sum_{k=0}^{y} p_t^y(x_j = k)k, \\ p(y, t) &= p(y+1, t+1)\gamma(y+1, t+1) + \lambda \sum_{j \in S} P_j(S) p(y-1, t+1) \\ &\quad + \left(1 - \lambda \sum_{j \in S} P_j(S) - \gamma(y, t+1) \right) p(y, t+1), \\ \gamma(y, t) &= \sum_{j=1}^{n} \sum_{k=0}^{y} p_t^y(x_j = k)\gamma_j^t(k), \\ c(y, t) &= \frac{\sum_{j=1}^{n} \sum_{k=0}^{y} p_t^y(x_j = k)\gamma_j^t(k) r_j(t)}{\gamma(y, t)}.\end{aligned}$$

Moreover, under the linear and equal cancellation rates Assumption 3.3 we have that

$$\begin{aligned}\gamma(y, t) &= \sum_{j=1}^{n} \gamma_j^t \bar{x}_j(y, t), \\ c(y, t) &= \frac{\sum_{j=1}^{n} r_j(t)\gamma_j^t \bar{x}_j(y, t)}{\gamma(y, t)}.\end{aligned}$$

Proof. See Appendix 3.A. □

We can use the equations from Lemma 3.4 and Lemma 3.5 to iteratively calculate the conditional cancellation probability $\gamma(y,t)$ and the expected costs $c(y,t)$ of a cancelled product in state (y,t) under a policy π.

The Independence of Cancellations algorithm iterates the following procedure until we hit a stopping criterion, for example after a fixed number of iterations. First, start with an arbitrary expected costs $\hat{c}(y,t)$ and cancellation probability $\hat{\gamma}(y,t)$ for each state (y,t). A good initial solution would be to use the linear and equal cancellation rates Assumption 3.3 and evaluate $\hat{c}(y,t)$ and $\hat{\gamma}(y,t)$ according to the obtained strategy. Then solve the problem using dynamic programming with Equation (3.5), where we use $\hat{c}(y,t)$ as an approximation of the expected costs $c(y,t)$ and $\hat{\gamma}(y,t)$ as an approximation of the conditional cancellation probability $\gamma(y,t)$. Then we find an optimal strategy π which can be used use to calculate the expected costs $\hat{c}(y,t)$ and the c cancellation probability $\hat{\gamma}(y,t)$ corresponding to π, using the equations in Lemma 3.4 and Lemma 3.5. See Algorithm 3.1 for a summary of this heuristic.

Algorithm 3.1. Independence of Cancellations Heuristic (IOC)

1. Start with initial expected costs $\hat{c}(y,t)$ and cancellation probability $\hat{\gamma}(y,t)$.
2. Find an optimal strategy π with dynamic programming on Equation (3.5) using $\hat{c}(y,t)$ as an approximation for $c(y,t)$ and $\hat{\gamma}(y,t)$ as an approximation for $\gamma(y,t)$.
3. Use π and the equations in Lemma 3.4 and Lemma 3.5 to calculate $\hat{c}(y,t)$ and $\hat{\gamma}(y,t)$.
4. Go to step 2 unless a stopping criterion is hit.

We now investigate the complexity of Step 2 of this algorithm.

Efficient Sets

Another elegant property of the Independence of Cancellations (IOC) algorithm is that only efficient offer sets are potential optimal sets, while the others are never optimal. Moreover, this set of potential offer sets is the same for all states (y,t). In Proposition 2 of Talluri & van Ryzin (2004a) it is shown that in the customer choice base model inefficient sets are never optimal. We show in Proposition 3.2 that inefficient sets are never optimal for the IOC algorithm and that the set of efficient sets is the same for all (y,t).

Proposition 3.2. An inefficient set is never an optimal solution to Equation (3.5). Moreover, the set of efficient sets is the same for all (x,t).

Proof. In Equation (3.5) we have that $\Delta H_j(y,t) = 0$ for all j and (y,t), so the set of efficient sets is the same for all (y,t). The proof that inefficient sets are never an optimal solution to Equation (3.5) is analogous to the proof of Proposition 3.1. □

3.3.4 Linear Cancellation Rates Heuristic (LCR)

The goal of the second heuristic is to be fast and efficient. To accomplish this we do the following:

- *Work under the linear cancellation rates Assumption 3.2.* This way the model is restricted, but computation time is gained.
- *Only keep track of the total number of reservations.* Information on the number of reservations per product is lost, but computation time is gained.
- *Only consider efficient sets in our policy.* Then the computation time is reduced because less solutions need to be considered, but some potential better solutions are lost.
- *Use the same set of efficient sets in each state.* Computation time is gained because the efficient sets do not have to be evaluated for each state, but some potential better solutions are lost.

In the second heuristic we work under the linear cancellation rates, i.e., assume that $\gamma_j(x_j) = \gamma_j x_j$ for all j. The heuristic can also be applied if we work under the linear and equal cancellation rates Assumption 3.3, since this assumption is a special case of the linear cancellation rates assumption. We solve a one-dimensional problem, where we only keep track of the total number of reservations and not of the reservations per product. We reformulate the problem such that the costs of cancellations are incorporated in the maximisation term of the Bellman equation, as described in Theorem 3.1. Only this time we only consider offer sets that are efficient. In particular, choose the efficient sets that result from taking $R(S,x,t) = 0$ for all (x,t). The motivation for choosing only these efficient sets is that the numerical results in Section 3.5 show that (a) leaving out inefficient solutions does not have a great impact on performance; and (b) choosing among the efficient sets is better than choosing among an equal number of arbitrary sets. The Bellman equation that has to be solved is given by

$$\tilde{W}_t(y) = \begin{cases} \max\limits_{\substack{S\subset N\\ \text{efficient}}} \Bigg\{ \lambda \sum\limits_{j\in S} P_j(S)\big(r_j - \Delta H_j(t) + \tilde{W}_{t-1}(y+1)\big) & \text{if } t > 0, \\ \qquad + \gamma y \tilde{W}_{t-1}(y-1) & \\ \qquad + \left(1 - \lambda \sum\limits_{j\in S} P_j(S) - \gamma y\right) \tilde{W}_{t-1}(y) \Bigg\} & \\ q(y) & \text{if } t = 0. \end{cases} \tag{3.6}$$

However, assume that the cancellation probabilities differ per product. We do not know the cancellation probability γ if we are in state y, since we do not know what the mix of current reservations is. This problem is similar to the problem we described in Section 3.3.3. We could use Lemma 3.5 to approximate γ iteratively, but this is computationally expensive. Moreover, numerical results show that simpler approximations outperform this method (see also Section 3.5). An example of a simple approximation of γ is taking the average of all γ's:

$$\gamma := \frac{1}{n} \sum_{j\in N} \gamma_j. \tag{3.7}$$

The advantage of using this method over the Independence of Cancellations algorithm is that it is computationally faster and performs better under the linear cancellation rates assumption. However, the downside of this method is that it can not be applied to general cancellation probabilities. See Algorithm 3.2 for an overview of the heuristic.

Algorithm 3.2. Linear Cancellation Rates Heuristic (LCR)

1. Estimate $\hat{\gamma}$ by using (γ_j), for example by Equation (3.7).
2. Find an optimal strategy of Equation (3.6).

3.3.5 Multiple Room Types Heuristic (MRT)

In practice hotels often have multiple room types. However, the assumption was made that the hotel has C rooms of the same type. Expanding the solution space with multiple room types leads to an intractable problem. To solve this problem, combine the Linear Cancellation Rates heuristic of Section 3.3.4 with *one-step improvement*, i.e., one step of the well known policy iteration (see also Puterman, 1994). In this approach, first find an approximation of the value function, and then use one maximisation step of policy iteration to find a better solution.

First consider the different room types separately. Let $\mathcal{I}$ be the set of room types. For each room type $i \in \mathcal{I}$ we have a certain capacity C_i. For each room type we only keep track of the total number of reservations for that room type. Under the linear and equal cancellation rates Assumption 3.3 this gives the exact solution; see Theorem 3.1. Let S_i be the offer set of products for room type i and let N_i be the set of products for room type i. Solve the problem per room type, as if it would be the only room type we have, using the Linear Cancellation Rates heuristic described in Section 3.3.4. The Bellman equation per room type $i \in \mathcal{I}$ is given by

$$\tilde{W}_t^i(y) = \begin{cases} \max\limits_{\substack{S_i \subset N_i \\ \text{efficient}}} \Bigg\{ \lambda \sum\limits_{j \in S_i} P_j(S_i)\big(r_j - \Delta H_j(t) + \tilde{W}_{t-1}^i(y+1)\big) & \text{if } t > 0, \\ \qquad +\gamma^i y \tilde{W}_{t-1}^i(y-1) & \\ \qquad + \left(1 - \lambda \sum\limits_{j \in S_i} P_j(S_i) - \gamma^i y\right) \tilde{W}_{t-1}^i(y) \Bigg\} & \\ q(y) & \text{if } t = 0. \end{cases} \tag{3.8}$$

Let u be the vector containing the number of reservations per room type, such that u_i represents the number of reservations for room type $i \in \mathcal{I}$. Then we want to solve the following recursive formula:

$$U_t(u) = \begin{cases} \max\limits_{S \subset N} \Bigg\{ \lambda \sum\limits_{i \in \mathcal{I}} \left[\sum\limits_{j \in S_i} P_j(S_i)\big(r_j - \Delta H_j(t) + U_{t-1}(u + e_i)\big) \right] & \text{if } t > 0, \\ \qquad + \sum\limits_{i \in \mathcal{I}} \left[\sum\limits_{j \in N_i} \gamma_j u_j U_{t-1}(u - e_i) \right] & \\ \qquad + \left(1 - \lambda \sum\limits_{j \in S} P_j(S) - \sum\limits_{j \in N} \gamma_j u_j \right) U_{t-1}(u) \Bigg\} & \\ q(u) & \text{if } t > 0. \end{cases}$$

However, this problem is still too large to solve, even if we use a small number of room types. Therefore, use *one-step improvement* to solve this problem. Suppose u is the current state and we wish to decide which products to offer. We approximate $U_{t-1}(u)$ by the following formula:

$$\tilde{U}_{t-1}(u) := \sum_{i \in \mathcal{I}} \tilde{W}^i_{t-1}(u_i). \tag{3.9}$$

The terms $\tilde{U}_{t-1}(u + e_i)$ and $\tilde{U}_{t-1}(u - e_i)$ can be approximated similarly. Then do one maximisation step to find an approximation $\bar{U}_t(u)$ for $U_t(u)$:

$$\bar{U}_t(u) = \begin{cases} \max_{S \in N} \Bigg\{ \lambda \sum_{i \in \mathcal{I}} \left[\sum_{j \in S_i} P_j(S_i)\big(r_j - \Delta H_j(t) + \tilde{U}_{t-1}(u + e_i)\big) \right] & \text{if } t > 0, \\ \qquad + \sum_{i \in \mathcal{I}} \left[\sum_{j \in N_i} \gamma_j x_j \tilde{U}_{t-1}(u - e_i) \right] & \\ \qquad + \left(1 - \lambda \sum_{j \in S} P_j(S) - \sum_{j \in N} \gamma_j x_j \right) \tilde{U}_{t-1}(u) \Bigg\} & \\ q(u) & \text{if } t = 0. \end{cases} \tag{3.10}$$

The multiple room type heuristic is summarised in Algorithm 3.3 below.

Algorithm 3.3. Multiple Room Types

1. Fix state u.
2. Evaluate $\tilde{W}^i_{t-1}(u_i)$, $\tilde{W}^i_{t-1}(u_i + 1)$, and $\tilde{W}^i_{t-1}(u_i - 1)$ using Equation (3.8) for all $i \in \mathcal{I}$.
3. Define $\tilde{U}_{t-1}(u)$, $\tilde{U}_{t-1}(u + e_i)$, and $\tilde{U}_{t-1}(u - e_i)$ using Equation (3.9).
4. Use Equation (3.10) to find an improved approximation $\bar{U}_t(u)$ of $U_t(u)$.

Remark. The action space of this problem is still large, even though we only use efficient sets per room type. The concept of efficient sets is only defined for single-resource problems. In the context of multiple room types it is worth looking into a generalisation of efficient sets to multiple-resource problems. However, this is beyond the scope of this chapter.

3.4 Estimating Parameters

In this section a method is proposed to estimate the parameters of the model from real data. Ideally an estimation model is used under perfect knowledge, i.e., assume we know (a) which products were offered at what time, (b) when customers arrived, (c) when they made a purchase, and (d) when they cancelled their reservation. A maximum-likelihood estimation model under perfect knowledge is described in Appendix 3.B. In practice not all information is available, unfortunately. For example, the no-purchase option is rarely observed. To overcome this problem we need to consider methods that deal with censored data. One method is the well-known expectation-maximisation algorithm, which is described in Appendix 3.B. However, the expectation-maximisation algorithm is unattractive for practical purposes for two reasons. First, the computation time is notoriously long. Second, the method estimates the parameters of the discretised model instead of the continuous model, which may lead to estimation errors. To overcome these problems we propose an estimation model that is based on the excellent estimation method proposed by Newman et alii (2014). The proposed method assumes the linear cancellation rates Assumption 3.2 and that the purchase probabilities are modelled by the *multinomial logit model* (MNL), which is briefly explained in Appendix 3.B.

In what follows the purchase probabilities and the offer set depend on the time period. Therefore, introduce the following notation. Let Z_{tj} be the variable containing values of the attributes of product j in time period t and let $Z_t = \{Z_{tj}\}_j$. Let S_t be the set of products that is offered in time period t. Let α be the utility of the no-purchase alternative. For convenience we introduce some notation. Let $P_{tj}(\alpha, \beta, S_t, Z_t)$ be the probability that product j is purchased in time period t under α, β, S_t, and Z_t, i.e.,

$$P_{tj}(\alpha, \beta, S_t, Z_t) := \frac{\exp(\beta^\top Z_{jt})}{\exp(\alpha) + \sum_{i \in S_t} \exp(\beta^\top Z_{it})}.$$

Define the probability $P_{t0}(\alpha, \beta, S_t, Z_t)$ that the no-purchase option is chosen as

$$P_{t0}(\alpha, \beta, S_t, Z_t) := \frac{\exp(\alpha)}{\exp(\alpha) + \sum_{i \in S_t} \exp(\beta^\top Z_{it})}.$$

It is convenient to use the probability $P_{t*}(\alpha, \beta, S_t, Z_t)$ that a product is purchased, i.e., the probability that the no-purchase alternative is not chosen. Hence $P_{t*}(\alpha, \beta, S_t, Z_t)$ is given by

$$\begin{aligned} P_{t*}(\alpha, \beta, S_t, Z_t) &:= \sum_{j \in S_t} P_{tj}(\alpha, \beta, S_t, Z_t) \\ &= \frac{\sum_{j \in S_t} \exp(\beta^\top Z_{jt})}{\exp(\alpha) + \sum_{i \in S_t} \exp(\beta^\top Z_{it})} \\ &= 1 - P_{t0}(\alpha, \beta, S_t, Z_t). \end{aligned}$$

Finally, let $P_{tj|*}(\beta, S_t, Z_t)$ be the probability that product j is purchased under the condition that the customer made a purchase, i.e.,

$$P_{tj|*}(\beta, S_t, Z_t) := \frac{\exp(\beta^\top Z_{jt})}{\sum_{i \in S_t} \exp(\beta^\top Z_{it})}.$$

We divide time into T time periods and assume that the parameters to be estimated are constant within these periods. Note that we allow more than one event to happen in one time period. The parameters that we need to estimate are: the no-purchase utility α; the weights β of the attributes of the MNL model; the parameters γ_j of the exponential distributions that model the cancellations; and the parameter λ of the Poisson process that models the arriving customers. We do this by using maximum likelihood estimation.

Assume that the offer set S_t does not change within this time period. Suppose we observe z_{jt} purchases for product j in time period t. Let $\zeta_j(y)$ be the total number of cancellations for product j when there were y reservations for product j, measured over all time periods. Suppose that we observe that for a period of $t_j(y)$ there were y reservations for product j, again measured over all time periods. Note that $t_j(y)$ might be 0.

The maximum likelihood function $L(z, \zeta | \alpha, \beta, \gamma, \lambda)$ is the probability that we observe z and ζ under given α, β, γ, and λ. Since the arrival process is independent from the cancellation process we can write

$$p(z, \zeta | \alpha, \beta, \gamma, \lambda) = p(z | \alpha, \beta, \gamma, \lambda) p(\zeta | \alpha, \beta, \gamma, \lambda),$$

with $p(z|\alpha,\beta,\gamma,\lambda)$ the probability that we observe z given α, β, γ, and λ; and $p(\zeta|\alpha,\beta,\gamma,\lambda)$ the probability that we observe ζ given α, β, γ, and λ. Since the cancellations are only dependent on γ we can even write

$$p(\zeta|\alpha,\beta,\gamma,\lambda) = p(\zeta|\gamma),$$

with $p(\zeta|\gamma)$ the probability that we observe ζ given γ. This also follows from the derivations below.

The cancellation processes are independent from each other, so we can write

$$p(\zeta|\gamma) = \prod_{j\in N} p(\zeta_j|\gamma_j),$$

with $p(\zeta_j|\gamma_j)$ the probability that we observe ζ_j given γ_j. Moreover, if we consider the time intervals $t_j(y)$ we can write

$$p(\zeta_j|\gamma_j) = \prod_{y=0}^{C_{\max}} p\big(\zeta_j(y)\big|\gamma_j\big).$$

Since the cancellations follow a Poisson process with parameter $\gamma_j y t_j(y)$ the probability $p(\zeta_j(y)|\gamma_j)$ that we observe $\zeta_j(y)$ cancellations in time period t conditional on γ_j is equal to

$$p\big(\zeta_j(y)\big|\gamma_j\big) = \big(\gamma_j y t_j(y)\big)^{\zeta_j(y)} \frac{e^{-\gamma_j y t_j(y)}}{\zeta_j(y)!}.$$

If we take the natural logarithm of $p(\zeta_j(y)|\gamma_j)$ we get

$$\log\big(p(\zeta_j(y)|\gamma_j)\big) = \zeta_j(y)\log(\gamma_j) + \zeta_j(y)\log\big(yt_j(y)\big) - \gamma_j y t_j(y) - \log\big(\zeta_j(y)!\big).$$

Combining all equations we find

$$\log\big(p(\zeta|\gamma)\big) = \sum_{j\in N}\sum_{y=0}^{C_{\max}} \big[\big\{\zeta_j(y)\log(\gamma_j) - \gamma_j y t_j(y)\big\} + \big\{\zeta_j(y)\log\big(yt_j(y)\big) - \log\big(\zeta_j(y)!\big)\big\}\big].$$

If we want to maximise this with respect to γ, then we can solve the following equations separately for each $j \in N$:

$$\sum_{y=0}^{C_{\max}} \big[\zeta_j(y)\log(\gamma_j) - \gamma_j y t_j(y)\big]. \tag{3.11}$$

If we take the first derivative with respect to γ_j we find

$$\frac{1}{\gamma_j}\sum_{y=0}^{C_{\max}} \zeta_j(y) - \sum_{y=0}^{C_{\max}} y t_j(y).$$

The optimum can be found by setting this equation to zero. This gives

$$\gamma_j = \frac{\sum_{y=0}^{C_{\max}} \zeta(y)}{\sum_{y=0}^{C_{\max}} y t_j(y)}.$$

Let $\zeta_j := \sum_{y=0}^{C_{\max}} \zeta(y)$, i.e., ζ_j is the total number of cancellations for product j. Then we have that

$$\gamma_j = \frac{\zeta_j}{\sum_{y=0}^{C_{\max}} y t_j(y)}. \tag{3.12}$$

By taking the second derivative of Equation (3.11) we find

$$-\frac{1}{\gamma_j^2} \sum_{y=0}^{C_{\max}} \zeta_j(y),$$

which is negative, such that the optimum found in Equation (3.12) is a maximum.

For completeness we below present the derivations of the results by Newman et alii (2014), which can directly be applied to our model.

The expression for $p(z|\alpha, \beta, \gamma, \lambda)$ does not depend on γ so we can write $p(z|\alpha, \beta, \lambda)$. Conditioning to the number of reservations m_t in time period t gives

$$p(z|\alpha, \beta, \lambda) = p(z|m, \beta, \lambda) p(m|\alpha, \beta, \lambda),$$

where $m = (m_1, \ldots, m_T)$. Note that $p(z|m, \beta, \lambda)$ does not depend on the no-purchase alternative parameter α. The probability $p(z|m, \beta, \lambda)$ that we observe z under the conditions that we observe m_t reservations with parameter β and λ is equal to

$$p(z|m, \beta, \lambda) = \prod_{t=1}^{T} \binom{m_t}{z_{1t}, \ldots, z_{nt}} \prod_{j \in N} \left[P_{tj|*}(\beta, S_t, Z_t) \right]^{z_{jt}}.$$

Note that since $\sum_{j \in N} z_{jt} = m_t$ and $z_{jt} = 0$ for $j \notin S_t$ we have that

$$\prod_{j \in N} \left[P_{tj|*}(\beta, S_t, Z_t) \right]^{z_{jt}} = \prod_{j \in S_t} \left[\frac{\exp(\beta^\top Z_{jt})}{\sum_{i \in S_t} \exp(\beta^\top Z_{it})} \right]^{z_{jt}} = \frac{\prod_{j \in S_t} \left[\exp(\beta^\top Z_{jt}) \right]^{z_{jt}}}{\left[\sum_{i \in S_t} \exp(\beta^\top Z_{it}) \right]^{m_t}}.$$

By taking the natural logarithm we find

$$\log \big(p(z|m, \alpha, \beta, \lambda) \big) = \sum_{t=1}^{T} \log \left(\frac{m_t!}{z_{1t}! \cdot \ldots \cdot z_{nt}!} \right) + \sum_{j \in S_t} z_{jt} \beta^\top Z_{jt} - m_t \log \left(\sum_{i \in S_t} \exp(\beta^\top Z_{it}) \right).$$

If we want to maximise this with respect to β we only have to consider

$$\sum_{t=1}^{T} \left[\sum_{j \in S_t} z_{jt} \beta^\top Z_{jt} - m_t \log \left(\sum_{i \in S_t} \exp(\beta^\top Z_{it}) \right) \right]. \tag{3.13}$$

On the other hand, m_t follows a Poisson distribution with parameter $\lambda \sum_{i \in S} P_{t*}(\alpha, \beta, S_t, Z_t)$. Therefore the probability $p(m|\alpha, \beta, \lambda)$ that we observe m_t reservations in time period t is given by

$$p(m|\alpha, \beta, \lambda) = \prod_{t=1}^{T} \frac{\left[\lambda \sum_{i \in S} P_{t*}(\alpha, \beta, S_t, Z_t) \right]^{m_t} e^{-\lambda P_{t*}(\alpha, \beta, S_t, Z_t)}}{m_t!}.$$

Taking the logarithm gives

$$\log\left(p(m|\alpha,\beta,\lambda)\right) = \sum_{t=1}^{T} m_t \left[\log\left(\lambda P_{t*}(\alpha,\beta,S_t,Z_t)\right)\right] - \lambda P_{t*}(\alpha,\beta,S_t,Z_t) - \log(m_t!). \tag{3.14}$$

Note that if we maximise Equation (3.14) with respect to α and λ we can ignore the last term. The first partial derivative with respect to λ equals

$$\frac{1}{\lambda}\sum_{t=1}^{T} m_t - \sum_{t=1}^{T} P_{t*}(\alpha,\beta,S_t,Z_t).$$

The optimum is attained at

$$\lambda = \frac{\sum_{t=1}^{T} m_t}{\sum_{t=1}^{T} P_{t*}(\alpha,\beta,S_t,Z_t)}, \tag{3.15}$$

where $\sum_{t=1}^{T} m_t$ is the total number of observed reservations. The second partial derivative with respect to λ equals

$$-\frac{1}{\lambda^2}\sum_{t=1}^{T} m_t,$$

which is negative for $\lambda \neq 0$, so the optimum from Equation (3.15) is a maximum. We can replace λ in Equation (3.14) by Equation (3.15), which gives

$$\log\left(p(m_t|\alpha,\beta,\lambda)\right) = \sum_{t=1}^{T} m_t \log\left(\frac{\sum_{t=1}^{T} m_t}{\sum_{t=1}^{T} P_{t*}(\alpha,\beta,S_t,Z_t)} P_{t*}(\alpha,\beta,S_t,Z_t)\right) - \frac{\sum_{t=1}^{T} m_t}{\sum_{t=1}^{T} P_{t*}(\alpha,\beta,S_t,Z_t)} P_{t*}(\alpha,\beta,S_t,Z_t) - \log(m_t!). \tag{3.16}$$

If β is given we can find α by maximising Equation (3.16) with respect to α. Then we can find λ by using Equation (3.15).

The estimation procedure is summarised as follows.

Algorithm 3.4. Three-Step Continuous Estimation

1. Find $\hat{\gamma}$ using Equation (3.12).
2. Find $\hat{\beta}$ using Equation (3.13).
3. Use the estimate $\hat{\beta}$ to find $\hat{\alpha}$ using Equation (3.16). Then use $\hat{\beta}$ and $\hat{\alpha}$ to find $\hat{\lambda}$ using Equation (3.15).

Theorem 3.2 below states that Algorithm 3.4 provides consistent estimators. Moreover, the estimator for the cancellation parameters γ is an unbiased estimator.

Theorem 3.2. As the sample size increases, Algorithm 3.4 leads to consistent estimators. Moreover, the estimator for γ found by Equation (3.12) is unbiased.

Proof. Theorems 1 and 2 in Newman et alii (2014) show that the estimators of β and (α, λ) are consistent, respectively. To show that the estimator $\hat{\gamma}$ of the cancellation parameter found by Equation (3.12) is a consistent estimator, we consider the cancellation parameters separately per product. Consider a product $j \in N$. Let γ_j^* be the true parameter and let γ_j^k be the parameter obtained using Equation (3.12) with sample size k. In other words,

$$\gamma_j^k = \frac{\sum_{y=0}^{C_{\max}} \zeta_j^k(y)}{\sum_{y=0}^{C_{\max}} y t_j^k(y)},$$

with $\zeta_j^k(y)$ the observed number of cancellations in state y and $t_j^k(y)$ the observed time spent in state y with k observations. We can think of one observation as an arrival day with all booking history. The expected number of cancellations when the state is y is equal to $\gamma_j^* y t_j^k(y)$. Note that $t_j^k(y)$ is the sum of k i.i.d. random variables $t_j(y)$, where $t_j(y)$ is the time spent in state y for one observation. As the sample size increases, we have that

$$\lim_{k \to \infty} \frac{y t_j^k(y)}{k} = y\mathbb{E}[t_j(y)].$$

On the other hand, $\zeta_j^k(y)$ is the sum of k i.i.d. random variables $\zeta_j(y)$, where $\zeta_j(y)$ is the number of cancellations for one observation. As the sample size increases, we have that

$$\lim_{k \to \infty} \frac{\zeta_j^k(y)}{k} = \mathbb{E}[\zeta_j(y)] = \mathbb{E}[\mathbb{E}[\zeta_j(y)|t_j(y)]] = \mathbb{E}[\gamma_j^* y t_j(y)] = \gamma_j^* y \mathbb{E}[t_j(y)].$$

Hence we have that

$$\lim_{k \to \infty} \gamma_j^k = \lim_{k \to \infty} \frac{\frac{1}{k}\sum_{y=0}^{C_{\max}} \zeta_j^k(y)}{\frac{1}{k}\sum_{y=0}^{C_{\max}} y t_j^k(y)} = \frac{\gamma_j^* \sum_{y=0}^{C_{\max}} y\mathbb{E}[t_j(y)]}{\sum_{y=0}^{C_{\max}} y\mathbb{E}[t_j(y)]} = \gamma_j^*,$$

so Equation (3.12) is a consistent estimator. Moreover, using conditional expectation it follows that

$$\mathbb{E}[\gamma_j^k] = \mathbb{E}\left[\mathbb{E}\left[\left.\frac{\sum_{y=0}^{C_{\max}} \zeta_j^k(y)}{\sum_{y=0}^{C_{\max}} y t_j^k(y)}\right| t_j^k(\cdot)\right]\right] = \mathbb{E}\left[\frac{\sum_{y=0}^{C_{\max}} \gamma_j^* y t_j^k(y)}{\sum_{y=0}^{C_{\max}} y t_j^k(y)}\right] = \gamma^*,$$

so Equation (3.12) provides an unbiased estimator. □

3.5 Numerical Results

In this section numerical results are provided to validate the customer choice cancellation model. The four main results that are provided are: 1) the fast computation time of the estimation method described in Section 3.4; 2) the fast computation time of the heuristics; 3) the good performance of the heuristics; and 4) the performance of the heuristics under estimated parameters. The parameters that are used are based on Simulation Example 2 in Talluri & van Ryzin (2004a). Let $n = 10$ be the number of products sold with corresponding price vector

$$r = (240, 220, 190, 160, 120, 112, 96, 80, 74, 70).$$

Demand, cancellation rates, and purchase probabilities are independent from the time period t. To allow the study of the effect of high volume in demand and low volume in demand the *load factor* l is introduced. In our studies we use $l \in \{0.6, 0.8, 1, 1.2, 1.4\}$. Demand λ per time unit is then defined as

$$\lambda(C, l, T) = \frac{Cl}{T}.$$

A low value of l implies a low demand relative to the capacity C and a high value of l implies a high demand relative to the capacity C.

Purchase probabilities are modelled by the MNL model. The only attribute that we consider is price, but we assume that there exist high price-sensitive and low price-sensitive customers which have different parameters $\beta_H = -0.005$ and $\beta_L = -0.0015$, respectively (as in Talluri & van Ryzin, 2004a). The no-purchase parameter is set to $\alpha = 0$ such that the MNL model is the same as in Talluri & van Ryzin (2004a). We allow overbooking up to 20% of the total capacity C. Cancellation rates are assumed linear and the parameters $\gamma \in \mathbb{R}^n$ depend on l and T in the following way:

$$\gamma = (18/25, 8/25, 14/25, 7/25, 1/5, 9/25, 4/25, 2/25, 1/125, 1/25)\frac{l}{T}.$$

The policies that follow from the heuristics need to be compared using the corresponding revenue. For small problem instances the exact solution method can be used, but for larger instances the problem instance is intractable. In that case simulation is used to estimate the revenue corresponding to a policy. The estimation errors are between 0.1% and 0.4%.

3.5.1 Computation Time of Estimation Methods

An estimation method only has value for practical applications if it can finish in reasonable time. The three-step continuous estimation method (3SC) described in Section 3.4 has this property.

In Table 3.3 the computation time of 3SC is compared with the computation time of Expectation-Maximisation (EM) for different sample sizes. The results show that the computation time of the EM-algorithm increases drastically when the sample size increases, while the computation time of the 3SC-method does not. This is only for a small example with $C = 20$ rooms, so the results will be worse for practical instances. Fortunately, the 3SC-method can be used also for larger instances, with more capacity. In Table 3.4 the computation time of 3SC is given for a hotel with $C = 200$ rooms for different sample sizes. The computation time increases but it stays reasonable.

3.5.2 Computation Time: Heuristics vs. Exact Solution

In the analysis of the cancellation model it was stated that the model suffers from the curse of dimensionality. The state space of the Markov decision process is too large, such

Observation days	Computation time EM	3SC
10	219.357	42.611
20	650.314	34.159
50	2869.260	39.311
100	9446.129	34.075

Table 3.3: Computation times (in seconds) for estimating parameters using EM and 3SC for $C = 20$, $T = 100$, $n = 10$, and different number of observation days.

Observation days	Computation time
10	50.47
20	55.39
50	64.52
100	81.22
200	112.98

Table 3.4: Computation time for 3SC estimation method for different number of observation days, using $C = 200$, $T = 1000$, $n = 10$.

that solving the exact problem takes too much time. In Table 3.5 the computation times of the solution methods are presented for a small example, where $C = 10$, $T = 20$, $l = 1$, and the number of products n varies from 2 to 10. The solution method where no cancellations are considered is used as a benchmark, denoted by TvR. Evaluating instances with $n > 6$ using the exact solution method causes system to be out of memory and the exact solution can therefore not be computed.

This example shows that when the number of products increases the computation time increases drastically, while the computation times for the heuristic stay more or less constant. In Table 3.6 below the computation times of the heuristics are presented for a larger instance, where $C = 200$, $T = 1000$, and $n = 10$. Even in this case the computation time is negligible.

n	Computation Time Exact	IOC	LCR	TvR
2	0.229	0.295	0.013	0.009
3	1.020	0.292	0.013	0.008
4	4.293	0.327	0.011	0.008
5	15.041	0.309	0.011	0.008
6	49.082	0.309	0.010	0.009
7	-	0.318	0.010	0.011
8	-	0.318	0.011	0.009
9	-	0.327	0.012	0.010
10	-	0.333	0.012	0.012

Table 3.5: Computation times heuristics vs. exact solution method, using $C = 10$, $T = 20$, load factor $l = 1$, $n = 2, \dots 10$.

3.5.3 Performance of Solution Methods

The cancellation model is only useful if the tractable heuristics outperform the base model. To this end, consider a large instance, with $C = 200$ rooms and $T = 1000$ time periods. The load factors l that are considered are $\{0.6, 0.8, 1, 1.2, 1.4\}$. In Table 3.6 the results are presented.

There are several noteworthy observations from Table 3.6. First and foremost, using the policies found by the IOC and LCR heuristics lead to solutions with *higher revenue* compared the TvR solution method (except for the case that $l = 0.6$, when LCR still

load factor	IOC	LCR	TvR	random	$\Delta_{\mathrm{IOC/TvR}}$	$\Delta_{\mathrm{LCR/TvR}}$	$\Delta_{\mathrm{LCR/random}}$
0.6	16,878	17,014	16,904	11,981	-0.15%	0.65%	42.00%
0.8	21,390	21,691	21,252	15,492	0.65%	2.06%	40.01%
1.0	25,297	25,949	24,169	18,818	4.67%	7.37%	37.90%
1.2	27,732	29,859	26,154	21,972	6.03%	14.17%	35.89%
1.4	29,386	33,225	27,848	24,969	5.52%	19.31%	33.06%
Comp. time	7.10	0.58	0.47				

Table 3.6: Performance of the heuristics over the base model, using $C = 200$, $T = 1000$, $n = 10$, and $l \in \{0.6, 0.8, 1, 1.2, 1.4\}$.

outperforms TvR, but IOC is slightly outperformed by TvR). Therefore it is reasonable to consider cancellations in the decision making process. Taking cancellation into account leads to a revenue gain of up to 20%.

Second, the IOC algorithm performs better than the TvR algorithm, but worse than the LCR algorithm. This is the motivation for choosing LCR over IOC if we work under the linear cancellation rates Assumption 3.2, but choosing IOC over TvR if we work with general probability rates under the independence of cancellations Assumption 3.1.

Third, the effects of taking cancellations into account increase when the arrival volume increases. This is intuitively clear: if the arrival rate is higher, the gap caused by a cancelled reservation will be filled quickly. If you do not take cancellations into account, the gap caused by the cancelled reservation will not be filled, which implies a loss in revenue.

3.5.4 Estimation Accuracy vs. Performance of Solution Methods

Estimation errors potentially lead to bad performance of the optimisation model. Depending on the sample size the estimation error might be small or large. In the next simulation the performance of the solution method using estimated parameters is measured. Again a large instance is considered, where $C = 200$, $T = 1000$, and $n = 10$. The load factor l is given different values, namely $\{0.6, 0.8, 1, 1.2, 1.4\}$, to consider the effect of different demand volumes. For each load factor a dataset of size 100 (arrival days) is created, from which the parameters are estimated. The estimated parameters are used by all heuristics to find the corresponding policy and revenue. In Table 3.8 the results of the estimation are presented. The true values as well as the estimated values are presented.

		Load factor				
		0.6	0.8	1	1.2	1.4
IOC	True	16,843	21,408	25,316	27,748	29,406
	Est.	17,168	20,198	24,277	26,764	30,039
LCR	True	16,996	21,685	25,964	29,905	33,186
	Est.	17,212	20,222	24,356	27,876	33,590
TvR	True	16,872	21,270	24,138	26,173	27,862
	Est.	17,162	20,080	23,151	24,758	28,432

Table 3.7: Performance of the heuristics under estimated values.

In Table 3.7 the revenues corresponding to load factor l, heuristic, and estimated or true parameter set is given. From this table several observations can be made. Firstly, the heuristics outperforms the TvR method in all cases: for all load factors, for all parameter sets, true or estimated. This fact suggests that choosing a heuristic that takes cancellations into account is preferred over the TvR model, even if the parameters are estimated and thus contain errors.

	load factor									
	0.6		0.8		1		1.2		1.4	
	True	Est.	True	Est.	True	Est.	True	Est.	True	Est.
α	0	-0.0241	0	0.3561	0	0.3096	0	0.2598	0	-0.0336
λ	0.12	0.1144	0.16	0.1666	0.2	0.2055	0.24	0.2447	0.28	0.2627
β_1	-0.005	-0.004551	-0.005	-0.004958	-0.005	-0.004717	-0.005	-0.005352	-0.005	-0.004794
β_2	-0.0015	-0.001392	-0.0015	-0.001480	-0.0015	-0.001513	-0.0015	-0.001659	-0.0015	-0.001466
γ_1	0.000432	0.0004539	0.000576	0.0005507	0.00072	0.0007590	0.000864	0.0008718	0.001008	0.0010052
γ_2	0.000192	0.0002130	0.000256	0.0002587	0.00032	0.0003509	0.000384	0.0004001	0.000448	0.0004650
γ_3	0.000336	0.0003534	0.000448	0.0004241	0.00056	0.0005843	0.000672	0.0006968	0.000784	0.0007589
γ_4	0.000168	0.0001300	0.000224	0.0002453	0.00028	0.0002947	0.000336	0.0003177	0.000392	0.0003684
γ_5	0.00012	0.0001004	0.00016	0.0001717	0.0002	0.0002049	0.00024	0.0002435	0.00028	0.0002671
γ_6	0.000216	0.0002182	0.000288	0.0002809	0.00036	0.0003480	0.000432	0.0003976	0.000504	0.0005031
γ_7	0.000096	0.0001021	0.000128	0.0001085	0.00016	0.0001625	0.000192	0.0001651	0.000224	0.0002265
γ_8	0.000048	0.00002455	0.000064	0.00005066	0.00008	0.00007288	0.000096	0.0001019	0.000112	0.0001115
γ_9	0.0000048	0.000003389	0.00000064	0.000005194	0.000008	0.000007687	0.0000096	0.00001236	0.0000112	0.000009344
γ_{10}	0.000024	0.00002257	0.0000032	0.00003663	0.00004	0.00005979	0.000048	0.00004427	0.000056	0.00005813

Table 3.8: Performance heuristics under estimated values. Variables: same as 1, except for load factor = 1. Use estimates from estimation method example.

Second, the model using the estimated parameters might outperform the model with the true parameters, or it might not. It seems that in this instance the extreme load factors ($l = 0.6$ and 1.4) lead to a better performance of the model under estimated parameters for all heuristics, while with the other load factors the model under the true parameters performs better.

3.6 Discussion

In this chapter we analysed the customer choice cancellation model in Section 3.2, including simplified versions that result from Assumptions 3.1, 3.2, and 3.3. Moreover, we introduced several solution methods in Section 3.3 and estimation methods in Appendix 3.B. In this section we briefly describe some guidelines for selecting a model, a solution method, and an estimation method. Moreover, we briefly discuss an extension of the cancellation model to a more realistic choice-based network cancellation model.

3.6.1 Selecting an Appropriate Model, Solution Method, and Estimation Method

In practice one needs to make a choice which model to use before deciding about the solution method and estimation method. To this end, a balance must be established for the trade-off between a good model, a good estimation method, and a good solution method. If little restriction is put on the cancellation probabilities, heuristics need to be used unless the problem is small. Subsequently, complex cancellation rate functions need to be estimated. On the other hand, under the linear and equal cancellation rates Assumption 3.3 the problem can be solved exactly in an efficient way, if the action space small. However, the assumption that cancellation probabilities are equal for all products is rather strong and not realistic in many cases.

For different model instances different solution methods and estimation methods may be appropriate. There are three dimensions that have to be considered when choosing a solution method: the model choice, the size of the state space, and the size of the action space. When choosing an estimation method we need to consider the model choice and in particular the model choice for the purchase probabilities and cancellation rates.

The efficient three-step estimation Algorithm 3.4 can only be applied to a model that satisfies Assumption 3.2 or 3.3 and the purchase probabilities are modelled by the multinomial logit model. In all other cases the expectation-maximisation algorithm described in Appendix 3.B can be used.

If none of the Assumptions 3.1, 3.2, or 3.3 is satisfied, we can only use the exact solution method to solve the problem. The computation time of this method may be prohibitively long if the state space or the action space is large.

Under the independence of cancellations Assumption 3.1 only (the reformulation of) the exact method described by Equation (3.3) or the Independence of Cancellations heuristic (IOC) described by Algorithm 3.1 are appropriate. If the state space or the action space is large we use the IOC heuristic and if the state space and the action space are small we can use the exact solution method.

Under the linear cancellation rates Assumption 3.2 we can apply the exact method, the IOC heuristic, and the Linear Cancellation Rates heuristic (LCR). However, numerical results show that LCR outperforms IOC under the linear cancellation rates assumption, so we will not consider IOC here. If both the action space and the state space are small, we will use the exact method. If however either the action space is large or the state space is large we prefer LCR because the exact method is intractable. The action space can be small when the state space is large if n is small but C is large.

Under the linear and equal cancellation rates assumption the problem can be solved exactly using Equation 3.4. However, if the action space is too large we have to use the LCR heuristic.

3.6.2 Future Directions: Choice-Based Network Cancellation Model

In practice it is common that fare products consist of multiple night stays/multiple flight legs, so it is worth looking at network RM problems. The theory described in this chapter is not sufficient to solve these complex problems. However, we provide an outline of a way to approach this problem. The choice-based network model by Liu & van Ryzin (2008) could be extended to include cancellations. The network consists of m legs and n fare products, where a fare product is now not only a combination of a price r_j and conditions, but also a combination of legs consumed by the product, e.g., several nights for a hotel stay.

As in the single-resource case we can write the dynamic programming formulation to solve the problem. In the network base model (without cancellations) we have to keep track of the available resources, which is at most a vector of size m. Liu & van Ryzin (2008) argue that this problem is intractable and provide a choice-based linear programming heuristic (CDLP) to solve the problem in reasonable time. Now if we also consider cancellations we have to keep track of the number of reservations per product. This is a vector of size at most $2^m - 1$, so we need to consider heuristics. One heuristic is to use state space reduction techniques similar to methods described in this chapter. Instead of keeping track of purchases of all products we can keep track of the availability of each resource, which reduces the state-space to an m-dimensional vector, as in the network base model. A heuristic similar to CDLP could be used to solve the remaining problem. However, cancellations contribute to the complexity of modelling and solving such heuristic. More research is necessary to investigate an efficient method to model and solve the choice-based network cancellation model. This topic would enrich the choice-modelling literature, but this is beyond the scope of this chapter.

3.7 Concluding Remarks

Cancellations have a big impact on revenue. Our model incorporates cancellations in the customer choice behaviour setting of RM. Policies that are optimal in a setting where cancellations are not considered can lead to a revenue loss of up to 20% compared to an optimal policy in our setting where cancellations are considered.

The exact solution can not be evaluated for all practical purposes because of the curse of dimensionality. To overcome this problem we introduced three heuristics: (1) the Independence of Cancellations heuristic 3.1, which reduces the state space and can be applied to models with general cancellation rates; (2) the Linear Cancellation Rates heuristic 3.2, which also reduces the state space and is appropriate under the linear cancellations Assumption 3.2; and (3) heuristic 3.3, which is an extension of the Linear Cancellation Rates algorithm and can be used for a hotel with several room types.

Two key results in this study ensure that the customer choice cancellation model can be an effective method for practitioners. First, the Linear Cancellation Rates heuristic from Section 3.3.4 is fast and gives a solution close to optimal. Second, the parameters of the model can be estimated in a consistent and fast way using Algorithm 3.4.

The results of this study can function as a foundation for several topics of further research. First of all, the effects of cancellation conditions and overbooking policies in combination with customer choice behaviour can now be studied with our customer choice cancellation model. Second, our customer choice cancellation model can be extended to include group bookings and multiple night stays/multiple flight legs, which are also very common in practice. Finally, cancellations can be embedded in other RM models to improve performance.

3.A List of Proofs

Proof of Proposition 3.1.

First we rewrite Equation (3.3) by taking all terms that do not depend on S out of the maximisation part. This leads to the following equation:

$$\begin{aligned}\tilde{W}_t(y) = &\max_{S\subset N} \{\lambda\big(R(S,y,t) - Q(S)\Delta\tilde{W}_{t-1}(y)\big)\} \\ &+ \gamma y \Delta\tilde{W}_{t-1}(y-1) + \tilde{W}_{t-1}(y),\end{aligned} \tag{3.17}$$

with $\Delta\tilde{W}_{t-1}(y) := \tilde{W}_{t-1}(y) - \tilde{W}_{t-1}(y+1)$.

Let $S^* \subset N$ be an inefficient set, i.e., there exist $\alpha(S) \geq 0$ for all $S \subset (N)$ with $\sum_{S\subset N} \alpha(S) = 1$ and

$$Q(S^*) \geq \sum_{S\subset N} \alpha(S)Q(S), \qquad R(S^*,y,t) < \sum_{S\subset N} \alpha(S)R(S,y,t).$$

Then we have that

$$\begin{aligned}&\lambda\left(R(S^*,y,t) - Q(S^*)\Delta\tilde{W}_{t-1}(y)\right) \\ &< \lambda\left(\sum_{S\subset N} \alpha(S)R(S,y,t) - \sum_{S\subset N} \alpha(S)Q(S)\Delta\tilde{W}_{t-1}(y)\right) \\ &= \sum_{S\subset N} \alpha(S)\lambda\big(R(S,y,t) - Q(S)\Delta\tilde{W}_{t-1}(y)\big).\end{aligned}$$

Since $\alpha(S)$ defines a probability, there is at least one $S \subset N$ such that

$$R(S^*,y,t) - Q(S^*)\Delta\tilde{W}_{t-1}(y) < R(S,y,t) - Q(S)\Delta\tilde{W}_{t-1}(y).$$

Hence S^* is not optimal. □

Proof of Lemma 3.4. The probability that $x_j \in \{0,\dots,y\}$ in state (y,t) is equal to 1. Therefore the probability $p_t^y(x_j = k)$ that $x_j = k$ for some $k \in \{0,\dots,y\}$ if we are in state (y,t) is equal to the probability $\hat{p}_t^y(x_j = k)$

that $x_j = k$ and we are in state y if we are at time t divided by the probability that we are in state y at time t. The latter is equal to

$$\sum_{k=0}^{y} \hat{p}_t^y(x_j = k),$$

so we only have to find $\hat{p}_t^y(x_j = k)$ for all $k \in \{0, \dots, y\}$.

There are three ways to get to state (y, t). First, one could come from state $(y+1, t+1)$ and a product was cancelled. The probability that we get to state $(y+1, t+1)$ is $p(y+1, t+1)$. There are two possible values for x_j at time $t+1$ that can lead to $x_j = k$ at time t, namely $x_j = k+1$ at time $t+1$ and product j is cancelled or $x_j = k$ at time $t+1$ and product $i \neq j$ is cancelled. The probability that the former event occurs is equal to

$$p(y+1, t+1)p_{t+1}^{y+1}(x_j = k+1)\gamma_j^{t+1}(k+1),$$

the probability that the latter event occurs is equal to

$$p(y+1, t+1)p_{t+1}^{y+1}(x_j = k)\sum_{i \neq j}\sum_{k'=0}^{y+1-k} p_{t+1}^{y+1}(x_i = k')\gamma_i^{t+1}(k').$$

The second possible state that one could come from to end up in state (y, t) is state $(y, t+1)$ and neither a product was cancelled nor a product was bought. The probability that we get to state $(y, t+1)$ is equal to $p(y, t+1)$. The only possibility to get to $x_j = k$ in state (y, t) if nothing happens is that $x_j = k$ in state $(y, t+1)$. The probability for this event is $p_{t+1}^y(x_j = k)$. The probability that nothing happens in state $(y, t+1)$ under the condition that $x_j = k$ is then equal to

$$1 - \lambda \sum_{j \in S} P_j(S) - \sum_{i \neq j}\sum_{k'=0}^{y-k} p_{t+1}^y(x_i = k')\gamma_i^{t+1}(k') - \gamma_j^{t+1}(k).$$

The third possible state that we could come from to end up in state (y, t) is state $(y-1, t+1)$ and a product was bought. The probability that we get to state $(y-1, t+1)$ is equal to $p(y-1, t+1)$. There are two possible values for x_j at time $t+1$ that can lead to $x_j = k$ at time t, namely $x_j = k-1$ at time $t+1$ and product j is bought or $x_j = k$ at time $t+1$ and product $i \neq j$ is bought. The probability that the former event occurs is equal to

$$p_{t+1}^{y-1}(x_j = k-1)\lambda P_j(S),$$

the probability that the latter event occurs is equal to

$$p_{t+1}^{y-1}(x_j = k)\lambda \sum_{i \in S \setminus \{j\}} P(S, i).$$

Combining all events shows that $\hat{p}_t^y(x_j = k)$ is equal to the stated formula for all k. $\square$

Proof of Lemma 3.5.

$\bar{x}_j(y,t)$ is the expected number of reservations for product j in state (y,t), which is equal to

$$\bar{x}_j(y,t) = \sum_{k=0}^{y} p_t^y(x_j = k)k.$$

$\gamma(y,t)$ is the probability that a product is cancelled in state (y,t). This is a combination of the cancellation probabilities of individual products. The cancellation probability of product j is equal to

$$\sum_{k=0}^{y} p_t^y(x_j = k)\gamma_j^t(k).$$

Therefore $\gamma(y,t)$ is equal to

$$\gamma(y,t) = \sum_{j=1}^{n}\sum_{k=0}^{y} p_t^y(x_j = k)\gamma_j^t(k).$$

Moreover, under the linear and equal cancellation rates Assumption 3.3 we have that

$$\gamma(y,t) = \sum_{j=1}^{n}\gamma_j^t\sum_{k=0}^{y} p_t^y(x_j = k)k = \sum_{j=1}^{n}\gamma_j^t\bar{x}_j(y,t).$$

$c(y,t)$ are the costs of a cancellation in state (y,t). If $x_j = k$, then the expected costs are $r_j(t)$. The expected costs for product j are thus

$$\sum_{k=0}^{y} p_t^y(x_j = k)\gamma_j^t(k)r_j(t).$$

Hence the expected costs for all products are

$$\sum_{j=1}^{n}\sum_{k=0}^{y} p_t^y(x_j = k)\gamma_j^t(k)r_j(t).$$

If we condition on the event that a cancellation occurs in state (y,t) we see that the expected costs for a cancellation are equal to

$$\begin{aligned} c(y,t) &= \frac{\sum_{j=1}^{n}\sum_{k=0}^{y} p_t^y(x_j = k)\gamma_j^t(k)r_j(t)}{\sum_{j=1}^{n}\sum_{k=0}^{y} p_t^y(x_j = k)\gamma_j^t(k)} \\ &= \frac{\sum_{j=1}^{n}\sum_{k=0}^{y} p_t^y(x_j = k)\gamma_j^t(k)r_j(t)}{\gamma(y,t)}. \end{aligned}$$

Moreover, under the linear and equal cancellation rates Assumption 3.3 we have that

$$c(y,t) = \frac{\sum_{j=1}^{n} r_j(t)\gamma_j^t\bar{x}_j(y,t)}{\gamma(y,t)}.$$

Finally, $p(y,t)$ is the probability that we are in state (y,t) at time t. First, we could come from state $(y+1,t+1)$ and we have a cancellation. The probability that this occurs is equal to

$$\begin{aligned} p(y+1,t+1)\sum_{j=1}^{n}\sum_{k=0}^{y+1} p_t^{y+1}(x_j=k)\gamma_j^t(k) \\ = p(y+1,t+1)\gamma(y+1,t+1) \end{aligned} \tag{3.18}$$

Second, we could come from state $(y,t+1)$ and there was no event. This occurs with probability

$$\left(1-\lambda\sum_{j\in S}P_j(S)-\gamma(y,t+1)\right)p(y,t+1). \tag{3.19}$$

Third, we could come from state $(y-1,t+1)$ and we have a reservation. This occurs with probability

$$\lambda\sum_{j\in S}P_j(S)p(y-1,t+1). \tag{3.20}$$

Combining Equations (3.18), (3.19), and (3.20) this gives

$$\begin{aligned} p(y,t) = p(y+1,t+1)\gamma(y+1,t+1) + \lambda\sum_{j\in S}P_j(S)p(y-1,t+1) \\ + \left(1-\lambda\sum_{j\in S}P_j(S)-\gamma(y,t+1)\right)p(y,t+1). \end{aligned}$$

□

3.B Estimating Parameters

3.B.1 Multinomial Logit Model

The *multinomial logit model* (MNL) is a method to describe a discrete probability distribution, given a set of independent variables. It is commonly used to predict the outcome of statistical classification problems, such as political elections (Dow & Endersby, 2004, see for example).

Assume that customers can choose among n products and the no-purchase alternative. Each product j provides a utility U_j for each customer, i.e., the value of product j to the customer. The no-purchase alternative provides a utility U_0. Each customer chooses alternative $j \in N \cup \{0\}$ such that his utility is maximised. The MNL assumes that the utilities are random variables of the form

$$U_j = u_j + \xi_j,$$

where u_j is the expected value of U_j and ξ_j is a random draw from the Gumbel distribution[2] with location equal to Euler's constant and scale equal to -1, such that the expected

[2]The cumulative distribution function of the *Gumbel distribution* is given by

$$F(x|\mu,\beta) = e^{-e^{-(x-\mu)/\beta}},$$

where μ is called the *location* and β is called the *scale* parameter. The expected value is equal to $\mu+\gamma\beta$, where $\gamma \approx 0.5772$ is Euler's constant.

value is zero. Under these assumptions it can be shown (see Ben-Akiva & Lerman, 1985) that the choice probabilities are equal to

$$P_j(S) = \frac{e^{u_j}}{\sum_{i \in S} e^{u_i} + e^{u_0}}.$$

The mean utility u_j of a product or the no-purchase option depends on $k \in \mathbb{N}$ different attributes. The values of the attributes are stored in the k-dimensional vector Z_j. The corresponding weights of each attribute are stored in the k-dimensional vector β. Hence the mean utility u_j for product j is given by

$$u_j = \beta^\top Z_j.$$

3.B.2 Complete Knowledge

Suppose that we have complete knowledge on arrivals, purchases and no-purchases, offer sets, and cancellations. Then we can estimate the parameters in the following way. Define $a_\lambda(t)$, $\{a_j(t)\}_{j \in N}$, and $a(t)$ by

$$\begin{aligned}
a_\lambda(t) &= \begin{cases} 1 & \text{arrival at time } t, \\ 0 & \text{otherwise,} \end{cases} \\
a_j(t) &= \begin{cases} 1 & \text{product } j \text{ cancelled at time } t, \\ 0 & \text{otherwise,} \end{cases} \\
a(t) &= \begin{cases} 1 & \text{no arrival or cancellation at time } t, \\ 0 & \text{otherwise,} \end{cases} \\
&= (1 - a_{\lambda(t)}) \prod_{j \in N} \big(1 - a_j(t)\big).
\end{aligned}$$

The parameters defined above are indicators to specify which event happened in time period t. Note that exactly one indicator is equal to 1 at time t and that this is in agreement with the assumption that only one event happens in each time period. In order to ensure a unique solution we set $\alpha = 0$ for convenience. Let $j(t)$ be the product that is purchased at time t under the condition that a customer arrived and $j(t) = 0$ if we observed an arrival but no purchase. Then $\lambda P_{tj(t)}(0, \beta, S_t, Z_t)$ is the probability that $j(t)$ is purchased at time t, or the probability that a customer arrived and purchased nothing if $j(t) = 0$. If product j was cancelled at time t then the probability that this occurs is $\gamma_j(x_j)$. If no event took place this occurs with probability $1 - \lambda - \sum_{j \in N} \gamma_j(x_j)$. Let D be the set of time periods. Then the likelihood function L is given by

$$L(\beta, \gamma, \lambda | a, x, S_t, Z_t) =$$

$$\prod_{t \in D} \left[\lambda P_{tj(t)}(0, \beta, S_t, Z_t)\right]^{a_\lambda(t)} \prod_{j \in N} \gamma_j(x_j)^{a_j(t)} \cdot \left[1 - \lambda - \sum_{j \in N} \gamma_j(x_j)\right]^{a(t)}.$$

The log-likelihood function LL is then

$$LL(\beta, \gamma, \lambda | a, x, S_t, Z_t) = \log(L) = \sum_{t \in D} \Bigg[a_\lambda(t) \log\big(P_{tj(t)}(\alpha, \beta, S_t, Z_t)\big) + a_\lambda(t) \log(\lambda)$$

$$+ \sum_{j \in N} a_j(t) \log\big(\gamma_j(x_j)\big) + a(t) \log\left(1 - \lambda - \sum_{j \in N} \gamma_j(x_j)\right) \Bigg].$$

Note that LL is separable in β and (γ, λ). Maximising the separated log-likelihood functions with respect to β and (γ, λ) gives estimates $\hat{\beta}$ and $(\hat{\gamma}, \hat{\lambda})$, respectively.

3.B.3 Incomplete Knowledge: Expectation Maximisation Algorithm

Suppose we can not observe the event that a customer arrives but does not make a purchase. Then we do not know the values of $a_\lambda(t)$ and $a(t)$. One way to overcome this problem is to apply the well known expectation maximisation algorithm (see Dempster et alii, 1977). Define $\mathcal{P}$ as the set of time periods where a purchase is observed and $\mathcal{C}_j$ the set of time periods where a cancellation for product j is observed. Then we can rewrite LL to

$$
\begin{aligned}
&LL(\beta,\gamma,\lambda|x,S_t,Z_t) = \\
&\sum_{t\in\mathcal{P}} \left[\log(\lambda) + \log\left(P_{tj(t)}(0,\beta,S_t,Z_t)\right)\right] + \sum_{j\in N}\sum_{t\in\mathcal{C}_j} \log\left(\gamma_j(x_j)\right) \\
&+ \sum_{t\notin\mathcal{P},t\notin\mathcal{C}_j} \log\left(1-\lambda-\sum_{j\in N}\gamma_j(x_j)\right) \\
&+ a_\lambda(t)\left[\log(\lambda) + \log\left(P_{t0}(0,\beta,S_t,Z_t)\right) - \log\left(1-\lambda-\sum_{j\in N}\gamma_j(x_j)\right)\right].
\end{aligned}
$$

We know the values for $a_\lambda(t)$ for $t\in\mathcal{P}$ and $t\in\mathcal{C}_j$ for all j. However, we do not know the values for $a_\lambda(t)$ for $t\notin P$ and $t\notin\mathcal{C}_j$ for all j. If we have estimates $\hat\beta$, $\hat\gamma$, and $\hat\lambda$ we can use Bayes theorem to calculate $\hat a_\lambda(t)$ (similar to Talluri & van Ryzin, 2004a):

$$
\begin{aligned}
\hat a_\lambda(t) &= \mathbb{E}[a_\lambda(t)|t\notin\mathcal{P},t\notin\mathcal{C}_j,\hat\beta,\hat\lambda,\hat\gamma] \\
&= \mathbb{P}[a_\lambda(t)=1|t\notin\mathcal{P},t\notin\mathcal{C}_j,\hat\beta,\hat\lambda,\hat\gamma] \\
&= \frac{\mathbb{P}[t\notin\mathcal{P},t\notin\mathcal{C}_j|a_\lambda(t)=1,\hat\lambda,\hat\beta,\hat\gamma]\cdot\mathbb{P}[a_\lambda(t)=1|\hat\lambda,\hat\beta,\hat\gamma]}{P[t\notin\mathcal{P},t\notin\mathcal{C}_j|\hat\lambda,\hat\beta,\hat\gamma]} \\
&= \frac{P_{t0}(\hat 0,\hat\beta,S_t,Z_t)\cdot\hat\lambda}{1-\hat\lambda P_{t*}(\hat 0,\hat\beta,S_t,Z)-\sum_{j\in N}\hat\gamma_j(x_j)}.
\end{aligned}
\tag{3.21}
$$

With $\hat a_\lambda$ we can approximate $a(t)$ by

$$
\begin{aligned}
\hat a(t) &= (1-\hat a_\lambda(t))\prod_{j\in N}\left(1-a_j(t)\right) \\
&= \begin{cases} 0 & \text{if a product was cancelled in time period } t, \\ 1-\hat a_\lambda(t) & \text{otherwise.}\end{cases}
\end{aligned}
$$

Under $\hat a_\lambda(t)$ the expected log-likelihood function is

$$
\begin{aligned}
&\mathbb{E}[LL|\hat\beta,\hat\lambda,\hat\gamma] = \\
&\sum_{t\in\mathcal{P}} \left[\log(\lambda) + \log\left(P_{tj(t)}(\beta,Z,S_t)\right)\right] + \sum_{j\in N}\sum_{t\in\mathcal{C}_j} \log\left(\gamma_j(x_j)\right) \\
&+ \sum_{t\notin\mathcal{P},t\notin\mathcal{C}_j} \log\left(1-\lambda-\sum_{j\in N}\gamma_j(x_j)\right) \\
&+ \hat a_\lambda(t)\left[\log(\lambda) + \log\left(P_{t0}(\beta,Z,S_t)\right) - \log\left(1-\lambda-\sum_{j\in N}\gamma_j(x_j)\right)\right].
\end{aligned}
\tag{3.22}
$$

Maximising this function with respect to (β, λ, γ) gives new estimates $(\hat{\beta}, \hat{\lambda}, \hat{\gamma})$. These values in turn can be used to find a new estimate $\hat{a}_\lambda(t)$, etc. The algorithm is summarised in Algorithm 3.5.

Algorithm 3.5. Expectation-Maximisation

1. Start with initial parameters $(\hat{\lambda}, \hat{\beta}, \hat{\gamma})$.
2. *Expectation* Estimate $\hat{a}_\lambda(t)$ using Equation (3.21).
3. *Maximisation* Estimate $(\hat{\lambda}, \hat{\beta}, \hat{\gamma})$ using $\hat{a}_\lambda(t)$ and Equation (3.22).
4. Go to step 2 unless we hit a stopping criterion.

CHAPTER 4

Single-Leg Choice-Based Revenue Management: a Robust Optimisation Approach

In practice, uncertainty in demand forecasts arises due to estimation errors and the stochastic nature of demand (see also Chapter 2). This chapter studies a *robust optimisation* approach to deal with the uncertainty of demand in the single-leg choice-based revenue management (RM) model from Chapter 3. To this extend, the probability vectors of general choice models are modelled by ϕ-divergence uncertainty sets. An important yet surprising result is that the robust solution method performs relatively better for smaller inventory than for larger inventory. Moreover, the robust solution method performs 2.5-3.25% better than the nominal solution when knowledge on cancellations is lacking.

4.1 Introduction

A popular trend in revenue management models is to capture the behaviour of customers that choose between different available products. The renowned paper by Talluri & van Ryzin (2004a) combined revenue management models with customer choice models. Many contributions to the body of literature on customer choice models have been made, e.g., better solution methods (Strauss & Talluri, 2012), network models that take into account multiple night stays/multiple flight legs (Liu & van Ryzin, 2008), and in Chapter 3 the model was extended to include cancellations, to which the model of this study is closely related. The solution methods used in these models often assume that the parameters of the model are known. However, in practice the true parameter values are unknown and have to be estimated from data. Estimating the parameters of a customer choice model requires more data than only sales data, which is not always available, so an estimation error is not unlikely. Newman et alii (2012) provide hotel data and describe the complexity of acquiring a proper data-set. Estimation errors lead to uncertainty in the parameters and very likely a misspecified model. Optimising a misspecified model leads to potentially suboptimal policies and revenue loss. A relatively recent field of optimisation that takes into account uncertainty in the optimisation procedure is robust optimisation. In this field, the values of the parameters are assumed to lie in an uncertainty set, rather than to be known exactly. Robust optimisation methods provide solutions where the worst-case

This chapter is based on Sierag & van der Mei (2016).

scenario is optimised, providing a trade-off between risk and average reward. See Ben-Tal et alii (2009) for an introduction of theories on robust optimization.

This chapter studies a robust optimisation approach to single-leg choice-based revenue management. In this model a single-leg revenue management problem is modelled as a Markov decision process and solved using dynamic programming. This study provides a general robust formulation of this model. In each step of the dynamic program a small maximin problem has to be solved. The minimisation problem can be formulated as a linear program. The uncertain parameters are probabilities, of which the uncertainty set is modelled using a ϕ-divergence measure. Tractable robust counterparts are presented for this problem. The methodologies that are used are based on Nilim & El Ghaoui (2005), who provide a robust formulation for general dynamic programming formulations, and the recent paper by Ben-Tal et alii (2013), who provide a novel formulation for robust counterparts for probabilities. The main contribution of this study is a tractable robust formulation for general choice models. Numerical results discussed in Section 4.4 below show that the robust solution method outperforms the nominal solution in many cases when using estimated parameters. Moreover, the robust solution method gives a relatively higher improvement in revenue for smaller hotels than for larger hotels. Related to this chapter is Rusmevichientong & Topaloglu (2012), where a robust formulation is provided of the multinomial logit model. In contrast to their paper, this chapter focuses on the *estimated probabilities* instead of the parameters of the chosen model. This way the robust solution method can be applied to any choice model, or any estimate of the choice probabilities.

Other revenue management problems have been solved using robust optimisation. Several static and dynamic single-leg revenue management problems that do not take into account customer choice behaviour have been studied. Ball & Queyranne (2009) provide a robust solution method for the single-leg revenue management problem that does not require demand information. Robust results for various policy classes are provided. Birbil et alii (2009) provide robust optimization methods for one static and one dynamic single-leg revenue management problem. Lan et alii (2008) provide a robust solution method for the single-leg revenue management problem under independent demand model when limited information is available on demand. Network revenue management problems have also been studied in robust optimisation context. Lai & Ng (2005) provide a stochastic programming solution to a network model for hotels, without using choice models. Perakis & Roels (2010) describe two robust solution methods for solving the network revenue management problem that does not take into account customer choice behaviour. One solution method solves the maximin problem and the other solves the minimax regret problem. As mentioned before, Rusmevichientong & Topaloglu (2012) provide a study of robust optimization applied to the assortment problem under the multinomial logit model. Farias et alii (2013) provide a robust non-parametric estimation method and a study on selecting the right choice model using sales data.

The remainder of this chapter is organised as follows. In Section 4.2 the single-leg customer choice model is described. The extension to cancellations is also presented. The model is reformulated in an equivalent formulation that is more convenient for this study. In Section 4.3 the robust counterpart of the nominal model is presented. First the formulation of the uncertainty sets is given. Then the robust dynamic program is formulated. The remainder of this section provides tractable reformulations for several ϕ-divergence measures. In Section 4.4 numerical results are presented to validate the model. In Section 4.5 some concluding remarks are given. See Appendix 4.A for references to robust linear optimisation. Moreover, the theory of ϕ-divergence uncertainty sets is explained, which is used in the theories of this chapter.

4.2 Model Description

In this section the set-up of the problem is presented in hotel-context. However, the general description applies to other areas, such as seats on a flight or tickets for a theatre performance. Consider a hotel with C identical rooms that wants to sell them in T time units, 0 being the arrival time. This arrival time is typically an arrival day, where the rooms for that night are offered for sale T days in advance. Overbooking is allowed up to $C_{\max}$ rooms. Each room can be sold using a *fare product* j, which is a combination of a room with a price r_j and certain conditions, such as the cancellation policy. Assume there is a finite number of fare products $N = \{1, \ldots, n\}$. At each moment in time the hotel manager decides which *offer set* $S \subset N$ of fare products to offer. Potential customers arrive according to a Poisson process with rate $\lambda > 0$. These customers show interest in the hotel, but their final decision is based on the offer set S displayed. The customer either buys one of the fare products $j \in S$, with probability $P_j(S)$, or leave and buy nothing at all, with probability $P_0(S)$. Customers are allowed to cancel their reservation, according to the cancellation policy. Assume that the cancellations of reservations happen independent from each other. This assumption is intuitively clear: clients are assumed to arrive independently, so their decision to cancel their reservation is not dependent on the decisions of other clients. If there are x_j reservations for fare product j, then the time to cancellations follows an exponential distribution with rate $\gamma_j x_j$, $\gamma_j \in \mathbb{R}^+$. See Figure 4.1 for an illustration of the model.

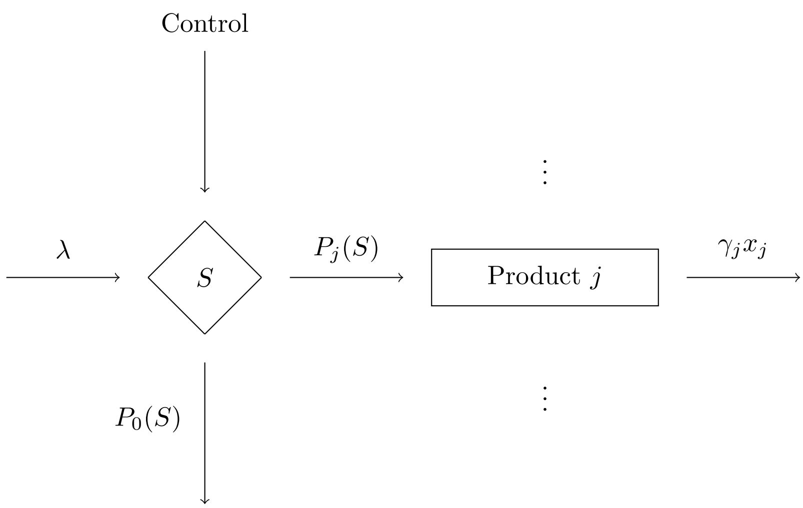

Figure 4.1: Illustration of the customer choice cancellation model, see also Chapter 3. Per arrival day the arrival process is Poisson distributed with parameter λ. The manager controls the offer set S. Under this offer set an arriving customer buys product $j \in S$ with probability $P_j(S)$, and with probability $P_0(S)$ the customer buys nothing. Finally, cancellations of product j follow an exponential distribution with parameter γ_j.

The problem for the manager is to decide which offer set $S \subset N$ to offer to maximise expected revenue. To solve this problem, the continuous time Markov decision process is discretised and set up as a dynamic programming formulation. The state space of this formulation keeps track of the number of reservations per product, and hence it grows exponentially with the number of products. In Chapter 3 a solution method was proposed that approximates the optimal solution by collapsing the state space (Algorithm 3.2). This is the only available tractable solution method that shows promising results. Therefore, this algorithm, which is called *nominal method* henceforth, is used as a benchmark and functions as a base for the robust solution method that is to follow. Another approximation described in Chapter 3 considers the model without cancellations. It is based on Talluri & van Ryzin (2004a), and is denoted by *Talluri and van Ryzin* (TvR) method henceforth. The TvR method is also included in the numerical studies.

The collapsed dynamic program only keeps track of the *total* number of reservations, leading to the state space $\{0, 1, \ldots, C_{\max}\}$. Time is divided into T time periods, where the length of the intervals is such that the probability that more than one event occurs is very small. Therefore it is assumed that only one event occurs per time period, where an event is either an arrival, a cancellation, or neither arrival nor cancellation. Denote with λ the probability that a customer arrives in a time period; and γy the probability that a product is cancelled in state y, $\gamma \in \mathbb{R}$. The single cancellation rate can be estimated from the individual cancellation rates, for example by the average, as is suggested in Chapter 3. The probability that no purchase occurs in a time period equals the sum of the probability that neither an arrival and nor a cancellation occurs, and the probability that an arrival occurs but the arriving customer makes no purchase. This probability is equal to

$$(1 - \lambda - \gamma y) + \lambda P_0(S) = 1 - \lambda \sum_{j \in S} P_j(S) - \gamma y.$$

In each time period the decision needs to be made which set S to offer. Note that time has to be scaled such that $\lambda + \gamma C_{\max} \leq 1$, otherwise the probabilities are not well defined.

Let $V_t(y)$ be the maximal expected revenue from time t to the arrival day in state y. Define $\Delta H_j(t)$ by

$$\Delta H_j(t) := \begin{cases} \gamma_j c_j(t) + (1 - \gamma_j)\Delta H_j(t-1) & \text{if } t > 1, \\ 0 & \text{if } t = 1, \end{cases}$$

for all $j \in N$. The Bellman equation corresponding to the discretised Markov chain is

$$V_t(y) = \begin{cases} \max\limits_{S \subset N} \Big\{ \lambda \sum\limits_{j \in S} P_j(S)\big[r_j - \Delta H_j(t) + V_{t-1}(y+1)\big] & \text{if } 0 < y < C_{\max}, \\ \quad +\gamma y V_{t-1}(y-1) & \\ \quad + \Big(1 - \lambda \sum\limits_{j \in S} P_j(S) - \gamma y\Big) V_{t-1}(y) \Big\} & \\ \gamma C_{\max} V_{t-1}(C_{\max} - 1) & \text{if } y = C_{\max}, \\ \quad + (1 - \gamma C_{\max})\, V_{t-1}(C_{\max}) & \\ \max\limits_{S \subset N} \Big\{ \lambda \sum\limits_{j \in S} P_j(S)\big[r_j - \Delta H_j(t) + V_{t-1}(1)\big] & \text{if } y = 0. \\ \quad + \Big(1 - \lambda \sum\limits_{j \in S} P_j(S)\Big) V_{t-1}(0) \Big\} & \end{cases} \tag{4.1}$$

Below a reformulation of Equation (4.1) is given, for convenience later on in the chapter. Note that λ, $P_j(S)$ and γ are probabilities. Define $\tilde{P}_j(S) := \lambda P_j(S)$ for all $j \in N$

($P_j(S) := 0$ if $j \notin S$), $\tilde{P}_{n+1}(S) := \gamma y$, and $\tilde{P}_0(S) = 1 - \sum_{j=1}^{n+1} \tilde{P}_j(S)$. Also define $\tilde{r}_j := r_j - \Delta H_j(t)$. Then Equation (4.1) can be rewritten as

$$V_t(y) = \max_{S \subset N} \left\{ \sum_{j \in S} \tilde{P}_j(S) \left[\tilde{r}_j + V_{t-1}(y+1) \right] + \tilde{P}_{n+1}(S) V_{t-1}(y-1) + \tilde{P}_0(S) V_{t-1}(y) \right\}. \tag{4.2}$$

4.3 Robust Reformulation

Finding accurate estimates for the problem at hand is challenging. Commonly used maximum likelihood estimation methods can be found in Newman et alii (2014), Sierag et alii (2015), or Talluri & van Ryzin (2004a). Estimating parameters from data often lead to estimation errors, as is the case with the mentioned methods. Incorrect estimates may lead to suboptimal decisions, and therefore suboptimal revenue. One way to deal with estimation errors in the optimisation process is *robust optimisation*. In this field, the values of the parameters are not known exactly, but are assumed to lie in an *uncertainty set*. The worst case scenario is optimised under the uncertainty set. The goal is to improve performance by using the robust solution method rather than the nominal solution method. In Appendix 4.A a brief introduction to robust linear optimisation is given.

Recently Ben-Tal et alii (2013) have provided tractable robust counterpart formulations for uncertainty sets that are based on ϕ-divergence measures. These uncertainty sets are used to model uncertainty in probabilities. The dynamic program in Equation (4.2) uses probabilities consisting of estimated parameters to find an optimal solution. The derivation of tractable robust counterparts for (4.2) under ϕ-divergence uncertainty sets is described below.

4.3.1 Reformulation

The uncertainty in the adjusted purchase probabilities $\tilde{P}_j(S)$ is assumed to deviate from the nominal value $\bar{p}$ according to a *ϕ-divergence measure*. Consider a function $\phi\colon \mathbb{R} \to \mathbb{R}$ that is twice differentiable on $[0, \infty)$, convex for $t \geq 0$, $\phi(1) = 0$, $0\phi(1/0) := a \lim_{t\to\infty} \phi(t)/t$ for $a > 0$, and $0\phi(0/0) := 0$. The ϕ-divergence measure $I_\phi(p, q)$ between two vectors p and q, $p, q \in \mathbb{R}^n$, is defined as

$$I_\phi(p, q) := \sum_{i=1}^{n} q_i \phi(p_i/q_i).$$

A ϕ-divergence measure is a metric for the distance between two (probability) vectors according to the divergence function ϕ. A popular choice for ϕ is the Cressie-Read divergence:

$$\phi^\theta(t) = \frac{1 - \theta + \theta t - t^\theta}{\theta(1-\theta)}, \qquad \theta \neq 0, 1, \quad t \geq 0.$$

The ϕ-divergence measure $I_\phi(p, q)$ is then given by

$$I_\phi(p, q) = \frac{1}{\theta(1-\theta)} \left(1 - \sum_i p_i^\theta q_i^{1-\theta} \right).$$

In Section 4.3.2 several other ϕ-divergence measures are discussed.

In Ben-Tal et alii (2013) it is shown that the robust counterpart of a linear constraint with ϕ-divergence uncertainty can be written in terms of the conjugate $\phi^*: \mathbb{R} \to \mathbb{R} \cup \infty$ of ϕ, which is defined as follows:

$$\phi^*(s) := \sup_{t \geq 0}\{st - \phi(t)\}. \tag{4.3}$$

Let X be a one-dimensional random variable with finite support $\{1, \ldots, m\}$. Denote $p = (p_1, \ldots, p_m)$ as the probability vector of X, such that $p_i = \mathbb{P}(X = i)$. Let $\hat{p}_0$ be the maximum likelihood estimator of p. From Ben-Tal et alii (2013) the uncertainty region for p is given by

$$\{p \in \mathbb{R}^m \,| I_\phi(p, \hat{p}_0) \leq \rho\},$$

with

$$\rho := \frac{\phi''(1)}{N}\chi^2_{m,1-\alpha},$$

with N the sample size, α the confidence level, and $\chi^2_{m,1-\alpha}$ is the $1-\alpha$ percentile of the χ^2_m-distribution. Then the uncertainty set Z_S for the purchase probabilities $\tilde{P}_j(S)$ is given by

$$Z_S = \left\{p \in \mathbb{R}^{n+2} \,| p \geq 0, Cp \leq d, I_\phi(p, \bar{p}) \leq \rho\right\},$$

with

$$C_{\bullet j} = \begin{cases} (1, -1) & \text{if } j \in S \cup \{0, n+1\}, \\ (0, 0) & \text{otherwise}, \end{cases}$$
$$d = (1, -1).$$

Nilim & El Ghaoui (2005) provide a robust formulation of dynamic programming with uncertainty in the probabilities. From their analysis it follows that the nominal dynamic program (4.1) under uncertainty sets Z_S can be solved using the following recursive formula:

$$V_t(y) = \max_{S \subset N} \Phi(S), \tag{4.4}$$

with

$$\Phi(S) = \min\left\{\sum_{j \in S} p_j[\tilde{r}_j + V_{t-1}(y+1)] + p_{n+1}V_{t-1}(y-1) + p_0 V_{t-1}(y) \,\middle|\, p \in Z_S\right\}. \tag{4.5}$$

The uncertainty problems $\Phi(S)$ can be solved for each S independently. When $\Phi(S)$ is known for all S, then optimisation problem (4.4) becomes a maximisation problem over a finite set of integers. The challenge is to evaluate $\Phi(S)$. For this purpose, it is notationally convenient to move the formula in the objective of (4.5) to the constraints, by setting the objective to minimise to $t \in \mathbb{R}$ and add the constraint

$$\sum_{j \in S} p_j[\tilde{r}_j + V_{t-1}(y+1)] + p_{n+1}V_{t-1}(y-1) + p_0 V_{t-1}(y) \leq t.$$

This leads to the following equivalent optimisation problem:

$$\min \left\{ t \in \mathbb{R} \;\middle|\; \begin{array}{r} \sum_{j \in S} p_j [\tilde{r}_j + V_{t-1}(y+1)] + p_{n+1} V_{t-1}(y-1) + p_0 V_{t-1}(y) - t \le 0, \\ \forall p \in Z_S \end{array} \right\}. \tag{4.6}$$

Define $x \in \mathbb{R}^2$, $a \in \mathbb{R}^2$, and $B \in \mathbb{R}^{2\times(n+2)}$ as follows:

$$\begin{aligned} x &= (t, x_0), \\ a &= (-1, 0), \\ B &= \begin{pmatrix} 0 & 0 & \cdots & 0 & 0 \\ V_{t-1}(y) & \tilde{r}_1 + V_{t-1}(y+1) & \cdots & \tilde{r}_n + V_{t-1}(y+1) & V_{t-1}(y-1) \end{pmatrix}. \end{aligned}$$

Then the constraint can be rewritten to

$$\begin{cases} 0 & \ge (a + Bp)^\top x, \;\; \forall p \in Z_S, \\ x_0 & = 1. \end{cases}$$

The following property from Ben-Tal et alii (2013) can now be applied.

Property 4.1. *Theorem 4.1 in Ben-Tal et alii (2013)*

Consider the linear constraint

$$(a + Bp)^\top x \le b, \qquad\qquad p \in Z, \tag{4.7}$$

where $x \in \mathbb{R}^n$ is the vector to be optimised, $a \in \mathbb{R}^n$, $B \in \mathbb{R}^{n\times m}$ and $b \in \mathbb{R}^n$ are given parameters, $p \in \mathbb{R}^m$ is the uncertain parameter, and

$$Z = \{p \in \mathbb{R}^m \,|p \ge 0, Cp \le d, I_\phi(p,q) \le \rho\},$$

is the uncertainty region of p with $q \in \mathbb{R}^m_+$, $\rho > 0$, $d \in \mathbb{R}^k$, and $C \in \mathbb{R}^{k\times m}$. Then a vector $x \in \mathbb{R}^n$ satisfies (4.7) if and only if there exist $\eta \in \mathbb{R}^k$ and $\xi \in \mathbb{R}$ such that (x, η, ξ) satisfies

$$\begin{cases} a^\top x + d^\top \eta + \rho\xi + \xi \sum_i q_i \phi^* \left(\dfrac{B_i^\top x - C_i^\top \eta}{\xi} \right) \le b, \\ \eta \ge 0, \xi \ge 0, \end{cases}$$

where B_i and C_i are the i-th columns of B and C, respectively, and ϕ^* is the conjugate function of ϕ.

The robust counterpart of (4.6) is then given by

$$\begin{aligned} a^\top x + d^\top \eta + \rho_P \xi + \xi \sum_{i=0}^{n} \bar{p}_i \phi^* \left(\frac{B_i^\top x - C_i^\top \eta}{\xi} \right) &\le 0, \\ \eta &\ge 0, \\ \xi &\ge 0. \end{aligned} \tag{4.8}$$

The robust counterparts for the boundaries (last two equations of (4.1)) are also given by Equation (4.8), where the parameters are given by

$$\begin{aligned} x &= (t, x_0), \\ a &= (-1, 0), \\ B &= \begin{pmatrix} 0 & 0 \\ V_{t-1}(y) & V_{t-1}(y-1) \end{pmatrix}, \\ p&(p_0, p_{n+1}), \end{aligned}$$

and

$$\begin{aligned} x &= (t, x_0), \\ a &= (-1, 0), \\ B &= \begin{pmatrix} 0 & 0 & \cdots & 0 \\ V_{t-1}(y) & \tilde{r}_1 + V_{t-1}(y+1) & \cdots & \tilde{r}_n + V_{t-1}(y+1) \end{pmatrix}, \\ p &= (p_0, p_1, \ldots, p_n), \end{aligned}$$

respectively.

A tractable reformulation of the robust counterpart under the Cressie-Read divergence is given by (Ben-Tal et alii, 2013, see appendix of)

$$\begin{cases} a^\top x + \eta_1 - \eta_2 + \xi\rho_P + \frac{\xi}{\theta}\sum_i \bar{p}_i \left((\frac{y_i}{\xi})^{\frac{\theta}{\theta-1}} - 1 \right) \leq 0 & \\ y_i = \xi - (1-\theta)(B_i^\top x - \eta_1 + \eta_2), & i = 1, \ldots, m, \\ \eta \geq 0, \xi \geq 0. & \end{cases}$$

This problem can solved using conic quadratic programming (CQP). Tractable robust counterparts for other popular ϕ-divergence measures can be found in Ben-Tal et alii (2013).

The parameters of the uncertainty sets need to be estimated from data. For some parameters the estimation procedure is trivial or already described. The parameter $\bar{p}$ follows directly from the maximum likelihood estimates, for example from the estimation procedure described in Section 3.4, and Newman et alii (2014) or Talluri & van Ryzin (2004a). The parameter ρ_P is defined by Equation (4.3).

4.3.2 Discussion of Selected ϕ-Divergence Measures

In this section, tractable reformulations of (4.8) are provided for popular choices of ϕ (Ben-Tal et alii, 2013). Each measure is appropriate in a different situation, and a motivation is given when to select which measure.

– **Kullback-Leibler** The Kullback-Leibler measure is

$$I_{\phi_{kl}}(p, \bar{p}) = \sum_{i=1}^{m} p_i \log(p_i/\bar{p}_i).$$

This asymmetric measure punishes upward deviation of p_i from $\bar{p}_i$ and rewards downward deviation. This measure is appropriate when the parameters are more likely to be smaller than the estimated parameters.

– **Burg Entropy** The Burg entropy measure is given by

$$I_{\phi_b}(p, \bar{p}) = \sum_{i=1}^{m} \bar{p}_i \log(p_i/\bar{p}_i).$$

Also for this asymmetric measure it holds that it punishes upward deviation of p_i from $\bar{p}_i$ and rewards downward deviation. The amplification of the punishment or reward is constant and dependent on the estimation of the nominal vector $\bar{p}$. The unfortunate property of this measure is that any value of p_i is feasible: p_i can be approximately close to zero, since it has a negative contribution to $I_{\phi_b}(p, \bar{p})$; and p_i can be greater than $\bar{p}_i$, up to 1, as long as p_j is small enough, for some $j \neq i$.

– **J-divergence** The J-divergence measure leads to a ϕ-divergence of

$$I_{\phi_j}(p, \bar{p}) = \sum_{i=1}^{m} (p_i - \bar{p}_i) \log(p_i/\bar{p}_i).$$

This asymmetric measure punishes both upward and downward deviation of p_i from $\bar{p}_i$. The measure is similar to Kullback-Leibler and Burg Entropy, yet the factor $p_i - \bar{p}_i$ ensures that all deviation is punished and none rewarded. Larger deviation is punished heavier than small deviations.

– **χ^2-distance** The ϕ-divergence of the χ^2-distance measure is given by

$$I_{\phi_c}(p, \bar{p}) = \sum_{i=1}^{m} \frac{(p_i - \bar{p}_i)^2}{p_i}.$$

This measure punishes larger deviations more heavily than smaller deviations. However, upward deviation is punished less than downward deviation, because of the $1/p_i$ term. The punishment is relative to the quantity of the vector p.

– **Modified χ^2-distance** The ϕ-divergence of the modified χ^2-distance is given by

$$I_{\phi_{mc}}(p, \bar{p}) = \sum_{i=1}^{m} \frac{(p_i - \bar{p}_i)^2}{\bar{p}_i}.$$

This measure also punishes larger deviations more heavily than smaller deviations. However, in this case upward deviation is punished equally strongly as downward deviation because of the $\bar{p}_i$. The punishment is relative to the quantity of the nominal value $\bar{p}_i$.

– **Hellinger distance** The ϕ-divergence of the Hellinger distance is given by

$$I_{\phi_h}(p, \bar{p}) = \sum_{i=1}^{m} (\sqrt{p_i} - \sqrt{\bar{p}_i})^2.$$

This is the l_2-norm applied to the vector $(\sqrt{p_1}, \ldots, \sqrt{p_n})$, which is a unit vector under the l_2-norm. This measure is more sensitive to changes in smaller probabilities than larger probabilities.

– **χ-divergence of order $\theta > 1$** The ϕ-divergence of the χ-distance of order $\theta > 1$ is given by

$$I_{\phi_{ca}}(p, \bar{p}) = \sum_{i=1}^{m} \bar{p}_i |1 - p_i/\bar{p}_i|^{\theta}.$$

This measure punishes larger deviations more heavily than small deviations. How severe the deviations are punished depends on the parameter θ. Larger θ punish deviations more than smaller θ.

– **Variation distance** The ϕ-divergence of the variation distance is given by

$$I_{\phi_v}(p, \bar{p}) = \sum_{i=1}^{m} |p_i - \bar{p}_i|.$$

This is the l_1 norm. Larger deviation is punished linearly more than smaller deviations.

– **Cressie and Read** The ϕ-divergence for the Cressie and Read measure with parameter θ is given by

$$I_{\phi_h}(p, \bar{p}) = \frac{1}{\theta(1-\theta)}(1 - \sum_{i=1}^{m} p_i^{\theta} \bar{p}_i^{1-\theta}).$$

This measure gives larger punishment to larger deviations if θ is large. For small θ the effect of large deviations is less severe.

4.4 Numerical Results

In this section numerical results are provided to validate the robust solution method for the single-leg revenue management model with cancellations and overbooking. The benchmark policy is given by the nominal solution method described in Section 4.2.

The model parameters that are used are based on Simulation Example 2 in Talluri & van Ryzin (2004a) and the example used in the numerical results in Section 3.5 of this dissertation. These parameters are realistic and based on observations in practice. Let $n = 10$ be the number of products sold with corresponding price vector

$$r = (240, 220, 190, 160, 120, 112, 96, 80, 74, 70).$$

Demand, cancellation rates, and purchase probabilities are independent from the time period t. To allow the study of the effect of high volume in demand and low volume in demand the *load factor* l is introduced. In our studies the values $l \in \{0.6, 0.8, 1, 1.2, 1.4\}$ are used. Demand λ per time unit is then defined as

$$\lambda(C, l, T) = \frac{Cl}{T}.$$

A low value of l implies a low demand relative to the capacity C and a high value of l implies a high demand relative to the capacity C.

Purchase probabilities are modelled by the multinomial logit (MNL) model. The only attribute that is considered is price, but it is assumed that there exist high price-sensitive and low price-sensitive customers which have different parameters $\beta_H = -0.005$ and $\beta_L = -0.0015$, respectively (as in Talluri & van Ryzin (2004a), and Section 3.5 above). The no-purchase parameter is set to $\alpha = 0$ such that the MNL model is the same as in Talluri & van Ryzin (2004a) and Section 3.5. Overbooking is allowed up to 20% of the total capacity C (this is more for computational reasons: the policies found will almost surely not open any product categories for sale as the overbooking is close to

20%). Cancellation rates are assumed linear and the parameters $\gamma \in \mathbb{R}^n$ depend on l and T in the following way:

$$\gamma = (18/25, 8/25, 14/25, 7/25, 1/5, 9/25, 4/25, 2/25, 1/125, 1/25)\frac{l}{T}.$$

To evaluate the performance of both robust and nominal policies for a particular model instance, 1000 data sets are simulated. From each dataset the parameters of the model are estimated, which are used by both solution methods. Simulation is used to estimate the revenue corresponding to a policy. The estimation errors are between 0.1% and 0.4%.

Hotel size

First the performances of the robust solution method for different hotel sizes are compared. Moreover, different ϕ-divergences are applied. The booking horizon is $T = 1000$ time units and the sample size is 100 arrival days. Table 4.1 shows the results.

	Hotel size									
	10		20		50		100		200	
ϕ-divergence metric	Rev	%	Rev	%	Rev	%	Rev	%	Rev	%
Nominal	1765		3549		8908		16711		29776	
TvR	1749	-0.91	3470	-2.23	8235	-7.56	15209	-8.99	26355	-11.49
Kullback-Leibler	1791	1.47	3620	2.00	8948	0.45	16793	0.49	29781	0.02
Burg Entropy	1791	1.47	3621	2.03	8945	0.42	16794	0.50	29779	0.01
J-divergence	1791	1.47	3621	2.03	8945	0.42	16792	0.48	29778	0.01
χ^2-distance	1792	1.53	3619	1.97	8946	0.43	16793	0.49	29784	0.03
Modified χ^2-distance	1791	1.47	3621	2.03	8946	0.43	16796	0.51	29780	0.01
Hellinger distance	1791	1.47	3620	2.00	8945	0.42	16794	0.50	29777	0.00
χ-div. of order $\theta = 0.5$	1792	1.53	3621	2.03	8946	0.43	16793	0.49	29781	0.02
Variation distance	1791	1.47	3619	1.97	8947	0.44	16795	0.50	29783	0.02
Cressie-Read	1791	1.47	3620	2.00	8946	0.43	16790	0.47	29782	0.02

Table 4.1: Performance of nominal and robust solution methods for various hotel sizes.

The results show that the robust solution methods perform better compared to the nominal solution method for small hotels. For larger hotels the difference is smaller. This might be the case because either the nominal solution performs better for larger hotels, or the robust solution methods perform worse for larger hotels. These results suggest that smaller hotels would relatively benefit more from using a robust solution method than larger hotels. This is an important observation, since most of the hotels are small and medium enterprises and have a relatively small number of rooms.

The results show no significant difference in performance between the robust solution methods. Using any robust solution method seems better than using none. For a fixed hotel size the divergences do not show much difference. Any difference might even be a caused by the small estimation errors. Also, whenever a divergence shows better performance for one hotel size, another divergence performs better for another hotel size. For example, for $C = 10$ the χ^2-divergence performs best, but for $C = 20$ the divergence performs worst.

Note that the TvR approximation does not perform well, and neither does it in the examples that follow. Also, no significant difference in performance of ϕ-divergences was found in the next examples. For convenience only the results of the Cressie-Read divergence are compared against the nominal method, and the TvR method is left out.

Load factor

The load factor influences the performance of the estimation method (see Section 3.5 above). Therefore the robust solution method might perform different under different load factors. The parameters are set to $C = 20$ and $C = 100$ rooms, $T = 100$ time periods, and the load factors that are used are $\{0.6, 0.8, 1, 1.2, 1.4\}$. The results are presented in Table 4.2.

	$C = 20$			$C = 100$		
Load factor	Nominal	Robust	%	Nominal	Robust	%
0.6	1934	1936	0.09	9107	9106	-0.02
0.8	2539	2553	0.56	11849	11853	0.03
1.0	3082	3125	1.39	14388	14443	0.38
1.2	3546	3619	2.03	16705	16781	0.46
1.4	3966	4039	1.82	18911	18982	0.37

Table 4.2: Performance of nominal and robust solution methods for various load factors.

The results show that under a small load factor the performance of the nominal solution and the robust solution are similar, while for larger load factors the robust solution outperforms the nominal solution. In practice, this means that the robust solution is preferred in popular areas or in high season, since it leads to significantly higher profits, while outside high season or at less popular locations the robust solution does not lead to higher profits.

Sample size

Now the behaviour of the solution methods according to different sample sizes is considered. One sample consists of the data collected for one arrival day. The uncertainty set parameter ρ is evaluated accordingly. The size of the hotel is $C = 20$ rooms and the booking horizon is $T = 100$ time periods.

	$C = 20$			$C = 100$		
Sample size	Nominal	Robust	%	Nominal	Robust	%
1	3434	3473	1.12	15510	15536	0.17
2	3286	3323	1.12	15472	15512	0.26
5	3482	3541	1.69	16442	16496	0.33
10	3513	3576	1.77	16646	16705	0.35
20	3542	3612	1.98	16665	16729	0.38
50	3545	3616	2.01	16690	16760	0.42
100	3546	3618	2.03	16708	16777	0.41

Table 4.3: Performance of nominal and robust solution methods for a hotels of size $C = 20$ and $C = 100$ and different sample sizes.

In Table 4.3 the performance of the nominal and robust solution methods is provided. The results show that the average performance of the robust solution method outperforms the nominal solution method by $1-2\%$ for $C = 20$, increasing as the sample size increases. For $C = 100$ the performance is lower, up to 0.42%.

Unknown cancellation behaviour

The robust solution method can be beneficial when the cancellation behaviour is not known. When the cancellation parameter is equal to zero, the robust solution method can be used using this input parameter. The robust solution method then takes into account an uncertainty around the zero vector, which may lead to an improved result. This has been applied as follows: the hotel size is $C = 20$ or $C = 100$ rooms, the booking horizon is $T = 100$ time periods, and the cancellation rate parameter is set to 0. The results are given in Table 4.4.

	$C = 20$			$C = 100$		
Sample size	Nominal	Robust	%	Nominal	Robust	%
1	3458	3544	2.48	15176	15651	3.13
2	3408	3484	2.22	15048	15500	3.00
5	3440	3524	2.44	15176	15660	3.19
10	3430	3515	2.47	15148	15619	3.10
20	3464	3559	2.74	15185	15678	3.25
50	3466	3559	2.70	15198	15699	3.30
100	3468	3561	2.68	15205	15709	3.31

Table 4.4: Performance of nominal and robust solution methods for various hotel sizes.

The results show that the average performance is increased by about 2.5% for $C = 20$, while for $C = 100$ the performance even increased by about 3.25%. This strongly suggests that the robust solution method can successfully be applied when no information about cancellation behaviour is known. This is an important observation since knowledge on cancellations is not always available in practice.

Number of products

The performance of the solution methods may be influenced by the number of products n. The next example shows the performance under a different number of products. The hotel size is $C = 20$ or $C = 100$ rooms, the load factor is $l = 1.2$, and the booking horizon is $T = 100$ time periods. The results are presented in Table 4.5.

	$C = 20$			$C = 100$		
Number of products	Nominal	Robust	%	Nominal	Robust	%
2	3472	3529	1.65	15943	15942	-0.01
3	3558	3628	1.97	16615	16684	0.41
4	3549	3620	2.01	16708	16776	0.40
5	3547	3621	2.08	16706	16778	0.43
6	3548	3619	1.99	16709	16779	0.42
7	3548	3620	2.03	16707	16777	0.42
8	3550	3622	2.04	16709	16778	0.41
9	3546	3621	2.12	16709	16781	0.43
10	3546	3619	2.03	16707	16774	0.40

Table 4.5: Performance of nominal and robust solution methods for different number of products.

The results suggest that the number of products has only a minor impact on the performance of the robust solution method. The robust solution method performs on average 1.65-2.12% better for $C = 20$, and 0.4% for $C = 100$. The results show highly robust revenues against different number of products, which is preferred.

4.5 Concluding Remarks

Estimating parameters of a revenue management model that takes into account customer choice behaviour is likely to lead to estimation errors. One method to improve performance is using robust optimisation to take into account the estimation error. In this chapter a robust solution method is described for the customer choice cancellation model introduced in Chapter 3.

The dynamic programming formulation is converted into a robust dynamic program. The uncertain parameters are modelled using ϕ-divergence uncertainty sets. In the dynamic program small minimisation problems containing the uncertain parameters have to be solved. Novel robust optimisation methodologies from Ben-Tal et alii (2013) lead to tractable solutions for several ϕ-divergence measures.

In the numerical studies it is shown that the robust solution method performs better for smaller hotels than for larger hotels. This can be explained by the results of Section 2.4, where it is shown that demand uncertainty is relatively higher for smaller hotels compared to larger hotels. As a consequence, the robust solution method is preferred for small and medium enterprise hotels, to which most hotels belong. Also, promising results are shown when cancellation behaviour is not known. The robust solution method outperforms the nominal solution method by up to 2% when using estimated parameters. The performance also shows to be robust under different number of products. The fact that the robust solution method is tractable and provides good results make it attractive for application in practice.

The results in this chapter can serve as a foundation for several topics further research. First, the methodology of this chapter can be extended to network revenue management models that use choice models. The effects of parameter estimation errors in these network models can be studied and solved using robust optimisation. Solving network problems is not straightforward, so an intense study is necessary. Second, an extensive study of the performance of different ϕ-divergence measures under different circumstances can provide promising insights. Each ϕ-divergence measure includes and excludes different distances from the estimated parameters, which may give insight to which measure is preferred in which situation.

4.A Introduction to Robust Linear Optimisation and ϕ-Divergence Uncertainty Sets

Consider a problem that can be modelled as a linear program:

$$\min\left\{c^\top x \,|Ax \leq b\right\},$$

with $c \in \mathbb{R}^n$ the cost vector, $x \in \mathbb{R}^n$ the vector of decision variables, $A \in \mathbb{R}^{m\times n}$ and $b \in \mathbb{R}^m$. The robust counterpart is given by

$$\min\left\{c^\top x \,|Ax \leq b, A \in Z\right\},$$

with $c \in \mathbb{R}^n$ the cost vector, $x \in \mathbb{R}^n$ the vector of decision variables, $A \in \mathbb{R}^{m\times n}$ the uncertain parameters in uncertainty set $Z \subset \mathbb{R}^{m\times n}$ and $b \in \mathbb{R}^m$ a constant vector. Ben-Tal et alii (2009) show that without loss of generality it may be assumed that only uncertainty in A exists, and not in the objective c or the right-hand-side b. Furthermore, Ben-Tal et alii (2009) show that the uncertainty can be approached constraint-wise, which is of the form

$$(a + B\zeta)^\top x \leq b, \quad \forall \zeta \in Z,$$

with the nominal value $a \in \mathbb{R}^n$ constant, $B \in \mathbb{R}^{n\times m}$ constant, $b \in \mathbb{R}$ constant, $\zeta \in \mathbb{R}^m$ uncertain, and Z the uncertainty region for ζ.

Tractable formulations of the robust counterpart for several standard uncertainty regions are provided in Ben-Tal et alii (2009).

CHAPTER 5

Pricing-Based Revenue Management for Flexible Products on a Network

This chapter proposes and analyses a pricing-based revenue management (RM) model that allows flexible products on a network, with a non-trivial extension to group reservations. Under stochastic demand the problem can be solved using dynamic programming, though it suffers from the curse of dimensionality. The solution under deterministic demand gives an upper bound on the stochastic problem, and serves as a basis for two heuristics, which are asymptotically optimal. Numerical studies, based on a problem instance from practice, show that the heuristics perform well, even under uncertainty in demand. Moreover, neglecting flexible products can lead to substantial revenue loss.

5.1 Introduction

Traditionally in RM, products are a combination of resources, e.g., a stay at a hotel (product) for two nights (resources). A recent development in constructing products is the concept of *flexible products* (Gallego et alii, 2004; Gönsch et alii, 2014; Mang et alii, 2012; Petrick et alii, 2010, 2012). In line with Gallego & Phillips (2004), a *flexible product* is a set of alternatives offered by a firm such that the customer is assigned to one of the alternatives at a later point in time. Offering flexible products, alongside traditional *specific products*, can achieve better price discrimination and potentially increase revenue (Gallego et alii, 2004; Gallego & Phillips, 2004; Mang et alii, 2012). For example, an airline serving an origin-destination pair with two flights a day could offer a flexible product on those two flights at a discounted price, alongside the specific products per flight. The airline will allocate customers that purchased flexible products to one of the flights shortly before departure in a balanced way, such that capacity is used more efficiently and new client segments can be reached without complete cannibalisation. Therefore, flexible products benefit both clients and company: more clients are able to purchase their desired products, and the company profits from better capacity utilisation and higher revenues.

Current literature on flexible products focusses on quantity-based RM (Gallego et alii, 2004; Gönsch et alii, 2014; Mang et alii, 2012; Petrick et alii, 2010, 2012), where resource

This chapter is based on Sierag (2016b).

capacities are optimally allocated over different classes of demand. This is typical for industries like hotels, airlines, and car rental companies, where it is relatively easy to remove products from the shelve in order to prevent sales and protect inventory for more profitable future sales. However, in many other industries, like retail, digital and TV advertising, and fast moving consumer goods (FMCG), it is not easy or common to remove products from the shelve. In this case, demand is matched with capacity through *dynamic pricing* rather than quantity-based RM. The mentioned industries are especially suitable for flexible products, since products can easily be replaced by one another. In particular, the digital and TV advertising industries, with 2014 global market volumes of US \$146.6 and \$189.4 billion, respectively, and growing (McKinsey&Company, 2015), sell their impressions in the form of flexible products. In order to use quantity-based methodologies and optimisation methods to solve dynamic pricing problems, one has to increase the number of products to approximate the continuity of the price range to the point that the problem unfortunately becomes highly intractable. Therefore, the quantity-based methods on flexible products in current literature are not sufficient for solving such dynamic pricing RM problems and an analysis of a *pricing-based* approach to flexible products is essential.

This study proposes and analyses a *pricing-based* RM model with capacity constraints that allows flexible products on a network. In addition, the model allows for group bookings (or bulk purchases), where a single customer is allowed to purchase more than one product at a time. Buying more than one product at once is a common phenomenon in retailing, online advertising, and FMCG. Due to the fact that flexible products are allocated to specific products at some point in time, the modelling of group reservations of flexible products differs from traditional group reservations for specific products. A stochastic model with multiple time periods is proposed that entails these characteristics, where demand is Poisson driven and time dependent. A dynamic programming formulation is given to solve the problem, though it suffers from the curse of dimensionality. The deterministic variant of the problem is tractable and provides an upper bound on the stochastic problem. Also, the deterministic pricing strategy serves as a basis for two asymptotically optimal heuristics.

This study builds upon two main research areas. The first area is pricing-based RM. Of particular interest to this study is the excellent study on network RM pricing by Gallego & van Ryzin (1997). The authors propose two heuristics to solve the intractable pricing model, on which the heuristics of this study are based. Improved solution methods that deal with the intractability of network models can be found in Maglaras & Meissner (2006), who introduced an action-space reduction algorithm, and, more recently, Zhang & Lu (2013), where a resource decomposition approach is considered. See Bitran & Caldentey (2003) and Soon (2011) for an overview of relevant pricing literature. The second research area studies flexible products. The concept of flexible products was introduced in the seminal paper by Gallego et alii (2004), who argue in favour of its applicability in many industries and analyse a model with one flexible product on two specific products. Subsequently Gallego et alii (2004) proposed and analysed a quantity-based network RM model with flexible products, together with an extension to choice-based demand. More recently, Petrick et alii (2010) provide an overview of different quantity-based RM methods; Petrick et alii (2012) study the effect of flexible products on revenue when demand is uncertain due to forecast errors; and Gönsch et alii (2014) provide an improved DLP approach to the flexible network problem. The empirical study by Mang et alii (2012) shows by means of a large field study of a low-cost airline that flexible products can increase profit by 5%. Related to flexible products is the study of probabilistic goods by Fay & Xie (2008), where capacity is allocated to the product immediately after purchase, rather than at a later moment in time.

The remainder of this chapter is organised as follows. First, in Section 5.2, the pricing-based network RM model under stochastic and deterministic demand is introduced. Also, the extension to group bookings is presented, as well as an upper bound on the stochastic problem. Second, Section 5.3 proposes two asymptotically optimal heuristics based on the deterministic problem to solve the intractable stochastic problem. Third, Section 5.4 provides numerical results the effects of flexible products, the performance of the heuristics, and the effects of uncertainty in demand on revenue. Finally, Section 5.5 concludes the chapter.

5.2 Model Description

In this section both a deterministic and a stochastic model for pricing flexible products on a network are introduced. The curse of dimensionality prevents the problem to be solved exactly. The solution to the deterministic problem is an upper bound on the stochastic solution, and is used in the heuristics presented in Section 5.3. Finally, an extension to a non-trivial model for group reservations is introduced.

5.2.1 Stochastic Model

Consider a firm that has m types of perishable *resources* available with capacity vector $C \in \mathbb{N}^m$, which can only be used in T time units. The firms sells the resources by offering n *specific products* $N = \{1, \dots, n\}$ and f *flexible products* $F = \{1, \dots, f\}$, which consume one or more resources. The resource consumption for the specific products is described by the incidence matrix $A = (a_{ij}) \in \mathbb{R}^{m \times n}$, where $a_{ij} = 1$ if product j consumes resource i. A flexible product $k \in F$ consists of a set of alternatives $F_k \subset N$ of specific products. At time T each customer that purchased a flexible product k is assigned to one of the $f_k = |F_k|$ alternatives $j \in F_k$.

The booking horizon is divided into T time periods in which products are offered. Demand for products in time period t is Poisson distributed with parameter $\lambda_j(p,t)$ for specific products $j \in N$ and parameter $\gamma_k(p,t)$ for flexible products $k \in F$, where $p \in \mathbb{R}_+^{n+f}$ is the price vector of the products. The demand functions $\lambda(p,t)$ and $\gamma(p,t)$ satisfy the following *regularity conditions*:

1. The demand function $(\lambda(p,t), \gamma(p,t))$ has an inverse $p(\lambda, \gamma, t)$. Hence the demand can and will function as the decision variables.
2. The *revenue* or *reward* at time t defined by $r(\lambda, \gamma, t) = (\lambda, \gamma)^\top p(\lambda, \gamma, t)$ is continuous, bounded, and concave.
3. The following asymptotic results hold for all finite (λ, γ) and for all $j \in N$ and $k \in F$:
$$\lim_{(\lambda,\gamma):\, \lambda_j \to 0} \lambda_j p_j(\lambda, \gamma, t) = 0,$$
$$\lim_{(\lambda,\gamma):\, \gamma_k \to 0} \gamma_k p_k(\lambda, \gamma, t) = 0.$$
4. The following ensures the revenue is bounded:
$$\sup_{j,t}[\arg \max_{(\lambda,\gamma):\, \lambda_j(t)} \lambda_j(t) p_j(\lambda, \gamma, t)] < \infty,$$
$$\sup_{k,t}[\arg \max_{(\lambda,\gamma):\, \gamma_k(t)} \gamma_k(t) p_k(\lambda, \gamma, t)] < \infty.$$

5. There exist *null prices* $p_j^\infty(t)$ and $p_k^\infty(t)$ such that:
   (a) if $\{p^i\}$ is a sequence for which $p_j^i \to p_j^\infty$, then $\lim_{i\to\infty} \lambda_j(p^i, t) = 0$,
   (b) if $\{p^i\}$ is a sequence for which $p_k^i \to p_k^\infty$, then $\lim_{i\to\infty} \gamma_k(p^i, t) = 0$.

A function that satisfies these conditions is called *regular*.

Let $v \in \mathbb{N}$ be the number of bookings for specific products. Let $s = C - Av$ be the capacity not yet committed. Let $y \in \mathbb{R}^f$ be the number of accepted requests for flexible products. Define (s, y) as the state of the network. Then (s, y) is feasible if and only if it satisfies the following system of linear (in)equalities

$$\begin{aligned} s &\geq 0, \\ \sum_{k=1}^{f} B_k z_k &\leq s, \\ y_k - \mathbf{1}^\top z_k &= 0, \quad k = 1, \ldots, f, \end{aligned} \tag{5.1}$$

where B_k is the submatrix of A containing columns corresponding to the products in F_k, and $z_k \in \mathbb{Z}^{f_k}$ represents the allocation of flexible products to specific products. Define $B = (B_1 \cdots B_f)$ and $z = (z_1, \ldots, z_f)$. Define $U \in \mathbb{R}^{f \times |z|}$ by $u_{ij} = 1$ if j corresponds to flexible product i. Let $N(t)$ and $M(t)$ be the random Poisson distributed demand under $\lambda(t)$ and $\gamma(t)$. The objective is to maximise the expected revenue:

$$\begin{aligned} E\left[\sum_{t=1}^{T} r(t, \lambda, \gamma)\right] =& E\left[\sum_{t=1}^{T} \big(N(t), M(t)\big)^\top p(t, \lambda, \gamma)\right] \\ =& E\left[\sum_{t=1}^{T}\left[\sum_{j\in N} N_j(t) p_j(t, \lambda, \gamma) + \sum_{k\in F} M_k(t) p_k(t, \lambda, \gamma)\right]\right]. \end{aligned} \tag{5.2}$$

Demand has to satisfy the feasibility constraints (5.1). The state (s, y) has to be feasible in each time period, but the flexible products can be distributed over their corresponding specific products at the end of the booking horizon. Therefore, it is not necessary to use a dummy variable vector z for each time period, but only one for the whole system. This leads to the following feasibility constraints:

$$\begin{aligned} \sum_{t=1}^{T} AN(t, \lambda(t), \gamma(t)) + Bz &\leq C, \\ \sum_{t=1}^{T} M(t, \lambda(t), \gamma(t)) - Uz &= 0. \end{aligned} \tag{5.3}$$

The stochastic optimisation problem is formulated as:

$$\max \left\{ E\left[\sum_{t=1}^{T} \big(N(t), M(t)\big)^\top p(t, \lambda, \gamma)\right] \;\middle|\; \begin{array}{ll} \sum_{t=1}^{T} AN(t) + Bz & \leq C \\ \sum_{t=1}^{T} M(t) - Uz & = 0 \\ \lambda, \gamma, z & \geq 0 \\ z \text{ integer} & \end{array} \right\}. \tag{5.4}$$

To deal with the stochasticity of demand, consider the following set-up to solve the problem. Assume that time periods are small enough that the probability that more than one purchase occurs in one time period is very small. The probability that no purchase occurs in time period t is then equal to

$$1 - \sum_{j \in N} \lambda_j(p,t) - \sum_{k \in F} \gamma_k(p,t).$$

The objective is to maximise revenue-to-go function $V_t(s,y)$. The Bellman equation equals

$$\begin{aligned} V_t(s,y) = \max_{p \in \mathbb{R}^{n+f}} \Bigg\{ & \sum_{j \in N} \lambda_j(p,t)[p_j + V_{t+1}(s - A_j, y)] \\ & + \sum_{k \in F} \gamma_k(p,t)[p_k + V_{t+1}(s, y + e_k)] \\ & + \left(1 - \sum_{j \in N} \lambda_j(p,t) - \sum_{k \in F} \gamma_k(p,t)\right) V_{t+1}(s,y) \Bigg\}. \end{aligned} \tag{5.5}$$

This problem is intractable for all practical purposes due to the curse of dimensionality. Therefore approximations or heuristics need to be considered. The next section introduces the deterministic variant of this problem, where stochastic demand is replaced by deterministic continuous demand. Subsequently, two heuristics, which both are asymptotically optimal, are presented that are based on the deterministic model.

5.2.2 Deterministic Model

Consider the stochastic model, but now assume that demand in time period t is deterministic with parameters $\lambda(p,t)$ and $\gamma(p,t)$, and relax the integrality constraints. Since the demand functions are regular, the demand variables can be used as decision variables instead of price variables, as will be the case. The objective function equals

$$E\left[\sum_{t=1}^{T} r(t,\lambda,\gamma)\right] = \sum_{t=1}^{T} r(t,\lambda,\gamma) = \sum_{t=1}^{T} \big(\lambda(t),\gamma(t)\big)^{\top} p(t,\lambda,\gamma). \tag{5.6}$$

The second constraint, regarding the distribution of flexible products over specific products, can be relaxed to an inequality without loss of generality. The optimisation problem is therefore given by

$$\max \left\{ \sum_{t=1}^{T} r(t, \lambda(t)\ \gamma(t)) \;\middle|\; \begin{array}{ll} \sum_{t=1}^{T} A\lambda(t) + Bz & \leq C \\ \sum_{t=1}^{T} \gamma(t) - Uz & \leq 0 \\ \lambda, \gamma, z & \geq 0 \end{array} \right\}. \tag{5.7}$$

To see that the second constraint can be relaxed to an inequality instead of an equality constraint, consider an optimal solution $(\lambda^*, \gamma^*, z^*)$ to (5.7) and assume that the second inequality constraint is strict, i.e., $\sum_{t=1}^{T} \gamma^*(t) < Uz^*$. Construct $\tilde{z}$ by subtracting a total amount of $(Uz^*)_k - \sum_{t=1}^{T} \gamma_k^*(t)$ from the z_i^*-s that correspond to product $k \in F$ in a distributive way, such that $\tilde{z} \geq 0$ and $\sum_i z_i^* - \tilde{z}_i = \mathbf{1}^{\top}\big(Uz^* - \sum_{t=1}^{T} \gamma^*(t)\big)$. Then

$U\tilde{z} = \sum_t \gamma^*(t)$ and $\sum_{t=1}^T A\lambda^*(t) + B\tilde{z} \le \sum_{t=1}^T A\lambda^*(t) + Bz^* \le C$, so $(\lambda^*, \gamma^*, \tilde{z})$ is an optimal solution to (5.7) where the second constraint is tight.

By the regularity assumptions on λ and γ both functions are concave, and therefore the following Karush-Kuhn-Tucker conditions are necessary and sufficient for optimality:

$$\begin{aligned}
\nabla_\lambda r_\lambda(\lambda, \gamma, t) &= A^\top \pi^*, \quad \forall t, \\
\nabla_\gamma r_\gamma(\lambda, \gamma, t) &= I\rho^*, \quad \forall t, \\
B^\top \pi^* - U^\top \rho^* &= 0, \\
(\pi^*)^\top \big(C - \sum_{t=1}^T A\lambda(t) - Bz\big) &= 0, \\
(\rho^*)^\top \big(Uz - \sum_{t=1}^T \gamma(t)\big) &= 0, \\
\pi^*, \rho^* &\ge 0.
\end{aligned} \tag{5.8}$$

Let $\{\lambda^d(t)\}_{t=1}^T$ and $\{\gamma^d(t)\}_{t=1}^T$ be an optimal solution to the deterministic problem with $\{p^d(t)\}_{t=1}^T$ the corresponding optimal prices and V^d the objective value. The objective function of (5.7) can be rewritten in a format that is useful in subsequent sections. Define

$$\begin{aligned}
\bar{p}_j &= \frac{\sum_{t=1}^T p_j^d(t)\lambda_j^d(t)}{\sum_{t=1}^T \lambda_j^d(t)}, & \alpha_j &= \sum_{t=1}^T \lambda_j^d(t), \\
\bar{p}_k &= \frac{\sum_{t=1}^T p_k^d(t)\gamma_k^d(t)}{\sum_{t=1}^T \gamma_k^d(t)}, & \alpha_k &= \sum_{t=1}^T \gamma_k^d(t).
\end{aligned}$$

Then the following holds:

$$V^d = \sum_{t=1}^T \big(\lambda^d(t), \gamma(t)\big)^\top p(t) = \sum_{j \in N} \alpha_j \bar{p}_j^d + \sum_{k \in F} \alpha_k \bar{p}_k^d. \tag{5.9}$$

Proposition 5.1 below shows that the revenue V^d of an optimal solution to the deterministic problem (5.7) is an upper bound on the objective value V^* of an optimal solution to the stochastic problem (5.4). The statement and proof are adjusted from (Liu & van Ryzin, 2008, Proposition 1) to match our model.

Proposition 5.1. Let V^* be the optimal objective value to (5.4) and let V^d be the optimal objective value to (5.7). Then $V^* \le V^d$.

Proof. Consider the Lagrange relaxation of Equation (5.4):

$$\begin{aligned}
\max_{\lambda,\gamma,z} \; E\Bigg[&\sum_{t=1}^T \big(N(t), M(t)\big)^\top p(t, \lambda, \gamma) + \pi^\top \left(C - \sum_{t=1}^T AN(t) - Bz \right) \\
&+ \rho^\top \left(Uz - \sum_{t=1}^T M(t) \right) \Bigg],
\end{aligned} \tag{5.10}$$

where $\pi \ge 0$. Consider an optimal solution (λ, γ, z) to the stochastic problem

(5.4). Then it holds that

$$C \geq \sum_{t=1}^{T} AN(t) + Bz \text{ (a.s.)},$$
$$0 \geq \sum_{t=1}^{T} M(t) - Uz \text{ (a.s.)}.$$

Since $\pi, \rho \geq 0$, the objective value V^L in the Lagrange relaxation (5.10) is larger than or equal to the objective value V^* of the stochastic problem (5.4). Therefore, the objective value $(V^L)^*$ of an optimal solution to the Lagrangian relaxation is larger than or equal to V^*:

$$V^* \leq (V^L)^*.$$

Note that Problem (5.10) is separable in t, and $E[N(t)] = \lambda(t)$ and $E[M(t)] = \gamma(t)$ hold. Therefore the problem is equivalent to maximising

$$\sum_{t=1}^{T} \left[r(t, \lambda, \gamma) - \pi^\top A\lambda(t) - \rho^\top \gamma(t)\right] + \pi^\top (C - Bz) + \rho^\top Uz.$$

The upper bound holds for all $\pi \geq 0$ and ρ, so also for π^* and ρ^*, the optimal dual prices from the deterministic problem, with corresponding deterministic solution $(\lambda^*, \gamma^*, z^*)$. Complementary slackness gives

$$(\pi^*)^\top (C - A\sum_{t=1}^{T} \lambda^*(t) + Bz^*) = 0,$$
$$Uz^* - \sum_{t=1}^{T} \gamma(t) = 0.$$

Therefore the objective value of the Lagrangian relaxation equals the deterministic objective value. Hence the deterministic solution V^d is an upper bound on the stochastic solution:

$$V^* \leq V^d.$$

□

5.2.3 Group Bookings

Many practical applications allow for *group bookings* or *bulk purchases*, where a customer purchases more than one item of the same product.[1] For example, a travel agency purchases a bulk of seats in airlines and rooms in hotels at the same time; and in online advertisement industry a publisher purchase a bulk of impressions. In the case of specific products, this can be modelled by introducing products with resource consumption lA_j, where $l \in \mathbb{N}$ is the size of the group booking. A flexible product on group products can

[1] Note that group bookings and bulk purchases do not refer to selling a stack of products to multiple customers, where each customer receives only one product.

then easily be defined by defining an appropriate set F_k. The allocation of the flexible product to specific products then proceeds by bulk: all l reservations will be allocated to one specific product j. However, it might be desirable to be able to allocate the l reservations to several *different* specific products. In order to do so, define $Q \in \mathbb{R}^{f \times f}$ such that Q_{ii} is the group size of product $i \in F$, and $Q_{ij} = 0$ if $i \neq j$. In the stochastic problem (5.4) the second constraint then changes to

$$Q \sum_{t=1}^{T} M(t) - Uz = 0, \tag{5.11}$$

and in the deterministic problem (5.7) the second constraint changes to

$$Q \sum_{t=1}^{T} \gamma(t) - Uz \leq 0. \tag{5.12}$$

The generalisation of the results in this chapter the group booking set-up is straightforward.

5.3 Solution Methods

This section proposes two heuristics to approximate the stochastic problem, called *make-to-stock* (MTS) and *make-to-order* (MTO). Both heuristics provide a *pricing and availability* strategy: the pricing strategy sets prices for products, and the availability strategy defines when products are open for sale or closed. In the case of MTS products are open until certain booking limits are reached, and with MTO a product is open until a sale is infeasible due to the capacity constraints. The intuition of both heuristics is based on the MTS and MTO heuristics by Gallego & van Ryzin (1997).

5.3.1 Heuristic One: Make to Stock (MTS)

Let $\{p^d(t)\}_{t=1}^T$ be the optimal deterministic prices, and let $\{\lambda^d(p^d(t), t)\}_{t=1}^T$ and $\{\gamma^d(p^d(t), t)\}_{t=1}^T$ be the corresponding intensities. Define the booking limits $b^\lambda \in \mathbb{N}^n$ and $b^\gamma \in \mathbb{N}^f$ by

$$b_j^\lambda := \lfloor \alpha_j \rfloor, \text{ and } b_k^\gamma := \lfloor \alpha_k \rfloor.$$

Price products according to $\{p^d(t)\}_{t=1}^T$ and offer the products until inventories are exhausted, booking limits are reached, or the deadline T is reached. Booking limits are beneficial for two reasons. First, the booking limits guarantee that the capacity constraints will be met. During the selling horizon it is therefore not necessary to keep track of the feasibility. Second, booking limits protect against *cannibalisation*, a reduction of sales of one product because another product is offered simultaneously. For example, train companies often offer a limited amount of seats at a discounted price, often under certain conditions like a no-refund policy, while simultaneously offering the full fare seats. In this case, cannibalisation occurs: valuable customers who are willing to buy at full fare but are also interested in the discounted price will buy discounted seats rather than full fare seats, which leads to a reduction in revenue. When the booking limits of the discounted seats are met, no more seats are offered at a discounted price, only the full fare. Now those valuable customers will buy seats at full fare rather than the discounted seats, such that no cannibalisation occurs. Proposition 5.2 below gives a bound on the performance of MTS. The statement and outline of the proof is taken from Gallego & van Ryzin (1997, Theorem 2). However, because our model is different, the proof has to be adjusted slightly.

Proposition 5.2. Define $u_j = \sup\left\{p_j^d(t)\,\middle|\,\lambda^d(t) > 0, \gamma^d(t) > 0, \forall t\right\}$ for all $j \in N \cup F$. Let $N(j)$ be Poisson distributed with parameter α_j, for all $j \in N \cup F$. Let V^{MTS} be the objective value of the strategy that follows from MTS evaluated in the stochastic problem (5.4). Then the following bound holds:

$$\frac{V^{MTS}}{V^*} \geq 1 - \frac{\sum\limits_{j \in N \cup F} u_j \left(\frac{\sqrt{\alpha_j}}{2} + 1\right)}{\sum\limits_{j \in N \cup F} \alpha_j \bar{p}_j^d}. \tag{5.13}$$

Proof. For each $j \in N \cup F$, denote with $\{T_j^k\}_{k=1}^{N(j)}$ the time periods where purchases occurred (one purchase per T_j^k). Let R^j be the revenue for product $j \in N \cup F$ under the deterministic solution evaluated in the stochastic problem. Then it holds that

$$\begin{aligned} E[R^j] &= E\left[\sum_{k=1}^{N(j)} p_j^d(T_j^k) - \sum_{k=b_j}^{N(j)} p_j^d(T_j^k)\right] \\ &\geq E\left[\sum_{k=1}^{N(j)} p_j^d(T_j^k)\right] - u_j E\left[(N(j) - b_j)^+\right]. \end{aligned} \tag{5.14}$$

From Wald's identity it follows that

$$E\left[\sum_{k=1}^{N(j)} p_j^d(T_j^k)\right] = E[N(j)]E[p_j^d(T_j^k)] = \alpha_j \bar{p}_j^d.$$

Similar to Gallego & van Ryzin (1997) the following inequality is used: for a random variable D with mean μ and standard deviation σ, and for any $d \in \mathbb{R}$, it holds that:

$$E[(D-d)^+] \leq \frac{\sqrt{\sigma^2 + (d-\mu)^2} - (d-\mu)}{2}. \tag{5.15}$$

From equation (5.15) it follows that

$$\begin{aligned} E\left[(N(j) - b_j)^+\right] &\leq \frac{\sqrt{\alpha_j + (b_j - \alpha_j)^2} - (b_j - \alpha_j)}{2} \\ &\leq \frac{\sqrt{\alpha_j}}{2} + |b_j - \alpha_j|. \end{aligned}$$

The fact that $b_j = \lfloor \alpha_j \rfloor$ leads to

$$E\left[(N(j) - b_j)^+\right] \leq \frac{\sqrt{\alpha_j}}{2} + 1,$$

for the second term of (5.14). Therefore the expected value of R^j satisfies

$$E[R^j] \geq \alpha_j \bar{p}_j^d - u_j(\sqrt{\alpha_j}/2 + 1).$$

Using the upper bound $V^* \leq V^d = \sum_{j \in N \cup F} \alpha_j p_j^d(t)$ completes the proof:

$$\frac{V^{MTS}}{V^*} \geq \frac{\sum\limits_{j \in N \cup F} \alpha_j \bar{p}_j^d - u_j(\frac{\sqrt{\alpha_j}}{2}+1)}{V^d} = 1 - \frac{\sum\limits_{j \in N \cup F} u_j \left(\frac{\sqrt{\alpha_j}}{2}+1\right)}{\sum\limits_{j \in N \cup F} \alpha_j \bar{p}_j^d}.$$

□

A consequence of Proposition 5.2 is that MTS is asymptotically optimal as demand and capacity go to infinity. To see this, consider the following set-up. For $h \in \mathbb{N}$ define $\lambda^h(p,t) = h\lambda(p,t)$, $\gamma^h(p,t) = h\gamma(p,t)$ and $C^h = hC$ for some fixed $\lambda(p,t)$, $\gamma(p,t)$, C, and T.

Corollary 5.1. Let V_h^{MTS} be the optimal objective value under deterministic model h. Then

$$\lim_{h \to \infty} V_h^{MTS} = V^*.$$

Proof. Let $\pi_h^* = h\pi^*$, $z_h^d = hz^d$, and $\rho_h^* = h\rho^*$ for $z^d(t)$ and shadow prices π^* and ρ^*. The claim is that $(h\lambda^d, h\gamma^d, hz)$ is an optimal solution to (5.8) for model h under prices p^d. Firstly, the revenue function in model h equals

$$\begin{aligned} R^h(\lambda^h, \gamma^h, z^h) &= \sum_{t=1}^{T} \left(\lambda^h(t), \gamma^h(t)\right)^\top p(t, \lambda, \gamma) \\ &= h \sum_{t=1}^{T} \left(\lambda(t), \gamma(t)\right)^\top p(t, \lambda, \gamma) \\ &= hR(\lambda, \gamma, z), \end{aligned}$$

hence $\bigtriangledown_{\lambda,\gamma,z} R^h(\lambda^h, \gamma^h, z^h) = h \bigtriangledown_{\lambda,\gamma,z} R(\lambda, \gamma, z)$. On the other hand

$$\begin{aligned} A^\top \pi_h^* &= hA^\top \pi^*, \\ I\rho_h^* &= hI\rho^*, \\ B^\top \pi_h^* - U^\top \rho_h^* &= h(B^\top \pi^* - U^\top \rho^*), \end{aligned}$$

hence the stationarity KKT conditions are satisfied. Next, observe that

$$\begin{aligned} C^h &= hC = h(\sum_{t=1}^{T} A\lambda(t) - Bz) = \sum_{t=1}^{T} A(h\lambda(t)) - B(hz), \\ U(hz) &= hUz = \sum_{t=1}^{T} h\gamma(t), \end{aligned}$$

hence the primal feasibility KKT conditions are satisfied. Note that

$$\alpha_j^h = h \sum_{t=1}^{T} \lambda_j^d(t) = h\alpha_j,$$

for all $j \in N \cup F$. The bound for model h now equals

$$\begin{aligned}
\frac{V_h^{MTS}}{V^*} &\geq 1 - \frac{\sum_{j \in N \cup F} u_j \left(\frac{\sqrt{h\alpha_j}}{2} + 1 \right)}{\sum_{j \in N \cup F} h\alpha_j \bar{p}_j^d} \\
&= 1 - \frac{\sum_{j \in N \cup F} u_j \left(\frac{\sqrt{\alpha_j}}{2} + 1 \right)}{\sqrt{h} \sum_{j \in N \cup F} \alpha_j \bar{p}_j^d} \\
&= 1 - O(h^{-1/2}),
\end{aligned}$$

which converges to 1 as $h \to \infty$. $\square$

5.3.2 Heuristic Two: Make to Order (MTO)

The MTO heuristic prices products according to $\{p^d(t)\}_{t=1}^T$. The difference from the MTS heuristic is that no booking limits are applied. Instead, products are offered until selling the product is infeasible. For specific products j this occurs when the state $(s - A_j, y)$ is infeasible, and for flexible products k when $(s, y+e_k)$ is infeasible. Like the MTS heuristic, MTO is asymptotically optimal and a bound for the performance of MTO is given. For convenience, introduce the following notation for resources i and products $j \in N \cup F$:

$$c_{ij} = \begin{cases} a_{ij} & \text{if } j \in N, \\ \max\{a_{ik} \,|k \in F_j\} & \text{if } j \in F. \end{cases}$$

The statement and outline of the proof are taken from Gallego & van Ryzin (1997, Theorem 3), but are adjusted to match our model.

Proposition 5.3. Let $S_i = \{j \in N \cup F \,|c_{ij} > 0\}$ and $\alpha(S_i) = \sum_{j \in S_i} \alpha_j$. Define $v_i = \max\{u_j \,|j \in S_i\}$, $\bar{v}_i = v_i \max\{c_{ij} \,|j \in S_i\}$, and $\hat{p}_i = \sum_{j \in S_i} \bar{p}_j \alpha_j / \alpha(S_i)$. Let V^{MTO} be the objective value of the strategy that follows from MTO evaluated in the stochastic problem (5.4). Then the following bound holds:

$$\frac{V^{MTO}}{V^*} \geq 1 - \frac{\sum_{i=1}^{m} \bar{v}_i \sqrt{\alpha(S_i)}}{2 \sum_{i=1}^{m} \hat{p}_i \alpha(S_i)}. \tag{5.16}$$

Proof. Consider a modified system where negative inventories are allowed, but a penalty of v_i is charged for every unit of resource i that is backlogged. In this system, if α_{ij} is strictly greater than the capacity left of i, backlogged revenue for product $j \in N \cup F$ is equal to

$$p_j(t) - v_i \leq p_j(t) - u_j \leq 0. \tag{5.17}$$

Backlogged products also consume resources, which cannot be used by non-backlogged products. Therefore $V^{mod} \leq V^{MTO}$ holds, where V^{mod} is the revenue corresponding to the modified model. In expectation this gives

$$\begin{aligned} V^{MTO} &\geq V^{mod} \\ &= \sum_{j \in N \cup F} E[\sum_{k=1}^{N(j)} p_j^d(T_j^k)] - \sum_{i=1}^{m} v^i E[(\sum_{j=1}^{n} c_{ij} N(j) - C_i)^+]. \end{aligned} \tag{5.18}$$

Like in the proof for MTS the first term equals $\sum_{j=1}^{n} \alpha_j \bar{p}_j$. For the second term, observe that if $d \geq \mu$ Equation (5.15) becomes (triangle inequality)

$$\begin{aligned} E[(D-d)^+] &\leq \frac{\sqrt{\sigma^2 + (d-\mu)^2} - (d-\mu)}{2} \\ &\leq \frac{\sigma + (d-\mu) - (d-\mu)}{2} = \frac{\sigma}{2}. \end{aligned} \tag{5.19}$$

Note that

$$E[\sum_{j \in N \cup F} c_{ij} N(j)] = \sum_{j \in N \cup F} c_{ij} \alpha_j \leq C_i.$$

Consequently,

$$E[(\sum_{j \in N \cup F} c_{ij} N(j) - C_i)^+] \leq \frac{\sigma_i}{2},$$

with

$$\begin{aligned} \sigma_i^2 = Var[\sum_{j \in N \cup F} c_{ij} N(j)] &= \sum_{j \in N \cup F} c_{ij}^2 Var(N(j)) = \sum_{j=1}^{n} c_{ij}^2 \alpha_j \\ &\leq \max\left\{ c_{ij}^2 \middle| j \in S_i \right\} \alpha(S_i). \end{aligned}$$

Hence

$$\begin{aligned} \sum_{i=1}^{m} v_i E[(\sum_{j \in N \cup F} c_{ij} N(j) - C_i)^+] &\leq \sum_{i=1}^{m} v_i \frac{\sigma_i}{2} \\ &\leq \frac{1}{2} \sum_{i=1}^{m} v_i \sqrt{\max\left\{ c_{ij}^2 \, | j \in S_i \right\} \alpha(S_i)} \\ &= \frac{1}{2} \sum_{i=1}^{m} v_i \max\left\{ c_{ij} \, | j \in S_i \right\} \sqrt{\alpha(S_i)} \\ &= \frac{1}{2} \sum_{i=1}^{m} \bar{v}_i \sqrt{\alpha(S_i)}. \end{aligned}$$

Hence:

$$\frac{V^{MTO}}{V^*} \geq \frac{V^d - \frac{1}{2} \sum_{i=1}^{m} \bar{v}_i \sqrt{\alpha(S_i)}}{V^d} = 1 - \frac{\sum_{i=1}^{m} \bar{v}_i \sqrt{\alpha(S_i)}}{2 \sum_{i=1}^{m} \hat{p}_i \alpha(S_i)}. \tag{5.20}$$

□

Corollary 5.2. Let V_h^{MTO} be the optimal objective value under deterministic model h. Then

$$\lim_{h\to\infty} V_h^{MTO} = V^*.$$

Proof. Similar to MTS, for the model it holds that $h\ \alpha_h(S_i) = h\alpha(S_i)$ (note that the h-s in $\bar{p}_i$ are cancelled). This gives

$$\frac{V^{MTO}}{V^*} \geq 1 - \frac{\sum_{i=1}^m \bar{v}_i \sqrt{h\alpha(S_i)}}{2\sum_{i=1}^m \hat{p}_i h\alpha(S_i)} = 1 - O(h^{-1/2}), \tag{5.21}$$

which goes to 1 as $h \to \infty$. □

5.4 Numerical Examples

This section provides numerical results to illustrate the model and the performance of the heuristics. The example that is used is based on a problem instance faced by a firm in the online advertisement industry, the problem that motivated the research of this study. In the online advertisement market publishers sell space on their website to advertisers. Consider a publisher who owns three websites A, B and C. On each website, two advertisement spots are available: a banner (1) on the top of the web page, and a box (2) a bit lower. See Figure 5.1 for a visualisation.

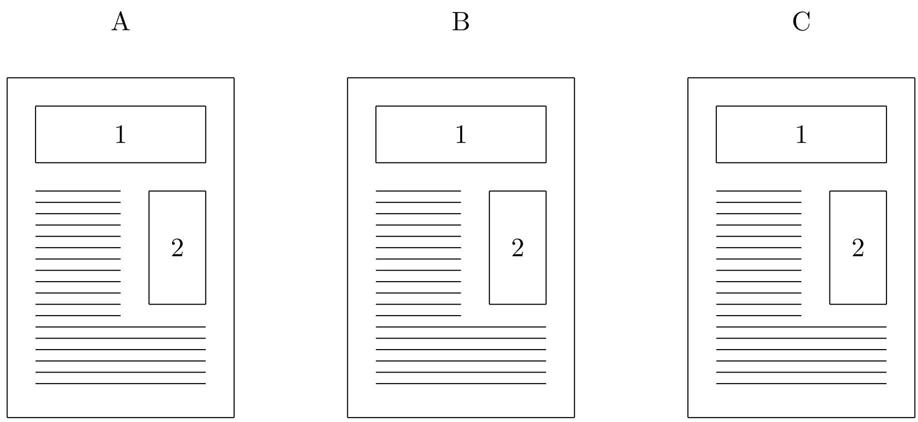

Figure 5.1: Visualisation of advertisement spots on three websites.

The capacity of the resources are $C = (100, 100, 70, 70, 30, 30)$, which are the views of the websites, respectively. The resources are consumed in $T = 10$ time period (say, weeks, or months), over which the publisher sells them in the form of products. The publisher sells the specific products where only one resource is consumed separately, or both banner and box of one website together. The incidence matrix A is therefore given by

$$A = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix}. \tag{5.22}$$

The publisher also uses flexible products:

$$F_1 = \{1,3\} \qquad F_4 = \{2,4,6\}$$
$$F_2 = \{1,3,5\} \qquad F_5 = \{7,8\}$$
$$F_3 = \{2,4\} \qquad F_6 = \{7,8,9\}.$$

The demand function is linear and independent from other products, i.e., demand for product i is equal to $\lambda_i(t) = a_i(t) + b_i(t)p_i(t)$. The price sensitivity does not change over time, so $b(t) = b$ for all t. However, the intercept $a(t)$ changes over time, and is given by

$$a(t) = \tilde{a}\frac{\sqrt{t}}{\sum_{s=1}^{10}\sqrt{s}}. \tag{5.23}$$

The parameters $\tilde{a}$ and b are given in Table 5.1.

		$\tilde{a}_i$	b_i
Specific products	1	259	-0.10
	2	211	-0.09
	3	222	-0.07
	4	114	-0.06
	5	103	-0.05
	6	81	-0.04
	7	168	-0.08
	8	97	-0.07
	9	81	-0.06
Flexible products	1	108.11	-0.04
	2	135	-0.05
	3	108	-0.04
	4	86	-0.06
	5	124	-0.05
	6	92	-0.06

Table 5.1: Parameter values for specific and flexible products.

5.4.1 MTS and MTO

First the performance of the MTS and MTO heuristics under the parameters described above is discussed. To capture the effect of the size of the problem, the problem instances are scaled by a factor $h \in \{0.1, 0.2, 0.5, 1, 2, 5, 10\}$. Furthermore, the heuristics are compared to two additional heuristics: MTS-NF and MTO-NF (no flexible products). MTS-NF and MTO-NF correspond to the MTS and MTO heuristics, except for the fact that no flexible products are offered. The actual revenue of the strategies is calculated by means of simulation. The errors are within 0.5% of the stated values, with 95% confidence. The results are presented in Table 5.2.

		MTS		MTO		MTS-NF		MTO-NF	
h	UB	Rev.	%UB	Rev.	%UB	Rev.	%UB	Rev.	%UB
0.1	7869	5372	68.27	6971	88.59	4622	58.74	5352	68.01
0.2	15737	12370	78.60	14548	92.44	10506	66.76	11296	71.78
0.5	39342	34628	88.02	37493	95.30	28471	72.37	29605	75.25
1	78685	72162	91.71	76098	96.71	59133	75.15	60558	76.96
2	157370	148093	94.11	153748	97.70	121330	77.10	123131	78.24
5	393425	379339	96.42	387734	98.55	309386	78.64	312131	79.34
10	786850	767164	97.50	778854	98.98	624898	79.42	628605	79.89

Table 5.2: Performance of MTS and MTO heuristic, together with MTS-NF and MTO-NF.

The results give two main insights. First, the MTO heuristic performs significantly better than the MTS heuristic in this example. The original problem with $h = 1$ gives an optimality gap of at most 3.29% for MTO and 8.29% for MTS. The gap could be smaller because the reference revenue is only an upper bound on the optimal revenue attainable. For smaller h the optimality gap of MTS increases substantially to up to 31.73% for $h = 0.1$, while for MTO the gap increases only up to 11.41%. As a result, protecting resources for certain products to prevent cannibalisation is not profitable in the case that demand is scarce. As the scale of the problem increases, the optimality gap becomes smaller. The size of the problem compensates the booking limits.
The second insight is that not offering flexible products leads to an enormous revenue loss compared to MTS or MTO: 19.3% for $h = 10$ to up to 23.2% for $h = 0.1$. Intuitively, flexible products provide the flexibility of deciding which resources are consumed at a later moment in time, when all demand is known. This flexibility and control over resource consumption leads to more efficient management of resource allocation and utilisation of capacity. This is in line with the findings of Gallego et alii (2004), who argue that offering flexible products alongside specific products leads to higher capacity utilisation and could attract additional customers without complete cannibalisation.

5.4.2 Estimation Error

Uncertainty in demand due to forecasting or estimation errors is an important issue and may lead to suboptimal policies. As an illustration, this example shows the impact of forecast/estimation errors on the performance of the different heuristics. In each simulation the parameters of the assumed demand are randomly drawn according to a normal distribution with mean μ, the parameter value of the true demand ($a(t)$ and $b(t)$), and standard deviation $1.96\sqrt{|\mu|}/l$. Here, l can be interpreted as the sample size, and the standard deviation follows from the confidence interval of the 'estimated' parameter μ. The results are shown in Table 5.3.

		MTS		MTO	
l	UB	Rev.	%UB	Rev.	%UB
∞	78685	72162	91.71	76098	96.71
10	78685	69936	88.88	72479	92.11
20	78685	70918	90.13	73978	94.02
50	78685	71537	90.92	75039	95.37
100	78685	71755	91.19	75469	95.91
200	78685	71875	91.34	75713	96.22
500	78685	72012	91.52	75923	96.49
1000	78685	72065	91.59	76000	96.59

Table 5.3: Results under estimation errors. The first row, where $l = \infty$, represents the case where demand is known and is the same as in Table 5.2.

As can be seen from the results, the performance does not suffer much from forecasting errors. The good performance under demand uncertainty is in line with the results of Petrick et alii (2012). Only for small datasets the revenue loss is 0.9-3.1% for the MTS heuristic and 1.4-4.8% for the MTO heuristic (compared to the MTS and MTO heuristics under true demand, respectively). Note that the revenue loss is higher for the MTO heuristic than for the MTS heuristic. This can be explained by the fact that MTS reserves resources for allocation of products. When products are priced too low due to forecasting errors, MTO will sell them until the resources are exhausted, such that the more expensive ones can not be sold any more. On the other hand, MTS will protect the resources for the more expensive products.

5.4.3 Group Bookings

In this final example the effect of group bookings is measured. Group bookings are common in many industry applications like hotels, airlines, and online advertisements. In this section the model from Section 5.2.3 that incorporates group bookings is used. Flexible products are allowed to be sold in groups of size 5 and 10. Demand is based on the parameters of Table 5.1, except for the fact that the demand parameter $a(t)$ is multiplied by 0.7, for groups of size 5 with 0.2, and groups of size 10 with 0.1. Table 5.4 shows the results. The MTS and MTO strategies take group bookings into account; MTS-NG and MTO-NG are the MTS and MTO strategies that do not take group bookings into account, but instead price group reservations the same as they do single-product reservations.

		MTS		MTO		MTS-NG		MTO-NG	
h	UB	Rev.	%UB	Rev.	%UB	Rev.	%UB	Rev.	%UB
0.1	4687	2981	63.59	4155	88.64	2934	62.59	4035	86.09
0.2	9374	7177	76.56	8662	92.40	6932	73.94	8403	89.64
0.5	23436	20002	85.35	22323	95.25	21083	89.96	21394	91.29
1	46872	42243	90.12	45328	96.71	42994	91.73	42832	91.38
2	93745	87444	93.28	91567	97.68	85729	91.45	85536	91.24
5	234362	224873	95.95	230919	98.53	213386	91.05	213326	91.02
10	468724	455440	97.17	463773	98.94	426261	90.94	426203	90.93

Table 5.4: Performance of MTS and MTO under group bookings, together with MTS-NG and MTO-NG.

Three observations can be made. First, there is a clear difference in performance between the MTS and MTS-NG heuristics on the one side and the MTO and MTO-NG heuristics on the other side. Similar to previous examples, the MTO heuristic performs better than the MTS heuristic: by 7.3% for $h = 1$ and up to 39.5% for smaller problem sizes. Hence selecting the wrong heuristic can have a dramatic impact on revenue. The fact that MTS reserves resources for products by rounding the deterministic α_j's may cause this significant loss in revenue.

The second observation is that not taking group bookings into account results in a substantial revenue loss of up to 8.1%, for the MTO heuristic. The gap is larger when h increases. MTS, however, does not always outperform MTS-NG: for $h = 0.5$ and $h = 1$ MTS-NG performs 5.4% and 1.8% better than MTS, respectively. According to these examples, MTS does not only perform worse, but also shows no monotonicity in performance, which makes it unreliable.

In conclusion, group reservations have a big impact on revenue, but forecasting and implementing a method that can effectively take group reservations into account is challenging.

5.5 Concluding Remarks

Selling inventory as flexible products in bulk is common practice in the online advertisement industry, but other industries (like retailing and FMCGs) could also profit from introducing flexible products. Flexible products give the company the flexibility to assign the customer close to consumption to a selection specific products, as capacity allows. Moreover, since flexible products give the company this flexibility, it can ask a lower price, which attracts new customer segments. Hence flexible products can lead to better capacity utilisation and higher revenues. The numerical studies endorse this by showing an increase in revenue of up to 20% when flexible products are offered alongside specific products.

This chapter introduces and analyses a pricing-based network RM model with flexible products, and a non-trivial extension to group bookings/bulk purchases. Two practical solution methods are examined that are efficient to solve and implement: the make-to-order (MTO) and make-to-stock (MTS) heuristics, which are based on the deterministic variant of the problem. Both heuristics set price product prices according to demand for discrete decision moments in time, such that prices can dynamically be updated. The MTS heuristic applies booking limits on products, with the goal of protecting inventory for profitable future sales. However, numerical results show that the MTO heuristic, that does not enforce booking limits, performs much better than MTS (up to 30%), especially when inventories and demand are small. As an explanation, the enhanced capacity utilisation that comes with assigning flexible products to specific products offsets the lost demand that is caused by booking limits.

The analytical results provide theoretical optimality bounds on the heuristics, i.e., the optimality gap is guaranteed to lie beneath a certain value, depending on the problem instance. Moreover, it is shown that the heuristics are asymptotically optimal as capacity and demand increase. The numerical examples, based on a problem instance from industry, support the analytical findings under different problem sizes: the optimality gap converges from 31.73% for a small instance to 1.02% for a large instance.

Besides the flexible nature of the products, the selling in bulk – or in group bookings – is a particular aspect of the problem that is dominating the industry. For example, publishers do not sell their inventory per impression, but rather per 1,000 or 10,000 impressions per customer. Group reservations of flexible products offer more flexibility in the assignment of flexible products to inventory, as opposed to specific group reservations. The simulation studies endorse the effectiveness of flexible group reservations: taking this effect into account leads to 7.3% increase in revenue for an average-size problem.

Another aspect that is important in practice but often overlooked in academics is the robustness of pricing strategies. The pricing methods of this study assume that demand and price sensitivity is known exactly, but in practice this is not the case and only forecasts are available. A desirable feature of the pricing strategy is that it is robust and stable under forecast errors. To this account, a numerical study was performed to measure the performance of the pricing strategies under forecast errors. The results show that for both heuristics the revenue loss is reasonable (from 0.9% for small errors to 4.8% for large errors), implying that the heuristics are quite robust and reliable for applications.

The results of this study have several implications that lead to topics for further research. A popular trend in RM is customer behaviour, such as stockpiling and consideration sets of customers. Solving the problem under choice-based demand is not straightforward, and might lead to promising insights. Also, an in-depth analysis of more heuristics could lead to improved results. Although the optimality gap is at most 3.29% in the case of $h = 1$ in Section 5.4.1, the potential increase in revenue can lead to substantial increase in profit. The additional revenue will only account for more profit, since no extra costs are incurred.

CHAPTER 6

Choice-Based Single-Leg Revenue Management under Online Reviews

This chapter proposes a revenue management model that integrates reviews and ratings. The dependency between reviews and revenue is two-fold: the content of a review depends on the product the customer purchases, and reviews impact the demand. A complicating factor in this model is that the effects of reviews are delayed, i.e., by sacrificing revenue now in order to get better reviews, long-term revenue can be increased. Because the full planning problem of finding an optimal strategy for the proposed model is intractable, a novel solution methodology is proposed to solve the problem approximately by restricting the space of possible solutions to *equilibrium strategies*. It is shown that equilibrium strategies for the full problem can be found by viewing the full problem as a series of multi-objective Markov decision processes subproblems, and aiming for a target review ratio in the subproblems while optimising revenue. Numerical studies show that taking reviews into account in this manner can lead to an increase in revenue of up to 11% compared to the case where the sole objective is revenue.

6.1 Introduction

A recent development that influences sales in various industries is the wide availability of reviews. For example, customers who attended a play in a theatre are given the opportunity to share their experience with other potential clients via websites like Tripadvisor. Research strongly suggests that reviews influence the buying behaviour of customers (Pan & Zhang, 2011; Park & Lee, 2009; Yoo & Gretzel, 2011). For instance, Sparks & Browning (2011) show that 'consumers seem to be more influenced by early negative information' and Pavlou & Dimoka (2006), Vermeulen & Seegers (2009), and Ye et alii (2011) show that more recent reviews have more impact. It is therefore in the interest of the theatre to get good reviews and avoid bad reviews. Conversely, customer reviews are influenced by the price/quality perception of the customer Zhou et alii (2014). This chapter presents a pioneering study that proposes a decision-theoretic model that explicitly includes reviews to maximise the revenue.

This chapter is based on Sierag & Roijers (2016).

This chapter studies the problem of optimising revenue over multiple performances when demand depends on reviews, and where cancellations, overbooking, and customer choice preferences are taken into account. Moreover, the model contains a feedback mechanism: on the one hand the creation of reviews depends on the purchases, and on the other hand demand depends on reviews. The focus is on recent reviews instead of rating, since rating tends to remain more or less constant, while customers are more influenced by recent reviews, either positive or negative. Demand for a certain performance depends on reservations of previous performances, due to the reviews that are released in the meantime. This makes it a hard problem, suffering heavily from the curse of dimensionality. The approach of this study is to maximise revenue while the ratio between positive and negative reviews remains constant, called *equilibrium strategies.*[1] Searching in the space of equilibrium strategies is an approximation, yet even under these circumstances finding an optimal solution is challenging. An elaborate though tractable solution method is provided, using a multi-objective Markov decision process formulation.

This study builds on three research areas. The first research area is the work on choice-based assortment problems, such as the work in Chapter 3 and Talluri & van Ryzin (2004a). We refer to Shen & Su (2007) for an overview of literature in this area. The basic idea is that customer preferences depend on the products that are offered, and the set of offered products should be optimised accordingly. The second area is the work on e-word-of-mouth (eWOM), in particular the effect of reviews on demand, and the effect of purchases on reviews. See for example the survey by Ye et alii (2011) for an overview of related literature. Studies often show that the impact of reviews is strong, though no connection is made to optimise assortments accordingly. Closest to the work in this chapter is the article by Dellarocas et alii (2004), where reviews are used to forecast box office revenues. However, the reviews are not used to optimise revenue, and the creation of new reviews is is not considered. The third research area is the work on multi-objective Markov decision processes (Roijers et alii, 2013). When it is not clear a priori what the relative importance of each objective is and how the objectives interact, an set of different trade-offs between the objectives is considered. Such a set is called a coverage set. In this study, a series of multi-objective problems approximates an otherwise intractable problem.

This study makes the following research contributions: 1) a model that incorporates reviews in the optimisation process (Section 6.2); 2) an analysis of the model indicating that the full problem is intractable, and an approximation as series of tractable multi-objective problems (Sections 6.3 and 6.4); and 3) numerical experiments indicate that taking reviews into account induces significantly higher long-term revenue than when solely revenue is optimised (Section 6.5). Finally, implications for research and practice are discussed in Section 6.6.

6.2 General Model Description

The model is framed in the context of theatres. However, we emphasise that the model is more general and can be applied to other single-leg experience-based products, like sports events or concerts, or the cinema.[2]

Consider a theatre with C identical seats that wants to sell them for one or more plays on multiple performances. In particular, assume a (possibly infinite) set $I \subset \mathbb{N}$ of fu-

[1] Note that this is not the only possible solution approach. However,this solution has the advantage that it is tractable and leads to large improvements in revenue as Section 6.5 indicates.

[2] In general, the approach can be used for any problem that can be modelled as a series of decision problems in which the output of an earlier problem (partially) parameterises the next problem in the series.

ture arrival dates, each with booking horizon of T time units (selling seats of a particular performance only starts at T units before the performance), 0 being the arrival time. Overbooking is allowed up to $C_{\max}$ seats.[3] Each customer has the opportunity to give a *review* of their experience, which is *positive* or *negative*, and corresponds to a performance $i \in I$ of the past. The set of past reviews is given by $\mathcal{R}$.

Demand depends on the reviews. Clients for performance $i \in I$ arrive according to a Poisson process $\lambda_i(\mathcal{R})$ dependent on the current set of reviews, starting T days before performance i. Each seat is sold as a *fare product* j, which is a combination of a seat with a price r_j and conditions, such as the cancellation policy. Moreover, customers are influenced by the fare product when writing a review. The probability that fare product j leads to a positive review is q_j^p and to a negative review it is q_j^n.

Assume that there is a finite number of fare products $N = \{1, \ldots, n\}$. At each moment in time, for each future performance separately that is at least T time units ahead, the theatre manager decides which *offer set* $S \subset N$ of fare products to offer. Depending on the offer set S displayed and the reviews $\mathcal{R}$, an arriving customer decides to either buy one of the fare products $j \in S$, with probability $P_j(\mathcal{R}, S)$, or leave and buy nothing at all, with probability $P_0(\mathcal{R}, S)$.

Reservations are allowed to be cancelled, where a potential refund depends on the cancellation conditions of the product. Assume that cancellations occur independently from each other. The cancellation rate $\gamma_i(\mathcal{R})$ is independent of the product, but depends on the reviews. If there are x_j^i reservations for fare product j for performance i, then cancellations occur with rate $\gamma_i(\mathcal{R})x_j^i$. The costs for the theatre of a cancelled reservation of product j at time t equals $c_j^i(t)$. When there are more reservations than capacity at the end of the performance, i.e., when $x_i > C_{\max}$, a penalty $q(x_i)$ is incurred. See Figure 6.1 for an overview of the model.

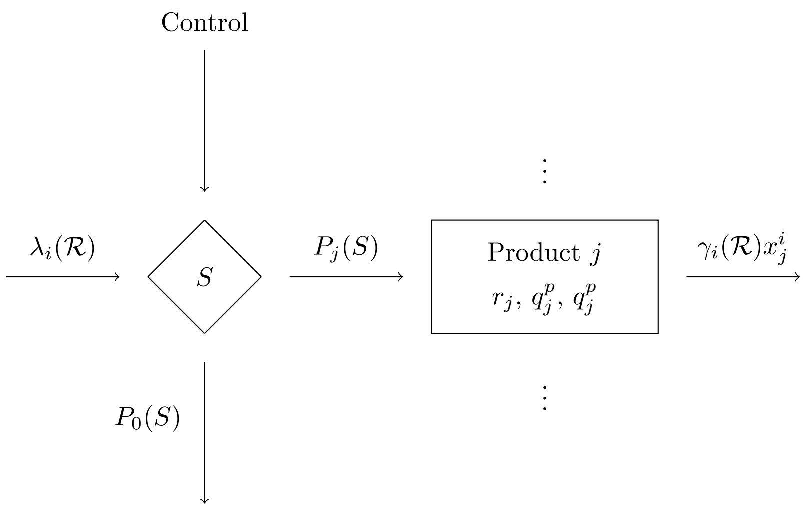

Figure 6.1: Illustration of the model. Per performance the arrival process is Poisson distributed with parameter $\lambda_i(\mathcal{R})$. The manager controls the offer set S. Under this offer set an arriving customer buys product $j \in S$ with probability $P_j(\mathcal{R}, S)$. With probability $P_0(\mathcal{R}, S)$ the customer buys nothing. Finally, cancellations of product j follow an exponential distribution with parameter $\gamma_i(\mathcal{R})$.

[3] Note that $C_{\max}$ can be equal to C, hence the model without overbooking is a special case of this model.

6.2.1 Modelling Demand and Cancellations as a Function of Reviews

Demand is influenced by reviews, either positive or negative (Pan & Zhang, 2011; Park & Lee, 2009; Sparks & Browning, 2011; Vermeulen & Seegers, 2009; Yoo & Gretzel, 2011). To capture this effect, consider the following set-up. Let Q_i^p and Q_i^n be the number of positive and negative reviews resulting from performance i. Recent reviews are more relevant than older reviews. Let $\alpha \in (0,1)$ be the discounting parameter for relevance of reviews and let M be the number of past performances that are relevant. Define the *discounted reviews* $\tilde{Q}_k^p$ and $\tilde{Q}_k^n$ on performance k by

$$\tilde{Q}_k^p := \sum_{i=k-M}^{k-1} \alpha^{k-i-1} Q_i^p, \qquad \tilde{Q}_k^n := \sum_{i=k-M}^{k-1} \alpha^{k-i-1} Q_i^n.$$

The most recent reviews of performance $k-1$ are not discounted, while the relevant reviews of arrival $k-M$ are discounted most, by factor α^{M-1}, $M \geq 2$.

Let $\bar{\lambda}_i, \bar{\gamma}_i \in \mathbb{R}$ be the base parameter of demand and cancellations, respectively, for performance $i \in I$. Let $\beta_p^\lambda, \beta_n^\lambda, \beta_p^\gamma, \beta_n^\gamma \in \mathbb{R}$. Define $\lambda_i(\mathcal{R})$ and $\gamma_i(\mathcal{R})$ by

$$\begin{aligned}
\lambda_i(\mathcal{R}) &:= \bar{\lambda}_i \exp\left(\beta_p^\lambda \tilde{Q}_k^p + \beta_n^\lambda \tilde{Q}_k^n\right), \\
\gamma_i(\mathcal{R}) &:= \bar{\gamma}_i \exp\left(\beta_p^\gamma \tilde{Q}_k^p + \beta_n^\gamma \tilde{Q}_k^n\right).
\end{aligned}$$

In accordance with literature (Vermeulen & Seegers, 2009) it is assumed that positive reviews have positive effect on demand and negative reviews have a negative effect on demand, which translates to $\beta_p^\lambda > 0$ and $\beta_n^\lambda < 0$. For cancellations the opposite is expected: positive reviews result in less cancellations and negative reviews result in more cancellations, i.e., $\beta_p^\gamma < 0$ and $\beta_n^\gamma > 0$. In line with Sparks & Browning (2011), the effect of negative reviews is larger than positive reviews: $\beta_p^\lambda < -\beta_n^\lambda$ for demand, and $\beta_p^\gamma < -\beta_n^\gamma$ for cancellations. The exponential function provides these intuitive features, as well as the fact that demand cannot be negative. Define the *review ratio* ρ by

$$\rho := \frac{\tilde{Q}_k^p}{\tilde{Q}_k^p + \tilde{Q}_k^n}.$$

Then an equivalent formulation for $\lambda_i(\mathcal{R})$ and $\gamma_i(\mathcal{R})$ is

$$\begin{aligned}
\lambda_i(\mathcal{R}) &:= \bar{\lambda}_i \exp\left(\tilde{\beta}_p^\lambda \rho + \tilde{\beta}_n^\lambda (1-\rho)\right), \\
\gamma_i(\mathcal{R}) &:= \bar{\gamma}_i \exp\left(\tilde{\beta}_p^\gamma \rho + \tilde{\beta}_n^\gamma (1-\rho)\right),
\end{aligned}$$

where $\tilde{\beta}_p^\lambda = \beta_p^\lambda(\tilde{Q}_k^p + \tilde{Q}_k^n)$, $\tilde{\beta}_n^\lambda = \beta_n^\lambda(\tilde{Q}_k^p + \tilde{Q}_k^n)$, $\tilde{\beta}_p^\gamma = \beta_p^\lambda(\tilde{Q}_k^p + \tilde{Q}_k^n)$, and $\tilde{\beta}_n^\gamma = \beta_n^\gamma(\tilde{Q}_k^p + \tilde{Q}_k^n)$.

6.2.2 Running Example

To illustrate the model, consider the following running example. To capture the effects of reviews under different circumstances, four instances are considered by adjusting the review probabilities and the review parameters of the demand. For both cases the effects are either large or small, leading to four scenario's: 1) large effect on both demand and review probabilities; 2) large effect on demand and small effect on review probabilities; 3) small effect on demand and large effect on review probabilities; and 4) small effect on both demand and review probabilities. The base demand is $\bar{\lambda}$ for all performances, with parameters $(\beta_p^\lambda, \beta_n^\lambda) = (1, -1.5)$, i.e., a large effect on demand, in cases 1) and 2);

and $(\beta_p^\lambda, \beta_n^\lambda) = (0.9, -1)$, i.e., a small effect on demand, in cases 3) and 4). The review probabilities are presented in Table 6.1 below. A difference is made between a reservation that is overbooked and one that is not. The difference is independent of the product, such that it can be incorporated in a similar way that overbooking costs are incurred. The values of the attributes are motivated by the findings of Zhou et alii (2014).

	Large effect (cases 1 and 3)				Small effect (cases 2 and 4)			
	No overbooking		Overbooking		No overbooking		Overbooking	
Product j	q_j^p	q_j^n	q_j^p	q_j^n	q_j^p	q_j^n	q_j^p	q_j^n
1	0.160	0.120	0.140	0.140	0.130	0.098	0.110	0.118
2	0.130	0.160	0.110	0.180	0.090	0.144	0.070	0.164
3	0.210	0.095	0.190	0.115	0.155	0.088	0.135	0.108
4	0.240	0.080	0.220	0.100	0.170	0.082	0.150	0.102
5	0.230	0.110	0.210	0.130	0.140	0.124	0.120	0.144
6	0.288	0.056	0.268	0.076	0.194	0.072	0.174	0.092
7	0.304	0.048	0.284	0.068	0.202	0.069	0.182	0.089
8	0.270	0.090	0.250	0.110	0.160	0.116	0.140	0.136
9	0.326	0.037	0.306	0.057	0.213	0.065	0.193	0.085
10	0.280	0.085	0.260	0.105	0.165	0.114	0.145	0.134

Table 6.1: Review probabilities for both large and small effect. In case of overbooking the probabilities differ.

In this example, the purchase probabilities are modelled according to the commonly used *multinomial logit* (MNL) model, where the sole attribute is price. The results of this paper are more general, and the MNL model is only used as an illustration. Utility plays a crucial role in the MNL model (Ben-Akiva & Lerman, 1985). Let $u_j = \beta r_j$ be the mean utility of product j, where $\beta \in \mathbb{R}$ is the corresponding weight. The purchase probabilities are then given by

$$P_j(S) = \begin{cases} \dfrac{e^{u_j}}{\sum_{i \in S} e^{u_i} + e^{u_0}} & \text{if } j \in S, \\ 0 & \text{otherwise,} \end{cases}$$

for all $j \in N$. Assume there are two types of customers: high price-sensitive and low price-sensitive. High price-sensitive customers have parameter $\beta^H = -0.005$, while low price-sensitive customers have parameter $\beta^L = -0.0015$.

6.3 A Tractable Model

In this section we formulate the optimisation problem of maximising revenue in the long run by taking into account the reviews. First, computational issues related to the size of the problem are discussed. Second, after concluding that the full problem is intractable, a model that approximates the large problem by breaking it down to a series of subsequently smaller problems is proposed. Within this (infinite) series, the proposed solutions are those that lead to stable review sets. Such solutions are referred to as *equilibrium strategies.* On the level of a smaller problem, the review sets must be stable for an equilibrium strategy, for two objectives: positive reviews, and negative reviews. These objectives supplement the (immediate) revenue objective. A dynamic programming formulation is given to find the resulting three-dimensional value function. In the long run, the review scores influence the long-term revenue. However, it is not clear at the level of a smaller problem how the review scores translate to long-term effects on the revenue. Therefore, in this section the focus lays on a three-objective model that incorporates the different trade-offs between these objectives. In Section 6.4 the computational method to find equilibrium strategies is discussed.

6.3.1 Computational Issues

Suppose that the sales manager needs to optimise revenue for multiple shows or performances. Then the creation of new reviews after a show influences the purchasing behaviour of customers in the future. Moreover, the reviews may be updated during the booking horizon of another performance. Hence, strategies for a particular performance need to take into account the strategies of other performances.

To illustrate this, consider the following example, with only three performances with a booking horizon of three days. Only the reviews from the last performance are used, i.e., $M = 1$. Denote S_t^i as the offer set that is offered t days beforehand for performance i. The offer sets of the first performance are not influenced by any previous decisions, as well as the first two offer sets of performance 2 and the first of performance 3. However, after the first performance, reviews are updated. The amount of positive and negative reviews depends on offer sets S_1^1, S_2^1, and S_3^1. Therefore, the decision which offer set S_1^2 to use on the last day before performance 2 depends on all offer sets of performance 1. Note that S_2^3 is also dependent on all offer sets of performance 1. Furthermore, after performance 2 the reviews are updated, such that S_1^3 depends on all offer sets of arrival 2, and hence also on all offer sets of performance 1. See Figure 6.2 for an overview of this example.

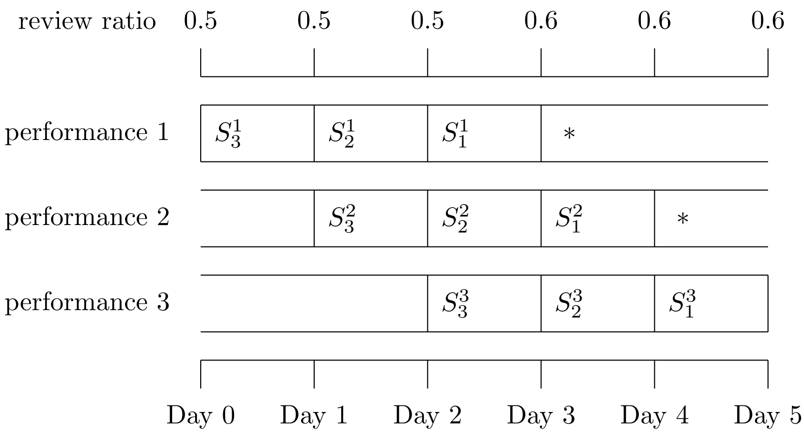

Figure 6.2: Illustration of example with $M = 1$, $K = 0$, and $T = 3$ over a five day period. A star $*$ denotes the end of a performance and the update of the reviews.

Let R^i be the revenue that results from performance i and denote by $S^i = \{S_T^i, \ldots, S_1^i\}$ the strategy of performance i. Then R^i is a function of the strategies of all performances up to day i: $R^i\big(S^1, \ldots, S^i\big)$. The objective $\phi^I(S)$ to maximise the finite horizon problem, with $I \in \mathbb{N}$ the horizon length, is therefore given by

$$\phi^I(S) = \sum_{i=1}^{I} R^i\big(S^1, \ldots, S^i\big),$$

while the a-discounted infinite horizon objective function $\phi^\infty(S)$ is given by

$$\phi^\infty(S) = \sum_{i=1}^{\infty} a^{i-1} R^i\big(S^1, \ldots, S^i\big).$$

At every moment in time products for up to T performances are in consideration to be for sale. Therefore, the state space for these T parallel processes has to be taken into account when solving the problem. Just using the Cartesian product of the T processes is not enough. Namely, this leads to a system that does not have the Markov property, because to make a decision it is necessary to take the review ratio into account. One option is to include the review ratio in the state space as a real number. However, this leads to a continuous state space, yielding an intractable problem. An option with finite state space is to keep track of the strategies of all past performances that are necessary for calculating the review ratio. This model has the Markov property. However, for both the finite horizon and the discounted infinite horizon problem the action space as well as the state space are still too large: the action space has size N^T and the state space has size up to N^I for the finite horizon, and converges to an infinite state space for the infinite horizon problem. Therefore, finding an optimal solution in the set of all feasible solutions is intractable because of the curse of dimensionality.

6.3.2 Equilibrium Strategies

Motivated by these deliberations, consider the set of *equilibrium strategies*, which consists of strategies that keep the expected review ratio constant and equal to the *target review ratio*. In equilibrium, i.e., when the target review ratio is achieved and remains fixed (in expectation), it is clear beforehand what the expected review ratios at all time periods will be for each performance. Also, the long-term revenue is no longer influenced by changing review ratios; R_i solely depends on S_i. Therefore, each equilibrium can be solved separately and the problem can be solved per target review ratio. In such cases, the long term revenue is equal to the number of times a *single-performance dynamics problem* is executed, times the revenue attained in this single-performance dynamics problem, given the fixed target review ratio. Such a single-performance dynamics problem is modelled as a multi-objective Markov decision process and is tractable.

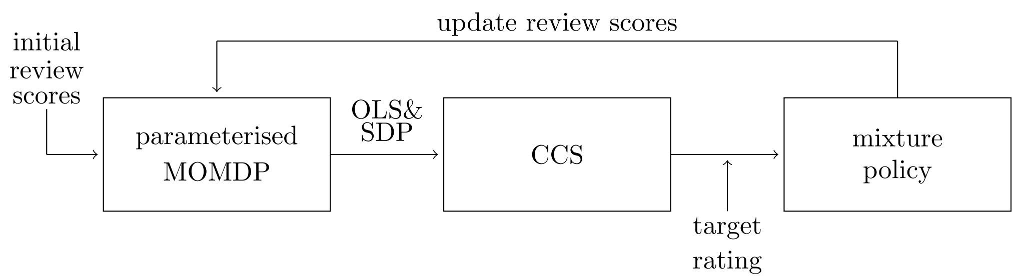

Figure 6.3: The process of finding an equilibrium policy. Starting from an initial parameterised MOMDP, a CCS is identified, from which the mixture policy with the intended target review ratio is chosen, of which the output in terms of positive and negative review scores are used reparameterise the MOMDP. This process repeats until the parameters of the MOMDP converge.

However, before the system is in equilibrium, it must be shown that the equilibrium, i.e., the right review ratio, can in fact be reached. To this end, the process of Figure 6.3 is followed. An initial MOMDP is constructed on the basis of initial review scores, i.e., the parameters for demand and review probability are calculated from the current reviews. Then, a *convex coverage set* (formally defined in Section 6.4.1) is computed, from which a stochastic mixture policy can be constructed that optimises the short-term revenue while guaranteeing that the target review ratio is met. Executing this stochastic mixture policy leads to new review scores. With these new review scores, the process is restarted until the parameters of the MOMDP converge. After that, the maximal revenue for the target review ratio is determined, as the value of the equilibrium strategy.

To illustrate the equilibrium strategies in the successive model, consider the example from earlier in this section. Suppose the review ratio on day 0 is 0.5 and the target review ratio is 0.6. Then the review ratios on days 1 to 5 equal 0.5, 0.5, 0.5, 0.6, and 0.6, respectively. All performances are then separately solved and revenue is optimised to target review ratio 0.6.

In the remainder of this section, a method is developed that optimises revenue to a target review ratio.

6.3.3 One Performance Dynamics

The one performance problem can be modelled as a finite-horizon continuous-time multi-objective Markov decision process over T time units. Define the state space by

$$X := \{0, \ldots, C_{\max}\},$$

where $x \in X$ represents the number of reservations. The *action space* $\mathcal{A}$ at each time step is the set of all possible offer sets $S \subset N$.

To solve the problem, time is discretised and divided into T time periods, where the length of the intervals is such that the probability that more than one event occurs is very small (following Chapters 3 and 4). To formulate the *transition probabilities*, it is assumed (following Chapters 3 and 4) that only one event occurs per time period, where an event is either an arrival, a cancellation, or neither arrival nor cancellation. Denote with λ the probability that a customer arrives in a time period; and γx the probability that a product is cancelled in state x. The probability that no purchase occurs in a time period equals the sum of the probability that neither an arrival and nor a cancellation occurs, and the probability that an arrival occurs but the arriving customer makes no purchase. This is equal to

$$(1 - \lambda - \gamma x) + \lambda P_0(S) = 1 - \lambda \sum_{j \in S} P_j(S) - \gamma x.$$

In each time period the decision needs to be made which set S to offer. Recall that time has to be scaled such that $\lambda + \gamma C \leq 1$ for the probabilities to be well defined. The transition probabilities from a given state $0 < x < C_{max}$, are thus:

$$\begin{aligned}
P(x' = x + 1 | x, S) &= \lambda \sum_{j \in S} P_j(S), \\
P(x' = x - 1 | x, S) &= \gamma x, \\
P(x' = x | x, S) &= 1 - \lambda \sum_{j \in S} P_j(S) - \gamma x,
\end{aligned}$$

and 0 otherwise. When $x = C_{max}$ it is not allowed to offer any fare products, and the only possible transitions are due to cancellations.

Now define the *reward function* $R \colon X \times \mathcal{A} \to \mathbb{R}^3$ as follows. Let $R^1(x, S)$ be the expected immediate revenue reward; $R^2(x, S)$ the expected good review reward; and $R^3(x, S)$ the expected bad review reward, in a state x for an action S. Define $r_{1j} = r_j$, $r_{2j} = q_j^p$, and $r_{3j} = -q_j^n$ for all $j \in N$. Define $c_{1j}(t) = c_j(t)$, $c_{2j}(t) = q_j^p$, and $c_{3j}(t) = -q_j^n$ for all $j \in N$ and for all t. The expected costs that follow from cancellations can be added to the immediate rewards when a product is purchased, since neither the booking system nor the manager has control over the cancellations of current reservations (see Section 3.3 for a detailed description of this derivation). By doing so, it is not necessary to account for the costs of cancellations of a product purchased at a time step later in the decision process, allowing the state x to be Markov.

The expected cancellation costs for objective $i \in \{1, 2, 3\}$ that follow from selling product j in time step t is equal to

$$\Delta H_{ij}(t) = \begin{cases} \gamma c_{ij}(t) + (1-\gamma)\Delta H_{ij}(t-1) & \text{if } t > 1, \\ 0 & \text{if } t = 1. \end{cases}$$

The expected reward that includes the expected costs of cancellations is therefore given by $r_{ij} - \Delta H_{ij}(t)$, for all $j \in N$, for $i = 1, 2, 3$, and for all t. The immediate reward for a cancellation,[4] as well as nothing happening is 0.

6.3.4 Policies and Value Vectors

A manager can interact with the model by executing a *policy* with a corresponding *value*. In the multi-objective (rather than single-objective) setting, stochastic policies can Pareto-dominate deterministic policies (Vamplew et alii, 2009; Wakuta, 1999). Therefore, define the policy space Π, as all possible mappings of states, time steps, and actions (i.e., fare-products) to a probability of taking that action: $\pi\colon X \times \{0, \dots, T\} \times \mathcal{A} \to [0, 1]$. For each state and time step, the probabilities for each action should be positive and should sum to 1.

A given policy, $\pi(x, t, S)$, induces a probability distributions over execution trajectories $(x_T, S_T, r^T, \dots\ x_1, S_1, r^1, x_0)$, where r^t is the reward vector corresponding to time step t. Each such a trajectory has an associated vector-valued *return*, i.e., the sum of the reward vectors in the trajectory. The value of a policy is its expected return. Because the returns are additive (Roijers et alii, 2013), the value of a policy π for a given time step and state, $V_t^\pi(x)$ can be expressed recursively with a Bellman equation in vector form, or, as separate Bellman equations per objective:

$$\begin{aligned} V_{i,t}^\pi(x) = \sum_{S \subset N} \pi(x, t, S) \Bigg\{ & \lambda \sum_{j \in S} P_j(S) \big[r_{ij} - \Delta H_{ij}(t) + V_{i,t-1}^\pi(x+1) \big] \\ & + \gamma x V_{i,t-1}^\pi(x-1) \\ & + \left(1 - \lambda \sum_{j \in S} P_j(S) - \gamma x \right) V_{i,t-1}^\pi(x) \Bigg\}. \end{aligned} \tag{6.1}$$

A *deterministic policy* is a policy π such that for all (x, t) it holds that $\pi(x, t, S(x, t)) = 1$ for some $S(x, t) \subset N$, and $\pi(x, t, S) = 0$ for all $S \neq S(x, t)$. A deterministic policy $\pi(x, t)$ may be seen as a mapping: $\pi\colon X \times \{0, \dots, T\} \to \mathcal{A}$. In this case, the marginalisation over actions drops out of the Bellman equation:

$$\begin{aligned} V_{i,t}^\pi(x) = & \lambda \sum_{j \in S} P_j(\pi(x,t)) \big[r_{ij} - \Delta H_{ij}(t) + V_{i,t-1}^\pi(x+1) \big] \\ & + \gamma x V_{i,t-1}^\pi(x-1) \\ & + \left(1 - \lambda \sum_{j \in S} P_j(\pi(x,t)) - \gamma x \right) V_{i,t-1}^\pi(x). \end{aligned} \tag{6.2}$$

[4] The immediate reward for cancellations is 0, because the expected costs are already accounted for in the purchase event rewards (see Section 3.3).

6.3.5 Scalarisation and Scalarised Solution Methods

In general there is no single policy that is best with respect to all objectives, because there are different policies that yield different trade-offs between the objectives. In the multi-objective decision making literature, a certain trade-off is typically preferred on the basis of the partially unknown preferences of a user, i.e., a human decision maker. These preferences can be expressed in terms of a *scalarisation function*, or *utility function*, f, which collapses the value vector of a policy, i.e., the vector of values of a policy in each objective, to a scalar:

$$V_w^\pi = f(V^\pi, w), \tag{6.3}$$

where V^π is the value vector of a policy π, and w is a vector parametrising f. In practice, only limited information is available about the function f. Therefore, the solution to a multi-objective decision problem is a set of value vectors and associated policies that *cover* all possible f and w. The human decision maker can choose her preferred policy from this set.

In our problem, there is no human decision maker. Instead, the function f represents the long-term revenue, as a result of optimally balancing the review scores and the immediate revenue. To determine the optimal long-term revenue exactly however, the intractable problem discussed at the beginning of this section has to be solved. Gladly, it not necessary to do that at this stage. Instead, assume that f is unknown but with some known constraints. This results in a tractable model that describes one-day-arrival dynamics.

In the problem at hand, i.e., the one-day-arrival dynamics multi-objective Markov decision process (MOMDP) model, the exact shapes of f and w are unknown. Therefore, the aim is to find a set of all possibly optimal solutions for the MOMDP, i.e., a set that *covers* all possible f and w that fit the known constraints about f and w (Roijers et alii, 2013):

> **Definition 6.1.** Let $\mathcal{F}$ be a set of scalarisation functions. Then a set of policies CS is called a *coverage set* if for every scalarisation function $f \in \mathcal{F}$ and every weight w there is a policy $\pi \in CS$ such that $V_w^\pi \geq V_w^{\pi'}$ for all $\pi' \in \Pi$, and, conversely, if for every policy $\pi \in CS$ there exist a scalarisation function $f \in \mathcal{F}$ and a weight w such that $V_w^\pi \geq V_w^{\pi'}$ for all $\pi' \in \Pi$.

In the next section, the CS is specified further, by imposing constraints on f that follow from the available information about how the value vectors attainable in the MOMDP model affect the long-term revenue. After that, the coverage set for the one-day-arrival dynamics MOMDP is used to determine approximately optimal strategies for the full problem.

6.4 Computing Coverage Sets for Single-performance Planning

In this section, coverage sets for one-day-arrival dynamics MOMDP planning are discussed. Firstly, the appropriate CS is derived from the available information about how immediate revenue and reviews affect the long-term revenue. Secondly, an algorithm that computes this CS is discussed. Finally, a way to employ the CS to improve long-term revenue is proposed.

6.4.1 Coverage Sets

Recall that the long-term revenue can be expressed in terms of a scalarisation function $f(\mathbf{V}^\pi, \mathbf{w})$. Which undominated policies need to be contained in the coverage set depends on what is known about f and w. Sometimes it is known beforehand that f has a particular shape. For example, f might be *linear* (White & Kim, 1980).

Definition 6.2. A *linear scalarisation function* is the inner product of a weight vector w and a value vector V^π:

$$V_w^\pi = w \cdot V^\pi. \tag{6.4}$$

Each element of w is greater than or equal to 0, and specifies how much one unit of value for the corresponding objective contributes to the scalarised value.

Linear scalarisation functions are both common and intuitive. The most common situation in which linear scalarisation applies is when the value vectors can be translated into monetary value. For example, consider a task in which objective corresponds to quantities of various resources that need to be bought or sold on a market. For revenue management this might therefore be the most intuitive way of scalarising. However, linearity is a strong assumption, that cannot be made in this context.

In fact, only the following information about f can safely be assumed:

1. immediate revenue always contributes positively to the long-term revenue,
2. positive reviews contribute positively to the long-term revenue, and
3. negative reviews always contribute negatively to the long-term revenue.

When a negative review objective is redefined as -1 times the number of negative reviews, an f is obtained that is *monotonically increasing* in all objectives.

Definition 6.3. Let V_i^π denote the value of policy π in the i-th objective. Let $\pi, \pi' \in \Pi$ be two policies. A scalarisation function f is *monotonically increasing* if for all $i \in \{1, \ldots, I\}$ such that $V_i^\pi \geq V_i^{\pi'}$, the following inequality holds for all weights w:

$$f(V^\pi, w) \geq f(V^{\pi'}, w). \tag{6.5}$$

The assumption that is f is *monotonically increasing* guarantees that if more of one objective is obtained while not losing anything in another objective, the utility cannot go down. Note that linear scalarisation functions (with non-zero positive weights) are included in the family of monotonically increasing functions. Monotonicity is therefore a less strict assumption than linearity.

When an optimal solution is needed with respect to all linear f, a coverage set is needed that contains an optimal solution for every possible weight vector w in Equation 6.4. The coverage with respect to linear f is called a *convex coverage set* (CCS).

Definition 6.4. A set of policies CCS is a *convex coverage set* if for all weights w there is a policy $\pi \in CCS$ such that $w \cdot V^{\pi} \geq w \cdot V^{\pi'}$ for all $\pi' \in \Pi$ and, conversely, if for all policies $\pi \in CCS$ there is a weight w such that $w \cdot V^{\pi} \geq w \cdot V^{\pi'}$ for all $\pi' \in \Pi$. Without loss of generality, assume that w adheres to the simplex constraints, i.e., all elements of w are positive and sum to one.

The coverage set for the family of monotonically increasing functions is called a *Pareto coverage set (PCS).*

Definition 6.5. Let $\succeq_P$ denote weak Pareto-dominance, i.e., $V^{\pi} \succeq_P V^{\pi'} \equiv \forall i \quad V_i^{\pi} \geq V_i^{\pi'}$. A set of policies PCS is a *Pareto coverage set* if for all policies $\pi \in \Pi$ there exists a policy $\pi \in PCS$ such that

$$V^{\pi} \succeq_P V^{\pi'}.$$

I.e., for every possible policy, there is a policy in the PCS, with at least equal value in all objectives.

Because monotonicity is a less strict assumption than linearity, the PCSs are typically much larger than the CCSs. Furthermore, because of the possibly non-linearity of the scalarisation function, policies that constitute a PCS are much harder to obtain than those that constitute a CCS. Nonetheless, it is not necessary to represent PCSs explicitly if it can be assumed that stochastic policies are allowed. This is due to a result by Vamplew et alii (2009) which states that for any MOMDP, a PCS of stochastic policies can be constructed from a given CCS of deterministic policies, by taking so-called *mixture policies* from policies in this CCS. A mixture policies is constructed by taking a subset of N policies from the CCS and assigning each a probability of being executed.

The difference between the CCS and the PCS with or without stochastic policies is illustrated in Figure 6.4, where all the possible values of deterministic policies for a two-objective MOMDP are denoted as points. The axes represent the different objectives. Note that the grey points are neither in any PCS or CCS, as they are dominated by one of the other points (i.e., there is another point with a higher value in all objectives). The points π_1, π_2, and π_3 represent a possible CCS of deterministic policies, while pi_4 represents a point that may, but is not necessarily in the CCS. Note that a CCS of deterministic policies is also a CCS for stochastic policies as there is always a deterministic policy that is optimal for any w (Howard, 1960; Roijers et alii, 2013). Point π_5 would be in a PCS if stochastic policies were not allowed as it is not dominated by another point. However, when stochastic policies (including mixture policies) are allowed, the PCS is represented by all the values on the black lines, on which there are policy values that dominate π_5.

To summarise, all the Pareto-optimal policy values of stochastic policies can be constructed by mixing deterministic policies from a CCS of deterministic policies. Therefore, this study focuses on finding methods for finding a CCS, and construct policies with values on the PCS of stochastic policies when necessary. Thus, define solving an MOMDP as *finding a convex coverage set.*

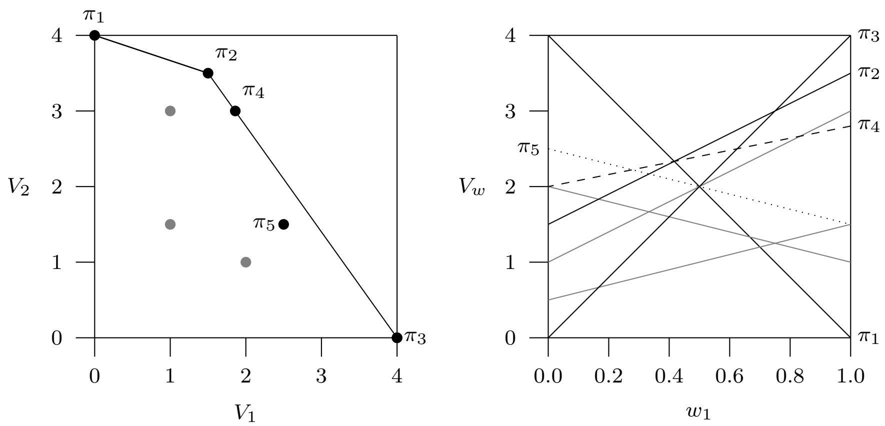

Figure 6.4: *(Left)* A stochastic PCS can be constructed from a deterministic CCS. *(Right)* The *scalarised value* as a function of the linear scalarisation weight.

6.4.2 Optimistic Linear Support

In order to compute the CCS for the single arrival model, the *Optimistic Linear Support* (OLS) algorithm (Roijers et alii, 2015) is used. OLS is a general framework for solving multi-objective decision problems (including MOMDPs, Roijers et alii, 2014). It takes a single-objective solver – such as dynamic programming – as a subroutine, and produces a CCS within a finite number of calls to this subroutine. In this section, the framework is described, and the specific implementation of OLS for multi-objective revenue management.

Because OLS computes the CCS, it can make use of linear scalarisation. Under this assumption, define the *scalarised value function* $V^*_{CCS}(w)$ that provides the maximal scalarised value given a linear scalarisation weight w:

$$V^*_{CCS}(w) = \max_{V^\pi \in CCS} w \cdot V^\pi.$$

Here, $V^*_{CCS}(w)$ is a piecewise linear and convex (PWLC) function, because each value function defines a (hyper)plane over the weight simplex, as illustrated in Figure 6.4 (right), and $V^*_{CCS}(w)$ maximises over these hyperplanes. That is, it consists of the convex upper surface of the lines in Figure 6.4 (right).

OLS builds up the CCS incrementally by solving a series of linearly scalarised instances of the MOMDP, for different w. The optimal policy π to an instance scalarised with w maximizes $V^\pi_w = w \cdot V^\pi$. When this π is identified, V^π is added to a partial CCS X, which converges to a CCS.

To select good w's for scalarisation, OLS exploits the observation that $V^*_X(w) = \max_{V^\pi \in X} w \cdot V^\pi$ is PWLC over the weight simplex. In particular, OLS selects only so-called *corner weights* that lie at the intersections of line segments of the PWLC function $V^*_X(w)$ that correspond to the value vectors found so far. For example, in Figure 6.5 on the left, $V^*_X(w)$ is indicated with bold line segments. There are two pay-off vectors in X, and there is one corner weight. The maximal potential error reduction that can be made by identifying a new pay-off vector $u(a)$ is at a corner weight Cheng (1988). The potential error reduction is denoted with dashed blue vertical lines, and is at present Δ. OLS scalarises the MOMDP at this corner weight and solves it using a single-objective solver, obtaining the optimal policy for that weight. The multi-objective

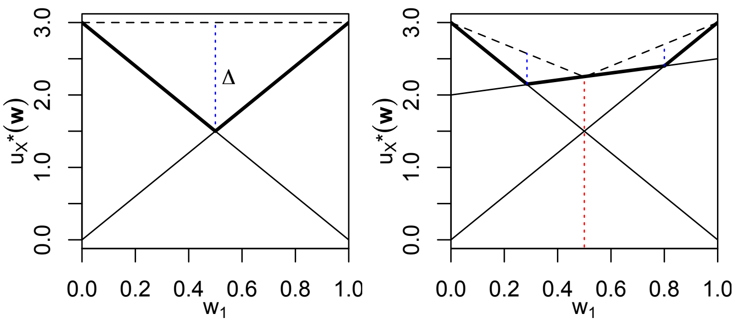

Figure 6.5: The scalarised value as a function of weights $V_X^*(w)$ (bold segments) for $X = \{(0,3),(3,0)\}$. There is one corner weight: $(0.5, 0.5)$ *(Left)*. Adding a new value vector, $(2.0, 2.5)$, to X, thereby improving $V_X^*(w)$ *(Right)*.

value of this policy, V^π, improves over $V_X^*(w)$ at that corner weight, as shown on the right-hand side of Figure 6.5 indicated by the red dashed line on the right. By continuing to find new corner weights and solving scalarised MOMDPs for corresponding to these corner weights, OLS is guaranteed to produce an exact CCS after solving a finite number of single-objective problems. That is, when no improvements are found at any remaining corner weights, the possible error reduction becomes 0, and a CCS is found. When there are multiple corner weights, OLS first solves the scalarised problem with the highest possible error reduction Δ.

After solving a scalarised MOMDP, for a given w, a policy w is obtained. If a standard implementation of dynamic programming would be used, the single-objective policy value V_w^π would also be obtained. Because OLS requires V^π however, this would mean that policy evaluation has to be used to obtain this multi-objective value. Therefore, in this implementation, an improvement is made over standard DP which is called *scalarised dynamic programming* (SDP). SDP keeps track of the multi-objective value vectors while maximising the value for w, thereby preventing having to perform separate policy evaluation steps.

6.4.3 Equilibrium Strategy

This study proposes to find equilibrium strategies that maximise expected revenue in the long run. For a given CCS with input review ratio ρ, an equilibrium strategy $\pi(\rho)$ that achieves ρ can be found as follows. Each deterministic stationary policy $\pi \in$ CCS has a corresponding value vector $\mathbf{V}^\pi$. Hence each policy can be represented in terms of a value vector. This leads to a three-dimensional value space CCS$'$ with revenue, positive reviews, and negative reviews as axes:

$$\text{CCS}' = \{\mathbf{V}^\pi \,|\pi \in \text{CCS}\}\,. \tag{6.6}$$

Consider Figure 6.6 below, where a sketch of CCS$'$ is given by the convex non-linear surface. This surface represents all potential optimal value vectors, covering all review ratios. However, not all value vectors need to be considered: the space CCS$'$ can be reduced to the value vectors $\mathbf{V}^\pi \in$ CCS$'$ for which the review ratio equals ρ. In order to do so, consider the hyperplane $\mathcal{H}$ for which the review ratio equal ρ:

$$\mathcal{H} = \left\{ x \in \mathbb{R}^3 \,\middle|\, \frac{x_2}{x_2 - x_3} = \rho \right\} = \left\{ x \in \mathbb{R}^3 \middle|\, x_2(\rho - 1) - \rho x_3 = 0 \right\}.$$

A sketch of the hyperplane is given in Figure 6.6. The reduced value vector space, consisting of value vectors in CCS$'$ for which the review ratio is equal to ρ, is then given by the intersection CCS$' \cap \mathcal{H}$. In Figure 6.6 the intersection CCS$'$ and $\mathcal{H}$ is emphasised by the black thick line.

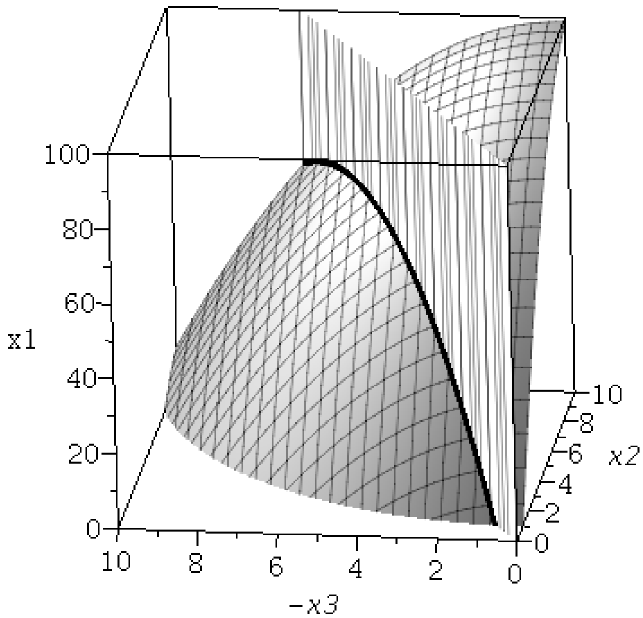

Figure 6.6: Visualisation of the set CCS$'$ of potentially optimal value vectors (convex non-linear surface), the hyperplane $\mathcal{H}$ of value vectors for which the review ratio is equal to ρ (hyperplane cutting CCS$'$), and their intersection CCS$' \cap \mathcal{H}$, the set of all potentially optimal value vectors for which the review ratio equals ρ (thick black line).

Hence the value vector $\mathbf{V}^\pi \in \text{CCS}' \cap \mathcal{H}$ with optimal revenue gives the solution to the problem at hand:

$$\pi(\rho) = \arg\max_{\pi \in \text{CCS}} \left\{ V_1^\pi \,\middle|\, \mathbf{V}^\pi \in \text{CCS}' \cap \mathcal{H} \right\}.$$

A closer look at the solution space provides insight in the evaluation of the optimal solution $\pi(\rho)$. First note that CCS, and therefore CCS$'$, is a convex set consisting of a finite number of faces. Therefore, CCS$' \cap \mathcal{H}$ is a line in $\mathbb{R}^3$ consisting of a finite number of line segments, corresponding to faces of CCS$'$ that intersect with $\mathcal{H}$. Hence $\pi(\rho)$ is one of the corners of those line segments.

Now that an equilibrium strategy $\pi(\rho)$ can be found for every ρ, the goal is to find the strategy π^*, with corresponding target review ratio ρ^*, that optimises revenue:

$$\pi(\rho^*) = \arg\max_{\pi(\rho)} V_1^{\pi(\rho)}.$$

Two computational challenges arise in these evaluations:

1. Not all target review ratios are attainable. The CCS with input ρ can attain the review ratios $[\underline{\rho}, \overline{\rho}]$. If $\rho \notin [\underline{\rho}, \overline{\rho}]$, then no equilibrium strategy exists and ρ is not

a feasible target review ratio. Only feasible target review ratios need to be considered. However, it is not clear beforehand which values of ρ are feasible.

2. Optimising over all target review ratios is challenging. The set of feasible target review ratios is continuous, and in order to find the optimal revenue corresponding to a certain target review ratio ρ the whole procedure described in this section has to be followed. Under these deliberations the authors propose two approaches to deal with the continuous variable ρ. The first approach descretises the set of attainable values of ρ, denoted by $\mathcal{P} \subset \mathbb{R}$, $|\mathcal{P}| < \infty$. Then, an optimal policy $\pi(\rho^*)$ over $\mathcal{P}$ is used:

$$\rho^* = \arg\max_{\rho \in \mathcal{P}} V_1^{\pi(\rho)}. \tag{6.7}$$

The second approach is iterative and of stochastic nature. Each iteration an arbitrary ρ is selected from the set of feasible values for ρ, according to some distribution. If $V_1^{\pi^*(\rho)} > V_1^{\pi^*(\rho')}$, with ρ' the best target review ratio found so far, then update ρ' with ρ. Continue until some stopping criterion is hit.

6.5 Numerical Examples

In this section numerical results are provided of an implementation of the multi-objective revenue management model. The two main goals of these examples are 1) to show how to interpret and use the convex coverage set; and 2) to numerically validate equilibrium strategies. Examples of realistic size are used.

All examples use the following set-up. Part of the parameters that are used are from the example of Section 6.2, and part is based on the example used in the numerical results in Section 3.5 and Talluri & van Ryzin (2004a). Let $n = 10$ be the number of products sold with corresponding price vector

$$r = (240, 220, 190, 160, 120, 112, 96, 80, 74, 70).$$

Overbooking is allowed up to 20% of the total capacity $C = 200$. Time is discretised to $T = 1000$. The demand, cancellation rate, and purchase probabilities are independent from the time period t. The base demand per time unit is equal to $\bar{\lambda} = 0.2$. The cancellation rate is assumed to be $\gamma = 0.0004$. The values from the example of Section 6.2 are used for both review probabilities and purchase probabilities, including the four different scenarios.

6.5.1 One Instance: Policy Analysis

This example examines the results for scenario 1 from the example of Section 6.2, where the current review ratio is $\rho = 0.6$. The resulting convex coverage set is presented in Table 6.2. Besides the total expected revenue, the positive reviews, and the negative reviews, also the resulting review ratio ρ is given. The solutions are ordered by revenue. The value space is three dimensional (revenue, positive reviews, and negative reviews). To provide a graphical representation of the value space, consider the three faces presented in Figures 6.7, 6.8, and 6.9. In each figure the convex coverage set that follows from revenue and review ratio is also given. Moreover, Figure 6.10 provides a visual representation of review ratio versus revenue.

Two observations of interest can be made from this table and accompanied figures. First, by sacrificing revenue the review ratio can be increased, as is expected. Optimising revenue results in total expected revenue of 27098.09 and review ratio of 0.57, while optimising the review ratio yields total expected revenue of only 7493.76, a 72% decline, and review ratio of 0.9, a 56% increase. Sacrificing revenue does not necessarily lead to a higher review ratio though. For example, the optimal solution corresponding to the sixth row in Table 6.2 leads to an expected revenue of 15137.33 and review ratio 0.84. The fifth row shows that by sacrificing revenue to 14252.8 the number of positive reviews can be increased, but the number of negative reviews is increased. In this case it leads to a decrease in review ratio to 0.81. Therefore, the procedure in Section 6.3 needs to be used to find an equilibrium strategy. In this case, The equilibrium of $\rho^* = 0.6$ has optimal revenue of 27027.18, a 3.18% increase with respect to solely optimising revenue (at $\rho = 0.57$).

Revenue	Positive reviews	Negative reviews	Rating ρ
0	0	0	-
7493.76	33.01	3.75	0.90
12513.99	42.88	9.28	0.82
12595.63	41.42	6.30	0.87
14252.80	42.48	9.77	0.81
15137.33	41.09	7.57	0.84
16021.48	41.67	10.38	0.80
17983.01	40.40	11.16	0.78
18693.24	39.78	10.61	0.79
19985.47	38.61	9.99	0.79
21719.66	36.79	11.74	0.76
23528.79	33.55	12.81	0.72
24686.86	26.26	12.34	0.68
25194.50	29.93	13.87	0.68
26943.83	25.33	15.07	0.63
27098.09	21.11	15.62	0.57

Table 6.2: Convex coverage set for $\rho = 0.6$.

A second observation is that the positive reviews tend to increase as revenue decreases, and negative reviews tend to decrease as revenue decreases. This is in concordance with the set-up of the scenarios, where higher prices lead to less positive and more negative reviews. It can also be seen in Figures 6.7 and 6.8. However, there is no strict increase or decrease. This can be explained by the trade-off between positive and negative reviews with respect to revenue. There is however a difference in behaviour between positive and negative reviews. At first positive reviews tend to increase as revenue decreases. Policies are selected that give slightly less revenue, but provide more positive reviews and less negative reviews. However, at some point revenue can only be increased more if at times no products are offered, to the point that no products are offered any more. This leads to the zero-solution: if no products are offered, no revenue is earned and no reviews are given. This in fact is an optimal policy when only negative reviews are the objective. A side-effect is that at some point, when revenue decreases, also positive reviews decrease, because of lack of offer. For example, consider the solution with 12513.99 revenue, 42.88 positive reviews, and 9.28 negative reviews; compared to revenue of 7493.76, 33.01 positive reviews, and 3.75 negative reviews. Note that the reviews score does increase from 0.82 to 0.90.

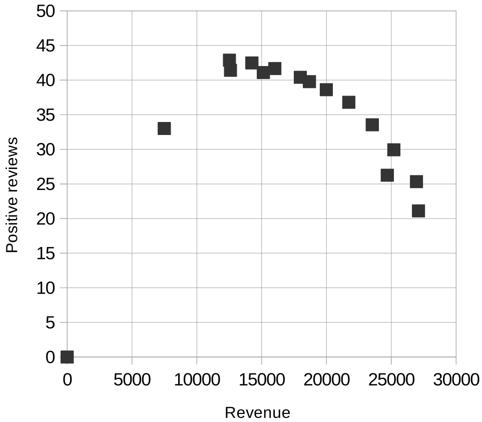

Figure 6.7: Convex coverage set plot of revenue and number of positive reviews.

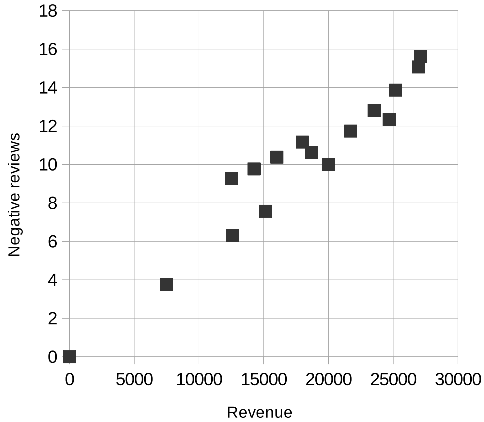

Figure 6.8: Convex coverage set plot of revenue and number of negative reviews.

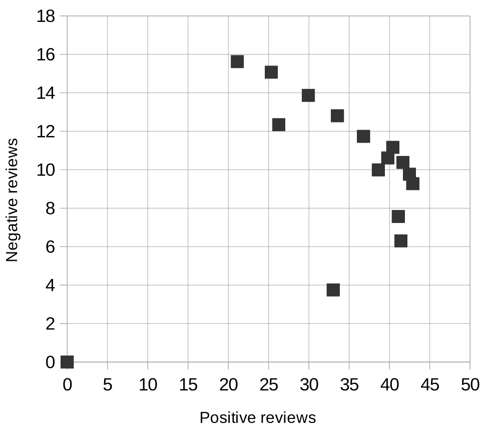

Figure 6.9: Convex coverage set plot of number of positive and negative reviews.

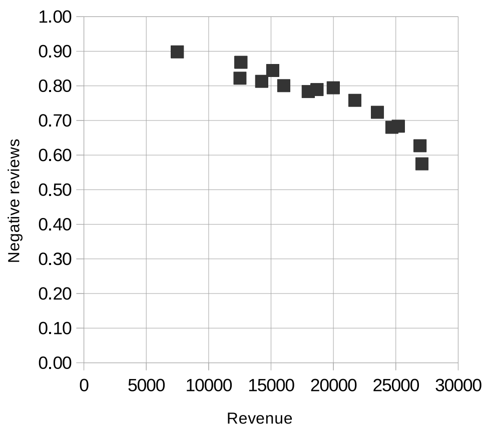

Figure 6.10: Plot of revenue and review ratio resulting from CCS.

6.5.2 Equilibria

An important application of the revenue-review trade-off model is to make decisions *now* that positively influence *future* purchase behaviour and revenues. By sacrificing revenue now in order to get better reviews, future revenues can be increased substantially. One way to do this is to consider the equilibrium strategies described in Section 6.3 (other approaches that are both tractable and indicate how long-term revenue can be increased are not known to the authors). Each of the four scenarios of the example of Section 6.2 is considered. The optimal revenue that results from the target review ratio can be calculated using stochastic mixture policies on the convex coverage set of revenue and reviews, as is shown in Section 6.3. For example, the equilibrium strategy for scenario 1 with $\rho = 0.6$ yields an expected revenue of 27027.18. Just optimising revenue, on the other hand, leads to an equilibrium policy that yields an expected revenue of 26194.23. By observing the policy space in Table 6.2, an initial revenue *loss* of 0.26% is incurred (from 27098.09 if solely revenue is optimised). However, on the long term this leads to a revenue *increase* of 3.18%.

The question is what target review ratio will lead to maximal revenue in the long term. Therefore the following experiment is conducted. For all four scenarios the total expected revenues corresponding to the equilibrium policy are evaluated for a range of target review ratios. For convenience, losses for some higher target reviews are omitted. The results are presented in Figure 6.11. Observe that the curves are not smooth, nor indeed is there a reason why they should be.

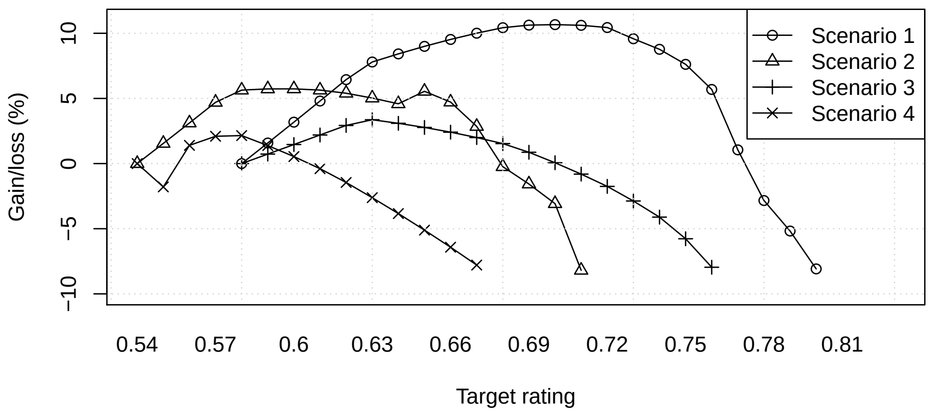

Figure 6.11: Revenues corresponding to equilibrium policies.

The four scenarios give different results, though the prevailing observation is that there is a lot to gain by optimising to target review ratio instead of solely optimising revenue. Scenario 1, where both the effect of review ratio on purchases and of purchases on review ratio are large, shows a structural increase in revenue of 11% when an optimal target review ratio is selected. If the demand effect is small, but the effect of review probabilities high, in scenario 2, then optimising to target review ratio leads to a 6% structural increase in revenue. The other scenarios 3 and 4, with the effect of review probabilities low, show a decent increase in revenue of about 2-3%.

Ye et alii (2011) found that a 10% increase in rating increases online bookings by more than 5%. In terms of revenue, our example gives similar results. Scenario 1 induces a revenue increase of 10.66% with an increase in review ratio of 20.69%. The increase in revenue is higher that in the results by Ye et alii (2011), but so is the increase in review ratio. In scenario 2 a similar observation is made: 5.73% increase in revenue and 11.11% increase in review ratio. Scenario 3 and 4, however, show slightly poorer results. Scenario 3 yields an increase of 3.36% in revenue with 8.62% increase in review ratio and scenario 4 an increase of 2.15% in revenue and 7.41% in review ratio. This can be explained by the low impact of review ratio on demand in these scenarios.

6.6 Concluding Remarks

In this chapter a novel revenue management model is introduced that captures the trade-off between revenue and reviews. Revenue management strategies influence the perception of customers, which results in a changing review ratio. On the other hand, review ratio influences buying behaviour. The formulated model captures the long-term effect of optimising revenue according to a target review ratio.

A new solution method to approach a problem of such complexity is introduced to optimise revenue in the long run. The methodology builds on recent developments in multi-objective Markov decision process theory, and contributes to this body of literature. Because the policy space is restricted to policies where revenue is optimised such that the target review ratio remains constant, the full problem can be reduced to a series of multi-objective Markov decision problems.

Our numerical studies show how to interpret the solution space, the convex coverage set, or a single multi-objective MDP in the series. Moreover, results of the equilibrium strategies of the successive model show that revenue improvements of up to 11% are achievable if reviews are taken into account in the optimisation process, instead of the sole objective of revenue. All results, featuring different scenarios, suggest revenue increases of at least 2%. In practical terms for theatre, this leads to a significant increase in revenue that can reach into the millions annually.

The results of this study have several implications that suggest topics for further research. First, the model can be used to identify the effect of improving facilities of the theatre to revenue and review ratios. Second, the model can be extended to a network setting for applications with such structure, like hotels. However, this is not straightforward, and it increases the dimensionality of the problem to the extend that is intractable. Next, Chapter 7 introduces and analyses a choice-based network model that includes reviews in the optimisation process.

CHAPTER 7

Choice-Based Network Revenue Management under Online Reviews

This chapter proposes a choice-based network revenue management (RM) model that integrates the effect of reviews. The dependency between reviews and revenue is two-fold: customers write reviews based on their price/quality perception, and reviews impact sales. A complicating factor in this model is that the effects of reviews are delayed, e.g., by sacrificing revenue now in order to get better reviews, long-term revenue can be increased. Faced by the complexity of the model, two heuristics are proposed, one of which uses robust optimisation techniques. Numerical results show a 3.5-5.2% improvement when reviews are taken into account. Moreover, the impact of reviews is greater under low demand intensity than under high demand intensity.

7.1 Introduction

A recent development that impacts sales of numerous industries is that customers are given the opportunity to share their experience with other potential clients via websites like Booking.com, Expedia, or Tripadvisor. Examples of industries where this is common today are hotels, airlines, travel agencies, car rentals, and short-term storage space leases. Evidence from literature shows that customers are highly influenced by reviews in their purchasing process (Pan & Zhang, 2011; Park & Lee, 2009; Yoo & Gretzel, 2011). In particular, negative information has relatively more impact than positive information (Sparks & Browning, 2011); positive reviews improve the confidence and willingness to buy at a hotel (Vermeulen & Seegers, 2009); and more recent reviews have more impact than older reviews (Pavlou & Dimoka, 2006; Vermeulen & Seegers, 2009; Ye et alii, 2011). In order to increase demand it is therefore beneficial for a company to get more positive reviews and try to avoid negative reviews. On the other hand, the price/quality perception of the customer impacts the reviews. This study proposes a decision model that maximises revenue in the long run in a network set-up, where reviews are implicitly incorporated.

This chapter is based on Sierag (2016a).

This study considers the problem of deciding which products should be offered at what time, with the objective of maximising revenue, in a network setting where the demand depends on past reviews and customer choice preferences are taken into account. In particular, the reviews are modelled as a feedback mechanism: on the one hand demand depends on reviews, and on the other hand, reviews depend on the price/quality perception of customers. This feedback mechanism was introduced in Chapter 6, where a a *single-leg* choice-based RM problem under reviews is considered. In line with Chapter 6, the focus is on reviews rather than ratings, since recent reviews impact sales and ratings tend to remain more or less constant. Demand of future arrivals in the planning horizon depends on past reviews as well as reviews that are released during the planning horizon. This highly complex stochastic problem is intractable, and this chapter proposes two heuristics to solve the problem. The first heuristic is a deterministic variant of the problem, which is shown to be an upper bound to the stochastic counterpart and converges to the optimal value when demand and capacity are scaled. The second heuristic uses robust optimisation techniques to deal with uncertainty in reviews. Numerical results in Section 7.5 show that taking reviews into account leads to a revenue improvement of 3.5-5.2%. Moreover, the results show that 1) reviews have more impact when the demand is low; and 2) that small hotels are more effected by the review mechanism than larger hotels.

This chapter builds on the work in Chapter 6, where we propose a similar feedback mechanism to deal with reviews in a single-leg environment. The solution methods of Chapter 6 cannot be extended to a network setting, and therefore this study is essential. This chapter makes the following research contributions: 1) a model that incorporates reviews in the optimisation process in a network setting (Section 7.2); 2) a deterministic variant of the intractable stochastic problem, which is shown to be an upper bound and converges asymptotically to the optimal revenue (Section 7.3); 3) Two heuristics to solve the stochastic problem: one based on the deterministic problem and one based on robust optimisation methods (Section 7.4); 4) numerical experiments indicating that including the review feedback mechanism lead to higher long-term revenue (Section 7.5). Finally, implications for research and practice are discussed in Section 7.6.

7.2 Stochastic Model

This section introduces the full stochastic problem of optimising long-term revenue in a network setting under customer choice behaviour and reviews. For clarity of presentation, this chapter is written in the context of hotels, since the combination of reviews and the network structure, in the form of multiple night stays, is very typical for this branch. However, the models and analysis are more general and apply to all network RM problems where demand depends on reviews and reviews depend on the experience of the customer.

7.2.1 Model Description

Consider a hotel with identical rooms. The hotel manager wants to sell the rooms for $m \in \mathbb{N}$ nights. The capacity C_i for night $i \in \{1, \ldots, m\}$ is allowed to differ per night, for example due to renovation of some rooms. The firm sells the rooms in the form of n *products*, where product $j \in N = \{1, \ldots, n\}$ is a combination of one or more rooms, possibly for multiple nights, a reward r_j, and certain conditions (like a cancellation policy). In the hotel context used in this chapter, the room consumption is determined by an arrival night i and a *length of stay* (LOS). The reward r_j may include the expected revenue from the room price and other sources of revenue, such as food and beverage, spa and fitness, and casino revenues. Let $A = (a_{ij}) \in \mathbb{R}^{m \times n}$ be the incidence matrix of the products and rooms, where a_{ij} is the number of rooms of night i that is consumed by product j. Furthermore, a customer that purchased product j has a probability of q_j^p that he will write a positive review about its purchase and an independent probability of q_j^n that he will write a negative review. Hence he will write both a positive and negative review with probability $q_j^p q_j^n$. Parameters are chosen such that $q_j^p + q_j^n + q_j^p q_j^n \leq 1$, so the probabilities are well defined. Note that in practice it is natural that customers write both positive and negative reviews at the same time, as a customer can write about both positive and negative aspects of the experience.

The products are sold continuously over T time units. Arrival nights occur also within the booking horizon, so products may perish before the end of the booking horizon is met. At certain moments in time T_d ($1 \leq d \leq D$, $D \in \mathbb{N}$) the reviews are updated: all products for which consumption ended after T_{d-1} but before or at T_d, denoted by the set $N_d \subset N$, are published at time T_d. The period between T_{d-1} and T_d is denoted by *period* d. Let Q_d^p and Q_d^n be the number of positive reviews and negative reviews, respectively, that are published at time T_d (with $T_0 = 0$, the start of the booking horizon). The initial positive and negative reviews are given by Q_0^p and Q_0^n, respectively.

Demand is influenced by reviews, either positive or negative (Pan & Zhang, 2011; Park & Lee, 2009; Sparks & Browning, 2011; Vermeulen & Seegers, 2009; Yoo & Gretzel, 2011). More recent reviews have more impact on demand than older reviews (Pavlou & Dimoka, 2006; Vermeulen & Seegers, 2009; Ye et alii, 2011). To capture this effect, in accordance with Chapter 6, the positive and negative reviews are discounted by a factor $\alpha \in (0, 1)$. The *discounted reviews* for time T_d are given by

$$\tilde{Q}_d^p := \sum_{d'=0}^{d} \alpha^{T_d - T_{d'}} Q_{d'}^p, \qquad \tilde{Q}_d^n := \sum_{d'=0}^{d} \alpha^{T_d - T_{d'}} Q_{d'}^n. \tag{7.1}$$

Assume that customers in period d arrive according to a Poisson process with rate λ_d, where λ_d depends on the the reviews:[1]

$$\lambda_d = (\bar{\lambda}_d + \beta^p \tilde{Q}_{d-1}^p + \beta^n \tilde{Q}_{d-1}^n)^+, \tag{7.2}$$

where $\bar{\lambda}_d \in \mathbb{R}$ is the base arrival rate and $\beta^p, \beta^n \in \mathbb{R}$. The operator $(x)^+$ is defined by $(x)^+ = \max\{0, x\}$, which enforces demand to be non-negative. In this set-up, the LOS and arrival night are included in the product $j \in N$. In accordance with Vermeulen & Seegers (2009) it is assumed that positive reviews have positive effect on demand and negative reviews have a negative effect on demand, which translates to $\beta_d^p > 0$ and $\beta_d^n < 0$. In line with findings in literature (Sparks & Browning, 2011), the effect of negative reviews is greater than positive reviews: $\beta_d^p < -\beta_d^n$.

The results of this chapter can be derived when the parameters β^p, β^n depend on time, but for clarity it is assumed that β^p, β^n are equal for all time periods. Let $q_j = \beta^p q_j^p + \beta^n q_j^n$, let $Q_d = \beta^p Q_d^p + \beta^n Q_d^n$, and let $\tilde{Q}_d = \beta^p \tilde{Q}_d^p + \beta^n \tilde{Q}_d^n$. Then the demand can be rewritten as

$$\lambda_d = \bar{\lambda}_d + \tilde{Q}_{d-1}.$$

[1] It is possible to let the arrival rate λ_d be time dependent. The analysis and results remain the same, but for ease of notation a constant arrival rate is used throughout this chapter.

Continuously in time the hotel manager decides which subset $S \subset N$ to offer to arriving clients. According to this *offer set* S an arriving client purchases product j with probability $P_j(S)$ or declines from purchasing anything with probability $P_0(S)$. In accordance with Markov decision process literature a feasible solution π to the problem at hand is called a *policy.* Let Π be the set of all policies, where non-deterministic policies are allowed. The policy $\pi^* \in \Pi$ that optimises expected revenue is called an *optimal policy.*

7.2.2 Problem Formulation

Let $N(S_\pi(t)) \in \mathbb{N}^n$ be the stochastic process of the vector of purchases at time t under policy π, where $S_\pi(t) \subset N$ is the offer set corresponding to π. The problem statement is then given by

$$\begin{aligned} &\max_{\pi \in \Pi} E\left[\int_0^T r^\top N(S_\pi(t))dt\right] \\ &\text{s.t.} \int_0^T AN(S_\pi(t))dt \le C, \\ &\quad S_\pi(t) \subset N, \qquad\qquad \forall t \in [0,T]. \end{aligned} \tag{7.3}$$

To solve this problem, time is discretised into T time units, such that the probability that more than one event occurs in one time unit is small. The probability that a customer arrives in time period t ($1 \le t \le T$) is given by λ_t. Define the state space by $X \times Y$, where X is the set of all possible occupation scenarios, i.e., $X = \{x \in \mathbb{N}^m \,|x \le C\}$, and Y is the set of all possible outcomes for reviews. Since there are only a finite number of decision moments, $X \times Y$ is finite, however large. Note that λ_t depends on the reviews y; therefore, denote the arrival probability by $\lambda_t(y)$, $y \in Y$. Denote the revenue to go at time t in state (x, y) by $V_t(x, y)$. Under these considerations an optimal policy can be found by solving the following Bellman equation:

$$\begin{aligned} V_t(x,y) = &\max_{S \subset N}\left\{\lambda_t(y)\sum_{j\in S} P_j(S)\left[r_j + V_{t+1}(x + A_j, y + q_j) - V_{t+1}(x,y)\right]\right\} \\ &+ V_{t+1}(x,y). \end{aligned} \tag{7.4}$$

The high dimensionality of the state space (which requires at least m dimensions to keep track of the consumed rooms per night) makes this stochastic problem intractable, similar to other choice-based network studies (Bront et alii, 2009; Hosseinalifam et alii, 2016; Liu & van Ryzin, 2008; Meissner & Strauss, 2012; Meissner et alii, 2013; Strauss & Talluri, 2012). Therefore, approximations have to be considered. The next sections are dedicated to performance measures, including an upper bound, and two heuristic approaches.

7.3 Deterministic Model

This section proposes a deterministic counterpart of the stochastic problem (7.3). The objective value is shown to be an upper bound to the stochastic problem. Furthermore, the upper bound is shown to be asymptotically tight when demand and capacity are scaled.

7.3.1 CRLP Formulation

Consider the deterministic variant of the stochastic problem, where demand and reviews are deterministic and continuous variables. Under deterministic demand a mixed integer linear program formulation is proposed, called *choice-based review linear program* (CRLP). In other studies concerning choice-based network RM a similar linear program is derived, called *choice-based deterministic linear program* (CDLP) (Gallego et alii, 2004; Liu & van Ryzin, 2008). However, the CRLP does not follow directly from the CDLP, since the demand parameter is not a constant any more, but depends on the past reviews. A crucial step in deriving the CRLP is the introduction of the decision variables $x(S,d)$, representing the total number of clients that are offered set S during period d. This contrasts earlier work (Gallego et alii, 2004; Liu & van Ryzin, 2008), where $t(S)$ is used as a decision variable, representing the *time* that set S is offered rather than the total number of clients that was offered set S. While using $t(S)$ is an appropriate decision variable in their studies, a direct implementation to the review model leads to a non-linear program.

The objective of CRLP is the expected reward. A customer who is offered set $S \subset N$ leads to an expected reward of $\sum_{j\in S} r_j P_j(S)$. With $x(S,d)$ as decision variables, the objective of CRLP is then given by

$$\sum_{d=1}^{D} \sum_{S\subset N} x(S,d) \sum_{j\in S} P_j(S) r_j. \tag{7.5}$$

When the offer set $S \subset N$ is offered to a customer, the expected resource consumption is equal to $AP(S)$. The capacity constraints are therefore given by

$$\sum_{d=1}^{D} \sum_{S\subset N} x(S,d) AP(S) \leq C. \tag{7.6}$$

To model the demand and time constraints, observe that the total demand over all sets $S \subset N$ offered during period d is upper bounded by the demand rate of period d times the length of period d:

$$\sum_{S\subset N} x(S,d) \leq (\bar{\lambda}_d + \tilde{Q}_{d-1})^{+}(T_d - T_{d-1}). \tag{7.7}$$

Note that the demand rate is non-negative. For ease of notation, assume that $T_d - T_{d-1} = 1$ for all $1 \leq d \leq D$. Using the definition of discounted reviews this then leads to[2]

$$\sum_{S\subset N} x(S,d) \leq \left(\bar{\lambda}_d + \alpha^{d-1} Q_0 + \sum_{d'=1}^{d-1} \alpha^{d-d'-1} \sum_{j\in N_{d'}} \sum_{S\subset N} \sum_{d''=1}^{d'} x(S,d'') P_j(S) q_j \right)^{+}, \tag{7.8}$$

for all $1 \leq d \leq D$. For notational convenience, let $\tilde{\lambda}_d = \bar{\lambda}_d + \alpha^{d-1} Q_0$, and define $\mu\colon \{S \subset N\} \times \{1,\ldots,D\} \times \{1\ldots,D\} \to \mathbb{R}$ by

$$\mu(S,d,d') := \sum_{j\in S} P_j(S) q_j \left(\sum_{d''=d'}^{d-1} \alpha^{d-d''-1} \mathbb{I}\{j \in N_{d''}\} \right), \tag{7.9}$$

for all $S \subset N$ and for all $1 \leq d, d' \leq D$. The constant $\mu(S,d,d')$ can be interpreted as the expected additional demand in period d per unit of demand in a previous period d' that

[2] Note that $\alpha^{d-d'}$ has to be replaced with $\alpha^{T_d - T_{d'}}$ in Equation (7.8) when $T_d - T_{d-1} = 1$ does not hold.

set S was offered. Hence $\mu(S,d,d')x(S,d')$ is the expected additional demand for period d that follows from offering set S in period d'. The demand constraints (7.8) can now be rewritten in a more convenient form:

$$\sum_{S\subset N} x(S,d) \le \left(\tilde{\lambda}_d + \sum_{S\subset N}\sum_{d'=1}^{d-1} x(S,d')\mu(S,d,d')\right)^+, \tag{7.10}$$

for all $1 \le d \le D$. The CRLP is then given by

$$\begin{aligned} \max_{\substack{x(S,d):\\ S\subset N, 1\le d\le D}} & \sum_{d=1}^{D}\sum_{S\subset N} x(S,d)\sum_{j\in S} P_j(S)r_j \\ \text{s.t.} & \sum_{d=1}^{D}\sum_{S\subset N} x(S,d)AP(S) \le C, \\ & \sum_{S\subset N} x(S,d) \le \left(\tilde{\lambda}_d + \sum_{S\subset N}\sum_{d'=1}^{d-1} x(S,d')\mu(S,d,d')\right)^+. \end{aligned} \tag{7.11}$$

The demand constraints contain a maximisation function, such that the problem can be solved as a mixed integer linear program (MILP). Mathematical software programs such as CPLEX and Gurobi are capable of handling quite large MILP instances.

7.3.2 Upper Bound and Asymptotic Optimality

Let V^{CRLP} be the optimal objective value of the CRLP (7.11) and let V^* be the optimal objective value of the stochastic problem (7.3). Proposition 7.1 shows that V^{CRLP} is an upper bound to V^* (similar to Gallego et alii, 2004, Proposition 2, and Liu & van Ryzin, 2008, Proposition 1).

Proposition 7.1. $V^{\text{CRLP}} \ge V^*$.

Proof. Let π^* be an optimal policy of the stochastic problem (7.3) and let $S_{\pi^*}(t,\mathcal{F}_t)$ be the stochastic process of sets offered at time t under π^* and $\mathcal{F}_t$ the history of the system up to time t. Since π^* is feasible to (7.3), it holds in a path wise fashion that

$$\int_0^T AN(S_{\pi^*}(t,\mathcal{F}_t))dt \le C,$$

and therefore the expectation is finite:

$$E\left[\int_0^T AN(S_{\pi^*}(t,\mathcal{F}_t))dt\right] \le C < \infty.$$

Note that $AN(S(t))$ is non-negative, so Fubini's Theorem applies, with the following result:

$$\int_0^T AE[N(S_{\pi^*}(t,\mathcal{F}_t))]dt = E\left[\int_0^T AN(S_{\pi^*}(t,\mathcal{F}_t))dt\right] \le C.$$

Next, define $x_{\pi^*}(S,d)$ by

$$x_{\pi^*}(S,d) = E\left[\lambda_d \int_{T_{d-1}}^{T_d} \mathbb{I}\{S_{\pi^*}(t,\mathcal{F}_t) = S\}dt\right],$$

for all $S \subset N$ and $1 \leq d \leq D$, with

$$\lambda_d = \left(\tilde{\lambda}_d + \sum_{S\subset N}\sum_{d'=1}^{d-1} x(S,d')\mu(S,d,d')\right)^+.$$

The expected number of clients that purchase product j is equal to the summation over S of the expected number of sales of product j, given the number of clients that were exposed to set S:

$$\sum_{d=1}^{D}\sum_{S\subset N} x_{\pi^*}(S,d)P_j(S) = \int_0^T E[N_j(S_{\pi^*}(t,\mathcal{F}_t))]dt \leq C.$$

Therefore, the constraints of CRLP (7.11) are satisfied, and π^* is a feasible solution to CRLP (7.11). Furthermore, the objective values of the stochastic problem (7.3) and CRLP coincide:

$$V^* = E\left[\int_0^T r^\top N(S_{\pi^*}(t,\mathcal{F}_t))dt\right] = \sum_{d=1}^{D}\sum_{S\subset N} x_{\pi^*}(S,d)\sum_{j\in S} P_j(S)r_j,$$

Since $x_{\pi^*}(S,d)$ is a feasible solution to CRLP, the corresponding objective value is bounded by the optimal objective value V^{CRLP}, which completes the proof. □

Consider the k-scaled problem instance, where demand, capacity, and initial reviews Q_0^p and Q_0^n are scaled by a factor $k \in \mathbb{N}$, i.e., the demand rate equals $k\lambda$, the capacity vector equals kC, and the initial reviews are given by kQ_0^p and kQ_0^n. Let V_k^* and V_k^{CRLP} be the optimal objective value of the k-scaled stochastic problem and CRLP, respectively. As $k \to \infty$, the objective values of the stochastic and deterministic values converge to V^{CRLP}, which is shown in Proposition 7.2 below (adjusted to our model from Liu & van Ryzin, 2008, Proposition 2).

Proposition 7.2. $\lim_{k\to\infty} \frac{1}{k}V_k^* = \lim_{k\to\infty} \frac{1}{k}V_k^{\text{CRLP}} = V^{\text{CRLP}}$.

Proof. Let $x^*(S,d)$ be an optimal solution to the unscaled CRLP problem (7.11). The objective value of the k-scaled CRLP is equal to k times the objective value of the unscaled CRLP. Also, the constraints of the k-scaled CRLP equal k times the constraints of the unscaled CRLP. Therefore, $kx^*(S,d)$ is an optimal solution to the k-scaled CRLP with optimal value kV^{CRLP}, so the second equality above holds.

Next, construct an optimal policy $\pi \in \Pi$ for the k-scaled stochastic problem from $x^*(S,d)$ as follows. In period d, offer set S a deterministic amount of time equal to

$$t_\pi(S,d) := x^*(S,d)/\lambda_d, \tag{7.12}$$

where λ_d is the deterministic demand rate given by

$$\lambda_d = \left(\tilde{\lambda}_d + \sum_{S\subset N}\sum_{d'=1}^{d-1} x(S,d')\mu(S,d,d')\right)^+. \tag{7.13}$$

This is the time that set S is offered during period d in the (k-scaled) CRLP. The order that the sets are offered is arbitrary. Let $D^k(S,d,t)$ be the random vector of product demand under set S over t time units in the k-scaled stochastic problem. So, for product j and offer-set S in period d, $D_j^k(S,d,t_\pi(S,d)$ is Poisson distributed with parameters $x^*(S,d)P_j(S)$. Under π not all demand is accepted, only the demand under $kx^*(S,d)P(S)$ is accepted. Let $N_\pi(S,d)$ be the accepted demand according to policy π:

$$N_\pi(S,d) := \min\{D^k(S,d,t_\pi(S,d)), kx^*(S,d)P(S)\}. \tag{7.14}$$

Due to the demand constraint of the k-scaled problem of the CRLP it holds that

$$\begin{aligned}\sum_{d=1}^{D}\sum_{S\subset N} AN_\pi(S,d) &= \sum_{d=1}^{D}\sum_{S\subset N} A\min\{D^k(S,d,t_\pi(S,d)), kx^*(S,d)P(S)\}\\ &\leq kC,\end{aligned}$$

so π is an admissible policy for the k-scaled stochastic problem. The objective value equals

$$\sum_{d=1}^{D}\sum_{S\subset N} r^\top N_\pi(S,d) = \sum_{d=1}^{D}\sum_{S\subset N} r^\top \min\{D^k(S,d,t_\pi(S,d)), kx^*(S,d)P(S)\}.$$

Letting $k\to\infty$ this leads to

$$\lim_{k\to\infty}\frac{1}{k}\sum_{d=1}^{D}\sum_{S\subset N} r^\top \min\{D^k(S,d,t_\pi(S,d)), kx^*(S,d)P(S)\} \tag{7.15}$$

$$= \lim_{k\to\infty}\sum_{d=1}^{D}\sum_{S\subset N} r^\top \min\{\frac{1}{k}D^k(S,d,t_\pi(S,d)), x^*(S,d)P(S)\}. \tag{7.16}$$

For all $j \in N$, $D_j^k(S,d,t_\pi(S,d))$ has the same distribution as $\sum_{y=1}^k D_{j,y}(S,d,t_\pi(S,d))$, with $D_{j,y}(S,d,t_\pi(S,d)) \sim \text{Pois}(x^*(S,d)P_j(S))$ i.i.d. for all $1 \leq y \leq k$. By the law of large numbers the sequence $\frac{1}{k}D^k(S,d,t_\pi(S,d))$ then converges to $x^*(S,d)P(S)$ a.s. as $k\to\infty$. Since the minimisation function is continuous, the continuous mapping theorem can be applied, which yields

$$\begin{aligned}&\lim_{k\to\infty}\sum_{d=1}^{D}\sum_{S\subset N} r^\top \min\{\frac{1}{k}D^k(S,d,t_\pi(S,d)), x^*(S,d)P(S)\}\\ &= \sum_{d=1}^{D}\sum_{S\subset N} r^\top \min\{x^*(S,d)P(S), x^*(S,d)P(S)\}\\ &= \sum_{d=1}^{D}\sum_{S\subset N} x^*(S,d)r^\top P(S)\\ &= V^{\text{CRLP}}.\end{aligned}$$

This completes the proof.

□

7.4 Solution Methods

In this section two solution methods are proposed to solve the stochastic problem (7.3). The first method, called CRLP, uses the solution to CRLP (7.11) as a strategy. The second method is motivated by the fact that the demand rate depends on past reviews, which entails uncertainty. A robust optimisation formulation is proposed to deal with this uncertainty, with the goal of providing an improved solution method. Finally, this section is concluded with a discussion on some computational challenges and opportunities.

7.4.1 CRLP Approximation

The first heuristic is a straightforward implementation of the outcome of CRLP. Let $x^*(S,d)$ be an optimal solution to CRLP. In term of time units, the optimal strategy $x^*(S,d)$ translates to offering set S for a total of $t(S,d) = x^*(S,d)/\lambda_d$ time units in period d. Since the strategy assumes deterministic demand, it dictates no specific order in which the sets are offered. Moreover, set S does not have to be offered continuously for $t(S,d)$ time units in period d: it may be offered a couple of different time segments within period d, as long as the total offer time equals $t(S,d)$.

However, in the stochastic model demand is stochastic while the capacity is limited. Each time a customer arrives, a purchase can only be accepted if the remaining capacity allows for it. By fixing the order in which sets S are offered in period d, for example by lexicographical ordering or ordering by expected revenue, products from sets S that are of higher order will be more likely to be purchased than lower order products, since by the time the lower order is reached, less capacity remains. This leads to a bias towards products of offer sets on top of the list and differs from core strategy of CRLP. An ordering might have a positive impact on performance, but this is not a straightforward process in general. Therefore, to overcome this problem and keep as close to the CRLP strategy as possible, a randomisation of offer sets is applied, in combination with splitting the period $t(S,d)$ in smaller parts. That is, the offer time $t(S,d)$ is split into K periods $t_k(S,d) = t(S,d)/K$ $(1 \le k \le K)$, such that the total time of offering set S in period d still equals $t(S,d)$. Then a random ordering is applied to all time segments $t_k(S,d)$ of all offer sets, per period d.

7.4.2 Robust CRLP

A crucial aspect of reviews is their impact on future demand. The CRLP assumes that the reviews are deterministic, and the CRLP approximation described in the previous section assumes that the demand rates are known. However, due to the stochasticity of the reviews, the actual demand rates may fluctuate. To overcome this problem, a robust version of the CRLP is proposed, where uncertainty is assumed in the outcome of the reviews. In particular, consider the demand constraint of CRLP (7.11):

$$\sum_{S\subset N} x(S,d) \le \left(\tilde{\lambda}_d + \sum_{S\subset N}\sum_{d'=1}^{d-1} x(S,d')\mu(S,d,d')\right)^+. \tag{7.17}$$

The demand in this constraint heavily depends on the outcome of the reviews, in the form of $\mu(S,d,d')$, which is stochastic. Small changes in the outcome of the reviews lead to different demand rates. To take this into account, assume that there is uncertainty in the parameters $\mu(S,d,d')$. For convenience, write $\mu(S,d,d')$ in terms of its nominal value $\bar{\mu}(S,d,d') \in \mathbb{R}$ and a *primitive factor* $\zeta(S,d,d') \in \mathbb{R}$:

$$\mu(S,d,d') = \bar{\mu}(S,d,d') + \zeta(S,d,d'). \tag{7.18}$$

The uncertain parameter $\zeta = \{\zeta(S,d,d')\}$ is assumed to lie in an uncertainty set

$$Z \subset \{S \subset N\} \times \{1,\ldots,D\} \times \{1,\ldots,D\}. \tag{7.19}$$

In Ben-Tal et alii (2009) it is shown that the uncertainty can be approached constraint-wise. The robust formulation of constraint (7.17) can then be reformulated as

$$\sum_{S\subset N} x(S,d) \leq \left(\tilde{\lambda}_d + \sum_{S\subset N}\sum_{d'=1}^{d-1} x(S,d')\left[\bar{\mu}(S,d,d') + \zeta(S,d,d')\right]\right)^+, \tag{7.20}$$

for all $\zeta \in Z$.

The choice of the uncertainty set Z impacts the tractability of the program. Tractable formulations of the robust counterpart for several standard uncertainty regions are provided by Ben-Tal et alii (2009). Although the robust counterparts of those standard uncertainty regions are denoted as tractable, some still involve solving non-linear programs, and often the number of constraints increases. In case of the CRLP, tractability is already an issue. Under these deliberations, consider the interval/box uncertainty set Z_∞, given by

$$Z_\infty = \{\zeta \mid \|\zeta\| \leq \rho\}, \tag{7.21}$$

which leads to the following robust counterpart (see Ben-Tal et alii, 2009):[3]

$$\sum_{S\subset N} x(S,d) \leq \left(\tilde{\lambda}_d + \sum_{S\subset N}\sum_{d'=1}^{d-1} x(S,d')\left[\bar{\mu}(S,d,d') - \rho\right]\right)^+. \tag{7.22}$$

The advantage of this uncertainty region is that the constraints remain linear and the number of constraints does not increase. However, using the same ρ for all variables is very conservative. Therefore we propose the following uncertainty set:

$$Z = \{\zeta \mid |\zeta(S,d,d')| \leq \rho(S,d,d')\}, \tag{7.23}$$

with $\rho(S,d,d') > 0$, $S \subset N$, $1 \leq d, d', D$. Under this uncertainty region the robust counterpart becomes

$$\sum_{S\subset N} x(S,d) \leq \left(\tilde{\lambda}_d + \sum_{S\subset N}\sum_{d'=1}^{d-1} x(S,d')\left[\bar{\mu}(S,d,d') - \rho(S,d,d')\right]\right)^+. \tag{7.24}$$

This adaptation of the interval/box uncertainty region is conservative but has two main advantages. First, it leads to a MILP which is tractable in the case of CRLP: no new integer variables are added and the number of constraints of the robust CRLP remains the same. Second, the nominal values $\bar{\mu}(S,d,d')$ may very from each other in orders of magnitude, but Z can take this into account by setting the bounds $\rho(S,d,d')$ relatively to the nominal value $\bar{\mu}(S,d,d')$.

[3] Note that $|x(S,d)| = x(S,d)$, since $x(S,d) \geq 0$.

7.4.3 Computational Challenges

The CRLP has $D2^n$ decision variables and $D+m$ constraints. With an exponential number of decision variables in n this problem is intractable, even for modest sizes. However, certain conditions on demand and branch-and-price techniques can be used to attempt to solve the problem. This paragraph examines both techniques in context of CRLP (7.11).

Segmentation under Disjoint Consideration Sets

An important step in improving the tractability of CRLP is to make assumptions on the demand that reduce the amount of decision variables. These assumptions may follow naturally from the problem at hand, for example from segmentation of customers. Segmentation is an important aspect of RM and one of the first steps of a RM implementation (Talluri & van Ryzin, 2004b, p. 579-585). Assume that customers are partitioned into $|L| \in \mathbb{N}$ segments, where L is the set of segments and each segment $l \in L$ has its own characteristics and preferences. According to these characteristics and preferences, the hotel can target a segment $l \in L$ by offering them certain products S_l. Customers may be segmented in such a way that products can be offered solely to one segment, without being exposed to another, through different booking channels. For example, transient clients might use booking websites, while group reservations are made through agencies directly with the hotel; and in Chapter 2 it is shown that business and leisure customers can be segmented by day of the week and time of the day: business clients tend to make a reservation during weekdays between 8:00 and 17:00, while leisure clients tend to make reservations during weekdays from 17:00 and in the weekend.

Now the assumption that reduces the number of decision variables is that segments have disjoint consideration sets N_l, $l \in L$. That is, a customer from segment $l \in L$ only considers a subset $N_l \subset N$ of the products, with $N_l \cap N_k = \emptyset$ for all $l, k \in L$. Although room consumption may be overlapping, the decision variables $x(S, d)$, for all $S \subset N$ and for all $1 \leq d \leq D$, can be replaced by $x(S_l, d)$, for all $S_l \subset N_l$, $l \in L$, and for all $1 \leq d \leq D$. This reduces the number of decision variables from $D2^n$ to $D \sum_{l \in L} 2^{|N_l|}$, a significant reduction. If the consideration sets N_l are of reasonable size, the number of decision variables can be contained and CRLP is tractable.

In other choice-based RM models, such as Kunnumkal & Topaloglu (2008), Liu & van Ryzin (2008), Meissner & Strauss (2012), and Vossen & Zhang (2015), a special case of segmentation under disjoint consideration sets is discussed and used to validate the models. In these works, the multinomial logit model is used as a choice model, and to build on these works the same choice model is used in the numerical section 7.5 of this chapter. However, we stress that the reduction of decision variables holds for any choice model in combination with segmentation under disjoint consideration sets.

Branch-and-Price

When the number of decision variables of CRLP is too large, branch-and-price techniques can be used to attempt to solve CRLP. Branch-and-price is a column generation technique combined with branching for MILPs and is strongly related to branch-and-cut methods. See Barnhart et alii (1998) for a general discussion of branch-and-price and Cóccola et alii (2015) for an application of branch-and-price to a ship routing and scheduling problem. In the following, a sketch of a branch-and-price application to CRLP is given.

The branch-and-price algorithm consists of two aspects: 1) a branching tree to deal with the integer variables, and 2) a column generation procedure, executed in each node of the tree, to determine whether to branch or to stop, according to some decision rules. The column generation procedure gives a *local upper bound* (LUB) for each node in the branching tree. If the strategy of this LUB is feasible to CRLP, then the LUB is compared with the *global lower bound* (GLB), the best feasible solution to CRLP found so far. If the LUB is lower than the GLB (whether it is feasible to CRLP or not), then no optimal solution will be found in this branch and it is explored no more. If the LUB is not feasible to CRLP and larger than the GLB, then this node is explored further by splitting it into two child nodes according to some decision rule.

In particular, the column generation procedure in a node of the branching tree is as follows. Instead of solving the CRLP *master problem* (MP) (7.11), a subset $N_d \subset N$, $d = 1, \ldots, D$, of offer sets is considered as decision variables (columns of CRLP). This *reduced master problem* (RMP) is tractable when $|N_d|$ is substantially smaller than $|N|$, for all $d = 1, \ldots, D$. The LP relaxation of the RMP leads to dual variables $\pi \in \mathbb{R}^m$ and $\rho \in \mathbb{R}^D$, corresponding to the capacity constraints and demand constraints, respectively. Next, consider the reduced costs that follow from the dual solution of the relaxed RMP, of all columns of the MP that are not in RMP. If any of those columns has positive reduced costs, one of those columns is selected and added to the RMP. To determine which column has positive reduced costs and can be added to the RMP, the following column generation *sub-problem* has to be solved:

$$\max_{\substack{S \subset N \\ 1 \leq d \leq D}} \left\{ (r^\top - \pi^\top A) P(S) - \rho_d + \sum_{d'=d+1}^{D} \rho_{d'} \mu(S, d', d) \right\}. \tag{7.25}$$

If the optimal value of sub-problem (7.25) is negative (i.e., if there is no column with positive reduced costs) and the relaxed RMP strategy is feasible to RMP (i.e, the relaxed integer variables are integers), then the strategy/LUB is compared with the GLB. If it is smaller than GLB, the exploration of this branch ends here.

Otherwise, if the optimal value of the sub-problem (7.25) is negative and the relaxed RMP solution is not feasible to RMP, branching is applied to the integer variables. If the LUB of the node is greater than GLB, then, according to some decision rule, an integer variable is selected to construct the next step in the branching tree. Otherwise the search in this branch ends here. This procedure is continued until an optimal strategy is found. The process always converges since the branching tree is finite.

Other papers on choice-based network RM also discuss column generation techniques (Gallego et alii, 2004; Liu & van Ryzin, 2008). However, they consider CDLP, which a linear program without any integer variables. Those procedures cannot be applied to CRLP, but the branch-and-price procedure described in this section can.

7.5 Numerical Results

This section provides some numerical experiments to illustrate the model and performance of the solution methods. The first experiment focuses on the performance of the robust solution method under different loads. The second experiment focuses on the weight of review probabilities. The third experiment considers different hotel sizes. The fourth and final experiment shows the performance under estimation errors of the demand parameters.

Suppose a manager wants to find a strategy for his hotel of size $C \in \mathbb{N}$ over an 8 week period. After each week reviews are released, so $D = 8$. Assume that rooms are sold over a 90-day booking horizon. For simplicity only two resource types per week are used, by grouping weekdays together: resource type 1 represents Monday-Friday, resource 2 represents Friday-Monday. See Figure 7.1 for an illustration.

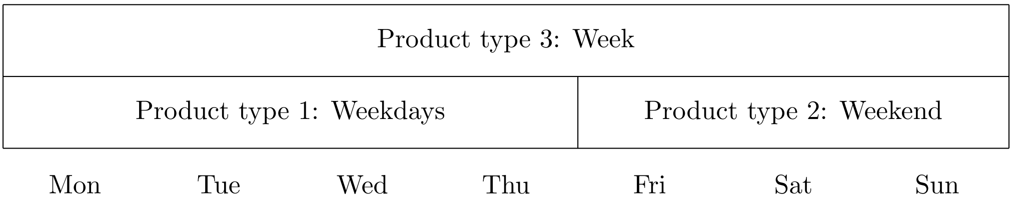

Figure 7.1: The three different product types per week.

On these resources, three product types are constructed: product type 1 consumes one unit of resource type 1, representing a midweek stay; product type 2 consumes one unit of resource type 2, representing a weekend stay; product type 3 consumes one unit of resource type 1 and one unit of resource type 2, representing a whole week stay. Each product type has three price levels, see Table 7.1.

		Product		
Product type		1	2	3
1	(weekdays)	360	480	600
2	(weekend)	250	340	400
3	(week)	500	750	900

Table 7.1: Prices per product per product type.

An arriving customer is only interested in one product type in one particular arrival week d ($1 \leq d \leq 8$). Hence customers can be segmented such that their consideration set is restricted to the particular arrival week and product type. The base arrival rate of customers for week d equals $\bar{\lambda}_d = 2C$. Within one week, the probability that an arriving customer is interested in product type 1, 2, and 3 is equal to 1/2, 1/3, and 1/6, respectively. Furthermore, assume that customers of each segment select a product based on the multinomial logit model, where the MNL weights are given by $(10, 7, 3, 4)$ (note that the last entry represents the no-purchase option).

Regarding the reviews, assume that the positive and negative review probabilities equal $q^p = (0.1, 0.05, 0.01)$ and $q^n = (0.01, 0.02, 0.05)$, respectively, for all product types. Let $\alpha = 0.9$ be the discount parameter and let $\beta = (1, 1.5)$ be the parameter that measures the impact of reviews on demand. The initial reviews are set to zero.

The solution methods that are used in the examples are CRLP and the robust CRLP. The uncertainty parameters $\rho(S, d)$ are defined by

$$\rho(S, d) = \rho|\mu(S, d)|, \tag{7.26}$$

where $\rho \in \mathbb{R}^+$ is fixed for all $S \subset N$ and for all $1 \leq d \leq D$. The robust CRLP method under parameter $\rho \in \mathbb{R}^+$ is denoted by CRLP$_\rho$. CDLP is used as a benchmark. The primary measure of performance is the expected revenue. A secondary objective is the variation of revenue in terms of the coefficient of variation c_v, which can measure the robustness of the heuristic.

To analyse the impact of different problem sizes some scaling parameters are used for hotel size, demand load, and review probabilities. These parameters are specified in the

l_f	Method	Revenue	c_v (%)	% Opt. gap
	UB	805698	0.00	0.00
	CRLP	803539	5.75	0.27
	CDLP	688235	5.73	14.58
0.6	$\text{CRLP}_{0.05}$	803353	5.72	0.29
	$\text{CRLP}_{0.1}$	802974	5.74	0.34
	$\text{CRLP}_{0.2}$	803626	5.85	0.26
	$\text{CRLP}_{0.5}$	798323	5.84	0.92
	CRLP_{1}	787477	5.98	2.26
	UB	1060989	0.00	0.00
	CRLP	1037750	3.53	2.19
	CDLP	916199	4.94	13.65
0.8	$\text{CRLP}_{0.05}$	1037682	3.43	2.20
	$\text{CRLP}_{0.1}$	1037283	3.28	2.23
	$\text{CRLP}_{0.2}$	1035870	3.26	2.37
	$\text{CRLP}_{0.5}$	1027387	3.04	3.17
	CRLP_{1}	1032296	3.79	2.70
	UB	1220873	0.00	0.00
	CRLP	1187402	2.90	2.74
	CDLP	1136718	3.72	6.89
1	$\text{CRLP}_{0.05}$	1188385	2.83	2.66
	$\text{CRLP}_{0.1}$	1188001	2.76	2.69
	$\text{CRLP}_{0.2}$	1190194	2.65	2.51
	$\text{CRLP}_{0.5}$	1187481	2.24	2.74
	CRLP_{1}	1175339	1.99	3.73
	UB	1326409	0.00	0.00
	CRLP	1282403	2.34	3.32
	CDLP	1251263	2.95	5.67
1.2	$\text{CRLP}_{0.05}$	1283714	2.25	3.22
	$\text{CRLP}_{0.1}$	1283909	2.13	3.20
	$\text{CRLP}_{0.2}$	1284633	2.15	3.15
	$\text{CRLP}_{0.5}$	1286872	1.95	2.98
	CRLP_{1}	1282403	1.71	3.32
	UB	1403950	0.00	0.00
	CRLP	1355459	2.13	3.45
	CDLP	1339087	2.68	4.62
1.4	$\text{CRLP}_{0.05}$	1356163	1.99	3.40
	$\text{CRLP}_{0.1}$	1357059	1.98	3.34
	$\text{CRLP}_{0.2}$	1356214	1.80	3.40
	$\text{CRLP}_{0.5}$	1352999	1.68	3.63
	CRLP_{1}	1345883	1.73	4.14

Table 7.2: Performance of different heuristics under various demand loads l_f.

appropriate examples. Simulations are used to approximate the expected revenues of the different heuristics. The errors are within 0.5% of the stated values, with 95% confidence.

7.5.1 Demand Load

The first experiment validates the performance of the different heuristics under various demand loads. The hotel size is set to $C = 200$ and demand is multiplied by a load factor $l_f \in \{0.6, 0.8, 1, 1.2, 1.4\}$. The applied heuristics are CRLP and CRLP_ρ for $\rho \in \{0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1\}$. The results are presented in Table 7.2.

The solution methods behave differently under the different demand loads. Surprisingly, an increase in demand load brings the optimality gaps of CRLP and CDLP closer together. One would assume that under higher loads the effect of reviews is much more evident, which translates to a better performance of CRLP. A plausible explanation is that the strategies of CRLP and CDLP are similar when the demand load is increased. If demand is high, then it makes less sense to sacrifice revenue now in order to get better reviews, and higher demand, in the future: there will be enough demand anyway. However, when demand is low, then it is worth sacrificing some revenue now in order to get higher demand in the future.

The robust solution methods perform quite well, in some cases even better than CRLP. Also, the coefficient of variation of CRLP$_\rho$ is generally lower than of CRLP. An exception is at $l_f = 0.6$, though the difference is small. Based on these observations, CRLP$_{0.5}$ is selected to be used in the other examples: it provides a lower c_v than CRLP, indicating that the solution is more stable; and the expected revenue does not deviate that much from CRLP, and for $l_f = 1$ it is equal (in the remaining example the demand load is fixed to $l_f = 1$).

7.5.2 Weight of Review Probabilities

This sections investigates the impact of the weight of the review probabilities. In an environment with $C = 200$ rooms and no scaled load (see the previous experiment), the review probabilities are multiplied with weight $w \in \{0.01, 0.1, 0.2, 0.5, 1, 2\}$. A low weight implies a small effect of review probabilities, while a high weight implies a large effect of review probabilities. The results are presented in Table 7.3.

		CRLP			CDLP			CRLP$_{0.5}$		
w	UB	Rev	c_v	%Gap	Rev	c_v	%Gap	Rev	c_v	%Gap
0.01	1165145	1154996	1.87	0.87	1154996	1.87	0.87	1154996	1.87	0.87
0.1	1170226	1148537	1.84	1.85	1153348	2.09	1.44	1149764	1.91	1.75
0.2	1178102	1153052	2.00	2.13	1151556	2.27	2.25	1153876	1.97	2.06
0.5	1197396	1168466	2.36	2.42	1144965	2.96	4.38	1165078	2.17	2.70
1	1220873	1187402	2.90	2.74	1136718	3.72	6.89	1187481	2.24	2.74
2	1281146	1232147	3.09	3.82	1115901	5.02	12.90	1223037	1.84	4.54

Table 7.3: Performance of different heuristics under various weights w of review probabilities.

The results show that the weight of the review probabilities impacts the performance. For the low values of $w = 0.01$ the performances of all heuristics give the same results. CRLP outperforms CDLP in all other cases, with the surprising exception of $w = 0.1$. CRLP$_{0.5}$ performs better than CRLP for small weights, but when the weight increases CRLP performs better in terms of expected revenue and worse in terms of variation.

7.5.3 Hotel Size

In this experiment the hotel size is varied as $C \in \{50, 100, 200, 500, 1000\}$, while the other parameters remain fixed. The results are presented in Table 7.4.

		CRLP			CDLP			CRLP$_{0.5}$		
C	UB	Rev	c_v (%)	%Gap	Rev	c_v (%)	%Gap	Rev	c_v (%)	%Gap
50	305218	287855	5.35	5.69	278054	6.44	8.90	288752	5.07	5.39
100	610436	586188	4.11	3.97	563560	4.98	7.68	587729	3.42	3.72
200	1220873	1187402	2.90	2.74	1136718	3.72	6.89	1187481	2.24	2.74
500	3052182	3000490	1.84	1.69	2859800	2.58	6.30	2991797	1.22	1.98
1000	6104363	6030491	1.28	1.21	5731378	1.88	6.11	5999069	0.75	1.72

Table 7.4: Performance of heuristics under different hotel sizes C.

We make two main observations, apart from the fact that CRLP and CRLP$_{0.5}$ outperform CDLP for all hotel sizes, both in terms of expected revenue (by 3.5-5.2 percent) and variation. First, the method that performs best depends on the hotel size: for small hotels, with $C = 50$ and $C = 100$, CRLP$_{0.5}$ outperforms CRLP in terms of both revenue and variation. On the other hand, for large hotels, with $c = 500$ and $C = 1000$, CRLP provides higher expected revenue but CRLP$_{0.5}$ provides lower variation. For small hotels it is therefore beneficial to use the robust CRLP method, while large hotels need to decide on the trade-off between revenue and variation. As a large hotel can bare more risk due to its large volume, the CRLP method will lead to higher revenues in the long run.

A second observation is that the optimality gap and variation decrease as the hotels size increases. There are two forces at work that could explain this behaviour. First, by Propositions 7.1 and 7.2, the upper bound is larger than the optimal value, and converges to the optimal value as the hotel size and demand increase. Second, the CRLP heuristic might perform better as hotel size and demand increase. In both cases the difference between the upper bound and CRLP gets tighter.

7.5.4 Estimation Error

In this experiment the performance of the heuristics is measured under estimation errors of demand. The strategies of the heuristics are evaluated using the current parameters, but the actual demand is assumed to deviate from $\bar{\lambda}_d$ by a factor of $u \in \mathbb{R}^+$. In each simulation, u is drawn from a uniform distribution: $u \in \text{Unif}[1 - \bar{u}, 1 + \bar{u}]$, with $\bar{u} \in \{0.01, 0.02, 0.05, 0.1, 0.2, 0.5\}$. The results are presented in Table 7.5.

	CRLP			CDLP			CRLP$_{0.5}$		
$\bar{u}$	Rev	c_v (%)	%Loss	Rev	c_v (%)	%Loss	Rev	c_v (%)	%Loss
∞	1187200	2.89	0.00	1136084	3.78	4.31	1187561	2.21	-0.03
0.01	1187586	2.84	-0.03	1135729	3.76	4.34	1187495	2.26	-0.02
0.02	1187304	2.95	-0.01	1136061	3.84	4.31	1187025	2.40	0.01
0.05	1184460	3.44	0.23	1134494	4.43	4.44	1186081	2.77	0.09
0.1	1178394	4.82	0.74	1131404	5.88	4.70	1181261	3.97	0.50
0.2	1158335	8.49	2.43	1117748	9.67	5.85	1162424	7.45	2.09
0.5	1078115	21.06	9.19	1059288	22.12	10.77	1090038	19.90	8.18

Table 7.5: Performance of heuristics under estimation errors.

The results show that CRLP is quite robust in the sense that the performance does not suffer too much in presence of estimation errors. For estimation errors of up to 5% the revenue loss compared to perfect knowledge does not deviate too much, and in fact increases a little bit (although this might be caused by simulation errors). For all values of $\bar{u}$ the variation is lowest for CRLP$_\rho$. This is expected, as the robust solution method takes uncertainty into account. However, when the uncertainty is too big, with $\bar{u} = 0.5$, all methods do not perform well.

7.6 Concluding Remarks

The main conclusion is that online reviews have a huge impact on revenue. This study introduces a choice-based network RM model that incorporates reviews in a feedback mechanism: on the one hand demand is impacted by reviews, and on the other hand, reviews depend on the price/quality perception of customers. Including reviews in the decision process can lead to a significant increase in revenue.

The full stochastic problem cannot be evaluated for practical purposes, a common problem in choice-based network RM, because of the network structure and because the problem has to keep track of the reviews. To this extent a deterministic variant of the problem is proposed. The deterministic variant serves as an upper bound to the stochastic problem, and converges to the optimal stochastic value when demand and capacity are scaled.

Two heuristics are proposed to solve the problem: one based on the deterministic variant and a robust optimisation approach, where the outcome of reviews is assumed to be uncertain. Both heuristics show a significant improvement of 3.5-5.2% over the benchmark, which does not take into account reviews. Two insightful results from the numerics in Section 7.5 are that considering reviews in the decision process has more impact when demand is low compared to when demand is high; and that small hotels are more effected by the review feedback mechanism than larger hotels. Two other core advantages make the heuristics an effective tool for practitioners: both heuristics can be evaluated efficiently, and the robust solution method reduces the risk that comes with estimation errors.

The remainder of this section addresses two important aspects of choice-based network RM under reviews. First, opportunities and challenges of various application areas are discussed. Second, as a topic for further research, the value of dynamic strategies is highlighted, as well as challenges to acquire such strategies.

7.6.1 Application Areas

The choice-based review model can be applied in many relevant areas. Hotels, airlines, travel agencies, rental cars, and short-term storage space leases have a network structure and sales are highly influenced by recent reviews. Websites like booking.com, hotels.com, tripadvisor, airlinequality.com, rentalcars.com, and Yelp assist customers in their decision process by offering reviews of the various options of companies and products.

The effect of reviews on demand differs per industry. In some industries, like hotels, the reviews impacts the demand of the *individual property* more than the brand itself. This is due to the fact that ratings differ per property of the same branch, and for a property only corresponding reviews are listed. In other industries, on the other hand, like airlines, *brands* are rated rather than individual flight legs.[4] The effect of a review of a flight leg therefore impacts the whole fleet. One way to deal with this is joining the whole fleet together in an optimisation problem, such that the effects of reviews over multiple flight legs is accounted for. However, this leads to intractability issues of biblical proportions. A tractable approach would be to optimise local fleets, where the effect of reviews from flight legs outside the local fleet are accounted for by forecasts.

Applications of reviews to RM that where a network structure is absent, or where booking horizons do not overlap, are treated in Chapter 6.

7.6.2 Further Research

As this is a pioneering study on choice-based network RM under reviews, investigating improved solution methods to reduce the optimality gap is a promising future direction. In particular, the heuristics described in this study are of static rather than dynamic nature. However, other choice-based RM studies have shown an improvement in performance when dynamic strategies are used rather than static strategies (e.g., see Liu & van Ryzin, 2008; Maglaras & Meissner, 2006). However, most studies use only the current state of reservations to decide on an offer set, and clearly not the current state of reviews. This is exactly why a straightforward implementation of such strategy most likely will fail: it

[4]As a hybrid, car rental sites rate the brands, but tend to show reviews related to the pickup location.

does not take into account the crucial effect of reviews, leading to a downwards spiral of decreasing reviews and decreasing demand.

Decomposition methods, like in Liu & van Ryzin (2008) and Talluri & van Ryzin (2004b, p. 100-108), solve the problem of the huge state space of the full dynamic problem by solving a number of single-night stay sub-problems. However, in review context, these sub-problems would still be intractable if reviews are to be taken into account. The main challenge is to develop a novel method that incorporates the reviews in the decision process without blowing up the state space.

Summary and Conclusion

This dissertation proposed and analysed several revenue management models, with focus on the effects of cancellations, overbooking, customer purchase behaviour, and reviews, in presence of flexible products and group reservations in a network setting. Chapter 2 provided exploratory data analysis of an independent Dutch hotel, demonstrating how one should analyse client segment mix, the nature of demand, and effect of cancellations and group reservations, amongst others. The results showed that a major part (21.71%) of the reservations are cancelled, which has a big impact on revenues. Motivated by this, in Chapter 3 we proposed a single-leg revenue management model that takes into account cancellations and overbooking along with purchasing behaviour of customers. A dynamic programming formulation to solve the discretised Markov process suffers from the curse of dimensionality, since it has to keep track of the purchases of different product types. Therefore, three heuristics were proposed, each appropriate under different assumptions. Numerical results revealed that not taking cancellations into account can lead to a revenue loss of up to 20%.

Another result from the data analysis is that demand follows a Poisson Process, which implies relatively more uncertainty in demand for small hotels than for larger hotels. Therefore, revenue optimisation models, which generally attempt to optimise the *expected* revenue, should take this uncertainty into account, e.g., by using robust optimisation techniques. Chapter 4 (and Section 7.4) discuss robust optimisation techniques for revenue management models. Chapter 4 is devoted to a robust solution method for the single-leg choice-based RM model of Chapter 3. The uncertainty in customer purchasing probabilities is modelled using a ϕ-divergence measure, and tractable reformulations of robust counterparts are presented. A numerical study implies that using the robust solution method to model uncertainty in demand, e.g., due to estimation errors or the distribution of demand, can lead to significantly higher revenues than when the nominal solution is used.

Chapter 5 presented a network RM model that allows flexible products, which are common practice in TV and online advertising, and show potential to increase revenues in retailing and fast-moving consumer goods as well. Flexible products give the company the flexibility to assign the customer close to consumption to a selection specific products, as capacity allows. Moreover, since flexible products give the company this flexibility, it can ask for a lower price, which attracts new customer segments. Hence flexible products can lead to better capacity utilisation and higher revenues. The numerical studies endorse this by showing an increase in revenue of up to 20% when flexible products are offered alongside specific products.

A recent development that impacts revenue is the wide availability of reviews and online ratings. Chapters 6 and 7 proposed RM methodologies that model the effect of reviews on demand and the effect of the price/quality perception of clients on writing reviews. This feedback mechanism complicates the model, because demand depends on past reviews. For instance, by sacrificing revenue now, *long-term* revenue can be increased. Chapter 6 proposes a single-leg model with a novel solution method, to model the effects of reviews for amongst other theatres, concerts, sport events, or cinemas. The elaborate solution method can unfortunately not be used in the extension to networks Therefore, in Chapter 7 two heuristics are proposed, one of which deals with uncertainty in demand by means of robust optimisation. Results show a significant improvement in long-term revenue of up to 11%. Two insightful results from the numerics in Section 7.5 are that considering reviews has more impact when demand is low than when demand is high, relative to the hotel size; and that small hotels are more effected by reviews than larger hotels.

The results in this paper motivate future studies where challenges and opportunities in RM, observed by analysing practical instances and data, are exploited. The wide availability of data and technological advances provide great opportunities to perform research on customer behaviour and price sensitivity, and to develop new product types to conquer different market segments. In particular I would like to point out the opportunities of research on the effect of reviews on long-term revenue in *collaborative consumption* platforms. *Collaborative consumption* is the coordination of people to share products amongst relative strangers by means of the online sharing economy (Belk, 2014). Prominent examples of shared products are *temporary housing* (e.g., Airbnb, Couchsurfing) and *transportation* (e.g., Blablacar, Lyft, and Uber).

To continue in the spirit of the hospitality industry, consider the popular online mediation platform Airbnb. Airbnb allows house owners to host travellers, by offering them (a part of) their home (e.g., the whole house or a room) to spend the night, which otherwise would accommodate nobody. For example, when a house owner travels he can rent the house during that period to relative strangers; or a house owner can offer their guest room to travellers. A cornerstone of Airbnb is their reputation system, where both house owners and guests write reviews about one another and the stay itself (Edelman & Luca, 2014), which builds trust between house owners and guests. Guests are more willing to stay at a stranger's house when the house owner has a good reputation, and, vice versa, house owners are more willing to accept a stranger in their home when he has a good reputation.

House owners can offer competitive prices compared to established means of temporary housing such as hotels and hostels. It is beneficial for Airbnb to increase their market share, and pricing is an effective tool to reach this goal. Each property can be described by a vector of perhaps hundreds of features, including the space, facilities, geographical location; and demand may depend on many factors including seasonality, competitor prices, and reviews (Javanmard & Nazerzadeh, 2016). With this vast amount of information at hand, the field of RM has a great opportunity to grow by developing high-dimensional optimisation and statistical techniques that use this big data for accurate demand forecasts, price sensitivity, and, finally, for optimising long-term expected revenues.

Bibliography

Self References

Sierag, D. D. (2016a). Choice-based network revenue management under online reviews. *To appear in Journal of Revenue and Pricing Management.*

Sierag, D. D. (2016b). Pricing-based revenue management for flexible products on a network. *To appear in Journal of Revenue and Pricing Management.*

Sierag, D. D., & Hanzon, B. (2016). Pricing derivatives on multiple assets: Recombining multinomial trees based on pascals simplex. *Submitted.*

Sierag, D. D., Koole, G. M., van der Mei, R. D., van der Rest, J. I., & Zwart, B. (2015). Revenue management under customer choice behaviour with cancellations and overbooking. *European Journal of Operational Research*, *246*(1), 170-185.

Sierag, D. D., & Roijers, D. M. (2016). Revenue management under reviews and online ratings. *Submitted.*

Sierag, D. D., & van der Mei, R. D. (2016). Single-leg choice-based revenue management: A robust optimisation approach. *To appear in Journal of Revenue and Pricing Management.*

Sierag, D. D., van der Rest, J. I., Koole, G. M., van der Mei, R. D., & Zwart, B. (2016). Exploratory data analysis in revenue management forecasting: a case study of a small and independent hotel in the netherlands. *To appear in Internation Journal of Revenue Management.*

Den Boer, A. V., & Sierag, D. D. (2016). Decision-based model selection. *Submitted.*

Koupriouchina, L., Van der Rest, J. I., Sierag, D. D., & Schwartz, Z. (2017). Forecasting in hotel revenue management: a critical review of scale-dependent and scale-independent accuracy measures. In M. Uysal, Z. Schwartz, & E. Sirakaya-Turk (Eds.), *Management science in hospitality and tourism: Theory, practice, and applications* (p. 269-304). Apple Academic Press/CRC Press/Taylor & Francis Group.

Van der Rest, J. I., Loosschilder, G., Schwartz, Z., Cordella, P., & Sierag, D. D. (2016). The impact of sorting and filtering on ota customers' choices, search results hotel's placement value and revenue. *To appear in Journal of Revenue and Pricing Management.*

References

Aghazadeh, S. M. (2007). Revenue forecasting models for hotel management. *The Journal of Business Forecasting*, *26*(3), 33-37.

Aydin, N., Ş. İ Birbil, Frenk, J. B. D., & Noyan, N. (2012). Single-leg airline revenue management with overbooking. *Transportation Science*, *47*(4), 560-583.

Baker, T., Murthy, N. N., & Jayaraman, V. (2002). Service package switching in hotel revenue management systems. *Decision Sciences*, *33*(1), 109-132.

Ball, M. O., & Queyranne, M. (2009). Toward robust revenue management: Competitive analysis of online booking. *Operations Research*, *57*(4), 950-963.

Barnhart, C., Johnson, E. L., Nemhauser, G. L., Savelsbergh, M. W., & Vance, P. H. (1998). Branch-and-price: Column generation for solving huge integer programs. *Operations Research*, *46*(3), 316-329.

Belk, R. (2014). You are what you can access: Sharing and collaborative consumption online. *Journal of Business Research*, *67*(8), 1595–1600.

Belobaba, P. P. (1987a). Airline yield management: An overview of seat inventory control. *Transportation Science*, *21*, 63-73.

Belobaba, P. P. (1987b). *Air Travel Demand and Airline Seat Inventory Management.* PhD Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Belobaba, P. P. (1989). Application of a probabilisitc decision model to airline seat inventory control. *Operations Research*, *37*(2), 183-197.

Ben-Akiva, M. E., & Lerman, S. R. (1985). *Discrete Choice Analysis: Theory and Application to Travel Demand* (Vol. 9). MIT press, Cambridge, MA.

Bendoly, E. (2013). Real-time feedback and booking behavior in the hospitality industry: Moderating the balance between imperfect judgment and imperfect prescription. *Journal of Operations Management*, *31*(1), 62-71.

Ben-Tal, A., den Hertog, D., De Waegenaere, A., Melenberg, B., & Rennen, G. (2013). Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, *59*(2), 341-357.

Ben-Tal, A., El Ghaoui, L., & Nemirovski, A. (2009). *Robust Optimization.* Princeton University Press.

Bermúdez, J. D., Corberán-Vallet, A., & Vercher, E. (2009). Multivariate exponential smoothing: A Bayesian forecast approach based on simulation. *Mathematics and Computers in Simulation*, *79*(5), 1761-1769.

Bertsimas, D., & Popescu, I. (2003). Revenue management in a dynamic network environment. *Transportation Science*, *37*(3), 257-277.

Birbil, S. I., Frenk, J. B. G., Gromicho, J. A., & Zhang, S. (2009). The role of robust optimization in single-leg airline revenue management. *Management Science*, *55*(1), 148-163.

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

Bodea, T., Ferguson, M., & Garrow, L. (2009). Data set-choice-based revenue management: Data from a major hotel chain. *Manufacturing & Service Operations Management*, *11*(2), 356-361.

Brännäs, K., Hellström, J., & Nordström, J. (2002). A new approach to modelling and forecasting monthly guest nights in hotels. *International Journal of Forecasting*, *18*(1), 19-30.

Bront, J. J. M., Méndez-Díaz, I., & Vulcano, G. (2009). A column generation algorithm for choice-based network revenue management. *Operations Research*, *57*(3), 769-784.

Chen, C., & Kachani, S. (2007). Forecasting and optimisation for hotel revenue management. *Journal of Revenue & Pricing Management*, *6*(3), 163-174.

Chen, C.-C., Schwartz, Z., & Vargas, P. (2011). The search for the best deal: How hotel cancellation policies affect the search and booking decisions of deal-seeking customers. *International Journal of Hospitality Management*, *30*(1), 129-135.

Chen, C.-C., & Xie, K. L. (2013). Differentiation of cancellation policies in the US hotel industry. *International Journal of Hospitality Management*, *34*, 66-72.

Cheng, H.-T. (1988). *Algorithms for partially observable Markov decision processes* (Unpublished doctoral dissertation). University of British Columbia, Vancouver.

Clarke, K. (2004). Case study: Implementing a revenue management system at Malaysia Airlines. *Journal of Revenue and Pricing Management*, *3*(1), 41-48.

Cóccola, M. E., Dondo, R., & Méndez, C. A. (2015). A MILP-based column generation strategy for managing large-scale maritime distribution problems. *Computers & Chemical Engineering*, *72*, 350-362.

Cranage, D. (2003). Practical time series forecasting for the hospitality manager. *International Journal of Contemporary Hospitality Management*, *15*(2), 86-93.

Cross, R. G. (1997). *Revenue Management: Hardcore Tactics for Market Domination.* Broadway Books, New York, NY.

Dellarocas, C., Awad, N., & Zhang, X. (2004). Exploring the value of online reviews to organizations: Implications for revenue forecasting and planning. *ICIS 2004 Proceedings*, *30*.

Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society Series B (Methodological)*, *39*(1), 1-38.

Dow, J. K., & Endersby, J. W. (2004). Multinomial probit and multinomial logit: a comparison of choice models for voting research. *Electoral Studies*, *23*, 107-122.

Edelman, B. G., & Luca, M. (2014). Digital discrimination: The case of airbnb.com. *Harvard Business School NOM Unit Working Paper*(14-054).

Erdelyi, A., & Topaloglu, H. (2010). A dynamic programming decomposition method for making overbooking decisions over an airline network. *INFORMS Journal on Computing*, *22*(3), 443-456.

Farias, V. F., Jagabathula, S., & Shah, D. (2013). A nonparametric approach to modeling choice with limited data. *Management Science*, *59*(2), 305-322.

Farouk El Gayar, N., Saleh, M., Atiya, A., El-Shishiny, H., Alkes Youhanna Fayez Zakhary, A., & Abdel Aziz Mohammed Habib, H. (2011). An integrated framework for advanced hotel revenue management. *International Journal of Contemporary Hospitality Management*, *23*(1), 84-98.

Fay, S., & Xie, J. (2008). Probabilistic goods: A creative way of selling products and services. *Marketing Science*, *27*(4), 674-690.

Ferguson, M., & Smith, S. (2014). The changing landscape of hotel revenue management and the role of the hotel revenue manager. *Journal of Revenue and Pricing Management*, *13*(3), 224-232.

Fildes, R., Nikolopoulos, K., Crone, S. F., & Syntetos, A. A. (2008). Forecasting and operational research: a review. *Journal of the Operational Research Society*, *59*(9), 1150-1172.

Fritz, R. G., Brandon, C., & Xander, J. (1984). Combining time-series and econometric forecast of tourism activity. *Annals of Tourism Research*, *11*(2), 219-229.

Gallego, G., Iyengar, G., Phillips, R., & Dubey, A. (2004). Managing flexible products on a network. *Working Paper*.

Gallego, G., & Phillips, R. (2004). Revenue management of flexible products. *Manufacturing & Service Operations Management*, *6*(4), 321-337.

Gallego, G., & van Ryzin, G. (1997). A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research*, *45*(1), 24-41.

Geraghty, M. K., & Johnson, E. (1997). Revenue management saves national car rental. *Interfaces*, *27*(1), 107-127.

Gönsch, J., Koch, S., & Steinhardt, C. (2014). Revenue management with flexible products: The value of flexibility and its incorporation into DLP-based approaches. *International Journal of Production Economics*, *153*, 280-294.

Guadix, J., Cortés, P., Onieva, L., & Muñuzuri, J. (2010). Technology revenue management system for customer groups in hotels. *Journal of Business Research*, *63*(5), 519-527.

Guerriero, F., Miglionico, G., & Olivito, F. (2014). Strategic and operational decisions in restaurant revenue management. *European Journal of Operational Research*, *237*(3), 1119–1132.

Haensel, A., & Koole, G. (2011a). Booking horizon forecasting with dynamic updating: A case study of hotel reservation data. *International Journal of Forecasting*, *27*(3), 942-960.

Haensel, A., & Koole, G. (2011b). Estimating unconstrained demand rate functions using customer choice sets. *Journal of Revenue & Pricing Management*, *10*(5), 438-454.

Holverson, S., & Revaz, F. (2006). Perceptions of european independent hoteliers: hard and soft branding choices. *International Journal of Contemporary Hospitality Management*, *18*(5), 398-413.

Hopman, D. (2013). Revenue management in small-scale hotels. *Working Paper*.

Hosseinalifam, M., Savard, G., & Marcotte, P. (2016). Computing booking limits under a non-parametric demand model: A mathematical programming approach. *Journal of Revenue & Pricing Management*, *15*(2), 170-184.

Howard, R. A. (1960). *Dynamic Programming and Markov Decision Processes.* MIT Press.

Iliescu, D. C., Garrow, L. A., & Parker, R. A. (2008). A hazard model of us airline passengers refund and exchange behavior. *Transportation Research Part B: Methodological*, *42*(3), 229-242.

Javanmard, A., & Nazerzadeh, H. (2016). Dynamic pricing in high-dimensions. *Working paper*.

Kimes, S. E. (1999). Group forecasting accuracy in hotels. *Journal of the Operational Research Society*, *50*(11), 1104-1110.

Kimes, S. E. (2000). Revenue management on the links: applying yield management to the golf-course industry. *The Cornell Hotel and Restaurant Administration Quarterly*, *41*(1), 120-127.

Kouprouchina, L., Van der Rest, J. I., Sierag, D. D., & Schwartz, Z. (2017). Forecasting in hotel revenue management: a critical review of scale-dependent and scale-independent accuracy measures. In M. Uysal, Z. Schwartz, & E. Sirakaya-Turk (Eds.), *Management science in hospitality and tourism: Theory, practice, and applications* (p. 269-304). Apple Academic Press/CRC Press/Taylor & Francis Group.

Koupriouchina, L., van der Rest, J.-P., & Schwartz, Z. (2014). On revenue management and the use of occupancy forecasting error measures. *International Journal of Hospitality Management*, *41*, 104-114.

Koushik, D., Higbie, J. A., & Eister, C. (2012). Retail price optimization at intercontinental hotels group. *Interfaces*, *42*(1), 45-57.

Kunnumkal, S., Talluri, K., & Topaloglu, H. (2012). A randomized linear programming method for network revenue management with product-specific no-shows. *Transportation Science*, *46*(1), 90-108.

Kunnumkal, S., & Topaloglu, H. (2008). A refined deterministic linear program for the network revenue management problem with customer choice behavior. *Naval Research Logistics*, *55*(6), 563-580.

Lai, K. K., & Ng, W. (2005). A stochastic approach to hotel revenue optimization. *Computers & Operations Research*, *32*(5), 1059-1072.

Lan, Y., Gao, H., Ball, M. O., & Karaesmen, I. (2008). Revenue management with limited demand information. *Management Science*, *54*(9), 1594-1609.

Langeveld, C. B. G. (2006). *Economie van het Theater: Vestiging, Prijsvorming en Economies of Scale in een Wereld met Bezieling.* Erasmus University Rotterdam.

Lautenbacher, C. J., Stidham, S., & Subramanian, J. (1999). Airline yield management with overbooking, cancellations, and no-shows. *Transportation Science*, *33*(2), 147-167.

Law, R. (1998). Room occupancy rate forecasting: a neural network approach. *International Journal of Contemporary Hospitality Management*, *10*(6), 234-239.

Law, R. (2004). Initially testing an improved extrapolative hotel room occupancy rate forecasting technique. *Journal of Travel & Tourism Marketing*, *16*(2-3), 71-77.

Lee-Ross, D., & Johns, N. (1997). Yield management in hospitality SMEs. *International Journal of Contemporary Hospitality Management*, *9*(2), 66-69.

Lemke, C., Riedel, S., & Gabrys, B. (2013). Evolving forecast combination structures for airline revenue management. *Journal of Revenue and Pricing Management*, *12*(3), 221-234.

Li, G., Song, H., & Witt, S. F. (2005). Recent developments in econometric modeling and forecasting. *Journal of Travel Research*, *44*(1), 82-99.

Li, G., Song, H., & Witt, S. F. (2006). Time varying parameter and fixed parameter linear aids: An application to tourism demand forecasting. *International Journal of Forecasting*, *22*(1), 57-71.

Li, G., Wong, K. K., Song, H., & Witt, S. F. (2006). Tourism demand forecasting: A time varying parameter error correction model. *Journal of Travel Research*, *45*(2), 175-185.

Lim, C., Chang, C., & McAleer, M. (2009). Forecasting h(m)otel guest nights in New Zealand. *International Journal of Hospitality Management*, *28*(2), 228-235.

Littlewood, K. (1972). Forecasting and control of passenger bookings. *AGIFORS 12-th Annual Symposium Proceedings*, 95-128.

Liu, Q., & van Ryzin, G. (2008). On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management*, *10*(2), 288-310.

Luciani, S. (1999). Implementing yield management in small and medium sized hotels: an investigation of obstacles and success factors in florence hotels. *International Journal of Hospitality Management*, *18*(2), 129-142.

Maddag, B., Moussawi-Haidar, L., El-Taha, M., & Rida, H. (2010). Dynamic cruise ship revenue management. *European Journal of Operational Research*, *207*(1), 445-555.

Maglaras, C., & Meissner, J. (2006). Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing & Service Operations Management*, *8*(2), 136-148.

Makridakis, S., Wheelwright, S., & Hyndman, R. J. (1998). *Forecasting: Methods and Applications.* John Wiley & Sons.

Mang, S., Post, D., & Spann, M. (2012). Pricing of flexible products. *Review of Managerial Science*, *6*(4), 361-374.

McGill, J. I., & van Ryzin, G. J. (1999). Revenue management: Research overview and prospects. *Transportation Science*, *33*(2), 233-256.

McKinsey&Company. (2015). *Global media report 2015: Global industry overview* (Tech. Rep.).

Meissner, J., & Strauss, A. (2012). Network revenue management with inventory-sensitive bid prices and customer choice. *European Journal of Operational Research*, *216*(2), 459-468.

Meissner, J., Strauss, A., & Talluri, K. (2013). An enhanced concave program relaxation for choice network revenue management. *Production and Operations Management*, *22*(1), 71-87.

Milla, S., & Shoemaker, S. (2008). Three decades of revenue management: What's next? *Journal of Revenue & Pricing Management*, *7*(1), 110-114.

Newman, J., Ferguson, M. E., & Garrow, L. A. (2012). Estimating discrete choice models with incomplete data. *Transportation Research Record 2302*, 130-137.

Newman, J. P., Garrow, L., Ferguson, M., Jacobs, T. L., & Purnomo, H. (2014). Estimation of choice-based models using sales data from a single firm. *Manufacturing & Service Operations Management*, *16*(2), 184-197.

Nilim, A., & El Ghaoui, L. (2005). Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, *53*(3), 780-798.

Pan, Y., & Zhang, J. (2011). Born unequal: a study of the helpfulness of user-generated product reviews. *Journal of Retailing*, *87*(4), 598-612.

Park, C., & Lee, T. (2009). Information direction, website reputation and eWOM effect: A moderating role of product type. *Journal of Business Research*, *62*(1), 61-67.

Pavlou, P. A., & Dimoka, A. (2006). The nature and role of feedback text comments in online marketplaces: Implications for trust building, price premiums, and seller differentiation. *Information Systems Research*, *17*(4), 392-414.

Perakis, G., & Roels, G. (2010). Robust controls for network revenue management. *Manufacturing & Service Operations Management*, *12*(1), 56-76.

Petrick, A., Gönsch, J., Steinhardt, C., & Klein, R. (2010). Dynamic control mechanisms for revenue management with flexible products. *Computers & Operations Research*, *37*(11), 2027-2039.

Petrick, A., Steinhardt, C., Gönsch, J., & Klein, R. (2012). Using flexible products to cope with demand uncertainty in revenue management. *OR Spectrum*, *34*(1), 215-242.

Phillips, R. L. (2005). *Pricing and Revenue Optimization.* Stanford University Press.

Puterman, M. (1994). *Markov Decision Processes.* Wiley.

Rajopadhye, M., Ghalia, M. B., Wang, P. P., Baker, T., & Eister, C. V. (2001). Forecasting uncertain hotel room demand. *Information sciences*, *132*(1), 1-11.

Van der Rest, J. I., Loosschilder, G., Schwartz, Z., Cordella, P., & Sierag, D. D. (2016). The impact of sorting and filtering on ota customers' choices, search results hotel's placement value and revenue. *To appear in Journal of Revenue and Pricing Management.*

Roijers, D., Vamplew, P., Whiteson, S., & Dazeley, R. (2013). A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, *48*, 67-113.

Roijers, D., Whiteson, S., & Oliehoek, F. (2015). Computing convex coverage sets for faster multi-objective coordination. *Journal of Artificial Intelligence Research*, *52*, 399-443.

Roijers, D. M., Scharpff, J., Spaan, M. T. J., Oliehoek, F. A., de Weerdt, M., & Whiteson, S. (2014). Bounded approximations for linear multi-objective planning under uncertainty. In *Icaps 2014: Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling* (p. 262-270).

Rusmevichientong, P., & Topaloglu, H. (2012). Robust assortment optimization in revenue management under the multinomial logit choice model. *Operations Research*, *60*(4), 865-882.

Schwartz, Z., & Cohen, E. (2004a). Hotel revenue-management forecasting evidence of expert-judgment bias. *Cornell Hotel and Restaurant Administration Quarterly*, *45*(1), 85-98.

Schwartz, Z., & Cohen, E. (2004b). Subjective estimates of occupancy forecast uncertainty by hotel revenue managers. *Journal of Travel & Tourism Marketing*, *16*(4), 59-66.

Schwartz, Z., & Hiemstra, S. (1997). Improving the accuracy of hotel reservations forecasting: curves similarity approach. *Journal of Travel Research*, *36*(1), 3-14.

Shen, Z. M., & Su, X. (2007). Customer behavior modeling in revenue management and auctions: A review and new research opportunities. *Production and Operations Management*, *16*(6), 713-728.

Sierag, D. D. (2016a). Choice-based network revenue management under online reviews. *Submitted*.

Sierag, D. D. (2016b). Pricing-based revenue management for flexible products on a network. *To appear in Journal of Revenue and Pricing Management*.

Sierag, D. D., Koole, G. M., van der Mei, R. D., van der Rest, J. I., & Zwart, B. (2015). Revenue management under customer choice behaviour with cancellations and overbooking. *European Journal of Operational Research*, *246*(1), 170-185.

Sierag, D. D., & Roijers, D. M. (2016). Revenue management under reviews and online ratings. *Submitted*.

Sierag, D. D., & van der Mei, R. D. (2016). Single-leg choice-based revenue management: A robust optimisation approach. *Journal of Revenue and Pricing Management*, 1–14.

Sierag, D. D., van der Rest, J. I., Koole, G. M., van der Mei, R. D., & Zwart, B. (2016). Exploratory data analysis in revenue management forecasting: a case study of a small and independent hotel in the netherlands. *To appear in Internation Journal of Revenue Management*.

Song, H., Gao, B. Z., & Lin, V. S. (2013). Combining statistical and judgmental forecasts via a web-based tourism demand forecasting system. *International Journal of Forecasting*, *29*(2), 295-310.

Song, H., & Li, G. (2008). Tourism demand modelling and forecasting – a review of recent research. *Tourism Management*, *29*(2), 203-220.

Soon, W. (2011). A review of multi-product pricing models. *Applied Mathematics and Computation*, *217*(21), 8149-8165.

Sparks, B. A., & Browning, V. (2011). The impact of online reviews on hotel booking intentions and perception of trust. *Tourism Management*, *32*(6), 1310-1323.

Strauss, A., & Talluri, K. (2012). A tractable consideration set structure for network revenue management. *Working Paper*.

Talluri, K. T., & van Ryzin, G. (2004a). Revenue management under a general choice model of consumer behavior. *Management Science*, *50*(1), 15-33.

Talluri, K. T., & van Ryzin, G. J. (2004b). *The Theory and Practice of Revenue Management.* Springer Science & Business Media.

Vamplew, P., Dazeley, R., Barker, E., & Kelarev, A. (2009). Constructing stochastic mixture policies for episodic multiobjective reinforcement learning tasks. In *Australasian joint conference on artificial intelligence* (pp. 340–349).

van Ryzin, G., & Vulcano, G. (2013). An expectation-maximization method to estimate a rank-based choice model of demand. *Technical Note*.

van Ryzin, G. J. (2005). Future of revenue management: models of demand. *Journal of Revenue and Pricing Management*, *4*(2), 204-210.

Vermeulen, I. E., & Seegers, D. (2009). Tried and tested: The impact of online hotel reviews on consumer consideration. *Tourism Management*, *30*(1), 123-127.

Vinod, B. (2004). Unlocking the value of revenue management in the hotel industry. *Journal of Revenue and Pricing Management*, *3*(2), 178-190.

Vossen, T. W., & Zhang, D. (2015). A dynamic disaggregation approach to approximate linear programs for network revenue management. *Production and Operations Management*, *24*(3), 469-487.

Wakuta, K. (1999). A note on the structure of value spaces in vector-valued Markov decision processes. *Mathematical Methods of Operations Research*, *49*(1), 77-85.

Weatherford, L. R., & Kimes, S. E. (2003). A comparison of forecasting methods for hotel revenue management. *International Journal of Forecasting*, *19*(3), 401-415.

Weatherford, L. R., Kimes, S. E., & Scott, D. A. (2001). Forecasting for hotel revenue management: Testing aggregation against disaggregation. *Cornell Hotel and Restaurant Administration Quarterly*, *42*(2), 53-64.

White, C. C., & Kim, K. M. (1980). Solution procedures for solving vector criterion Markov decision processes. *Large Scale Systems*, *1*, 129-140.

Xiang, Z., Schwartz, Z., Gerdes, J. H., & Uysal, M. (2015). What can big data and text analytics tell us about hotel guest experience and satisfaction? *International Journal of Hospitality Management*, *44*, 120-130.

Yang, Y., Pan, B., & Song, H. (2014). Predicting hotel demand using destination marketing organizations web traffic data. *Journal of Travel Research*, *53*(4), 433-447.

Ye, Q., Law, R., Gu, B., & Chen, W. (2011). The influence of user-generated content on traveler behavior: An empirical investigation on the effects of e-word-of-mouth to hotel online bookings. *Computers in Human Behavior*, *27*(2), 634-639.

Yoo, K., & Gretzel, U. (2011). Influence of personality on travel-related consumer-generated media creation. *Computers in Human Behavior*, *27*(2), 609-621.

Yüksel, S. (2007). An integrated forecasting approach to hotel demand. *Mathematical and Computer Modelling*, *46*(7), 1063-1070.

Zakhary, A., Atiya, A. F., El-Shishiny, H., & Gayar, N. E. (2011). Forecasting hotel arrivals and occupancy using monte carlo simulation. *Journal of Revenue & Pricing Management*, *10*(4), 344-366.

Zhang, D., & Lu, Z. (2013). Assessing the value of dynamic pricing in network revenue management. *INFORMS Journal on Computing*, *25*(1), 102-115.

Zhou, L., Ye, S., Pearce, P., & Wu, M. (2014). Refreshing hotel satisfaction studies by reconfiguring customer review data. *International Journal of Hospitality Management*, *38*, 1-10.