

Demand-Driven Re-Fleeting in a Dynamic Pricing Environment

Zur Erlangung des akademischen Grades eines
Doktors der Wirtschaftswissenschaften

(Dr. rer. pol.)

von der Fakultät für
Wirtschaftswissenschaften
des Karlsruher Instituts für Technologie

genehmigte

DISSERTATION

von

Dipl.-Wi.-Ing. Matthias Viehmann

Tag der mündlichen Prüfung: 28. Juni 2012

Referent: Prof. Dr. Karl-Heinz Waldmann

Korreferent: Prof. Dr. Kay Mitusch

(2012) Karlsruhe

Eigenständigkeitserklärung

Ich versichere wahrheitsgemäß, die Dissertation bis auf die in der Abhandlung angegebene Hilfe selbstständig angefertigt, alle benutzten Hilfsmittel vollständig und genau angegeben und genau kenntlich gemacht zu haben, was aus Arbeiten anderer und aus eigenen Veröffentlichungen unverändert oder mit Abänderungen entnommen wurde.

Matthias Viehmann

Karlsruhe, 23. Januar 2012

Contents

1	Introduction	10
1.1	Contribution	16
1.2	Outline	17
2	Price Discrimination	18
2.1	Requirements	19
2.2	History and Classification	20
2.3	Practical Implications	23
2.4	Discrimination in a competitive environment	24
2.5	Welfare effects and legislation	25
3	Markov Decision Processes	27
3.1	Model and Notation	28
3.2	Decision Criteria	29
3.3	Backward Induction	30
3.4	Structured Policies	31
4	Revenue Management	36
4.1	Definition and History	36

4.2	Requirements and general assumptions	37
4.3	Revenue Management Process	41
4.4	Data collection and Forecasting	42
4.5	Capacity Control	44
4.5.1	Booking Controls	45
4.5.2	Static Models	46
4.5.3	Dynamic Models	53
4.6	Dynamic Pricing	56
5	Demand-Driven Re-Fleeting	70
5.1	Airline Planning Process	70
5.2	Fleet Assignment	73
5.3	Demand-Driven Re-Fleeting	76
5.4	Fleet Assignment Swaps	80
6	Demand-Driven Re-Fleeting and Dynamic Pricing	83
6.1	Two-leg model	85
6.2	Extension to Sequences	100
6.3	Heuristic Approaches	110
6.3.1	Period Re-Evaluation of the Fleet Assignment	110
6.3.2	Limited Lookahead Policies	112
6.3.3	Dynamic Pricing Linear Program	113
6.3.4	Demand-Driven Re-Fleeting Dynamic Pricing Heuristic	115
6.4	Numerical Studies	116
6.4.1	Simulation Environment	117
6.4.2	General Settings	119
6.4.3	Simulation Results	121
6.4.4	Summarized Results	130
7	Conclusion	132
7.1	Summary	133

7.2	Future Research	135
A	Additional Proofs	136
A.1	Proof of Theorem 6.2.1	136
A.2	Proof of Theorem 6.2.2	137
A.3	Proof of Lemma 6.2.3	139
A.4	Proof of Theorem 6.2.4	139
A.5	Proof of Theorem 6.2.6	140
A.6	Proof of Theorem 6.2.7	140
A.7	Proof of Theorem 6.2.10	140
B	Additional Algorithms	141
B.1	Two-leg Dynamic Pricing	141
B.2	Demand-Driven Re-Fleeting Heuristic	146

List of Figures

2.1	Monopoly prices and output quantities	21
2.2	Agglomerative price discrimination	22
2.3	Costs and benefits of segmentation	23
3.1	Examples of switching curves	35
4.1	Revenue Management Process	42
4.2	Relationships of different booking controls	47
4.3	Example of different booking controls	50
5.1	Airline planning process	72
5.2	Example of a time-expanded flight network	74
5.3	Example of Demand-Driven Re-Fleeting	77
5.4	Demand-Driven Re-Fleeting Process	78
5.5	Examples of two types of swapping opportunities	81
6.1	Example of a sub-network with two legs that are eligible for a swap	85
6.2	Example of a switching curve policy for the assignment action	95
6.3	Example of two flight sequences that are eligible for a swap	101

6.4	Process of Demand-Driven Re-Fleeting Heuristic 1	111
6.5	Influence of the time distance between re-optimization points	111
6.6	Process of Demand-Driven Re-Fleeting Heuristic 2	116
6.7	Arrival and thinning probabilities in the base-case scenario	120
6.8	Confidence intervals of the revenue increases in the base-case scenario .	122
6.9	Example of a strategy applying the dynamic programming model . . .	123
6.10	Example of a strategy applying Heuristic 1	123
6.11	Example of a strategy applying Heuristic 2	124
6.12	Example of a realization of the booking process	125
6.13	Mean relative revenue increases for varying demand factors	125
6.14	Mean relative revenue increases for varying capacities	126
6.15	Mean relative revenue increases for varying arrival probabilities	127
6.16	Mean relative revenue increases for varying demand sensitivities	128
6.17	Mean relative revenue increases for varying LAT	129
6.18	Mean relative revenue increases for varying variations of arrivals	130

List of Tables

1.1	Comparison of major LCC characteristics	16
2.1	Market Segmentation Example	21
4.1	Example of decreasing prices in time: Sales Probabilities	65
4.2	Example of decreasing prices in time: Values and optimal prices	67
6.2	Scenario Overview	120
6.3	Revenue increase and load factor results.	121

List of Abbreviations

CAB	Civil Aeronautics Board
CV	Coefficient of Variation
DCM	Dynamic Capacity Management (Demand-Driven Re-Fleeting)
DDD	Demand-Driven Dispatch (Demand-Driven Re-Fleeting)
D^3	Demand-Driven Dispatch (Demand-Driven Re-Fleeting)
DDR	Demand-Driven Re-Fleeting
DP	Dynamic Pricing
EMSR	Expected Marginal Seat Revenue
FAM	Fleet Assignment (Model)
GDS	Global Distribution System
LAT	Latest Assignment Time
LCC	Low-Cost Carrier
LLP	Limited Lookahead Policy
LP	Linear Program
MC	Marginal Cost
MDP	Markov Decision Process
MIP	Mixed Integer Program
MR	Marginal Revenue
RM	Revenue Management
TTD	Time to Departure
W.l.o.g.	Without loss of generality
WTP	Willingness-to-pay

1

Introduction

During the phase of the deregulation of the American airline industry in the late 70s and early 80s, established carriers were severely threatened by low-cost competitors entering the market. These offered aggressive low prices by highly utilizing their resources and by eliminating service features. The legacy carriers were losing market share, but could not match prices profitably because of their higher cost structures.

At the time, flights were only about half booked. As with all services, empty seats cannot be stored and perish at departure. In the airline industry, compared to the high fixed costs variable costs are negligible. American Airlines realized that with these characteristics empty seats had a marginal cost close to zero and could be sold at highly discounted prices.

The successful response of American Airlines was to engage in price discrimination by selling discount tickets to fill empty seats and match the aggressive prices. While attracting price-sensitive customers, the discounted tickets were restricted such that they did not appeal to the less price-sensitive business travelers that continued to purchase full-fare tickets.

However, price-sensitive leisure travelers usually book far in advance. Additionally, the discount tickets had an advance-purchase requirement. Hence, with price-sensitive demand arriving before high-value clients, the number of discount sales had to be limited to ensure no high-value demand was displaced. American at the time already had a computerized reservation system in place which they utilized in order to forecast demand and optimize the availability of discount tickets. Revenue Management (RM) was born and Bob Crandall, former chairman of American Airlines, is usually credited with its invention today.

American Airlines was able to match or undercut competitors' prices. Within months, the low-cost pioneer People Express that had experienced an unseen upsurge and was operating profitably was at the verge of bankruptcy. Don Burr, former CEO, attributes the failure to react to RM policies adequately for the downfall of the company (Cross, 1997).

In light of the success of American Airlines competing low-cost competition, all major American carriers followed and established RM policies. Smith et al. (1992) estimate that by utilizing RM policies American Airlines generated additional \$1.4 billion in revenues over a three-year period. Today, RM practices are a crucial factor for airlines worldwide (Pölt, 2002). Gains are estimated to be around 4-5%, which is comparable to an airline's profit in a good year (Talluri and Van Ryzin, 2004b).

RM has since seen much attention by researchers and practitioners and has grown to a mature business practice with sophisticated systems in place. While airlines have long remained the innovators in the field, many other industries today have introduced and elaborated RM practices. Examples are hotels, restaurants, car rentals companies, TV stations selling airtime, or energy suppliers.

One crucial assumption of RM models is the limited and fixed capacity in the short term. If capacity was flexible in the short term, all profitable demand could be accepted and there would be no need to restrict availability by RM measures. However, the assumption is overly restrictive in certain industries like the airline business, where a company sells several interchangeable capacities at the same time and only the overall capacity is fixed. An airline offers a number of flights and operates a whole fleet of aircraft. While the fleet and its overall capacity unquestionably is fixed in the short term, different planes with different capacities are usually available. Hence, the capacity of a single flight might be changed even shortly before departure. Realized demand and better forecasts closer to departure allow to better match demand and capacity. Cost savings flying with a smaller aircraft might well exceed additional revenues from selling seats at bargain prices. Similarly, bus companies operate a flexible fleet and car rental companies drive cars from one station to another to meet demand more effectively.

Other applications undoubtedly face a fixed capacity. A theater is not able to offer more seats when demand is high. But even for hotels, one might argue that a customer might be accommodated with a room from a competitor or a different hotel of the same chain. When a hotel engages in overbooking and sells more rooms than there are available in order to balance cancellations, this is usual practice to compensate booked customers that cannot be accommodated (Ivanov, 2006, Zhechev and Todorov, 2010).

In practice, airlines have been swapping aircraft between flights for a long time — a practice called Demand-Driven Re-Fleeting or Inclose Re-Fleeting. Berge and Hopper-

stad (1993) report manual processes at KLM and Australian. Pastor (1999) describes the (manual) process at Continental that is supported by simulations of different demand scenarios that help assess the benefits of swapping assigned equipment. At ANA, Oba (2007) reports that an automated system evaluates and proposes swaps to a revenue manager who then decides and executes the change. The only fully automated process is outlined in Zhao et al. (2007) at United Airlines. Zhao et al. (2007) report yearly benefits of \$5-\$10 million by swaps within a limited part of the fleet. Oba (2007) reports \$1.2 million per year at ANA.

While research has evolved and provided manifold extensions to basic RM models and important insights, the assumption of fixed capacities has remained unquestioned with few exceptions. Berge and Hopperstad (1993) provide a first analysis of systematic changes of capacities in simulation studies. Their results show a revenue potential of 1-5% increases. They apply a process that can be seen as industry practice today: After an initial equipment allocation, the booking process starts assuming fixed capacities. At certain planning points, the assignments are re-optimized restricted by the received bookings and using the updated demand forecasts. The process then continues with the updated capacities. A more recent study applying similar methods and using real-world data is the work by Frank et al. (2006). They find revenue gains of up to 2%.

The process formalized by Berge and Hopperstad (1993) allows to automate the evaluation and execution of equipment changes. However, in their framework RM optimizers do not anticipate possible changes in the future and still assume the current capacity as being fixed. Hence, depending on the final capacity allocation, controls might be overly or insufficiently restrictive in regards to low-value demand. Consequently, the benefits are strongly dependent on the time span between the re-optimization points of the assignments.

Few researchers have developed models to overcome this limitation and to anticipate possible swaps when optimizing RM controls. De Boer (2004) extends the popular EMSR-b heuristic. Wang and Regan (2006) develop a dynamic capacity control model. Both works build on given probabilities for the assignment of a certain capacity. The specification of adequate probabilities is itself challenging and problematic. Both works approximate the probabilities by the demand probabilities. Given two flights and two planes, the probability to assign the larger capacity to a flight is the probability of its demand exceeding the other flight's demand. Different valuations of the demand and additional costs of a swap (i.e. mainly changed fixed costs of the assignment) are ignored limiting the application to flights with the same origin and destination. Even then, valuing the demand equally is questionable because customers value the same flight at various times differently.

Only Wang and Meng (2008) propose a heuristic that is able to handle various demand structures with different valuations of flights. They extend a linear program used regularly by heuristics in capacity control settings. However, their model is not capable of including any costs that occur when changing the initial assignment either. Applications are hence again limited to flights with the same origin and destination.

Our motivation is to overcome the limitations of the current research and to provide optimization models that anticipate possible future changes of the capacity and that allow for different demand structures. Also allowing costs associated with a change, our models are not limited to flights with the same origin and destination or at least a similar cost structure.

Additionally, none of the existing publications analyze the proposed frameworks in comparison to an optimal policy¹. Our models are used to derive optimal policies and numerical structures found are used to find efficient algorithms and when developing heuristic approaches. In simulations, we compare the benefits of the policies derived through heuristics and the optimal policies.

Traditionally, RM has focused on so called capacity control models that ration supply for various products based on the same scarce capacity. Especially during the last decade with the gain of importance of low-cost carriers and the use of RM in new industries a different approach has attracted increasing attention: dynamic pricing. In a dynamic pricing setting, usually a single identical product is sold without artificial product versioning in order to segment demand. The price of the product is set dynamically over the booking horizon. Demand is stipulated by setting a lower price or discouraged by a higher price in order to match demand and capacity. The same restrictions normally apply to tickets sold by low-cost carriers regardless of the price the customer pays. In other industries, e.g. fashion retailing, versioning of the product is impossible once an order has been manufactured making capacity control policies inapplicable.

To be able to apply a dynamic pricing policy, naturally, the company must have flexibility in setting the price (Talluri and Van Ryzin, 2004b). If e.g. competition limits the ability to set prices freely, a company is more likely to apply a capacity control approach. The requirement for price flexibility also subsumes the ability to change prices dynamically in time at a reasonable or negligible cost. Before the Internet, setting individual prices was possible but usually expensive (e.g. customized mailings or frequent relabeling). Today, especially using the Internet as the main distribution channel, setting prices dynamically is possible at virtually no cost, which explains the

¹An optimal policy does not describe an ex-post optimal strategy, but we speak of an optimal policy if it maximizes the expected revenues. Refer to chapter 3 for more details on optimal policies.

increased acceptance and importance of dynamic pricing in the market. Given the choice, i.e. when all requirements are met for capacity control or dynamic pricing, the later is preferable (Talluri and Van Ryzin, 2004b).

Although, there has been a significant number of works dealing with dynamic pricing problems, still the vast majority of the RM literature concentrates on capacity control settings. The publications listed above dealing with possible changes of the limited capacity exclusively cover capacity control problems. Hence, another ambition of this work is to analyze the benefits of systematically considering changing capacity assignments in a dynamic pricing environment.

In the airline industry, dynamic pricing is applied mainly by low-cost carriers while legacy carriers sustain their established capacity control systems. While traditional carriers usually operate a heterogeneous fleet with a number of different aircraft types to meet requirements of various routes and their network structure, low-cost carriers often concentrate on a common fleet with few or even a single type. One extreme example is Ryanair that operates a completely homogeneous fleet of 275 Boeing 737-800 (Ryanair, 2011a). Their point-to-point network does not comprise connecting flights and flights are not meant to feed the demand for other legs. Not the fleet is chosen according to the requirements of the network, but the fleet characteristics are considered when analyzing potential new routes.

Even though capacity changes are not possible operating a common fleet, we are convinced that considering possible swaps in a dynamic pricing setting is worthwhile. There are other industries often applying dynamic pricing that have a limited flexibility in their capacities, e.g. train or bus companies. In Germany train transport has been liberalized and new competitors have only recently announced to soon enter the market (Schubert and Mrusek, 2009). Bus transportation still is heavily restricted, but legislation has pronounced its intention to liberalize the market (FOCUS Online, 2011). Increased competition will stipulate RM activity in these markets. American bus and train companies have been applying RM successfully for years. Additionally, even in the airline industry, we believe that dynamic pricing will continue to gain importance in practice and low-cost carriers might refrain from their single-type policies converging to more service orientated carriers.

The growth experienced by low-cost carriers in Europe has started to decline in recent years because the market starts to saturate (Binggeli and Pompeo, 2005). The North American market has seen similar signs of maturity. Attracting more demand is difficult because the no-frills services do not appeal to less price-sensitive segments. Adding cost-efficient service features is one option to attract different demand segments. Growth by offering new routes is also challenging because few short-haul routes

remain that have not been tapped by low-cost airlines and still yield profitable demand. In fact, low-cost carriers have increasingly been adding longer flights to their network. Offering long-distance routes, however, challenges the traditional low-cost business model that builds upon a high utilization of aircraft with short-haul traffic. Geographic expansion is additionally limited by fleet characteristics that determine efficient ranges and loads. To grow further, carriers might need to add different aircraft and might be forced to offer connections with short flights feeding long-distance capacities. Customers will perceive the increased number of destinations by allowing connections as an increased service level (Dunn and Dunning-Mitchell, 2011).

Hence, we expect the European market to experience some degree of assimilation in the near future. Low-cost carriers will need to move towards traditional airlines and to offer an increased level of service. They might offer long-distance destinations and abandon their point-to-point network and common-fleet strategy in order to facilitate future growth. Legacy airlines on the other side have already experienced the pressure imposed by low-cost competition and continue to react by cutting cost and offering bargain prices. Large carriers have established low-cost subsidiaries and integrate their offers and networks (airliners.de, 2011). An increasing number of network carriers today charges separately for service features such as additional baggage or credit card payment. The trend already becomes evident looking at airlines that offer low fares by cutting costs but at the same time offer higher service levels than usual low-cost carriers. Examples of these hybrid carriers are Air Berlin or JetBlue. Air Berlin, Germany's second-largest airline, offers long-distance flights, connecting itineraries, a loyalty program, and free snacks and drinks. Air Berlin is also a designated Oneworld alliance member and offers code-shared flights (Oneworld, 2011). Many low-cost airlines have started to sell tickets through Global Distribution Systems (GDS), which they had avoided in the past to save transaction fees. Other examples are Air Asia X or Tiger Airways in Asia that offer long-distance destinations. Other carriers like Ryanair ponder over similar offers (Clark, 2009). Table 1.1 highlights strategy differences for a number of exemplary low-cost airlines.

In the meantime, an option to facilitate operating different capacities in order match demand and supply more flexible while continuously exploiting the synergies of a nearly homogeneous fleet are crew-compatible fleet families such as the Airbus A320 or Boeing 737 families. Pilots are licensed to operate any plane of the family with a single type rating saving training and simulator costs and reducing standby requirements. Additionally, costs of spare parts and maintenance are reduced and the airline might be able to negotiate larger discounts when purchasing aircraft.

Table 1.1: Comparison of major LCC characteristics (updated from Belobaba (2009a))

	Southwest	JetBlue	AirTran	easyJet	Ryanair
Single aircraft type or single family of aircraft	✓	✗	✗	✗	✓
Point-to-point ticketing, no connecting hubs	✗	✗	✗	✓	✓
No labor unions, lower wage rates	✗	✓	✗	✗	✓
Single cabin service, no premium class	✓	✓	✗	✓	✓
No seat assignments	✓	✗	✗	✓	✓
Reduced frills for on-board service (vs. legacy)	✗	✗	✗	✓	✓
No frequent-flyer loyalty program	✗	✗	✗	✓	✓
Avoid global distribution systems (GDS)	✗	✗	✗	✗	✓

1.1 Contribution

The main contribution of this work are two dynamic-programming models that allow to set prices of several products simultaneously that use scarce capacities that might be interchanged. One model considers two flights only, while the second extends the scope to an arbitrary number of flight legs. We derive structures that allow for efficient calculations of optimal policies to maximize expected revenues. Costs of allocating the capacities are considered as well as differences of the flights regarding their demand, willingness-to-pay, and feasible price points.

Additionally, we develop two heuristics to consider capacity changes in a dynamic pricing environment. One conveys straightforward from common practice when applying capacity control strategies. Possible changes are not anticipated in the price policy and the capacity allocation is periodically re-optimized. The other extends a heuristic often applied in dynamic pricing that approximates the problem by its deterministic counterpart. Possible equipment changes do influence price policies derived with the second heuristic. Both heuristics allow for different demand structures and different allocation costs.

We simulate the booking process of an airline in different scenarios and test the different policies against each other. We analyze the benefits of swapping aircraft assignments

without and with prices reflecting possible future changes and the influence of various parameters.

1.2 Outline

Revenue Management refers to strategies and tactics that apply price discrimination policies. In chapter 2 we briefly introduce the principles and economic effects of price discrimination in general. In the following chapter we outline the theory and methodology of Markov Decision Processes which are the theoretic basis for the developed models.

Revenue Management itself is introduced in chapter 4. We start by outlining the general principles, processes, and controls. Then, static and dynamic capacity control models are described in detail before a one-leg dynamic pricing model is developed and structures are derived.

Chapter 5 introduces the airline planning process. The fleet assignment problem is considered in more detail in order to outline the concept, existing models, and algorithms of Demand-Driven Re-Fleeting.

The main contribution is presented in chapter 6. We develop a two-leg Demand-Driven Re-Fleeting model in a dynamic pricing context and then generalize the model to an arbitrary number of flights. Subsequently, two heuristics are developed. In simulation studies the benefits of applying Demand-Driven Re-Fleeting are analyzed and the policies derived from the heuristics are compared to the optimal policies.

We conclude with chapter 7 by summarizing the results. We additionally provide future research opportunities and other examples than the airline industry where the developed models and heuristics might be valuable.

Price Discrimination

In this chapter, we introduce the basic principles of price discrimination and their economic effects. Price discrimination describes a pricing policy whereby a company charges different prices to different consumer groups for the same product under the same circumstances (Fehl, 1981, Cole, 2008). Many authors narrow the definition to price differences not justified by a proportional difference in cost (e.g. Armstrong and Vickers, 2001).

Discrimination is usually motivated by profit maximization and achieved by the partial transfer of the consumer surplus to the producer's profit. Other motives might be to squeeze out competition or to enter a new market. Various policies can be applied including volume discounts, peak-load pricing, price-skimming and other forms of intertemporal pricing, bundling, price matching, or two-part pricing.

A producer engaging in price discrimination is usually only interested in his own benefits and consumers perceive the pricing strategy as benefiting only the producer. However, price discrimination in general might also bear positive welfare effects, e.g. if the overall output is increased or capacities can be utilized more efficiently.

We start by presenting general requirements to engage in discrimination in the next section. We then describe two different classifications of discrimination. Practical implications and problems are outlined, before we discuss the possibilities and effects in a competitive environment. A section briefly covering overall welfare effects and legislative limitations concludes the chapter.

2.1 Requirements

To introduce a discriminating pricing policy, the producer must have near monopoly control over the supply of the product or service. Under perfect competition, a purely price taking producer cannot discriminate prices. If the producer offers a premium price above the market's equilibrium, he loses all demand to his competitors. At a discounted price he does not maximize profits. However, price discrimination policies are still possible and found in practice under competition (c.f. e.g. Stole, 2007, Armstrong and Vickers, 2001). Competition might restrict the producer to set prices freely and limit him to a certain interval not to lose demand. This can also be valid for a monopolist facing the risk of potential competition. In a competitive environment discrimination might also be completely unprofitable even if the means to discriminate were available to the producer (see discussion below).

In order to convert the consumer surplus into the producer's profit, the demand function needs to show a negative slope. Also, the demand function must be stable to the introduction of price discrimination. Ott (1959) gives an example, where demand is dependent on the number of prices set. The demand function becomes steeper, the more segments are established. In the example, discrimination is not profitable and the monopolist maximizes profits by offering only the monopoly price.

The market segments in which different prices are charged need to be mutually independent. Customers' preferences and demand prices need to be different in the markets. The more heterogeneous customer behavior and preferences are, the more potential there is to exploit this heterogeneity (Talluri and Van Ryzin, 2004b). Arbitrage sales or demand change-over from one segment to the other must be impossible or unprofitable. In the case of personalized services, resales are inherently not possible. In other cases, contractual or other measures need to prevent resales. Examples are high changing fees for a personalized airline ticket or digital rights management that restricts usage to the consumer's computer.

The producer needs to have the means to identify and target the consumers in each segment separately. He must also be able to gather all information needed for the pricing decision, e.g. to estimate the price elasticities in the submarkets. Institutional or legal restrictions might constrain the manufacturer's ability to engage in price discrimination.

2.2 History and Classification

Discriminating prices have a long history and first publications date back to the late 19th century. Noteworthy, the first works studied price discrimination in the transportation service industry: railway tariffs (Dupuit, 1849, Lardner, 1850). Pigou (1920) established the theoretic economic foundation in his compulsory work. He classifies three types of price discrimination:

- *First degree* or *perfect* price discrimination is defined as the complete transfer of the consumer surplus to the producer by charging different prices for each unit of the product. The price for each unit is the maximum price the client is willing to accept. Hence, no consumer surplus remains and the marginal revenue curve coincides with the demand function. For example air cargo rates are usually individually negotiated with freight forwarders.
- In *second degree* price discrimination, the producer groups demand into several segments by individual preferences and charges each group a different price. The demand prices are not overlapping, and hence, the groups can be ordered by their demand prices. In the first group all clients with a demand price higher than f_1 are charged f_1 . The second group consists of all buyers with a demand price f such that $f_1 > f \geq f_2$ and f_2 is charged and so forth. An example for second degree discrimination is the pricing policy for Microsoft Windows. Various versions are sold appealing to different consumer groups with varying demand prices.
- In *third degree* discrimination, demand is clustered into groups by certain criteria, e.g. regional factors or by marketing channel. The demand prices within the groups can be overlapping. Of two customers in different groups with identical willingness-to-pay one might be able to purchase while the other is declined. Discounted student or senior fares are used to distinguish the lower-price segment from the less elastic segments.

Figure 2.1 gives an example of how a monopolist decides on his price and output in the case of linear demand and cost functions. In 2.1(a), he does not discriminate prices. Then his optimal output is located at the intersection of marginal cost (MC) and marginal revenue (MR). The shaded triangle shows the consumer surplus. In 2.1(b), the supplier applies second degree discrimination and groups demand into three segments. The total output increases and part of the consumer surplus is transferred to the supplier. The marginal revenue increases step-wise between the segments.

Figure 2.1(c) shows the situation when the monopolist enters first degree discrimination. He charges a different price for each unit and always charges the maximal price the customer is willing to accept. Hence, the consumer surplus is completely transferred to the producer's revenue. Note that the output in this case is equal to the output under perfect competition at the intersection of marginal cost and the competitive price. In comparison, only the producer's profit increases. First degree discrimination is the limiting case of second degree discrimination with one segment for each unit.

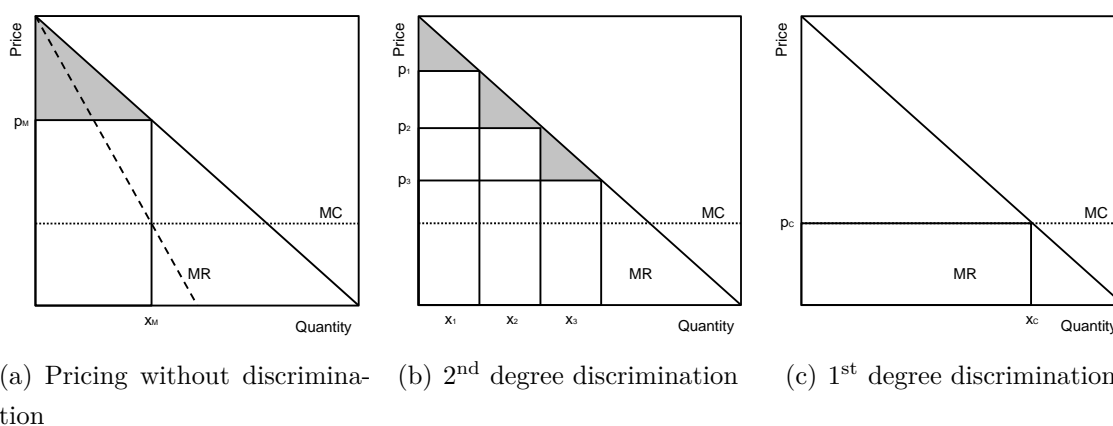


Figure 2.1: Monopoly prices and output quantities with and without price discrimination. The gray shaded areas mark the consumer surplus.

To illustrate the difference between second and third degree discrimination, assume we have a market with four buyers as given in Table 2.1. The supplier wants to sell four units and each buyer has a demand of one unit only. For second degree discrimination we segment by the demand price. We get two segments: client A and B are charged \$100 and client C and D are charged \$250. Total revenue earned is \$700 and a consumer surplus of \$300 remains.

If the producer segments by region, i.e. he applies third degree discrimination, he charges client A and C \$100 and B and D \$150 each. The revenue decreases to \$500 and the consumer surplus is now \$500.

Table 2.1: Market Segmentation Example

Buyer	Demand Price	Region
A	\$100	West
B	\$150	East
C	\$250	West
D	\$500	East

Weber (1956) constitutes another classification, which is non-overlapping with the classification by Pigou. He distinguishes *agglomerative* and *deglomerative* discrimination. While in the former separate markets are exogenously given, in the latter the manufacturer artificially splits one market into different segments. Deglomerative discrimination occurs in first, second, or third degree discrimination as outlined above. Third degree discrimination, on the other hand, can be either agglomerative or deglomerative. The distinction by Weber is important because the way prices and outputs are determined differ.

In agglomerative discrimination, the producer sets prices and his output completely separately. Each market's output is found at the intersection of marginal revenue and marginal cost as in a simple monopoly (see Figures 2.2(a) and 2.2(b)). Production cost and hence marginal cost are equal for both markets. Assuming linear demand functions, the total output is the same if the monopolist aggregated both markets to one (see Figure 2.2(c)). The less elastic market will be charged a higher price, while the price decreases in the market with more elastic demand. If both markets have the same elasticity, discrimination will not change prices or outputs. In the case of non-linear demand functions, in a monopolistic setting, the effect of discrimination on output and prices depends on the concavity or convexity of the demand (Varian, 1989).

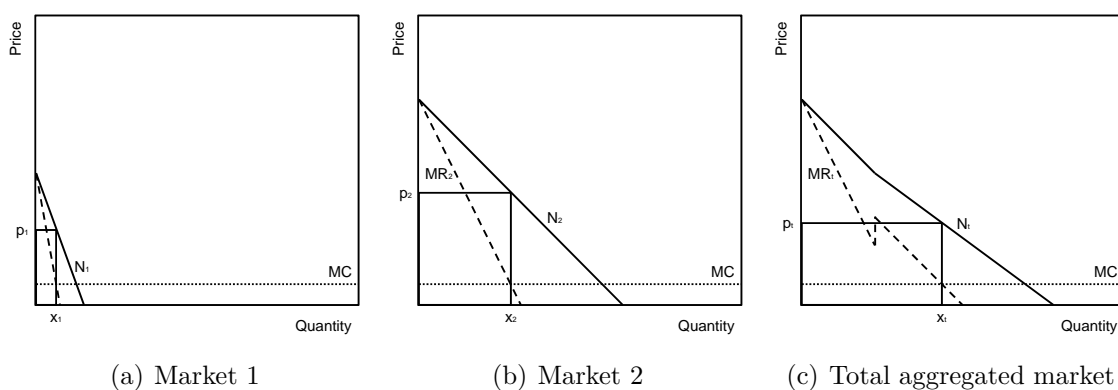


Figure 2.2: Agglomerative price discrimination. In comparison to a non-discriminating policy, the price is lowered on the more elastic market and raised on the less elastic market while the total output remains equal.

In practice, however, Weber's classification is blurred and can rather be seen as different phases of the discrimination process (Fehl, 1981). Usually, a monopolist seeks to split a market to enter discrimination based on differences of the willingness-to-pay, i.e. he is engaging in deglomerative discrimination. In practice, he needs to rely on some operative criteria that constitute the segments. Now, the prices actually need to be based upon agglomerative discrimination as the markets are already set by the previous phase. For that reason, especially in the English literature, usually third degree discrimination is used synonymously with agglomerative discrimination while

deglomerative discrimination is limited to first and second degree discrimination.

2.3 Practical Implications

While Pigou assumes the number and constitution of segments are previously determined, the major challenge in deglomerative discrimination is indeed to find the optimal number and composition. Ott (1959) and Lovell (1978) extend their models to endogenously determine the number of segments. They include costs of segmentation in their optimization models. These costs include the cost of information, the cost to target segments independently, and the cost to keep them separate. For each segment detailed information is required to identify the customers and to estimate their aggregated demand function. If resales are not inherently infeasible, the markets costly need to be separated and kept apart e.g. by versioning. Because the increase in revenue by adding another segment is declining in the number of segments, if there are segmentation cost, the trade-off between costs and benefits needs to determine the optimal number of submarkets. Figure 2.3 illustrates the cost-side and benefit-side approaches to find the optimum.

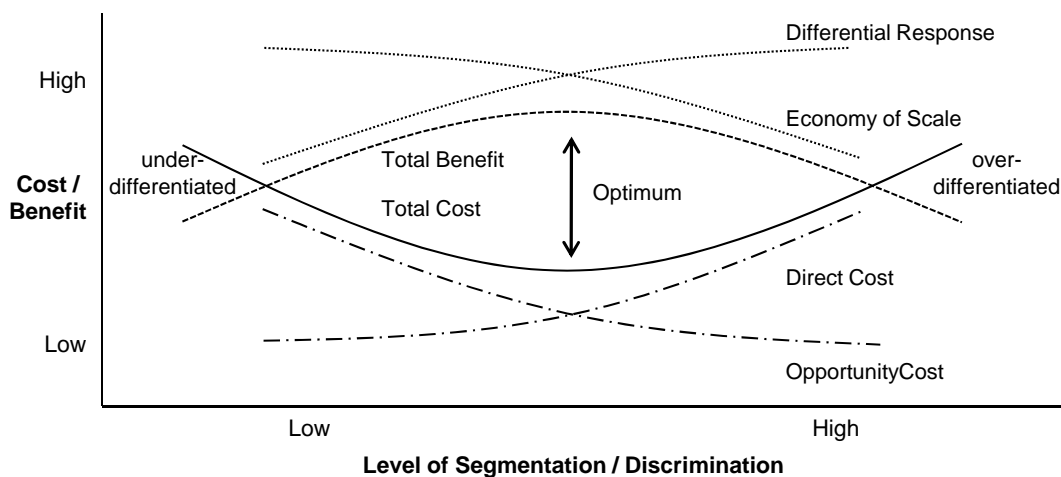


Figure 2.3: Costs and benefits of segmentation. The optimal level of segmentation can be found using either costs including opportunity costs or using benefits including economies of scale. Reprinted from Cui and Choudhury (2002).

A big portion of the consumer surplus is attained with a relatively small number of segments. Given linear demand, $n - 1/n$ portions of the consumer surplus are gained by establishing n submarkets. As a result, discrimination is usually performed with a limited number of segments in practice. For that reason and practical limitations gathering information and targeting individual buyers, first degree discrimination was long seen as rather academic and of limited importance in practice. Additionally, Ott

(1959) argues that even if first degree discrimination was optimal and practicable, then, the monopolist faces a bilateral monopoly. Therefore, the monopolist will not voluntarily weaken his position by entering the bilateral monopoly and negotiating prices separately with each customer.

Since more than a decade now, e-commerce has greatly improved companies' ability to discriminate prices. In the past, targeting individual customers e.g. by personalized physical catalogues was costly. Price updates were intricate to realize and to communicate. Today, offering individual prices through an online platform as well as price updates can be realized at negligible cost. First degree discrimination has become feasible, and hence has gained increased attention. Prices can also be updated frequently. Brynjolfsson and Smith (2000) find that online retailers adjust prices much more often and in much finer increments than conventional retailers. Airlines and hotels today adjust fares in near real-time. E-commerce eased not only the means to discriminate, but also the collection of required information. Through a combination of log-in data and loyalty programs, tremendous data are being collected. Using data-mining techniques, individual preferences can be deducted and used to discriminate.

If individual consumers cannot be identified and classified, one way facilitating discrimination in practice is *self-selection*. Self-selection describes the situation in which individuals are given the incentive to select themselves into a segment. The product or service is altered in a way that each variation appeals to one customer segment and the change to another variant, i.e. segment, is discouraged. Discrimination by self-selection needs to be distinguished from true product differentiation. The differences in cost for each variation do not justify the price differences. An example is selling a product with a slightly modified look under a different brand name for a different price. Airlines sell tickets with different restrictions in the same compartment for different prices.

2.4 Discrimination in a competitive environment

Normally, a monopolist benefits from a discriminatory policy in terms of profits. In a competitive environment, the basic principles of discriminatory prices remain the same as outlined above. However, effects on prices and thus benefits for the supplier and the consumers and on the overall welfare are ambiguous. Generally, if the demand is symmetric in the sense that companies rank segments equally by their price elasticities, the consequences are similar (Corts, 1998). In particular, discrimination leads to higher prices for the less elastic market and respectively lower prices on the more elastic market. In such an environment, companies benefit from discrimination. Examples for symmetric rankings are student fares or airline ticket restrictions tailored to leisure

travelers. Both segments are usually conceived as low-price groups.

If firms conceive the segments differently, one might price aggressively in a market another company considers high-value provoking competitive responses. In that case, discrimination might lead to *all-out competition* characterized by lower prices for all consumers (Corts, 1998). Companies then would be better off applying a uniform pricing strategy. However, if a company can increase profits by discrimination given fixed competitors' pricing policies, firms will be locked-in in a prisoner's dilemma where discrimination is a dominant strategy.

Ulph and Vulkan (2000, 2001) distinguish two major effects: the *enhanced surplus extraction effect* arises because firms can charge prices closer to the customers' reservation prices which positively affects revenues and profits. The *intensified competition effect* contrarily effects profits and originates from the fact that firms compete for consumers in every segment that each constitutes a single market. They analyze competitive effects under first and second degree discrimination (in their second paper in conjunction with mass customization). The results agree with Corts (1998). In many situations, especially when tastes are similar and buyers are not extremely loyal, the latter effect dominates the former and discrimination is not advantageous for the companies. Armstrong and Vickers (2001) get similar results for third degree discrimination.

Varian (2004) criticizes that Ulph and Vulkan (2000, 2001) assume full information about the consumers at the firms' disposal. He argues that a long-time supplier knows the customers' preferences and habits better than a potential competitor. Hence, the long-time supplier can offer superior personalized products and services than a competitor could. Then, the intensified competition effect is reduced and the enhanced surplus extraction effect might prevail.

For a detailed survey of price discrimination under competition the reader is referred to Stole (2007).

2.5 Welfare effects and legislation

As outlined in some remarks, overall welfare effects of discrimination are ambiguous. A necessary condition for increased welfare is an enlarged output (Varian, 1985). Especially when a monopolist serves markets that he would not supply under uniform pricing, social welfare benefits by cross-subsidization. Popular examples are sales of medicine in developing countries or subsidized train routes. This effect is well known as the *output enhancing effect*. In many cases the effect dominates the welfare loss by quality distortion for versioning. However, even when there is a positive output effect, the optimal allocation is deteriorated. In other cases overall welfare is reduced. For a

more comprehensive discussion of welfare implications we refer to Varian (1985, 2000). Especially in the context of revenue management or peak-load pricing, another positive effect is the efficient utilization of capacities. A demand shift away from peak times is encouraged by lower prices in off-peak periods. Hence, overall, less capacity suffices to satisfy demand allowing for lower average prices and possibly reducing negative external effects necessary to build up or maintain capacities.

In numerous countries price discrimination is restricted by legislation. In the United States, the Robinson-Patman Act¹ enacted in 1936 prohibits volume discounts not justified by cost reductions. Other forms of discriminatory prices are illegal if the motive is to reduce competition. In the European Union, *applying dissimilar conditions to equivalent transactions with other trading parties, thereby placing them at a competitive disadvantage*² is prohibited. Legal intervention against price discrimination is often criticized for restraining the freedom of contract and competition in favor of small and medium sized companies (Fehl, 1981).

¹U.S. Code, Title 15: Commerce and Trade, Chapter 1: Monopolies and Combinations in Restraint of Trade, Section 13

²Consolidated versions of the Treaty on European Union and the Treaty on the Functioning of the European Union, Article 102 (c), March 30, 2010

Markov Decision Processes

This chapter provides a brief introduction to the theory and methodology used in the subsequent chapters. We focus on finite-horizon, discrete-time Markov Decision Processes (MDPs) that are sufficient for our applications. For a comprehensive review of MDPs in general, the reader is referred to White (1993), Puterman (1994), or Hu and Yue (2008). Hinderer (1970) and Schäl (1975) provide insights to a general framework of MDPs.

In complex business environments, a decision maker is often faced with situations that require sequential decisions rather than a single decision that influence the overall outcome. Often one decision influences the subsequent state of the system and decision options available in the future. The decision maker is able to influence the subsequent development of the system by choosing actions. Hence, the decision maker tries to find a sequence of actions such that the outcome of the process is optimal in regards to a certain predefined criterion. These sequential decision problems are the subject of the dynamic programming theory.

Stochastic decision processes are used to model systems that do not evolve deterministically, but rather are subject to a probabilistic influence. A special case are Markov Decision Processes introduced by Bellman (1957). They feature the *Markov property*, i.e. the future state of the system does not depend on the entire history of states and actions of the process, but only the last observed state and the current action influence the evolution of the system. We focus on a special case, namely MDPs in discrete time with a finite planning horizon. We limit our discussion to finite state and action spaces and do not consider discounting.

3.1 Model and Notation

A stochastic system is observed at discrete points in time $t = T, T - 1, \dots, 0$ over the finite planning horizon T . Note that we count time backwards and the process starts in $t = T$. At the beginning of each time epoch, the system state $i_t \in I_t$ is observed and a decision maker selects one action $a_t \in A_t(i_t)$ from a set of feasible actions. The set of admissible actions $A_t(i_t)$ depends on the state of the system and the time of the decision. The decision maker receives a one-stage reward $r_t(i_t, a_t)$ before the system then evolves stochastically to the subsequent state $i_{t-1} \in I_{t-1}$. The reward r_t and the probability $\mathbf{p}_t(i_t, a_t, i_{t-1})$ for an evolution to a subsequent state depend on the current state and the action selected. The decision maker thus can influence the future evolution of the system by selecting an appropriate action. At the end of the planning horizon $t = 0$ the decision maker receives a terminal reward $V_0(i_0)$ depending on the final state $i_0 \in I_0$ of the system.

Such a system is modeled by a MDP with a finite planning horizon, which is the tuple $(T, I, A, \mathbf{p}, r, V_0)$ with

- (i) the planning horizon $T \in \mathbb{N}$,
- (ii) the non-empty, finite state space $I_t \subseteq I$ at each time $t = T, \dots, 0$,
- (iii) the set of actions A and the non-empty, finite set of admissible actions $A_t(i) \subseteq A$ in each state $i \in I_t$ at $t = T, \dots, 1$. We additionally define $\mathfrak{D}_t := \{(i, a) \in I_t \times A : a \in A_t(i)\}$.
- (iv) The transition law $\mathbf{p}_t : \mathfrak{D}_t \times I_{t-1} \rightarrow [0, 1]$ specifies the counting densities $\mathbf{p}_t(i, a, \cdot)$ for a transition to $j \in I_{t-1}$.
- (v) The one-stage reward is given by the functions $r_t : \mathfrak{D}_t \rightarrow \mathbb{R}, t = T, \dots, 1$ and
- (vi) the terminal reward by the function $V_0 : I_0 \rightarrow \mathbb{R}$.

To account for our needs in later chapters, we model an *inhomogeneous* system with the transition probabilities and one-stage rewards being dependent on time. Note that an inhomogeneous system can be easily transformed into a homogeneous system by adding a time dimension to the state space. Also, we model the state space I_t and the set of admissible actions $A_t(i)$ in state i dependent on time. In later applications this becomes necessary to accurately model the constraints. Where sufficient, we assume the state space and the set of admissible actions to be independent of time.

A map $\mathbf{f}_t : I_t \rightarrow A$ satisfying $\mathbf{f}_t(i) \in A_t(i)$ for all $i \in I_t$ is called a *decision rule*. Applying decision rule \mathbf{f}_t means that at time t given state i , the action $\mathbf{f}_t(i)$ is chosen.

Let Γ_t denote the set of all decision rules at time t . A sequence of decision rules $\phi = (\mathbf{f}_T, \mathbf{f}_{T-1}, \dots, \mathbf{f}_1)$ with $\mathbf{f}_t \in \Gamma_t$ for $t = T, \dots, 1$ is called a *policy* or *strategy*. It specifies the decision rule to use at each decision time. We denote the set of all (deterministic Markovian) policies by $\Phi = \times_{t=T, \dots, 1} \Gamma_t$.

Let J_t be the state of the system at time $t = T, \dots, 0$. Then, given a policy $\phi \in \Phi$, the sequence of states forms a stochastic process $\{J_t, t = T, \dots, 0\}$ which is well defined and is indeed a Markov chain (Hinderer, 1970). The probability for a realization j_T, j_{T-1}, \dots, j_0 is given by the product measure on $\mathfrak{J} := \times_{t=T, \dots, 0} I_t$

$$P_\phi(J_T = j_T, J_{T-1} = j_{T-1}, \dots, J_0 = j_0) = \\ P(J_T = j_T) \cdot \mathbf{p}_T(j_T, \mathbf{f}_T(j_T), j_{T-1}) \cdot \mathbf{p}_{T-1}(j_{T-1}, \mathbf{f}_{T-1}(j_{T-1}), j_{T-2}) \cdot \dots \cdot \mathbf{p}_1(j_1, \mathbf{f}_1(j_1), j_0)$$

under policy ϕ . Let E_ϕ denote the expectation with respect to P_ϕ . Usually, the initial state j_T of a process will be fixed. We then use the conditional expectation $E_\phi(\cdot | J_T = j_T)$.

At each time $t = T, T-1, \dots, 1$ the decision maker receives a reward $r_t(J_t, \mathbf{f}_t(J_t))$ and at the end of the planning horizon a terminal reward $V_0(J_0)$. All rewards depend on the realization of the process and as such are random variables under any policy. To compare different policies and finally to find an optimal strategy, we first need to define a performance measure. In later chapters we apply the expected total reward criterion, which is introduced in the next section.

3.2 Decision Criteria

The decision maker tries to find the best strategy in regards to his objective. In general our objective will be to maximize the total reward which is discussed in detail below. Where suitable, costs can be modeled as negative rewards. Other possible criteria include the maximization of the discounted total reward, the maximization of utility, or to maximize the probability to achieve a particular target. For details on other objectives we refer the interested reader to White (1993) and Bouakiz and Kebir (1995).

The total reward $R^T := \sum_{t=1}^T r(j_t, a_t) + V_0(j_0)$ depends on the realization of the process j_T, j_{T-1}, \dots, j_0 , which can in part be influenced by the decision maker by applying a policy $\phi \in \Phi$, i.e. by choosing appropriate actions a_T, a_{T-1}, \dots, a_1 in states j_T, j_{T-1}, \dots, j_1 . Applying ϕ , the total reward

$$R_\phi^T = \sum_{t=1}^T r(J_t, \mathbf{f}_t(J_t)) + V_0(J_0)$$

is itself a random variable. We can compare two policies by their conditional expectation of the total reward $E_\phi(R_\phi^T | J_T = j_T) =: V_{T,\phi}(j_T)$.

For $t = T, \dots, 1$, the *value function*

$$V_t(j_t) := \max_{\phi \in \Phi} V_{t,\phi}(j_t) = \max_{\phi \in \Phi} E_\phi(R_\phi^t | J_t = j_t) \quad , j_t \in I_t$$

is the maximum total expected reward until the end of the planning horizon starting in state j_t . A policy $\phi^* \in \Phi$ is called *optimal* if $V_T(j_T) = V_{T,\phi^*}(j_T)$ for all $j_T \in I_T$.

Since we assume a finite state and action space, there exists a deterministic Markovian policy that is optimal (Puterman, 1994, Proposition 4.4.3).

In the next section we show how such an optimal policy can be obtained.

3.3 Backward Induction

To determine the value functions and to obtain an optimal policy, instead of enumerating and evaluating all strategies, we use the more efficient method of *backward induction*. We make use of the following Theorem (Büning et al., 2000, Theorem 9.2.):

Theorem 3.3.1 (Optimality Equation).

For $t = 1, \dots, T - 1, T$

(i) V_t is the unique solution of the equation

$$V_t(i) = \max_{a \in A_t(i)} \left\{ r_t(i, a) + \sum_{j \in I_{t-1}} \mathbf{p}_t(i, a, j) V_{t-1}(j) \right\} \quad , i \in I_t. \quad (3.1)$$

(ii) Every policy $\phi = (\mathbf{f}_T, \mathbf{f}_{T-1}, \dots, \mathbf{f}_1)$ formed by actions $\mathbf{f}_t(i)$ maximizing the right-hand side of (3.1) is optimal.

(3.1) is also known as the *Bellman Equation*. If there are t periods left, the maximal reward is the sum of the one-stage reward of the action that maximizes the right-hand side of (3.1) and the expected remaining reward using an optimal policy. Optimal actions hence balance short- and long-term rewards. A maximizing action needs not to be unique. Often we will therefore add rules such as to use the largest maximizing argument in case of more than one optimal action.

To obtain an optimal policy, initially using the terminal value V_0 , we determine an action maximizing the right-hand side of (3.1) iteratively for each $t = 1, \dots, T$ using Algorithm 3.3.1.

Often only the total expected reward for a fixed initial state $V_T(i_T), i_T \in I_T$ and an optimal policy are of interest to the decision maker. Then, to save memory, only the

Algorithm 3.3.1 Algorithm to determine an optimal policy and all expected values.

Input: MDP $(T, I, A, \mathbf{p}, r, V_0)$

Output: Optimal policy $(\mathbf{f}_T^*, \mathbf{f}_{T-1}^*, \dots, \mathbf{f}_1^*)$

Expected values $(V_T, V_{T-1}, \dots, V_0)$

1. $t = 0$
 2. **while** $t < T$
 3. $t = t + 1$
 4. $I' = I_t$
 5. **while** $I' \neq \emptyset$
 6. Take an element $i \in I'$
 7. $V_t(i) = \max_{a \in A_t(i)} \left\{ r_t(i, a) + \sum_{j \in I_{t-1}} \mathbf{p}_t(i, a, j) V_{t-1}(j) \right\}$
 8. $\mathbf{f}_t^* = \arg \max_{a \in A_t(i)} \left\{ r_t(i, a) + \sum_{j \in I_{t-1}} \mathbf{p}_t(i, a, j) V_{t-1}(j) \right\}$
 9. $I' = I' \setminus \{i\}$
-

values of (3.1) of the current and the last iteration need to be saved. Algorithm 3.3.2 provides such an implementation with reduced memory requirements.

The algorithms provided both iterate the decision times, all states at the time, and all feasible actions for each state. Depending on the length of the planning horizon and the size of the state and action spaces, finding an optimal policy is computationally demanding. Many applications yield structures that allow to develop more efficient algorithms. We provide an overview of these structured policies in the next section.

3.4 Structured Policies

Many applications yield structures that convey to optimal decision rules and can be used to efficiently calculate and store optimal policies for an application. Some policies can even be completely characterized using few parameters. We start with a simple example of an application with such a structured optimal policy.

Example: Selling an asset (A stop problem) At times $t = T, \dots, 1$ a seller receives an offer x_t for an asset, which is a realization of a random variable $X_t \in \{0, \dots, M\}$. Let X_T, \dots, X_1 be independently identically distributed with $P(X_t = x) = q(x)$. The seller can accept or reject the offer. In case he accepts the offer, he receives x_t and the process stops. If he declines the offer, he cannot come back to the offer at a later time and the sales process continues in the next period $t - 1$. At the end of the sales period, if the asset has not been sold, he receives nothing and the sales process is stopped.

Algorithm 3.3.2 Algorithm to determine an optimal policy and the total expected value.

Input: MDP $(T, I, A, \mathbf{p}, r, V_0)$

Output: Optimal policy $(f_T^*, f_{T-1}^*, \dots, f_1^*)$

Total expected value V_T

1. $t = 0$
 2. $I' = I_0$
 3. **while** $I' \neq \emptyset$
 4. Take an element $i \in I'$
 5. $V''(i) = V_0(i)$
 6. $I' = I' \setminus \{i\}$
 7. **while** $t < T$
 8. $I' = I_t$
 9. **while** $I' \neq \emptyset$
 10. Take an element $i \in I'$
 11. $V'(i) = V''(i)$
 12. $I' = I' \setminus \{i\}$
 13. $I' = I_t$
 14. $t = t + 1$
 15. $I'' = I_t$
 16. **while** $I'' \neq \emptyset$
 17. Take an element $i \in I''$
 18. $V''(i) = \max_{a \in A_t(i)} \left\{ r_t(i, a) + \sum_{j \in I'} \mathbf{p}_t(i, a, j) V'(j) \right\}$
 19. $f_t^* = \arg \max_{a \in A_t(i)} \left\{ r_t(i, a) + \sum_{j \in I'} \mathbf{p}_t(i, a, j) V'(j) \right\}$
 20. $I'' = I'' \setminus \{i\}$
 21. $I' = I_t$
 22. **while** $I' \neq \emptyset$
 23. Take an element $i \in I'$
 24. $V_T(i) = V''(i)$
 25. $I' = I' \setminus \{i\}$
-

Let the state of the system $I_t = \{0, \dots, M\} \cup \{\infty\}$ represent the current offer at t . A state $i_t = \infty$ denotes the case, when the asset has been sold. The action space is $A = \{0, 1\}$, where $a = 1$ describes the acceptance of the current offer while declining an offer is represented by $a = 0$. In case the asset has been sold, an action does not have an impact at all. Hence, in all states, $A_t(i_t) = A$.

Considering $V_t(\infty) = 0$ for all $t = T, \dots, 1$ and the boundary condition $V_0 \equiv 0$, we reduce the optimality equation (3.1) to

$$V_t(i) = \max(i, \sum_{x=0}^M q(x)V_{t-1}(x)) \quad (3.2)$$

for $t = T, \dots, 1$.

The optimal policy $(f_T, f_{T-1}, \dots, f_1)$ built from decision rules

$$f_t(i) = \begin{cases} 0 & i < i_t^* \\ 1 & i \geq i_t^* \end{cases}$$

with $i_t^* = \sum_{x=0}^M q(x)V_{t-1}(x)$ is obvious from (3.2). We can further calculate the critical values i_t^* efficiently by

$$i_t^* = \begin{cases} 0 & t = 1 \\ \sum_{x=0}^M q(x) \max(x, i_{t-1}^*) & t = T, \dots, 2. \end{cases}$$

In the example, the optimal decision rule at each time can be characterized by only one parameter. Instead of storing a table for each decision time with the state and the respective optimal action, only the critical values need to be saved. Further, the threshold increases in time, which is an intuitive result since the decision maker reduces his minimum price as the sales horizon elapses. Exploiting the structure, an algorithm can efficiently calculate the optimal policy by recursively calculating the critical values $i_t^*, t = T, \dots, 2$. The example is discussed in more detail in Büning et al. (2000), where several other simple examples are presented.

The example shows how a structured policy can greatly reduce memory requirements to store an optimal policy and the computational effort to calculate it. In addition, a structured policy is easily understood and implemented by end users which again increases the acceptance of the strategy. Finding structured optimal policies that can be computed efficiently and which are intuitive and exercisable in practice is one of the central challenges of dynamic optimization. Powell (2007) highlights the importance of identifying structured optimal policies as *one of the most dramatic success stories from the study of Markov decision processes*.

In general, applications are more complex and even if models yield structured optimal policies, to find these and to prove their optimality is difficult and technically

demanding. In order to show structures, we usually need to prove that the value function is monotone or concave. Additionally, optimal policies normally cannot be stated explicitly as in the example and need to be calculated numerically for each application.

Policies $(f_T, f_{T-1}, \dots, f_1)$ with decision rules of the form

$$f_t(i) = \begin{cases} a_1 & i < i_t^* \\ a_2 & i \geq i_t^* \end{cases}$$

are called *threshold* or *control-limit* policies. $a_1, a_2 \in A_t(i)$ are distinct actions which are optimal depending on if the state $i \in I_t$ is above or below the critical value i_t^* . If we know a threshold policy is optimal, the problem of finding an optimal policy reduces to finding the threshold value at each decision time. Note that we implicitly assumed that the states can be completely ordered. Usually, the states have an intuitive interpretation providing a natural order, e.g. the monetary interpretation in the example described above or when the state represents a stock level as in our later applications.

Threshold policies are a special type of *monotone* policies. A monotone policy features decision rules that are non-increasing or non-decreasing in the system state. An optimal monotone policy might greatly reduce computational effort because during the numerical evaluation of the maximum in (3.1), some actions might not need to be considered at all. Examples of monotone decision rules can be found in chapter 4.

In chapter 6, we develop models with a multidimensional state space. A threshold policy in that case conveys to a *switching curve*, i.e. the threshold value in one dimension is a function of the other dimensions of the current state. The state space is separated into domains where a certain action is optimal. Figure 3.1 depicts two examples of switching curves. Note that switching curves need not necessarily be monotone functions.




Figure 3.1: Examples of switching curves. Switching curves divide the state space into multiple domains with a certain action being optimal within a domain.

Revenue Management

In this chapter, we introduce the basic concept of Revenue Management (RM). We start by presenting definitions and briefly outlining the history and development in the next section. We then discuss assumptions and requirements and sketch the general RM process. In section 4.4 interrelations and problems with collecting data and forecasting are outlined. Section 4.5 and section 4.6 supply a detailed discussion of the respective approaches and optimizations models applied. We build on these basic models in the subsequent chapters to develop more advanced models and heuristics.

4.1 Definition and History

Revenue Management refers to strategies and tactics used to predict and influence consumer demand in order to maximize revenue from constrained resources (Rosenberg, 2010). A simpler, more vivid, and popular definition originates from American Airlines in 1987: *Selling the right seats to the right customers at the right prices* (Smith et al., 1992). Today, some authors extend it by *at the right time and through the right distribution channel* (Pölt, 2002). RM is also known as Yield Management. Note that the definitions subsume the two main control strategies *capacity control* and *dynamic pricing*, which are introduced in the following sections, as well as other approaches such as overbooking. Contrarily, Revenue Management is often used only to describe capacity control problems (e.g. Phillips, 2005).

Before deregulation of the U.S. airline industry, market entry, routes, schedules, and fares were tightly regulated by the Civil Aeronautics Board (CAB). Hence, to optimize revenue, airlines were only able to engage in overbooking strategies to countervail losses by *no-shows*. No-shows are booked passengers that do not appear for boarding and

thus leaving seats empty on a flight. In 1961, the twelve largest American airlines faced no-shows of up to 10% of the total demand (Rothstein, 1971). In the early sixties, airlines started to *overbook* flights, i.e. to sell tickets in excess of the actual capacities. Airlines must balance the revenue gain from filling empty seats against the risk of denied boardings when more booked passengers show up for the flight than seats are available. Customers denied boarding need to be compensated and transferred to other flights. Denied boarding costs also include goodwill losses involved.

When the U.S. Congress passed the Airline Deregulation Act in 1978, restrictions on the domestic market were gradually phased out until 1983. Especially People Express, a no-frills carrier founded in 1981, entered the market with aggressive low prices. With fares up to 70% below those of the established carriers, it initially focused on untapped price-sensitive markets and showed a tremendous growth over the next years. The low fares imposed a major threat to the legacy airlines. Matching prices would not cover their cost, while only with low prices customers could be retained. The solution was price discrimination in the form of restricted discounted tickets matching the low prices. In early 1985, American Airlines introduced their *Ultimate Super Saver* fares that were subject to an advance-purchase restriction of two weeks and required a stay over a Saturday night. The fares did not appeal to the high-value business segment, but allowed to target the price-sensitive leisure market. The number of available discounted tickets was limited to ensure to be able to satisfy all high-value demand. Within one year, People Express was at the verge of bankruptcy and was eventually sold to Continental Airlines. In the light of the great success of American Airlines competing the low-cost competition, all major U.S. carriers introduced discounted fares and controls were increasingly computerized. Revenue Management was born and has since evolved to a highly sophisticated disciplinary today used by airlines worldwide. The impact of RM is huge: in 1992 American Airlines estimated an increase of revenue of \$1.4 billion over three years (Vinod, 2009). Pölt (2002) states that today, to any airline, RM is a crucial factor and airlines would not survive without it. RM has been adopted by many other industries, such as car rentals, hotels, broadcasting media, freight, theaters, retail, or manufacturing. Airlines, for many years, have remained the innovators regarding strategies and systems. Today, research is also partially driven by other industries' particular needs.

4.2 Requirements and general assumptions

To apply RM, services or products need to show the following characteristics, some of which are direct requirements for price discrimination (c.f. chapter 2):

- **Heterogeneous Demand.** Demand needs to be heterogeneous in terms of the willingness-to-pay and preferences on product features. Various demand groups need to be distinguishable into different segments. In the airline industry, different market segments, mainly leisure and business travel, clearly show different preferences and price sensitivities.
- **Arbitrage Prevention.** Resale must be impossible or uneconomical in order to prevent arbitrage sales. Airlines usually apply high fees for changing discounted tickets and tickets are personalized.
- **Advance Purchases.** Customers need to be able to purchase tickets in advance. The need for RM arises because customers with different valuations of the product arrive sequentially in time. High-value demand usually arrives later than demand for lower fare classes. The problem is to accept as many passengers early in the booking period to fully utilize the available capacity, but at the same time only as many such that high-value demand is not displaced. If tickets cannot be purchased in advance or customers do not arrive sequentially in the order of their willingness-to-pay, RM is useless or unnecessary.
- **Perishable Product.** The product must be perishable or storable only at significant cost. If products are storable at reasonable prices, other means to balance demand variations are more effective and efficient, e.g. producing and storing seasonal goods. Likewise, demand must not be storable and cannot be satisfied at a later time.
- **Limited Capacity.** The available capacity is limited and replenishment is impossible or costly in the short run. If short-term capacity upgrades are possible at modest cost, capacities do not need to be reserved for high-value demand and all profitable demand can be satisfied. Then, only the question of setting prices remains. While other forms of intertemporal pricing might be reasonable to consider, the dynamic nature of a RM strategy in regards to the observed remaining inventory is unnecessary and useless.

The fleet composition of an airline is a long-term decision and cannot be changed quickly because of long delivery times. However, the common assumption of a fixed capacity allocated to a specific product or resource is too strict in many applications. Short-term capacity adjustments are possible by changing the assignment of an aircraft from one flight to another. The main contribution of this work are optimization models and heuristics to obtain pricing strategies under the weaker assumption of the overall capacity being fixed, i.e. when the network capacity is fixed rather than the capacity assigned to each leg. These approaches

are introduced in chapter 6. In this chapter we introduce basic RM models and retain the assumption of a fixed capacity on a product or leg basis.

- **Stochastic Demand.** If demand varies stochastically over time and capacities cannot be adjusted in the short term, demand needs to be managed to match capacity. Otherwise the capacity could be adjusted to match demand in the long run.

At first glance, RM approaches might be considered even if demand is deterministic but fluctuating. However, similarly to the capacity requirement, the dynamic nature of a RM strategy is then unnecessary. Other forms of price discrimination, e.g. peak-load pricing, should be considered instead.

- **Marginal Costs.** Marginal costs are nil or negligible. In the light of the capacity requirement, fixed or quasi-fixed costs are usually high with low variable costs. For example, the costs of an additional passenger, i.e. e.g. airport fees or an additional menu, are marginal in comparison to the total costs of flying the airplane. Then, maximizing revenue approximately maximizes profits. If marginal costs are significant, contribution rather than revenue needs to be optimized. RM models in manufacturing usually optimize contribution (e.g. Wiggershaus, 2008). Hence, negligible marginal costs are not a necessary requirement. However, products or services with insignificant marginal costs that fulfill the other requirements naturally lend themselves to RM policies.
- **Market Acceptance.** The market needs to accept the pricing policy. Deutsche Bahn, the state-owned German railway operator, changed its traditional single-price policy late 2002 and introduced discounted fares with advance-purchase requirements and high fees to change reservations. At the same time, it stopped offering a 50% discount card which had been used by frequent travelers. Customers initially did not accept the new policies and the number of passengers declined by 10.6 % during the first months. Under public pressure Deutsche Bahn had to change its prices again after only seven months (DER SPIEGEL, 2003). The discount card has again been offered since and change fees were cut by 2/3. The advance-purchase discounts have been maintained and slowly have become accepted. The example shows that introducing RM strategies must be done carefully in order not to discourage clients.
Similarly, if price is a signal for quality, e.g. for luxury goods, discriminating policies without versioning are likely to have a negative impact on sales and revenues (Talluri and Van Ryzin, 2004b).

For a successful application of RM policies, in addition to the product requirements listed above, technological and managerial support are essential. Managers need confi-

dence in science and new technologies to accept automated decision systems. At least the pricing and the RM department are affected by and need to collaborate on the RM process. Sales and marketing departments also need to support systems and decisions. Technology needs to be in place to gather data, to accurately forecast demand, and to automatically apply controls. Online sales or sales through distribution systems naturally qualify for RM. If technical systems are not in place, to install these might itself be a challenging and costly proposition.

Additionally to the requirements in regards to the service or product, throughout the remainder, we adhere to the following common assumptions:

- Myopic customers. We will assume costumers do not behave strategically. They purchase as soon as the price is below their reservation price, i.e. their valuation of the product. Customers do not optimize the buying decision, e.g. postpone the purchase in the hope of lower prices. Considering strategic clients often makes models intractable.

Depending on the application, the assumption is also less harmful than it might seem. Often, especially when shopping for non-durable goods, consumers are spontaneous buyers and also have too little information to act strategically. The future price of an airline ticket depends on current bookings and the remaining inventory for sale, the time left for sales, and the strategy of the airline that reflects its expectations of future demand. In general, a consumer does not have any information about current bookings or the airline's expectations.¹ Hence, the customer cannot foresee future prices and act strategically. With high-value demand arriving late, ticket prices usually increase towards the departure of a flight. Then, a customer does not benefit from postponing his purchase in the hope for lower prices. If necessary, an airline can further take measures to prevent strategical behaviour e.g. by pursuing and communicating a strategy that prohibits decreasing prices.

Additionally, the demand forecasts, while assuming myopic customers, still indirectly include previous buying strategies. The forecasts are based on past purchases that reflect the strategic behavior. For example, if consumers wait for markdowns, the forecasts of the reservation prices are lower towards the end of the season.

¹Today, some airlines do reveal limited information about their inventories towards the end of the sales process, i.e. they post a price together with the number of seats left for sale or left for sale at the current price. However, the airline voluntarily releases the information only to encourage an immediate consumer response. The airline would refrain from publishing inventories if it did not expect to benefit from it.

Hence, models assuming myopic buyers are widely used in practice. For references concerning models considering strategic buyers we refer to Talluri and Van Ryzin (2004b).

- **Infinite Market.** We do not consider durable goods and hence assume that purchase probabilities are independent of past purchases. The assumption is reasonable because the market size is large compared to the offered capacity. Although one buyer is unlikely to purchase the same ticket again, the overall demand probabilities remain unchanged.
- **Monopolistic environment.** The observed demand is assumed to depend only on the offered price at a time and not on competitive actions. Similarly to the assumption of myopic customers, forecasts based on historical data indirectly include competitive strategies. Shugan (2002) suggests that assuming a monopolistic environment might be better than adopting complex competitive models. The later assume that competitors have followed and continue to follow optimal strategies which are reflected in the empirical data. The assumption is questionable in practice as well and the increased complexity is not offset by better results. If competitors change their strategies, either model is likely to perform poorly.

In the travel industry, depending on the origin and destination, routes may be highly competitive while others show little competition (Cole, 2008). However, even in a competitive environment, the market is inherently inhomogeneous if e.g. departure times or service levels differ. Hence, monopolistic models seem reasonable and gain important insights. They are widely found in practice. Nevertheless, recent studies analyze competitive effects in simplified frameworks and deduct important strategical insights (e.g. Isler and Imhof, 2008, Gallego and Hu, 2009, Talluri and Martínez de Albéniz, 2010).

4.3 Revenue Management Process

Talluri and Van Ryzin (2004b) in general distinguish *quantity-based* and *price-based* RM. Legacy carriers usually apply quantity-based capacity controls, i.e. they ration the availability of various predefined discounted fares. These *booking classes* or *fare classes* are well differentiated by restrictions such as advance-purchase requirements in order to encourage self-selection. The practice of restricting fares in such manner is called *fencing*. Contrarily, low-cost carriers often follow a price-based approach called *dynamic pricing*. They offer a single product with equal restrictions. The price for a ticket changes dynamically in time depending on the realized demand, the time

until departure, and updated expectations. Either strategy can be combined with overbooking.

Quantity- and price-based RM share a common conceptual process that is visualized in Figure 4.1. The following steps are continuously repeated (c.f. Talluri and Van Ryzin, 2004b):

- **Data Collection.** To be able to understand factors that influence demand and to forecast the demand historical data is needed. Ticketing data collected at airports and from distribution systems is readily available allowing airlines to also include competitive data in their models. As the booking process evolves to departure, data on realized demand is collected and included in the forecasts.
- **Forecasting.** The parameters of the demand model need to be forecasted. A major issue is that the input data is constrained. Only data on realized sales and not on rejected requests is available. However, unconstrained demand needs to be forecasted. Forecasts are updated repeatedly throughout the booking process.
- **Optimization.** Based on the forecasted demand booking limits or prices are optimized. As demand realizes and forecasts are updated, the controls are periodically re-optimized.
- **Booking Control.** The optimized controls are then applied to the arriving demand. This includes updating availability and prices through various distribution channels.

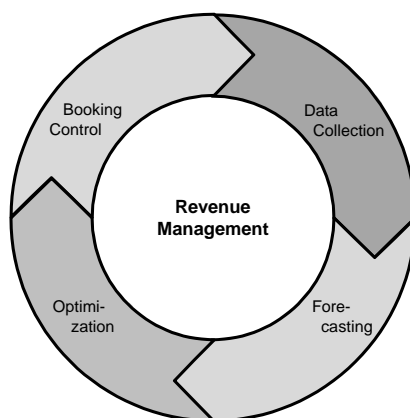


Figure 4.1: Revenue Management Process.

4.4 Data collection and Forecasting

Booking controls are optimized using forecasts that are based on the collected data. Hence, the forecasting quality greatly influences the quality of the resulting policy.

Lee (1990) estimates that a 10% increase in forecast accuracy yields revenue increases of up to 3%. A deep discussion of forecasting methods is beyond the scope of this work and we only sketch problems below and provide references for further reading. In the remainder we assume forecasts of all input data to the optimization models are available.

Most problems are not unique to RM applications and vast literature exists on methods that can be applied. Talluri and Van Ryzin (2004b) provide an introduction focused on RM and further references. A more detailed summary of current methods at the time is presented by Zeni (2001a). Surveys by McGill and Van Ryzin (1999) and Boyd and Bilegan (2003) provide further references.

In general, demand, cancellations, and no-shows must be forecasted. The actual parameters to be forecasted differ greatly depending on the RM approach and the actual models in use. Also, the level of granularity varies for different models. For example, static models require the estimation of the total demand per fare class. Contrarily, dynamic models need forecasts of the demand in each time period over the booking horizon. Depending on the parameters and the needed granularity, the sheer amount of values needed might present a problem. The number of forecasts is multiplied by the number of updates for each value. Lufthansa produces several million forecast values per day (Pölt, 2002). The amount of input data requires sophisticated data processing technologies. Somewhat surprisingly in regards to the law of large numbers, Weatherford et al. (2001) find that disaggregated forecasts strongly outperform aggregated forecasts in the hotel industry. The required brake-down process deteriorates the accuracy gain of forecasting aggregated values. The results are found to be similar for airlines in Weatherford and Kimes (2003).

Historical data is available from various sources such as distribution systems, airports, passenger records, shopping data, or accounting data (Mishra et al., 2005). Choosing sources and possibly combining data from various sources is a major issue. In general, historic data must be relevant for future predictions. Changes to the product, the economy, or the competitive environment need to be considered. Events and seasonalities likewise need to be included in the forecasts. Outliers need to be eliminated. Usually, a user interface allows analysts to adjust data according to external factors and some systems directly link data to external factors (Talluri and Van Ryzin, 2004b).

While historical data on bookings is readily available the data does not reflect demand. Sales are constrained by capacities and booking controls. If censored sales data is used as demand, forecasts tend to be lower than the actual demand. As a result, not enough capacity is reserved for high-value customers and revenue is lost. Subsequent forecasts are based on even more censored data which amplifies the revenue dilution. The pat-

tern repeats itself and is known as *spiral-down* effect. Cooper et al. (2006) model and analyze the effect and provide conditions for it to occur. The censored data can be unconstrained using statistical methods. A comprehensive discussion of unconstraining methods can be found in Zeni (2001a). Zeni (2001b) gives a short overview of different methods and their performance. He finds Expectation-Maximization algorithms to outperform other approaches. Another simulation study by Weatherford and Pölt (2002) yields the same results. The authors estimate the revenue gain of sophisticated methods to be in the range of 2% to 12%. More recently, customer-choice-based approaches have been developed and show good performances in tests with real-world constrained data (Sasse and Beleliev, 2010, Cote, 2010).

4.5 Capacity Control

In this section, basic capacity control models are introduced. These are RM models that control demand by rationing the supply in a revenue optimal manner. In contrast to dynamic pricing models (c.f. chapter 4.6), capacity control models assume given prices for different products or booking classes. Typically, in airline settings, these are determined by a separate pricing department before the sales process starts. Then, during the sales process, only the availability of the different products is controlled applying a RM system.

We first present common assumptions that we adhere to throughout the remainder. We then introduce different forms of booking controls before we describe different models in detail and present optimal policies and structures. A brief review of more advanced models relaxing some of the assumptions concludes the section.

For capacity control models, we adhere to the following common assumptions:

- There are no cancellations, no-shows, or go-shows. Each accepted reservation cannot be rejected later without significant costs. Consequently, overbooking is not considered.
- A denied passenger request is lost to the airline. Rejected demand cannot be recaptured on other flights and customers only request a certain booking class. They do not purchase a lower (buy-down) or higher (sell-up) fare class if the requested class is closed. Hence, demand for different fare classes is mutually independent and independent of any controls applied.
- Demand for different legs is mutually independent. Hence, we consider leg-based rather than network models. Network models are often intractable in practice and

single-leg models are still widely used. Additionally, many heuristics decompose the network problem into a collection of single-leg problems.

- Group arrivals can be partially accepted.
- The decision maker is risk-neutral. Hence, the *expected* revenues are maximized. Optimization based on expectations seems reasonable in the light of the number of departures every day. The decision process is repeated independently for each flight and every day of departure. Because of the large number of problem instances, the impact of a single actual realization is low and long-term average revenues are maximized using a risk-neutral approach.

4.5.1 Booking Controls

Booking Limits

A *booking limit* is the maximum number of units of the capacity available for sale to a class at a specific point in time. *Partitioned* booking limits divide the capacity into exclusive blocks that are only available to the respective class. For example, if a capacity of 30 is allocated to two classes, class 1 might have a booking limit of 20. Class 2 then has a booking limit of 10. We assume a higher fare in class 1. If all seats reserved for a class have been sold, further purchases are not possible, even if the requested class yields higher revenues than other open fare classes. In the example, if 20 seats of class 1 and less than 10 tickets of class 2 have been sold, a request for class 1 is declined even though the fare is higher than the fare of class 2.

To avoid suboptimal allocations in such manner, *nested* booking limits reserve capacity for classes with higher fares. A nested booking limit is the maximum number of units available for sale to the class and all classes with lower fares. The booking limit of the class with the highest fare is equal to the total capacity. In the described situation, requests for class 1 are accepted as long as seats are available. At the same time, a maximum of 10 seats is sold to class 2.

Protection Levels

A *protection level* describes the number of capacity units that are reserved for a booking class. The concept of partitioned and nested booking limits conveys directly to protection levels. Partitioned protection levels simply equal the booking limits. A nested protection level is the number of seats reserved for a class and all classes with

a higher fare. The nested protection level y_j for class j is

$$y_j = C - b_{j+1} ,$$

where C is the total capacity and b_{j+1} is the booking limit of the next class with a lower fare. The protection level of the lowest fare class, i.e. the capacity reserved for all classes altogether, is equal to the total capacity ($y_n = C$).

Bid Prices

A *bid price* is a control threshold. Only requests yielding revenues higher than the bid price are accepted. The major difference to booking limits or protection levels is obviously the revenue orientation, while the former are based on capacity units. Bid prices cannot allocate capacity in a partitioned way. All requests with higher fares are accepted. Talluri and Van Ryzin (2004b) note that bid prices need to be dependent on the remaining capacity to guarantee that sales do not exceed the capacity. A single static bid price might result in sales exceeding the available capacity. If the bid price is a function of the current capacity, an optimal policy is equivalent to a policy using booking limits or protection levels. Figure 4.2 illustrates the equivalence of the different controls.

In practice, the price of a booking class might change over time (Curry, 1990). Then, bid-price controls are preferable because the same booking class can be rejected or accepted depending on the current fare. Protection levels and booking limits are optimized using an average fare of the booking class. When a class is open, all requests are accepted regardless of the actual fare.

4.5.2 Static Models

Static models assume that the arrival process can be divided into non-overlapping time periods by fare class. During a time interval, only demand for one fare class is observed. Hence, demand for each fare class can be viewed as aggregated demand. Usually, demand is assumed to arrive sequentially ordered by fare, i.e. low-fare demand arrives before high-fare demand. This simplification is usually justified by advance-purchase requirements on discount fares and is widely used in practice.

Littlewood's Model

The first single-resource capacity control model was developed by Littlewood (1972) (McGill and Van Ryzin, 1999, Littlewood, 2005). He assumes two classes with fares

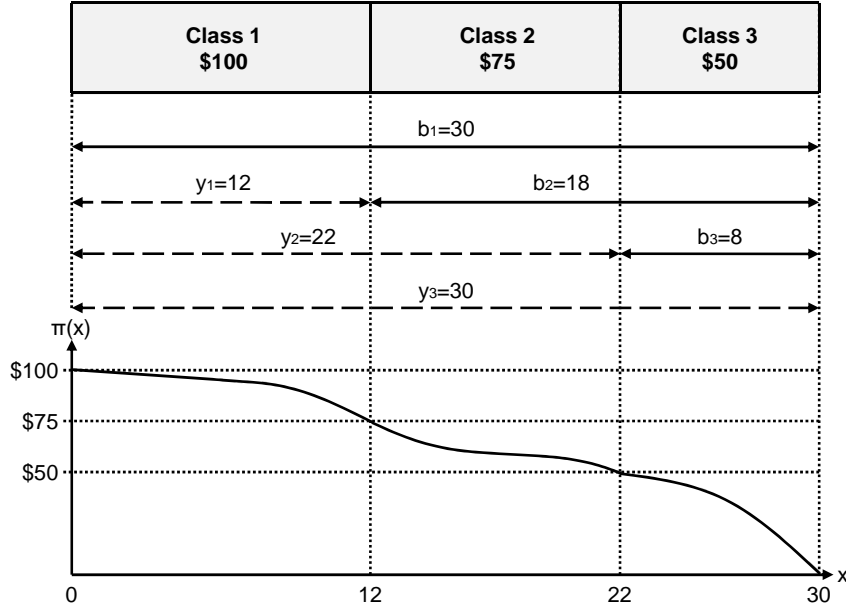


Figure 4.2: The relationships between booking limits b_j , protection levels y_j , and bid prices $\pi(x)$ dependent on a remaining capacity x . Reprinted from Talluri and Van Ryzin (2004b).

$f_1 > f_2 \geq 0$ and demand $D_j \in \mathbb{N}_0, j = 1, 2$. $P(D_1 \geq y_1)$ denotes the probability that class 1 demand exceeds the protection level y_1 , i.e. the maximum risk that a high-yield customer needs to be declined because of having accepted too many class 2 clients. Assessed with the fare of class 1 it yields the expected value $f_1 P(D_1 \geq y_1)$ of increasing the protection level by one unit to $y_1 + 1$. The expected value of increasing the optimal protection level y_1^* by another unit needs to be offset by the collected fare f_2 . Hence, the optimal protection level y_1^* needs to satisfy

$$f_2 < f_1 P(D_1 \geq y_1^*)$$

and

$$f_2 \geq f_1 P(D_1 \geq y_1^* + 1) .$$

This is known as *Littlewood's Rule*. Note that in a nested control setting, we implicitly assumed that the lower fare class sells out. If it did not sell out, because high-fare demand has access to all available seats, the displacement cost would be lower than f_2 .

Static n -class Model

Curry (1990), Wollmer (1992), and Brumelle and McGill (1993) independently extended Littlewood's model to $n > 2$ classes. Curry (1990) and Brumelle and McGill (1993) assume continuous demand. We will briefly sketch the model by Wollmer (1992) below. He uses discrete demand distributions. Li and Oum (2002) show the equiva-

lence of the optimality conditions of the three models. A short overview of all three models is given in Lautenbacher and Stidham (1999).

Let $D_j \in \mathbb{N}_0, j = 1, \dots, n$ be the demand for class j . According to the assumptions, $D_j, j = 1, \dots, n$ are distributed mutually independent and independent of the controls applied. W.l.o.g. we assume fares $f_1 > f_2 > \dots > f_n \geq 0$. We define a Markov Decision Process $(n, S, A, \mathbf{p}, r, V_0)$ with

- (i) the horizon $n \in \mathbb{N}, n > 2$. We count time backwards, and hence, based on the assumption of demand arriving in order from low-fare to high-fare, demand of class j materializes in period j . At $j = 0$ all remaining units perish, i.e. the flight departs and empty seats cannot be sold any longer.
- (ii) The state space is $S = \{0, \dots, C\}$, where C is the total capacity available. The state $c \in S$ refers to the remaining capacity at the current period.
- (iii) The action space is $A = \{0, \dots, C\}$. $a \in A$ specifies the maximum number of clients to accept. In state $c \in S$, the set of feasible actions $A(c) = \{0, \dots, c\}$ is limited by the remaining capacity.
- (iv) The transition law is

$$\mathbf{p}_j(c, a, c') = \begin{cases} P(D_j = d) & c' = c - d, d < a, \\ P(D_j \geq a) & c' = c - a, \\ 0 & \text{otherwise,} \end{cases}$$

for $j = 1, \dots, n$ with $c, c' \in S, a \in A(c)$.

- (v) The one-stage reward function is

$$r_j(c, a) = f_j \left[aP(D_j \geq a) + \sum_{d=0}^{a-1} dP(D_j = d) \right] = \sum_{d=0}^{a-1} f_j P(D_j \geq d + 1)$$

for $j = 1, \dots, n$ with $c \in S, a \in A(c)$.

- (vi) The terminal reward function is $V_0 \equiv 0$.

For any function u , we define $\Delta u(x) := u(x) - u(x - 1)$. Where necessary, we use a subscript to indicate the variable.

The value function for $j = 1, \dots, n$ and $c \in S$ is given by

$$V_j(c) = \max_{0 \leq a \leq c} \left\{ \sum_{d=0}^{a-1} \left[f_j P(D_j \geq d + 1) + P(D_j = d) V_{j-1}(c - d) \right] + P(D_j \geq a) V_{j-1}(c - a) \right\} \quad (4.1)$$

$$= \max_{0 \leq a \leq c} \left\{ V_{j-1}(c-a) + \sum_{d=0}^{a-1} \left[f_j P(D_j \geq d+1) + \Delta V_{j-1}(c-d) P(D_j \leq d) \right] \right\} \quad (4.2)$$

$$= \max_{0 \leq a \leq c} \left\{ \sum_{d=0}^{a-1} \left[[f_j - \Delta V_{j-1}(c-d)] P(D_j > d) \right] + V_{j-1}(c) \right\} \quad (4.3)$$

$$= \max_{0 \leq a \leq c} \left\{ \sum_{d=0}^{a-1} \left[[f_j - \Delta V_{j-1}(c-d)] P(D_j > d) \right] \right\} + V_{j-1}(c) . \quad (4.4)$$

We explicitly state (4.2) to highlight the equivalence to the model in Wollmer (1992). As the maximum is taken over a finite set for each j and c , there is an action a maximizing the inner expression. We refer to these a as being optimal. The value function shows structures as stated in the following Proposition:

Proposition 4.5.1. The function V_j defined in (4.1) has the following properties:

- (i) $V_j(c)$ is concave in c for all $j = 1, \dots, n$.
- (ii) $\Delta V_j(c)$ is non-decreasing in j for all $c \in S, c > 0$.

Proof. We only give proof of (ii). For a proof of (i), we refer to Wollmer (1992), Theorem 2. Fix $0 < c \leq C$ and let $a^* \in A(c-1)$ be optimal in $(c-1) \in S$ at stage $j \in \{2, \dots, n\}$. Then, using (i),

$$\begin{aligned} \Delta V_j(c) &\geq \sum_{d=0}^{a^*-1} \left[[f_j - \Delta V_{j-1}(c-d)] P(D_j > d) \right] + V_{j-1}(c) \\ &\quad - \sum_{d=0}^{a^*-1} \left[[f_j - \Delta V_{j-1}(c-1-d)] P(D_j > d) \right] - V_{j-1}(c-1) \\ \Leftrightarrow \Delta V_j(c) - \Delta V_{j-1}(c) &\geq \sum_{d=0}^{a^*-1} \left[[\Delta V_{j-1}(c-1-d) - \Delta V_{j-1}(c-d)] P(D_j > d) \right] \\ &\geq 0, \end{aligned}$$

which completes the proof. \square

It is easy to see that $V_j(c) \geq 0$ for all $c \in S, j = 0, \dots, n$. Using Proposition 4.5.1 (i), an optimal policy $\tau^* = (a_1^*, \dots, a_n^*)$ is given by

$$a_j^*(c) = \min \{ a \in \{0, \dots, c\} : f_j < \Delta V_{j-1}(c-a) \} \quad , \text{ for } j = 1, \dots, n, c \in S$$

with $\min \emptyset := c$. τ^* can be rewritten as a threshold policy:

Theorem 4.5.2. *There exists an optimal policy $\tau^* = (a_1^*, \dots, a_n^*)$ such that*

$$a_j^*(c) = \begin{cases} c - y_{j-1}^* & , c \geq y_{j-1}^* , \\ 0 & , \text{ otherwise,} \end{cases}$$

for $j = 1, \dots, n, c \in S$ with the optimal protection levels

$$y_j^* := \max \{c \in \{0, \dots, C\} : f_{j+1} < \Delta V_j(c)\} \quad , \text{ for } j = 1, \dots, n-1 ,$$

where $\max \emptyset := 0$, $y_0^* := 0$, and, by convention, $y_n^* := C$.

Note that the optimal protection levels are not dependent on the current capacity. Figure 4.3 illustrates an example of optimal controls.




Figure 4.3: Example of different booking controls. Assume the total capacity of $C = 30$. The resulting optimal protection levels are $y_1^* = 12$ and $y_2^* = 30$. Then, the booking limits are $b_1^* = 30$ and $b_2^* = C - y_1^* = 18$. At a current capacity of $c = 22$, 8 seats have been sold. We accept class 2 requests up to $a_2^*(c) = c - y_1^* = 10$.

Proposition 4.5.1 (ii) and the assumption of increasing fares $f_1 > \dots > f_n$ imply a nested structure:

Corollary 4.5.3. The optimal protection levels y_j^* are increasing in j .

EMSR Heuristics

The most popular capacity control heuristics are the *expected marginal seat revenue* heuristics version a (EMSR-a) (Belobaba, 1987a,b, 1989) and the slightly modified version b (EMSR-b) (Belobaba and Weatherford, 1996). Both are based on Littlewood's rule and extend it to $n > 2$ fare classes. The heuristics are widely found in practice, although computing optimal controls is not difficult. This is mainly due to their early implementation into RM systems at a time, when research lagged behind practice and computational power was highly expensive. The heuristics are easy to implement and run quickly, while their performance is usually quite close to the optimum (Wollmer, 1992, Robinson, 1995).

EMSR-a The EMSR-a computes protection levels by adding up protection levels generated using Littlewood's rule relative to all higher-fare classes. At stage $j \in$

$\{2, \dots, n\}$, one computes protection levels $y_{j-1,i}^*$, $i = j-1, \dots, 1$ such that

$$f_j < f_i P(D_i \geq y_{j-1,i}^*)$$

and

$$f_j \geq f_i P(D_i \geq y_{j-1,i}^* + 1) .$$

To obtain the protection level for all higher classes, i.e. the number of units to protect at stage j , the individual protection levels are added up:

$$y_{j-1}^a = \sum_{i=1}^{j-1} y_{j-1,i}^* .$$

EMSR-a was shortly believed to be optimal, but except for the highest fare class, EMSR-a does not yield optimal controls. Counter examples can be found e.g. in Wollmer (1992), Brumelle and McGill (1993), or Talluri and Van Ryzin (2004b). Controls are usually lower than the optimal controls because the statistical averaging effect is ignored. In general, however, controls can be overly or insufficiently protective (Brumelle and McGill, 1993). Although the controls gained by EMSR-a differ significantly from the optimal ones, revenue performance is usually quite close to the optimum (Wollmer, 1992, Robinson, 1995). EMSR-a performs significantly worse, when there is a large number of booking classes with fares close to each other (Talluri and Van Ryzin, 2004b).

EMSR-b The EMSR-b is also based upon Littlewood's rule, but instead of summing protection levels, demand of all higher fare classes is aggregated. The aggregated demand is treated as an artificial booking class with its fare computed from the weighted average fare of the aggregated classes. Hence, at stage $j \in \{2, \dots, n\}$, future demand is aggregated to

$$\bar{S}_{j-1} = \sum_{i=1}^{j-1} D_i .$$

The demand \bar{S}_{j-1} and the fare

$$\bar{f}_{j-1} = \frac{\sum_{k=1}^{j-1} f_k E(D_k)}{\sum_{k=1}^{j-1} E(D_k)}$$

of the artificial class are then used to determine its protection level by

$$f_j < \bar{f}_{j-1} P(\bar{S}_{j-1} \geq y_{j-1}^b)$$

and

$$f_j \geq \bar{f}_{j-1} P(\bar{S}_{j-1} \geq y_{j-1}^b + 1) .$$

$E(D_k)$ denotes the expectation of the demand D_k for class k .

The heuristic ignores effects of protection levels at future stages. Hence, revenue from all future demand *accepted* is approximated by the revenue from *total* future demand. However, it is likely that not all future demand will be accepted because of applied booking controls. EMSR-b is more common in practice and implemented more frequently than EMSR-a. Talluri and Van Ryzin (2004b) state that it usually performs better than EMSR-a, although they cite a study by Pölt (1999) with real-life data in which neither heuristic is found to be dominating the other.

Extensions to the static model

Pfeifer (1989) and Brumelle et al. (1990) consider the two class static model relaxing the assumption of independent demand. This includes generally correlated demand of the fares and correlation due to booking controls resulting in sell-ups, i.e. low-fare customers purchasing a full fare ticket when the discount class is closed. Bodily and Weatherford (1995) develop a heuristic for $n > 2$ classes with dependent demand and extend the two class model to integrate overbooking.

Belobaba and Weatherford (1996) provide a version of EMSR-b to include up-sell effects. Hopperstad (2000) finds that their model ignores up-sell effects in the underlying demand distribution and suggests an iterative approach to find the true demand distribution. Gallego et al. (2009) generalize the model in Brumelle et al. (1990). They also adapt the EMSR-b heuristic to include customer choice for $n > 2$ fare classes. In simulations the proposed model outperforms previous models for legs with high demand. For modest load factors the revenue gain is insignificant. Walczak et al. (2010) use data transformation to incorporate customer choice using the traditional EBSR-b heuristic.

Robinson (1995) relaxes the assumption of arrivals in increasing fare order and derives optimal policies.

Barz and Waldmann (2007) adapt the classic model to a risk-averse decision maker that maximizes expected utility. Weatherford (2004) develops an *expected marginal seat utility* heuristic based on EMSR for risk-averse decision makers. The heuristic is also discussed in Barz and Waldmann (2007) and Barz (2007).

Van Ryzin and McGill (2000) develop a heuristic that requires no separate forecasting or uncensoring and no cyclical reoptimization using an adaptive stochastic approximation method.

Curry (1990) considers network effects on optimal controls. Williamson (1992) provides a detailed overview of network effects on RM and a good literature review at the time. She develops and investigates the *prorated* EMSR scheme based on EMSR-a and

network adjusted revenues. A more recent detailed overview of network methodologies and controls is given in Talluri and Van Ryzin (2004b).

Belobaba (1987a) extends EMSR-a to include overbooking. A detailed and more recent discussion of cancellations and overbooking in a static context is presented in Chi (1995).

4.5.3 Dynamic Models

Dynamic capacity control models relax the assumption on the arrival process and allow requests to arrive in an arbitrary order. Hence, demand cannot be aggregated by fare class as in static models. All other assumptions are retained. Additionally, to make the models tractable, demand is assumed to be Markovian, i.e. the number of arrivals in a time window follows a Poisson distribution. This assumption is often criticized because it restricts the variance to equal the mean. Greater levels of variability as often found in practice cannot be captured. However, Walczak (2006) uses a compound Poisson process to amplify variance within the existing framework. We will retain the assumption of demand being Poisson-distributed throughout the remainder.

For convenience, time is usually finely discretized such that the probability of more than one request arriving in each epoch is negligible. One exception is the model by Lautenbacher and Stidham (1999). Throughout the text, we will adhere to the assumption of at most one customer arriving during each time epoch. Note that these periods need not be of equal length. At the beginning of the booking horizon, when few request arrive, periods of several days might be used. Close to departure epochs might represent time intervals of less than one hour.

Because of the use of disaggregated demand and possibly short time epochs, forecasting is more challenging for dynamic models than for static models. Hence, the availability of adequate forecasts is a major decisive factor for the model choice.

Dynamic n -class Model

We briefly outline the model and structural results as introduced by Lee and Hersh (1993). Note that this model differs from many other texts, which assume the class of the arriving request to be known *before* making the decision to decline or accept the request (e.g. Lautenbacher and Stidham, 1999, Barz, 2007). Both ways of modeling yield equivalent results in terms of optimal policies. For further discussion on the equivalence of the two approaches, we refer to Talluri and Van Ryzin (2004b) or Barz (2007). As a result, in our formulation, the decision variable is a vector constituted of the accept or deny decisions for each booking class. Contrarily, models assuming

known request classes only need scalar decision variables denoting the decision for the arriving booking class.

Again, we assume ordered fares $f_1 > f_2 > \dots > f_n \geq 0$ w.l.o.g. We denote a booking class by $j = 1, \dots, n$. Let λ_t^j be the probability for a request of class j at time t . We define a MDP $(T, S, A, \mathbf{p}, r, V_0)$ with

- (i) the planning horizon $T \in \mathbb{N}$. The length of the booking period is divided into periods with at most one customer arriving. Note that we count time backwards starting the booking period at $t = T$. The booking period ends at $t = 0$ with the departure.
- (ii) The state space is $S = \{0, \dots, C\}$, where C is the total capacity available. The state $c \in S$ refers to the remaining capacity at the current period.
- (iii) The action space is $A = \{0, 1\}^n$. For $a \in A$, $a(j)$ specifies the accept ($a(j) = 1$) or deny ($a(j) = 0$) decision for fare class j . In states $c \in S, c > 0$, the set of feasible actions is $A(c) = A$. In state $c = 0$, no capacity remains and no further requests can be accepted. Hence, $A(0) = \{(0, \dots, 0)'\}$.
- (iv) The time-dependent transition law is

$$\mathbf{p}_t(c, a, c') = \begin{cases} \sum_{j=1}^n a(j)\lambda_t^j & c' = c - 1, \\ 1 - \sum_{j=1}^n a(j)\lambda_t^j & c' = c, \\ 0 & \text{otherwise,} \end{cases}$$

for $t = T, \dots, 1$ with $c, c' \in S, a \in A(c)$.

- (v) The one-stage reward function is

$$r_t(c, a) = \sum_{j=1}^n a(j)\lambda_t^j f_j$$

for $t = T, \dots, 1$ with $c \in S, a \in A(c)$.

- (vi) The terminal reward function is $V_0 \equiv 0$.

The value function for $t = T, \dots, 1$ and $c > 0 \in S$ is given by

$$V_t(c) = \max_{a \in \{0,1\}^n} \left\{ \sum_{j=1}^n \left[a(j)\lambda_t^j [f_j + V_{t-1}(c-1)] \right] + \left[1 - \sum_{j=1}^n a(j)\lambda_t^j \right] V_{t-1}(c) \right\}. \quad (4.5)$$

For $c = 0$ the only feasible action $(0, \dots, 0)'$ is trivially optimal. Hence, with $V_0 \equiv 0$, $V_t(0) = 0$ for all $t = T, \dots, 1$.

(4.5) can be rewritten as

$$V_t(c) = V_{t-1}(c) + \sum_{j=1}^n \max_{a(j) \in \{0,1\}} \{a(j) \lambda_t^j [f_j - \Delta V_{t-1}(c)]\} . \quad (4.6)$$

The basic model has important characteristics and structures that we briefly present in the remainder of the section.

Proposition 4.5.4. The value function V_t has the following properties:

- (i) $V_t(c)$ is concave in c for all $t = T, \dots, 0$.
- (ii) $\Delta V_t(c)$ is non-decreasing in t for all $c \in S, c > 0$.

Proof. For a proof we refer to Lee and Hersh (1993) Theorem 1 and Theorem 2. \square

As a consequence of Proposition 4.5.4, there exists a set of critical time-dependent capacities $\{c_{t,j}^*\}$, such that a request of fare j at time t is accepted, if and only if $c \geq c_{t,j}^*$. Hence, an optimal policy can again be stated through nested protection levels. In contrast to static models, the controls are now dependent on time.

Theorem 4.5.5. *There exists an optimal policy $\tau^* = (a_T^*, \dots, a_1^*)$ such that*

$$a_t^*(j, c) = \begin{cases} 1 & , c \geq y_{t,j-1}^* \\ 0 & , \text{otherwise,} \end{cases}$$

for $j = 1, \dots, n, t = T, \dots, 1, c \in S$ with the optimal protection levels

$$y_{t,j}^* := \max \{c \in \{0, \dots, C\} : f_{j+1} < \Delta V_{t-1}(c)\} \quad , \text{ for } j = 1, \dots, n-1, t = T, \dots, 1 ,$$

where $\max \emptyset := 0, y_{t,0}^* := 0$ for all t , and, by convention, $y_{t,n}^* = C$ for all t .

The protection levels are monotone in time and in the booking classes as stated in the following corollary which follows from Proposition 4.5.4 and Theorem 4.5.4:

Corollary 4.5.6. The optimal protection levels $y_{t,j}^*$ are

- (i) non-decreasing in booking class j for fixed $t = T, \dots, 1$.
- (ii) non-decreasing in time t for fixed booking class $j = 1, \dots, n$.

Extensions to the dynamic model

Lee and Hersh (1993) formulate a dynamic program that considers groups arrivals which cannot be accepted partially. Note that Brumelle and Walczak (2003) provide a counterexample to their claim of a decreasing incremental value of a number of seats in time for a fixed available capacity. Van Slyke and Young (2000) formulate a stochastic knapsack problem for a capacity control application allowing inseparable group bookings.

Talluri and Van Ryzin (2004a) provide a dynamic model incorporating customer choice. Talluri and Van Ryzin (1998) consider network models. A more detailed introduction is presented in Talluri and Van Ryzin (2004b). Barz (2007) considers a risk-averse decision maker maximizing expected utility. Chi (1995) and Subramanian et al. (1999) consider overbooking, cancellations, and no-shows.

4.6 Dynamic Pricing

Dynamic pricing (DP) is another approach to optimize expected revenues. The obvious difference to capacity control problems is that the decision variable to manage demand is the price itself. However, the distinction is not sharp because closing a booking class can be considered rising the price of a ticket. The substantial difference is that demand is explicitly modeled dependent on price. Additionally, versioning or fencing, i.e. restricting different booking classes to encourage self-selection, is usually not applied when an airline uses DP. Instead, the same restrictions apply to all tickets.

To adopt DP, the company needs to have flexibility in price, e.g. a monopolist is able to influence demand by varying the price. Under perfect competition, a single company cannot influence prices. To optimize revenues the company can only adopt capacity control strategies. Additionally, price changes need to be possible quickly and at a reasonable or even negligible price.

Legacy airlines traditionally commit to prices in advance to the booking process. Hence, they are not flexible in price and apply capacity control models rather than dynamic pricing. Contrarily, many low-cost carriers (LCC) distribute tickets only online and have the ability to quickly adapt prices at virtually no cost. Many LCCs optimize revenues using DP models and sell tickets at changing prices throughout the booking horizon. Other industries such as fashion apparel might not be flexible in quantity but have flexibility in price. Then, DP naturally lends itself to maximize revenues.

The practice of dynamic pricing is probably as old as commerce itself. Merchants have always adapted prices in response to demand and vast literature on dynamically set-

ting prices exists from an economic, marketing, and operations research perspective. However, earlier publications often cover medium- to long-term price changes in the product life cycle. Only during the last decades, when technological advances facilitated frequent and inexpensive price changes, sophisticated decision support models have emerged for tactical dynamic pricing. Retailers have been the innovators in applying DP, initially in the form of markdowns. Today, other industries such as travel, e-commerce, and even manufacturing apply DP models and promote scientific and practical sophistication (Coy, 2000, Talluri and Van Ryzin, 2004b).

As mentioned, many early publications consider strategical pricing decisions, often in conjunction with production or replenishment decisions (e.g. Zabel, 1972, refer to Rajan et al. (1992) for a compact literature overview). Contrarily, consistent with the RM assumptions, we assume the capacity to be fixed and for now do not consider an option to increase the capacity in the short term. One early work on perishable products with a limited capacity was published by Kincaid and Darling (1963). In their model, prices are set continuously to maximize expected revenues. Time and the set of allowable prices are continuous and arrivals are assumed to follow a homogeneous Poisson process. Customers' reservation prices are assumed mutually independent and to follow a known time-dependent distribution. They derive optimality criteria and provide an optimal policy in closed form for an exponential demand function. Gallego and Van Ryzin (1994) formulate an equivalent problem and extend the closed-form optimal policy to a more general exponential family. They further extend the model to arrivals that follow an inhomogeneous or compound Poisson process. Additionally, they consider discrete sets of allowable prices and overbooking. They provide upper bounds and develop asymptotically optimal heuristic policies. Zhao and Zheng (2000) allow for inhomogeneous Poisson arrivals and derive a sufficient condition for a time-monotone optimal policy.

Our model formulation is similar to that in Bitran and Mondschein (1997). We also divide the booking horizon into discrete time periods with at most one arrival. In our formulation, however, time periods need not be of equal length. Arrivals are assumed to follow an inhomogeneous Poisson process. Additionally, we assume a discrete set of allowable prices that might arise from a strategical decision such as using specific price points.

We adhere to some common assumptions that are similar to those in section 4.5:

- There are no cancellations, no-shows, or go-shows. Each accepted reservation cannot be rejected later without significant costs. Overbooking is not considered. Below some references are presented that do consider overbooking.
- A denied passenger request cannot be recaptured on other flights and is lost to

the airline.

- Group arrivals are not considered, i.e. groups can be partially accepted.
- Demand for different legs is mutually independent. In that case, the multi-product problem reduces to a set of independent single-product problems. Hence, we consider leg-based rather than network models. Single-leg models are important components for many network heuristics as network models are often intractable in practice.
- The decision maker is risk-neutral. Hence, the *expected* revenues are maximized. Optimization based on expectations seems reasonable in the light of the daily repeated decision process for numerous departures.
- Demand is assumed to be distributed according to a Poisson distribution. Time is discretized finely such that the probability of two or more arrivals in a time period is negligible.

DP models require a demand model that captures how the price influences demand. Throughout the remainder, we assume a given willingness-to-pay (WTP) function for potential buyers. The WTP is expressed through a time-dependent distribution of the *reservation price* R_t , i.e. the maximum price a customer is willing to accept. We assume the reservation prices to be mutually independent and independent of an arrival. $q_t(f) := P(R_t \geq f)$ denotes the purchase probability when a customer arrives and fare $f \geq 0$ is offered. With rational consumers, the purchase probability for a price is at least as high as the purchase probability for a higher price. Hence, we assume $q_t(f) \geq q_t(f')$ for $f < f'$.

We now discuss the basic DP model for a single-leg flight. The results are used in the subsequent chapters.

Let $\mathcal{P} = \{f_0, f_1, \dots, f_k\}$ with $f_0 > f_1 > \dots > f_k \geq 0$ be the discrete set of allowable prices. The set must contain the *nullprice* f_0 with $q_t(f_0) = 0$ for all t . The nullprice is used if all requests are to be rejected, e.g. when there is no more remaining capacity and no more tickets can be sold. We define a MDP $(T, S, A, \mathbf{p}, r, V_0)$ with

- (i) the planning horizon $T \in \mathbb{N}$. The length of the booking period is divided into periods with at most one customer arriving. The probability for an arrival is denoted by λ_t , where t denotes the time left to sell any capacity. Because of the assumption of an infinite market, the arrival probabilities are mutually independent. Note that time is counted backwards and the booking period starts at $t = T$. The flight departs at $t = 0$ and the last decision is made at $t = 1$.

- (ii) The state space is $S = \{0, \dots, C\}$, where C denotes the total capacity available. The state $c \in S$ refers to the remaining capacity at the current period.
- (iii) The action space is $A = \mathcal{P}$. At each decision point, one price is selected from the set of feasible actions $A(c) = \mathcal{P}$ for $c > 0$. In state $c = 0$, no capacity remains and no further requests can be accepted. Hence, $A(0) = \{f_0\}$.
- (iv) The time-dependent transition law is

$$\mathbf{p}_t(c, a, c') = \begin{cases} \lambda_t q_t(a) & c' = c - 1, \\ 1 - \lambda_t q_t(a) & c' = c, \\ 0 & \text{otherwise,} \end{cases}$$

for $t = T, \dots, 1$ with $c, c' \in S, a \in A(c)$.

- (v) The one-stage reward function is

$$r_t(c, a) = \lambda_t q_t(a) a$$

for $t = T, \dots, 1$ with $c \in S, a \in A(c)$.

- (vi) The terminal reward function is $V_0 \equiv 0$.

The value $V_t(c)$ denotes the maximum aggregated expected revenue of the remaining time periods $t, t-1, \dots, 1$. The value function in each period $t+1$ and for remaining capacity $c \in S$ is given by the optimality equation

$$V_{t+1}(c) = \max_{a \in A(c)} \left\{ \lambda_t q_t(a) [a + V_t(c-1)] + [1 - \lambda_t q_t(a)] V_t(c) \right\}. \quad (4.7)$$

The following lemma will be of use when deriving structures of V_t :

Lemma 4.6.1. Let $g : X \rightarrow \mathbb{R}^+$, $X = \{0, \dots, n\}$ be non-decreasing and concave and $p_a \in [0, 1]$ with $p_a > p_{a'}$ if $a' > a$. Let $D(x)$ be a finite set of admissible a with at least one $a' \in D(x)$ with $p_{a'} = 0$. Further, let $D(x+1) = D(x)$. Then,

- (i) the largest maximizer $\psi(x) \in D(x)$ of the function

$$L_a g(x) := p_a [a + g(x-1)] + (1 - p_a) g(x) \quad , x > 0$$

is non-increasing in x .

- (ii) The function $h : X \rightarrow \mathbb{R}$ with

$$h(x) = \max_{a \in D(x)} \{ p_a [a + g(x-1)] + (1 - p_a) g(x) \} \quad , x > 0$$

and $h(0) = 0$ is non-decreasing and concave in x .

Proof.

(i) Let $x \in X \setminus \{0\}$, $a^* = \psi(x)$, and $a > a^*$. Then

$$L_{a^*}g(x) > L_ag(x)$$

and thus

$$\begin{aligned} L_{a^*}g(x+1) - L_ag(x+1) &= L_{a^*}g(x) + p_{a^*}\Delta g(x) + (1 - p_{a^*})\Delta g(x+1) \\ &\quad - L_ag(x) - p_a\Delta g(x) - (1 - p_a)\Delta g(x+1) \\ &> (p_{a^*} - p_a)(\Delta g(x) - \Delta g(x+1)) \\ &\geq 0. \end{aligned}$$

Hence, all actions $a > a^*$ are not optimal in $x+1$ and $\psi(x+1) \leq a^* = \psi(x)$.

(ii) For $x = 1$,

$$\begin{aligned} \Delta h(x) &= \max_{a \in D(x)} \{p_a[a + g(x-1)] + (1 - p_a)g(x)\} \\ &= \max_{a \in D(x)} \{p_a[a - \Delta g(x-1)]\} + g(x) \\ &\geq g(x), \end{aligned}$$

where we used that $p_{a'} = 0$ for at least one $a' \in D(x)$.

Now, let $x \in X, x > 1$ and $a^* = \psi(x-1)$. Then using the monotonicity and concavity of g

$$\begin{aligned} \Delta h(x) &\geq L_{a^*}g(x) - L_{a^*}g(x-1) \\ &= p_{a^*}\Delta g(x-1) + (1 - p_{a^*})\Delta g(x) \\ &\geq 0. \end{aligned}$$

Thus, h is non-decreasing in x .

For $x = 2$, let $a^* = \psi(x)$. Then, using that $p_{a'} = 0$ for at least one $a' \in D(x)$,

$$\begin{aligned} \Delta h(x) - \Delta h(x-1) &\leq p_{a^*}\Delta g(x-1) + (1 - p_{a^*})\Delta g(x) - g(x-1) \\ &\leq p_{a^*}\Delta g(x-1) + (1 - p_{a^*})\Delta g(x) - \Delta g(x-1) \\ &= (1 - p_{a^*})(\Delta g(x) - \Delta g(x-1)) \\ &\leq 0, \end{aligned}$$

where we used $g(x) \geq 0$ and the concavity of g .

Now, let $x \in X, x > 2$, $a^* = \psi(x)$, and $a^{**} = \psi(x - 2)$, then

$$\begin{aligned}
& \Delta h(x) - \Delta h(x - 1) \\
& \leq L_{a^*} g(x) - L_{a^*} g(x - 1) - L_{a^{**}} g(x - 1) + L_{a^{**}} g(x - 2) \\
& = p_{a^*} \Delta g(x - 1) + (1 - p_{a^*}) \Delta g(x) - p_{a^{**}} \Delta g(x - 2) - (1 - p_{a^{**}}) \Delta g(x - 1) \\
& = (1 - p_{a^*}) (\Delta g(x) - \Delta g(x - 1)) + p_{a^{**}} (\Delta g(x - 1) - \Delta g(x - 2)) \\
& \leq 0,
\end{aligned}$$

where we have used that g is concave in the last inequality. Thus, h is concave. \square

Using Lemma 4.6.1, we derive structures in the capacity:

Theorem 4.6.2. *In the DP model with the value function specified in (4.7)*

- (i) $V_t(c)$ is non-decreasing and concave in c for any given $t \in \{T, \dots, 0\}$.
- (ii) For a fixed t , the largest fare $\psi_t(c)$ maximizing the function

$$a \mapsto \lambda_t q_t(a) [a + V_{t-1}(c - 1)] + [1 - \lambda_t q_t(a)] V_{t-1}(c)$$

is non-increasing in c .

Proof. Using Lemma 4.6.1, the proof is trivial: Let $V_t(c)$ be positive, non-decreasing, and concave for some t . Then, applying Lemma 4.6.1 (ii), $V_{t+1}(c)$ is non-decreasing and concave. We easily show that $V_{t+1}(c) \geq 0$ using that $L_{f_0} V_t(c) = V_t(c)$ and $V_0 \equiv 0$. At the start of induction, at $t = 0$, the claim is true because $V_0 \equiv 0$. Hence, (i) is true. (ii) follows immediately from Lemma 4.6.1 (i) when considering that the largest feasible price is the nullprice that is also the only admissible action for $c = 0$. \square

Theorem 4.6.2 is frequently used throughout the remainder. It is also of interest when determining the optimal initial inventory, i.e. when assigning aircraft to flights during fleet assignment (c.f. chapter 5.4). In applications where the cost of the capacity is linear or convex, the expected profit is a concave function of the initial capacity. The maximum expected profit is attained where the marginal expected revenue equals the marginal cost (subject to discretization).

Fixed cost do not affect the location of the optimum. Clearly, fixed cost need to be offset by the expected profit. Otherwise it is optimal not to offer any capacity at all.

We now consider structures in the time t that are used to prove structures of the optimal policy in time.

Theorem 4.6.3. *In the DP model with the value function specified in (4.7)*

(i) $V_t(c)$ is non-decreasing in t for a fixed $c \in S$.

(ii) $V_t(c) - V_t(c-1)$ is non-decreasing in t for a fixed $c \in S, c > 0$.

Proof. $V_t(0) \equiv 0$ because the only feasible action is the nullprice f_0 . For $c > 0$,

$$\begin{aligned} V_{t+1}(c) &= \max_{a \in A(c)} \{ \lambda_{t+1} q_{t+1}(a) [a + V_t(c-1)] + [1 - \lambda_{t+1} q_{t+1}(a)] V_t(c) \} \\ &= \max_{a \in A(c)} \{ \lambda_{t+1} q_{t+1}(a) [a - \Delta V_t(c)] \} + V_t(c) \\ \Leftrightarrow V_{t+1}(c) - V_t(c) &= \max_{a \in A(c)} \{ \lambda_{t+1} q_{t+1}(a) [a - \Delta V_t(c)] \} \\ &\geq \lambda_{t+1} q_{t+1}(f_0) [f_0 - \Delta V_t(c)] \\ &= 0 . \end{aligned}$$

Hence, (i) is true. To show (ii), we first consider $c = 1$. Then, using (i),

$$V_t(1) - V_t(0) - V_{t-1}(1) + V_{t-1}(0) = V_t(1) - V_{t-1}(1) \geq 0 .$$

Now, for $c > 1$, let $a^* = \psi_t(c-1)$ be the optimal action for capacity $c-1$ at time t . Then

$$\begin{aligned} &V_t(c) - V_t(c-1) - V_{t-1}(c) + V_{t-1}(c-1) \\ &\geq \lambda_t q_t(a^*) [a^* + V_{t-1}(c-1)] + [1 - \lambda_t q_t(a^*)] V_{t-1}(c) \\ &\quad - \lambda_t q_t(a^*) [a^* + V_{t-1}(c-2)] - [1 - \lambda_t q_t(a^*)] V_{t-1}(c-1) \\ &\quad - V_{t-1}(c) + V_{t-1}(c-1) \\ &= -\lambda_t q_t(a^*) [\Delta V_{t-1}(c) - \Delta V_{t-1}(c-1)] \\ &\geq 0 , \end{aligned}$$

where we have used that V_{t-1} is concave in c in the last inequality. \square

Before we continue with finding a lower bound of the largest optimal action, we present the following lemma providing an upper bound for ΔV_t .

Lemma 4.6.4. For all $t = T, \dots, 0$ and all $c \in S$, the increase of the value function

$$\Delta V_t(c) \leq f_0 .$$

Proof. The proof follows by induction. Because $V_0 \equiv 0$ and $f_0 \geq 0$, the assertion holds for $t = 0$. Now, suppose there exists an $a \in A(c) \setminus \{f_0\}$ such that $a \geq \Delta V_t(c)$. Then, using Theorem 4.6.3 (i),

$$\begin{aligned} \Delta V_{t+1}(c) &= V_{t+1}(c) - V_{t+1}(c-1) \\ &\leq V_{t+1}(c) - V_t(c-1) \end{aligned}$$

$$\begin{aligned}
&= \max_{a \in A(c)} \{ \lambda_{t+1} q_{t+1}(a) [a - \Delta V_t(c)] \} + \Delta V_t(c) \\
&\leq \max_{a \in A(c)} \{ a - \Delta V_t(c) \} + \Delta V_t(c) \\
&= f_0 .
\end{aligned}$$

If $a < \Delta V_t(c)$ for all $a \in A(c) \setminus \{f_0\}$. Then, because $f_0 \in A(c)$,

$$\begin{aligned}
\Delta V_{t+1}(c) &= V_{t+1}(c) - V_{t+1}(c-1) \\
&\leq V_{t+1}(c) - V_t(c-1) \\
&= \max_{a \in A(c)} \{ \lambda_{t+1} q_{t+1}(a) [a - \Delta V_t(c)] \} + \Delta V_t(c) \\
&= \Delta V_t(c) \\
&\leq f_0 ,
\end{aligned}$$

which concludes the proof. \square

Next, we prove that $\Delta V_t(c)$ is a lower bound for the optimal action at $t+1$ in state $c \in S$. Note that during backward induction, at stage $t+1$, $\Delta V_t(c)$ is known. The bound is non-decreasing in time t (Theorem 4.6.3) and non-increasing in the capacity c (Theorem 4.6.2).

Theorem 4.6.5. *Let $c \in S$ and $t = T-1, T-2, \dots, 0$. Further, let $\psi_{t+1}(c) \in A(c)$ be the largest argument maximizing the function*

$$a \mapsto \lambda_{t+1} q_{t+1}(a) [a + V_t(c-1)] + [1 - \lambda_{t+1} q_{t+1}(a)] V_t(c) .$$

Then, $f_0 \geq \psi_{t+1}(c) \geq \min\{f_0, \Delta V_t(c)\}$.

Proof. Fix $c \in S$ and let $a \in A(c) \setminus \{f_0\}$ such that $a \leq \Delta V_t(c)$. Then

$$L_a V_t(c) = \lambda_{t+1} q_{t+1}(a) [a - \Delta V_t(c)] + V_t(c) \leq V_t(c) = L_{f_0} V_t(c) ,$$

which entails $f_0 \geq \psi_{t+1}(c) \geq \min\{f_0, \Delta V_t(c)\}$. If $a > \Delta V_t(c)$ holds for all $a \in A(c)$, the result is trivial. \square

In certain settings, the optimal pricing policy is monotone in time. We provide a sufficient condition by transferring the results in Zhao and Zheng (2000) to our setting. For $a, a' \in \mathcal{P}$, $a > a'$, we define

$$\mathbf{q}_t(a, a') := \begin{cases} q_t(a)/q_t(a') & \text{if } q_t(a') \neq 0, \\ 1 & \text{if } q_t(a') = 0. \end{cases}$$

Note that $q_t(a') = 0$ implies $q_t(a) = 0$ for $a > a'$.

Theorem 4.6.6. *If $q_t(a, a')$ is non-decreasing in t for $a, a' \in \mathcal{P}, a > a'$, then, for fixed $c \in S$, the largest maximizer $\psi_t(c)$ of the function*

$$a \mapsto \lambda_t q_t(a)[a + V_{t-1}(c - 1)] + [1 - \lambda_t q_t(a)]V_{t-1}(c)$$

is non-decreasing in t .

Proof. Let $a^* = \psi_t(c)$ be the largest optimal action in state $c \in S$ at t and let $a \in \mathcal{P}, a < a^*$. Let $q_t(a^*) \neq q_t(a)$, which is equivalent to $q_t(a^*) < q_t(a)$. Further, let $q_{t+1}(a^*) \neq q_{t+1}(a)$. Then

$$\begin{aligned} & L_{a^*}V_{t-1}(c) - L_aV_{t-1}(c) \\ &= \lambda_t q_t(a^*)[a^* - \Delta V_{t-1}(c)] - \lambda_t q_t(a)[a - \Delta V_{t-1}(c)] \geq 0 \\ \Leftrightarrow \quad \Delta V_{t-1}(c) &\geq \frac{q_t(a)a - q_t(a^*)a^*}{q_t(a) - q_t(a^*)} \\ &= \frac{a - [q_t(a^*)/q_t(a)]a^*}{1 - [q_t(a^*)/q_t(a)]} \\ &= a^* - \frac{a^* - a}{1 - [q_t(a^*)/q_t(a)]} \end{aligned}$$

Using the monotonicity of $\Delta V_t(c)$, we get

$$\begin{aligned} \Delta V_t(c) &\geq \Delta V_{t-1}(c) \\ &\geq a^* - \frac{a^* - a}{1 - [q_t(a^*)/q_t(a)]} \\ &\geq a^* - \frac{a^* - a}{1 - [q_{t+1}(a^*)/q_{t+1}(a)]}, \end{aligned}$$

which implies $L_{a^*}V_t(c) - L_aV_t(c) \geq 0$.

Now let $q_{t+1}(a^*) = q_{t+1}(a)$. Then

$$\begin{aligned} & L_{a^*}V_t(c) - L_aV_t(c) \\ &= \lambda_{t+1}q_{t+1}(a^*)[a^* - \Delta V_t(c)] - \lambda_{t+1}q_{t+1}(a)[a - \Delta V_t(c)] \\ &= \lambda_{t+1}q_{t+1}(a^*)[a^* - a] \\ &\geq 0. \end{aligned}$$

Note that $q_t(a^*) = q_t(a)$ implies $q_{t+1}(a^*) = q_{t+1}(a)$, and hence, the same argument works. \square

The assumption that $q_t(a, a')$ is non-decreasing in t for $a, a' \in \mathcal{P}, a > a'$ has the interpretation that an arriving customer is more likely willing to pay a premium over a

price a' the more time remains for sales. That might be a valid assumption in applications such as retailing, where early buyers accept paying a premium for exclusiveness. Later in the sales period, products become less popular and are often marked-down. However, the opposite is usually true in the airline industry. The closer to departure, the more time-sensitive customers are. A business traveler might purchase a ticket on short-notice to attend an appointment. That traveler is certainly willing to pay a premium to purchase a ticket in comparison to a leisure customer that books well in advance and is less time-sensitive. Hence, in the airline industry, the assumption is usually not valid. However, many authors assume that probabilities are constant in time. The case of constant probabilities is subsumed in the assumption of non-decreasing $q_t(a, a')$. Then, optimal prices are non-decreasing in time.

In general, the monotonicity in time does not hold. We give an example where the optimal price actually decreases in time. We consider only the last two decision times $t = 1, 2$ remaining in a booking process. Earlier decisions and sales do not influence the strategy and hence are not considered. The flight departs at $t = 0$. The sales probabilities are given in Table 4.1. These combine the arrival probabilities with the willingness-to-pay of the arriving customers. Assuming $V_0 \equiv 0$, i.e. not considering overbooking, the expected values and the optimal actions are presented in Table 4.2. For any positive capacity c , the optimal price in period $t = 1$ is $\psi_1^*(c) = 20$ and decreases to $\psi_2^*(c) = 10$ at $t = 2$. Note that the example does not depend on the

Table 4.1: Example of an optimal policy with decreasing prices in time. Sales Probabilities.

Price	$t = 1$	$t = 2$
10	0.4	0.2
20	0.3	0.05

exact value of the sales probability for $f = 10$ in $t = 1$ within the interval $[0, 0.6]$.

Algorithm 4.6.1 provides a pseudo-code implementation to calculate an optimal pricing policy using backward induction. The structures expressed in Theorem 4.6.2 and Theorem 4.6.5 are exploited to exclude suboptimal actions and increase efficiency. The monotonicity in time (Theorem 4.6.6) is not used because it does not hold in general. Note that Algorithm 4.6.1 returns the expected values of all decision times, i.e. the total expected revenue of the remaining sales period. To find an optimal policy and the expected value only at the beginning of the booking process, only the values of the current and the previous iteration need to be saved. Hence, the algorithm can be further optimized to reduce memory requirements, if only the policy and the initial expected values are of interest (c.f. chapter 3).

Algorithm 4.6.1 Algorithm to determine an optimal DP policy.

Input: Fares $f_0 > f_1 > \dots > f_k \geq 0$,
 arrival probabilities $(\lambda_T, \lambda_{T-1}, \dots, \lambda_1)$,
 purchase probabilities $(q_T, q_{T-1}, \dots, q_1)$,
 total initial capacity C ,
 length of the booking horizon T

Output: Optimal policy $(a_T^*, a_{T-1}^*, \dots, a_1^*)$,
 Expected values $(V_T, V_{T-1}, \dots, V_0)$

1. $t = 0$
 2. $c = 0$
 3. **while** $c \leq C$
 4. $V_t(c) = 0$
 5. $c = c + 1$
 6. **while** $t < T$
 7. $t = t + 1$
 8. $V_t(0) = 0$
 9. $a_t^*(0) = f_0$
 10. $c = 1$
 11. **while** $c \leq C$
 12. $i = k$
 13. **while** $f_i < \Delta V_{t-1}(c)$ **and** $i > 0$
 14. (* Using lower bound on price (Theorem 4.6.5) *)
 15. $i = i - 1$
 16. $a_t^*(c) = f_i$
 17. $V_t(c) = \lambda_t q_t(f_i)[f_i + V_{t-1}(c - 1)] + [1 - \lambda_t q_t(f_i)]V_{t-1}(c)$
 18. **while** $i > 0$ **and** $f_{i-1} \leq a_t^*(c - 1)$
 19. (* Using upper bound on price (Theorem 4.6.2) *)
 20. $i = i - 1$
 21. $L_{f_i} V_{t-1}(c) = \lambda_t q_t(f_i)[f_i + V_{t-1}(c - 1)] + [1 - \lambda_t q_t(f_i)]V_{t-1}(c)$
 22. **if** $L_{f_i} V_{t-1}(c) \geq V_t(c)$
 23. $V_t(c) = L_{f_i} V_{t-1}(c)$
 24. $a_t^*(c) = f_i$
 25. $c = c + 1$
-

Table 4.2: Example of an optimal policy with decreasing prices in time. Values and optimal prices.

	t=1			t=2		
	Capacity c			Capacity c		
Price f	0	1	> 1	0	1	> 1
10	-	4	4	-	6.8	8
20	-	6	6	-	6.7	7
f_0	0	0	0	0	6	6
Maximum	0	6	6	0	6.8	8
Optimal Price ψ^*	f_0	20	20	f_0	10	10

Extensions and related literature

Smith and Achabal (1998) study pricing for a single product when the demand depends on the price and the inventory level. They assume a known and constant demand intensity. Especially in a retail setting, demand dependency on inventory levels is a reasonable assumption. Chatwin (2000) considers multiple products with demand dependent on price and inventory. He restricts prices to a discrete set of allowable prices. He extends his results to time-dependent demand, policies dependent on time and inventory, and to allow for replenishment.

Feng and Gallego (1995) assume two fixed prices for a product and determine the optimal time to switch. Feng and Xiao (1999) extend the model to a risk-sensitive decision maker. Feng and Xiao (2000b) consider multiple predetermined prices and derive optimal switching times when a monotone pricing policy (i.e. only mark-ups or mark-downs) is required. Feng and Xiao (2000a) generalize these results to non-monotone policies. Feng and Gallego (2000) consider fares and arrival intensities that may depend on the time and the inventory.

Bitran and Mondschein (1997) study periodic pricing where prices can be set only at a finite set of decision times. They include a constraint to only allow prices decreasing in time as often used in retailing. In numerical studies, they find that the markdown constraint does not significantly decrease the overall revenue. The model is extended by Bitran et al. (1998) to consider different locations with separate demands and inventories. The same price is required at each location.

Gallego and Van Ryzin (1997) extend the results in Gallego and Van Ryzin (1994) to consider network problems. In both works, bounds and heuristics based on the deterministic problem are developed and the models are extended to include cancellations

and overbooking. Maglaras and Meissner (2006) provide a common general formulation of the dynamic pricing and capacity control problems. They consider multiple products sharing a common resource. They provide and numerically analyze several heuristics. Bitran et al. (2005) study optimal dynamic pricing in a multi-product setting with demand substitution due to price differences and stock-outs. They deploy a demand model called *Walrasian Choice model* that allows to model customer choice in conjunction with the substitution effects as well as facilitating a ranking among products.

Zhang and Cooper (2005) study a choice model to consider parallel flights that serve the same route at different times during the day. Customers make their choice between the flights depending on the set of all prices. To deal with the curse of dimensionality they analyze various pooling heuristics.

Gallego and Hu (2009) formulate a stochastic dynamic pricing game with multiple competitors selling substitutable perishable products. The probabilities of choosing the substitutes are based on a multinomial logit model. They derive asymptotically optimal heuristics based on the corresponding deterministic differential game. In a discrete-time version of a similar setting, Lin and Sibdari (2009) show the existence of a Nash equilibrium and characterize the price and the expected revenue when all real-time capacities are known to all competitors. They also propose a heuristic to apply when only the initial inventories of the competitors are known.

Using a Bayesian learning approach, Bitran and Wadhwa (1996) extend the model by Bitran and Mondschein (1997) to demand learning. They assume uncertainty in one parameter of the reservation-price distribution. They also consider non-homogeneous arrivals and a time-dependent willingness-to-pay. Additionally, they present and test an expected value heuristic. Besbes and Zeevi (2006) develop a two-step learning procedure for the multi-product pricing problem. First, prices are varied in a short learning phase. Then, with a non-parametric estimation of the demand function, using a deterministic heuristic, a fixed-price policy is adopted for the remainder of the sales horizon. Lin (2006) employs Bayesian learning to the arrival process. He assumes that the arrivals follow a Poisson process with an unknown parameter. Gallego and Talebian (2010) propose a demand learning process for multiple versions of one product and assume unknown arrival rates and valuations. They also provide a compact literature overview on demand learning and pricing.

Li and Zhuang (2009) model a single-leg dynamic pricing problem with homogeneous arrival rates and a constant willingness-to-pay function. They model a risk-sensitive decision maker and show that the monotone structures of the expected values and the optimal prices in time and capacity are preserved for exponential atemporal and general

additive utility functions. Optimal prices decrease with the degree of risk-aversion for additive and atemporal exponential utility functions. Levin et al. (2008) take another approach to incorporate risk into the decision. They augment the objective function with a penalty term for the probability to fall below a certain revenue level. They show that after the minimum revenue target has been reached, the risk-neutral and risk-averse policies coincide. They propose and numerically study a fixed-price heuristic as well as a heuristic allowing only a fixed number of price changes.

For more detailed surveys with additional references on specific problems in dynamic pricing we refer the reader to Bitran and Caldentey (2003) and Elmaghraby and Keskinocak (2003).

Demand-Driven Re-Fleetting

In this chapter, we introduce the concept of Demand-Driven Re-Fleetting. In order to explain the benefits we first outline the typical planning process of an airline and the timing of the various planning steps. The fleet assignment problem and its complexity are explained in more detail in section 5.2. In the next section we present the general idea of Demand-Driven Re-Fleetting and summarize the existing literature. We conclude by presenting an algorithm that facilitates finding swapping opportunities without affecting other flights of the network.

5.1 Airline Planning Process

The planning process of an airline comprises several complex problems that are usually solved sequentially as depicted in Figure 5.1. Among the most important long-term strategic decisions is the *fleet planning* that determines the composition of the fleet in the future. The airline decides which and how many new aircraft to purchase and which of the current aircraft to retire. The planning is based on traffic forecasts and market requirements that determine the needed technical and performance characteristics such as capacity and range of an aircraft. Economic and financial impacts need to be assessed as well as environmental and marketing aspects. Mainly due to the long delivery times, the fleet planning problem is solved years in advance to operations. It is a continuous process that greatly influences an airline's options to serve different markets and that influences its financial position and future crew and maintenance costs.

Subsequently, routes, frequencies, and times are fixed in *schedule design* planning usually more than one year prior to the start of the schedule. This step is again split into consecutive problems that involve strategic decisions on which markets (i.e. origins and

destinations) to serve and what times and frequencies to offer in order to satisfy customer needs. The offered schedule is a major determinant of demand and market share and hence the overall competitive position. The schedule is restricted by the available fleet in range, capacities, and airport requirements, although there is a feedback loop with fleet planning. Strategic considerations concerning routes and frequencies mainly drive fleet planning decisions.

During the next planning step, the *fleet assignment*, aircraft types are assigned to the legs in the schedule. Based on demand and revenue forecasts and cost estimates capacities are assigned to match demand in the most profitable way. Capacities for all scheduled legs are fixed. The fleet assignment greatly impacts profits as too small aircraft result in lost customers that cannot be accommodated (*spill*). Potential revenue and market share is lost due to insufficient capacities. In contrast, an assignment with excess capacity results in *spoilage*, i.e. empty seats. The larger airplane could possibly be used more efficiently on other flights with higher demand and exhibits higher operating cost than needed to satisfy the available demand. The fleet assignment problem is described in more detail in section 5.2.

The capacities determined are an important input to RM systems. Standard RM problems as described in chapter 4 assume a fixed and exogenously given initial capacity that is sold in a revenue maximizing way by selecting prices or booking limits. As the supply of seats is fixed, the demand is influenced by price or rationed by booking controls. The booking process typically starts about one year to departure and requires a finalized fleet assignment.

Crew and maintenance planning take place parallel to the sales process several weeks before departure. Legal and contractual requirements such as rest times or union agreements need to be satisfied. Often, crew plans must be finalized several weeks in advance to departures. Maintenance requirements are tightly regulated and some airlines adopt additional standards.

The planning steps are highly interdependent but need to be decided on different time scales. Each planning step is restricted by all preceding decisions made. The disposable fleet restricts the schedule design and the fleet assignment. RM and crew and maintenance planning require the fleeted schedule as an input. However, strategic decisions on routes and the schedule influence fleet planning. RM influences the realization of demand and hence the optimal fleet assignment at the time of departure. Fleet assignment planning also shows feedback loops with crew and maintenance planning.

Problems in regards to an integration arise because of the different time scales of the planning steps and the different optimization models used. The early planning steps need to be completed at times when demand is still highly uncertain. According to

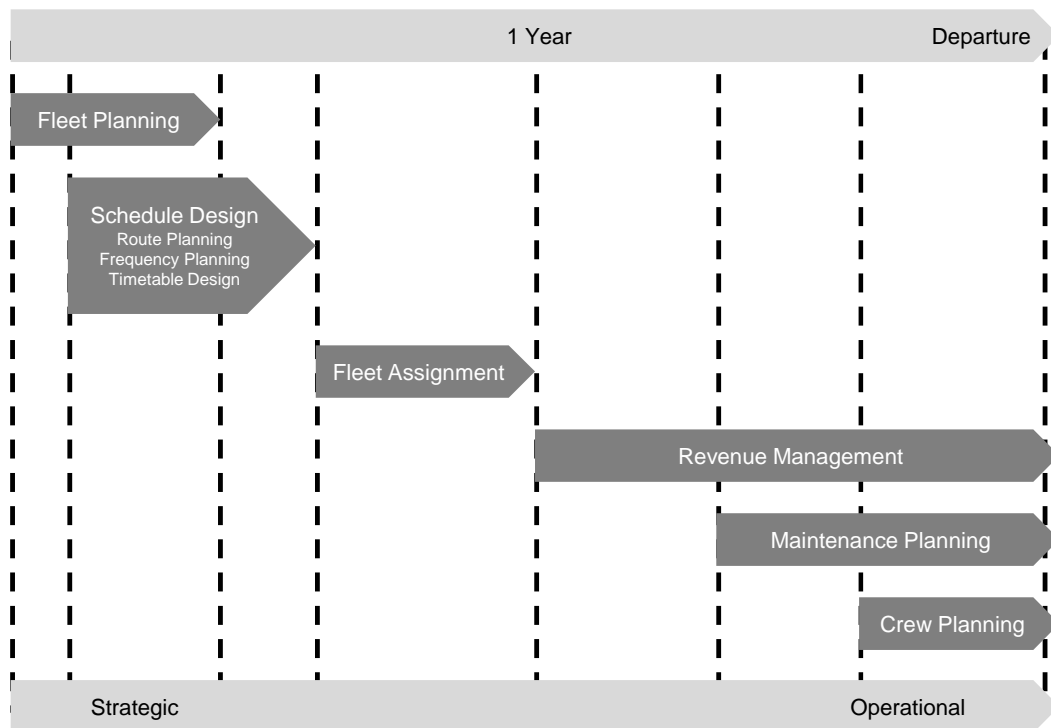


Figure 5.1: A typical sequential airline planning process.

Swan (2002) variations of 20%-50% of the mean are typical during the fleet assignment. Closer to departure, forecasts become more accurate and more granular. Because capacities need to be fixed based on highly volatile forecasts, airlines have concentrated on managing demand through RM to finely match it with the roughly planned supply. Forecasts for the fleet assignment are highly aggregated and averaged over typical time periods such as an average day. Different fares paid are mostly averaged as well. RM systems require much more granular demand forecasts and consider different prices paid by customers. Often, different departments of an airline are responsible for the different planning steps. Data from different sources is used and often not shared between different planners. Only recently, planning systems have been introduced that allow to exchange data and integrate different forecasts (c.f. e.g. Lufthansa Systems, 2010).

As a result of the stochastic nature of the demand and early forecasts used, even with sophisticated RM systems in place today, overall load factors are only in the range of 78%-82% (Lufthansa, 2011, Ryanair, 2011b). Lowering the capacity is not an option as even with low load factors significant spill still exists (Berge and Hopperstad, 1993).

The fleet assignment problem is described in more detail in section 5.2. RM techniques have been discussed in chapter 4. For further details on the other planning steps, the reader is referred to Belobaba (2009b).

5.2 Fleet Assignment

After the schedule has been fixed, fleet types are allocated to specific flight legs during the fleet assignment (FAM). Fleet types (e.g. Airbus A319, Boeing 737-800) rather than individual aircraft (*tails*) are assigned and capacities on each leg are fixed for subsequent planning steps. Individual planes are assigned in later planning phases under consideration of maintenance requirements.

For our purposes, we restrict ourselves to a brief description of the basic leg-based fleet assignment model as proposed by Subramanian et al. (1994) and Hane et al. (1995). Earlier models such as the ones by Abara (1989) or Daskin and Panayotopoulos (1989) are limited and have become obsolete today. More complex models consider recapture of spilled passengers (e.g. Lohatepanont, 2002, Lohatepanont and Barnhart, 2004), changes to the original schedule (Desaulniers et al., 1997, Rexing et al., 2000, Bélanger et al., 2006), or the integration of maintenance routing (Clarke et al., 1996, Barnhart et al., 1998, Haouari et al., 2009), crew planning (Barnhart et al., 2002b, Gao et al., 2009), or both (Sandhu and Klabjan, 2007, Papadakos, 2009). For a comprehensive review of the fleet assignment problem we refer to Lohatepanont (2002) or Grothklags (2006). The later work also contains an extensive discussion of solution algorithms.

The fleet assignment problem is usually modeled as a time-space network flow problem based on a multi-commodity flow problem. Each node corresponds to a specific departure or arrival at a certain airport. The associated departure times are simply those of the scheduled flights. The time of an arrival node is the *ready time*, i.e. the time when an aircraft is again ready for take-off after landing, disembarking, cleaning, fueling, boarding, etc. Flight legs are represented by flight arcs that connect nodes at different airports. Ground arcs model airplanes on ground and connect nodes at the same airport. A special type of ground arcs are overnight arcs that connect the last with the first arrival or departure of the planning horizon at each station. Overnight arcs guarantee that the same number of planes of each type originates and ends at each station to facilitate repeating a rolling plan. Figure 5.2 shows an exemplary flight network with only two stations.

Note that in the basic network illustrated flight durations and turn times (i.e. the time to prepare the aircraft for the next take-off) are assumed to be equal for all fleet types. In practice, flight and turn times vary by fleet type which can be modeled by separate copies of the network for each type. The length of the flight arcs are adapted to the type-specific flight and turn times. There are more nodes in the network, and hence, more restrictions are needed and the problem becomes less tractable with the number of copies.




Figure 5.2: Example of a time-expanded flight network for two airports. The flight arcs are copied for each fleet type.

We additionally assume that any leg can be flown with any fleet type. In practice, separate problems are solved for short- and long-distance flights. Hence, the assumption is not critical.

Let \mathcal{L} denote the set of all scheduled flight legs, \mathfrak{F} the set of the available fleet types, and \mathfrak{V} the set of all nodes. Further, let \mathfrak{G} be the set of all ground arcs and let $\mathcal{C}_{\mathcal{L}} \subseteq \mathcal{L}$ and $\mathcal{C}_{\mathfrak{G}} \subseteq \mathfrak{G}$ be the subsets of legs and ground arcs that cross the end of the planning interval, i.e. start before the end and arrive after the beginning of the period. Hence, $\mathcal{C}_{\mathcal{L}}$ and $\mathcal{C}_{\mathfrak{G}}$ can be used to count all aircraft used. The total number of aircraft of type $k \in \mathfrak{F}$ is denoted \mathfrak{N}_k . The set of flight and ground arcs beginning and ending at node $\mathfrak{v} \in \mathfrak{V}$ is denoted $I(\mathfrak{v})$ and $O(\mathfrak{v})$, respectively.

The binary decision variable x_{ik} describes the decision to fly leg $i \in \mathcal{L}$ with fleet type $k \in \mathfrak{F}$ ($x_{ik} = 1$) or not ($x_{ik} = 0$). The variable y_{ik} counts the number of aircraft of type $k \in \mathfrak{F}$ on ground arc $i \in \mathfrak{G}$. The basic FAM is then solved using the following mixed integer program:

(MIP 5.2.1)

$$\min \sum_{i \in \mathcal{L}} \sum_{k \in \mathfrak{F}} c_{ik} x_{ik} \quad (5.1)$$

$$\sum_{k \in \mathfrak{F}} x_{ik} = 1 \quad \text{for all } i \in \mathcal{L} \quad (5.2)$$

$$\sum_{i \in \mathcal{C}_{\mathcal{L}}} x_{ik} + \sum_{j \in \mathcal{C}_{\mathfrak{G}}} y_{jk} \leq \mathfrak{N}_k \quad \text{for all } k \in \mathfrak{F} \quad (5.3)$$

$$\sum_{i \in I(\mathfrak{v}) \cap \mathcal{L}} x_{ik} + \sum_{i \in I(\mathfrak{v}) \cap \mathfrak{G}} y_{ik} = \sum_{i \in O(\mathfrak{v}) \cap \mathcal{L}} x_{ik} + \sum_{i \in O(\mathfrak{v}) \cap \mathfrak{G}} y_{ik} \quad \text{for all } \mathfrak{v} \in \mathfrak{V}, k \in \mathfrak{F} \quad (5.4)$$

$$x_{ik} \in \{0, 1\} \quad \text{for all } i \in \mathcal{L}, k \in \mathfrak{F} \quad (5.5)$$

$$y_{ik} \geq 0 \quad \text{for all } i \in \mathfrak{G}, k \in \mathfrak{F} \quad (5.6)$$

The objective (5.1) minimizes costs. The coefficients c_{ij} denote the costs of flying leg $i \in \mathcal{L}$ with type $k \in \mathfrak{F}$. These costs not only include the operating costs, but also the costs of spilled passengers when the capacity is not sufficient. Part of the spilled passengers might be recaptured on other flights of the airline. These recaptured profits

should be excluded from the spill cost estimates. However, the basic model does not explicitly model passenger flows (i.e. spill and recapture rates). Hence, spill costs and recaptured profits can only be roughly approximated.

Alternatively, profit maximization might be used as the objective (e.g. Berge and Hopperstad, 1993). Profit coefficients need to be estimated for each leg and fleet type including the revenue generated and the costs of flying. In either model, the demand is assumed to be homogeneous and an average fare is used. If spilled passengers are valued using the average fare, both objectives yield the same optimal assignment. However, spill costs per passenger might often include an added loss of goodwill. In general, revenues and spill costs are hard to estimate, especially in regards of RM systems in place that deliberately spill certain passengers.

The *cover constraints* (5.2) in conjunction with the binary constraints (5.5) guarantee that each leg is served by exactly one type. The *count constraints* (5.3) ensure that the total number of aircraft is not exceeded while the *balance constraints* (5.4) require that the number of aircraft arriving and departing at each node is equal.

While solutions are usually integral, MIP 5.2.1 does not ensure the integrality of the number of aircraft on the ground arcs y_{ik} . With (5.4) and (5.5), a non-integer solution can only be attained if there are spare planes available, i.e. aircraft that are not needed to fly the schedule and remain unutilized on ground. The constraints (5.4) and (5.5) require that all ground arcs at one station are either integral or have the same non-integral residue. The possible residues at different stations need to add up to an integer because of the count constraints (5.3) and the implicitly assumed integrality of $\mathfrak{N}_k, k \in \mathfrak{F}$. Hence, a non-integer solution can easily be transferred into a feasible assignment by rearranging spare aircraft and we only ensure non-negativity by (5.6). If an integral solution is required, constraining one ground arc per station (e.g. all overnight arcs) suffices because then the balance and binary constraints (5.4) and (5.5) ensure the integrality of all other ground arcs as well (c.f. e.g. Gopalan and Talluri, 1998).

The constraints in (5.2), (5.3), and (5.4) are the most basic and most important constraints to find a feasible solution for the given schedule. Where necessary, other constraints might be added such as crew and maintenance requirements, noise constraints, and gate restrictions at the airport (Belobaba, 2009b).

Through preprocessing techniques and efficient algorithms Subramanian et al. (1994) and Hane et al. (1995) solve large real-world problems in reasonable time. These include the creation of subproblems based on aircraft type characteristics such as range (e.g. a wide-body aircraft sub-network), the consolidation of nodes, and the elimination of decision variables and flight and ground arcs. Belobaba (2009b) provides an example

with 2,044 legs and 9 fleet types that is solved in about 16 minutes. Still, the basic fleet assignment problem and its feasibility problem (i.e. if a feasible solution exists) are NP-complete for more than 2 fleet types (Gu et al., 1994, Grothklags, 2006).

One shortcoming of the basic FAM problem is the lack of integration with other planning steps despite their interdependencies. The references provided above include approaches to integrate these. Other shortcomings of the basic FAM are mainly due to the complexity that gives rise to approximations. Network effects, i.e. dependent leg demand and itinerary fares, are not considered. Spilled and recaptured demand is often ignored and can at best only be approximated. More complex network models such as the works by Barnhart et al. (2002a) and Jacobs et al. (2008) explicitly model the passenger flow through the network, i.e. booked passengers on each itinerary. The benefits are found to be significant in implementations at American Eagle Airlines and United Airlines.

A further shortcoming is the aggregation of the demand with the use of average fares on a leg or itinerary basis. Especially with sophisticated RM systems in place that protect capacity for high-value demand by spilling low-value demand, revenue and spill cost estimates based on average fares might be greatly distorted. Itinerary fares are usually prorated to legs equally or based on relative distance. However, capacity decisions only of those legs with scarce capacities impact the total expected revenue from the itinerary. As long as legs show large capacity buffers relative to the demand, these do not influence revenue. Hence, proration models that allocate fares only to scarce resources outperform traditional models (e.g. Barnhart et al., 2009).

The FAM planning horizon is usually one day because most airlines operate the same schedule Monday through Friday. These are often adjusted for weekends or schedules are considered separately from working days. Hence, the demand is assumed to be static, although in reality, the demand varies significantly by day and also by season. Airlines accept this shortcoming because they prefer uniform daily schedules for operational reasons (e.g. for crew, maintenance, or gate planning).

5.3 Demand-Driven Re-Fleeting

Changing the initial fleet assignment in response to demand is the concept of *Demand-Driven Re-Fleeting* (DDR), which is also known as Demand-Driven Dispatch (DDD or D^3), or Dynamic Capacity Management (DCM). During the booking process, as demand realizes and forecasts improve, certain assignments are changed. Changes are heavily restricted by crew and maintenance plans. Plans are bound to fulfill all maintenance requirements for each individual aircraft. Union agreements and legal con-

straints in many countries require settled crew schedules several weeks prior operations. Changes are both difficult and expensive. The introduction of cockpit-compatible fleet families (e.g. Airbus A319, A320, A321) has facilitated re-fleeting options without changing crew planning. Crews are certified to operate all aircraft in a family with a single type rating. Hence, planes with different capacities can be exchanged to better match demand without influencing crew planning. In spite of these restrictions, parts of an assignment can be revised as departure approaches to improve capacity utilization. This flexibility is exploited by engaging in systematic Demand-Driven Re-Fleeting.

To illustrate the idea, assume flight A and B depart from the same airport at about the same time and have the same pricing and cost structures. Thus, flight A is assigned a larger capacity if its expected demand is higher than that of flight B during the fleet assignment (see Figure 5.3(a)). As demand realizes, more information becomes available on the total demand of both flights. At some point during the booking process the fleet assignment is reevaluated and now flight B has the larger expected (unconstrained) demand. More expected revenue can be realized if the planes are switched and more demand can be accommodated (see Figure 5.3(b)).

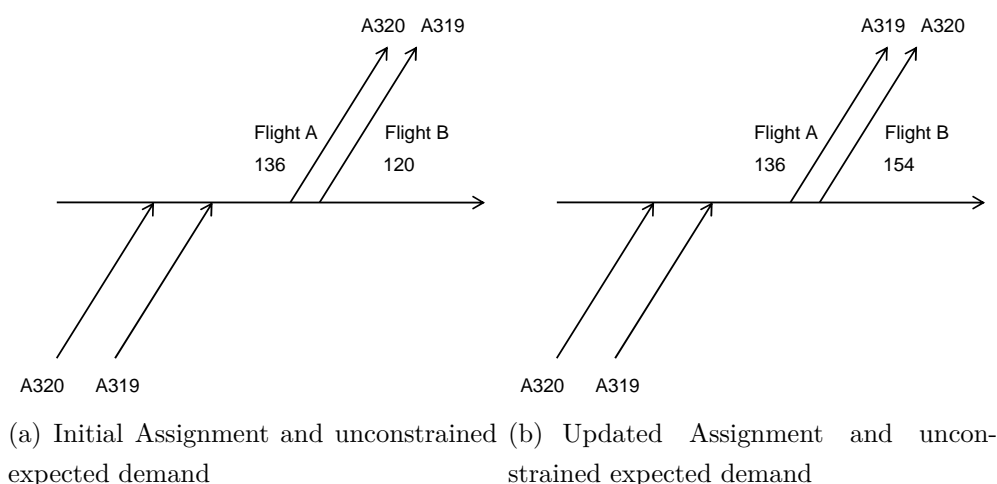


Figure 5.3: Example of Demand-Driven Re-Fleeting. The swap is induced by the updated demands on each leg.

While both, fleet assignment problems and Revenue Management, have drawn much attention from researchers and practitioners, literature on DDR is limited and has largely focused on recovery strategies from operational disruptions. The first scientific analysis of systematic changes in response to demand is the work of Berge and Hopperstad (1993). Their study shows a potential of 1%-5% in profits due to spill reductions and the use of smaller aircraft. Their results are affirmed by a more recent study by Frank et al. (2006) based on real-world data.

In both studies, booking limits are derived using the EMSR heuristic (c.f. section 4.5.2)

after an initial fleet assignment. At certain planning points during the booking process, forecasts are updated including the current bookings-in-hand. A fleet assignment problem (FAM) limited to the sub-network of exchangeable crew-compatible fleets is solved and booking policies are updated with the new capacity assignments. A feasible assignment is restricted to accommodate all current bookings. This process is repeated until the fleet assignment needs to be finalized. Frank et al. (2006) analyze various times for the latest assignment. Berge and Hopperstad (1993) do not restrict the latest assignment time. The process is depicted in Figure 5.3.

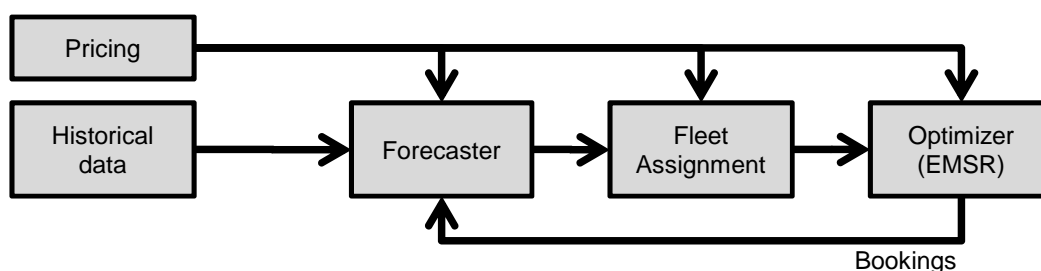


Figure 5.4: Demand-Driven Re-Fleeting process as proposed by Berge and Hopperstad (1993).

Simulation results show that the profit improvements are substantially dependent on the time spans between the planning points (c.f. section 6.3.1). Hence, the assignment needs to be reevaluated frequently. To reduce computational costs, Berge and Hopperstad (1993) propose two heuristics as alternatives to the FAM program described in section 5.2. An efficient and more effective algorithm is presented in Talluri (1996) which is described in detail in section 5.4.

Several practitioners have reported on implementations of DDR during the last two decades. Early approaches comprised a completely manual process for reviewing assignments close to departure supported by reservations and RM systems and their forecasts (e.g. at KLM and Australian Airlines as reported in Berge and Hopperstad (1993)). Pastor (1999) reports that at Continental Airlines, 60 and 14 days before departure, swaps are optimized and benefits are simulated to support manual reviews and decisions. Today, the process as proposed by Berge and Hopperstad (1993) represents the industry practice. However, the degree of manual interaction varies from manual decisions to a fully automated process (Oba, 2007, Zhao et al., 2007). Also, many airlines limit re-fleeting to cockpit-compatible families. All reports agree in that DDR yields substantial profit benefits due to lower operating cost, higher utilization, and revenue increases (lower spill) (see also Jacobs et al., 2001).

Despite the revenue improvements, the aforementioned approaches to DDR make use of traditional RM techniques. These assume the currently assigned capacity as de-

terministic. Hence, the effect of possible changes on the RM policies are neglected and the benefits of DDR are not completely exploited. A RM policy not capturing a possible upgrade or downgrade of capacity will be overly restrictive towards low-value demand or insufficiently restrictive, respectively. De Boer (2004) extends the EMSR-b heuristic to an uncertain capacity. He roughly approximates fleet assignment probabilities, i.e. probabilities that specific capacities are assigned to a flight. Using these approximated probabilities, the probability $\mathcal{Q}_j(b)$ that the aggregated demand of all higher fare classes $j - 1, \dots, 1$ exceeds the reserved capacity is calculated conditioned on the applied booking limit b . Similarly to EMSR-b, the booking limit b_j^d for class j is then determined by

$$b_j^d = \min\{0 \leq b \leq C : f_j < \bar{f}_{j-1} \mathcal{Q}_j(b)\} \quad (5.7)$$

with $\min \emptyset := C$. Note that the booking limits are bound by the currently assigned capacity C to guarantee a feasible solution. Further, (5.7) relaxes the integrality of the controls. The minimum is taken to break ties of different solutions.

Similarly, Wang and Regan (2006) develop a dynamic capacity control problem as described in section 4.5.3 when the capacity is uncertain. Consistent with De Boer (2004), their simulation studies show a revenue potential of up to 1.6% when incorporating the possible change into the RM policy.

The works by De Boer (2004) and Wang and Regan (2006) both assume a given set of flights that are subject to swaps that do not affect other assignments in the network. Such a set can be found efficiently using the algorithm introduced in section 5.4. Both approaches suffer from the requirement to approximate the probabilities for various fleet assignments. By assuming equal fares and assignment costs for all flights under consideration, the task is greatly simplified. Then, instead of evaluating profits, demand distributions can be used to calculate the probabilities. The simplification limits an application to flights serving the same markets with close departure times. Even then, the assumption of equal fares remains questionable.

The only DDR model that anticipates possible future equipment changes in the control policy and that incorporates assignment costs is the work by Wang and Meng (2008). They study a dynamic network capacity control problem in continuous time deriving a threshold policy that determines times to open or close a booking class. A set of feasible network assignments is needed as an input and is determined in a preceding optimization step subject to the basic constraints and possibly additional airline specific constraints such as allowing swaps only within cockpit-compatible fleet families. As bookings are accepted those assignments that become infeasible are removed from the set. Their solution procedure includes several steps involving feasibility checks of the fleet assignments. These are NP-complete problems and can be computationally

demanding. Hence, they propose a heuristic and do not further analyze their exact model. Unfortunately, the heuristic proposed is again limited to assignments involving equal costs which greatly restrains the scope of practical applications.

5.4 Fleet Assignment Swaps

If an assigned type is to be changed only for a certain leg in the schedule, aircraft can be swapped against each other with other legs while maintaining other assignments in the schedule. Berge and Hopperstad (1993) and Talluri (1996) provide algorithms that find swapping opportunities and preserve the cover (5.2), count (5.3), and balance (5.4) constraints in the schedule. We outline the later approach below because, in contrast to the former, the algorithm finds an opportunity if one exists (Talluri, 1996, Claim 1).

A change might become necessary during the planning phase if a planner makes manual corrections to the schedule or because of locked rotations, i.e. when the schedule does not allow for an acceptable maintenance routing. After the planning phase, operational disruptions such as delays or breakdowns might make a swap necessary.

The same algorithms can be used for Demand-Driven Re-Fleeting. If a manual re-fleeting process is adopted by an airline, the planner might use the algorithm to find an appropriate sequence of legs for the equipment swap. Models intended for an automated re-fleeting process such as the models proposed by De Boer (2004) or Wang and Regan (2006) assume a given set of legs that are subject to a possible swap. Following that idea, we will assume such a set as given in chapter 6. Using the described algorithm such a set of flights can be identified and then controlled applying DDR models.

Figure 5.5 illustrates examples of swapping opportunities. Note that the flight network might be reduced by consolidating nodes with feasible connections. Ground arcs are then only used to accurately model possible turns, e.g. when departures are timed such that some arriving flights can turn to all departures and others only to certain later departures. Figure 5.5 shows such a compact flight network facilitating shorter algorithm runs. Legs in the schedule apart from the shown sub-networks are not affected by an equipment change and the assignment constraints are maintained.

Depending on the objective, different costs are used to weight flight and ground arcs. If the number of legs affected is to be minimized, unit costs are applied. To minimize costs, swapping costs are added to differences of the assignment costs to the initial fleet assignment. If negative costs occur, the initial assignment is not optimal. In that case, adding a constant such that all costs are positive prevents cycles. To find a set of legs subject to a possible swap for Demand-Driven Re-Fleeting, demand correlations might be used. The correlations have to be rescaled such that all weights are positive.




Figure 5.5: Examples of two types of swapping opportunities. The dashed and solid arcs are assigned the same type. The assignments might be swapped without affecting other legs in the network. Reprinted from Talluri (1996).

A penalty for the number of equipment changes can be added when using assignment costs or correlations.

The fledged schedule and the cost matrix are used in Algorithm 5.4.1 to find an optimal swapping opportunity if one exists.

Algorithm 5.4.1 Algorithm to find a same-day swap opportunity (Talluri, 1996).

Input: Fledged flight network,
 cost matrix,
 leg l currently assigned type A,
 fleet type B to be assigned to leg l

Output: Set of legs with assignments to be swapped

1. Remove all overnight arcs and all arcs not assigned either type A or type B from the fledged flight network.
 2. Reverse the direction of all arcs with an assignment of type B in the flight network.
 3. Find a path from the head to the tail of leg l using a shortest-path algorithm with the provided cost matrix.
-

Algorithm 5.4.1 is restricted to swapping two types of airplanes. Swaps involving three or more types can be handled by re-solving the FAM problem for part of the flight network. Yet, Algorithm 5.4.1 can be applied sequentially to find an approximate solution. Talluri (1996) also provides details on how to implement the algorithm in case of diverse turn times and in case swaps are not restricted to retain overnight allocations.

In the algorithm, the shortest-path routine can be replaced by a procedure to rank a number of shortest paths (e.g. Martins and Pascoal, 2003). These can then be presented

to a planner as alternatives in a manual re-fleeting process.

Instead of solving a FAM instance reduced to a subset of legs the proposed algorithm builds on finding a shortest path in a directed network. For this problem, many efficient algorithms have been developed. The well-known algorithm by Dijkstra (1959) runs in $O(v^2)$ for a network with v nodes and has been improved several times since. A comprehensive review of efficient algorithms is given in Schultes (2008). In case of unit costs the run time decreases to $O(m)$ for m arcs (Ahuja et al., 1993).

Demand-Driven Re-Fleeting and Dynamic Pricing

In this chapter, we introduce dynamic pricing (DP) models incorporating Demand-Driven Re-Fleeting (DDR). As before, we adhere to the general assumptions listed in section 4.2. While we still consider the capacity as limited and fixed in the short term, we now refer to the overall network capacity rather than the capacity for each leg. Hence, capacities for individual legs might be changed by swapping aircraft with another leg.

Additionally, we make the following common assumptions:

- There are no cancellations, no-shows, or go-shows. Each accepted reservation cannot be rejected later without significant cost. Consequently, overbooking is not considered.
- Group arrivals can be partially accepted.
- Demand for different legs is mutually independent. The price offered for one leg does not affect the demand for other legs. Then, the network problem naturally reduces to separate leg instances. Hence, we only consider those legs in the network that are affected by a possible swap of aircraft.
- The decision maker is risk-neutral. Hence, the *expected* revenues are maximized. Optimization based on expectations seems reasonable in the light of the daily repeated decision process for numerous departures.
- The aggregated demand for all legs, i.e. the arrivals for any flight, is assumed to be distributed according to a Poisson distribution. Time is discretized finely such that the probability of more than one arrival in a time period is negligible.

- Operational, legal, or contractual requirements restrict changes to the fleet assignment to the time period until $\bar{T} > 0$ periods before departure. The assignment then needs to be finalized.

Our motivation is primarily to provide DP models that incorporate the re-fleeting decisions and that consider the possibility of changing aircraft in the pricing policy. In numerical studies we analyze the impact on revenues of DDR and the additional benefits of incorporating the possible swap into the pricing policy.

Additionally, we want to overcome the limitations found in existing capacity control models that develop single-leg policies conditioned on exogenously given probabilities for certain capacities (De Boer, 2004, Wang and Regan, 2006). The approach inherently yields some serious drawbacks for practical applications. The probabilities for each leg and capacity need to be estimated which is a challenging task itself. Hence, both works simplify by assuming equal fare, demand, and cost structures for legs involved in a swap. These simplifications limit the application to flights with the same origin and destination and with close departure times. However, even then, the assumption of equal fares and demands is questionable. Customers value flights differently at different times. We provide models that are capable of considering different demand structures and fares.

If fleet assignments are changed, the assignment costs usually change. These costs, which are assumed to be sunk in traditional RM models, need to be taken into account when evaluating a possible swap. The models proposed are capable of considering differences in assignment costs as well as other swapping costs that might arise (e.g. additional crew costs).

We take an approach of accounting for possible changes in the strategy while limiting the number of possible changes to one swap of the assignment. The idea is to start the booking process with the lower capacity for all flight legs. There is an option to assign the larger aircraft to one set of the flight legs that share an equal assignment. Once the option has been executed, no further changes are possible and the assignment is finalized. While this seems to be quite restrictive at first, it is reasonable to assume. As long as the possible swap is being considered in the pricing strategy, the actual decision can be postponed. Only when no further bookings could be accepted without executing the option, a decision needs to be made which is naturally final. To prove the point, consider two flights without overbooking. Assume one of the flights has no capacity left and the option is executed allocating the extra capacity to that flight. If we accept only one further booking request, the smaller capacity is not sufficient to satisfy the bookings-in-hand. Without overbooking, we now cannot switch back the assignment, it is final.

We start by modeling two legs subject to a swap in section 6.1. We derive structural results and use them to develop an efficient algorithm to find an optimal pricing policy. The model is extended in section 6.2 to sequences of several legs to overcome limitations of the two-leg model. Heuristic approaches are developed in section 6.3. We conclude by analyzing the performance of the various strategies in simulation studies in section 6.4. The benefits of applying DDR in general and the gain by considering a possible future swap in the pricing policies are assessed.

6.1 Two-leg model

We consider two aircraft with capacities C_1 and C_2 , respectively, and w.l.o.g., we assume $C_2 > C_1$. The planes are assigned to legs that are considered for an equipment swap. We assume that the swap does not affect the aircraft routings on other legs. For example, take two flights between the same origin and destination with close departure and arrival times as illustrated in Figure 6.1. Another application could be connections to two different hubs or bases where the aircraft can be switched back with backup aircraft or complementary swaps.




Figure 6.1: Example of a sub-network with two legs (Flight A and B) that are eligible for a swap with the same origin and destination.

We define a Markov Decision Process $(T, S, A, \mathbf{p}, r, V_0)$ as follows:

- (i) The planning horizon is $T \in \mathbb{N}$ which is the length of the booking period divided into periods with at most one customer arriving. Note that the periods need not be of the same length. Time is counted backwards and $t = T, T-1, \dots, 0$ denotes the number of periods until departure, i.e. the remaining periods for ticket sales.
- (ii) The joint state space is

$$\begin{aligned}
 S = & \{0, \dots, C_1\} \times \{0, \dots, C_1\} \times \{1\} \\
 & \cup \{0, \dots, C_2\} \times \{0, \dots, C_1\} \times \{0\} \\
 & \cup \{0, \dots, C_1\} \times \{0, \dots, C_2\} \times \{0\} .
 \end{aligned}$$

For $s = (s_1, s_2, s_3) \in S$, s_1 denotes the (remaining) capacity available on flight 1. s_2 is defined analogously for flight 2. If the option to assign the larger aircraft is still available $s_3 = 1$ and $s_3 = 0$, if it has been executed, i.e. the capacity difference $d := C_2 - C_1$ has already been added to one of the flights.

The assignment needs to be finalized the latest at $t = \bar{T}$. Hence, $S_t = S$ for $t = T, T - 1, \dots, \bar{T}$ and

$$S_t = S \setminus \{0, \dots, C_1\} \times \{0, \dots, C_1\} \times \{1\}$$

for $t = \bar{T} - 1, \bar{T} - 2, \dots, 0$.

- (iii) The action space is $A = \{(a_1, a_2, a_3) : a_1 \in \mathcal{P}_1, a_2 \in \mathcal{P}_2, a_3 \in \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}), (\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}), (\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})\}\}$, where $\mathcal{P}_i = \{f_0^i, f_1^i, \dots, f_{k_i}^i\}, k_i \in \mathbb{N}, i \in \{1, 2\}$, is the set of allowable prices for flight i including the nullprice f_0^i . The decision to execute the option on flight i is $a_3(i) = 1$, and $a_3(1) = a_3(2) = 0$ denotes the decision not to execute the option.

For all $t = T, T - 1, \dots, 1$ in state $s \in S_t$ with $s_3 = 0$ the sets of feasible actions $A_t(s) \subseteq A$ are

$$\begin{aligned} A_t(0, 0, 0) &= \{f_0^1\} \times \{f_0^2\} \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})\}, \\ A_t(0, s_2, 0) &= \{f_0^1\} \times \mathcal{P}_2 \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})\}, s_2 > 0, \\ A_t(s_1, 0, 0) &= \mathcal{P}_1 \times \{f_0^2\} \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})\}, s_1 > 0, \\ A_t(s_1, s_2, 0) &= \mathcal{P}_1 \times \mathcal{P}_2 \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})\}, s_1, s_2 > 0. \end{aligned}$$

For $t = T, T - 1, \dots, \bar{T} + 1$ in state $s \in S_t$ with $s_3 = 1$, we additionally have

$$\begin{aligned} A_t(0, 0, 1) &= \mathcal{A}_1 := \{f_0^1\} \times \mathcal{P}_2 \times \{(\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})\} \cup \mathcal{P}_1 \times \{f_0^2\} \times \{(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix})\} \cup \{f_0^1\} \times \{f_0^2\} \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})\}, \\ A_t(0, s_2, 1) &= \mathcal{A}_2 := \mathcal{P}_1 \times \mathcal{P}_2 \times \{(\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})\} \cup \{f_0^1\} \times \mathcal{P}_2 \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}), (\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})\}, s_2 > 0, \\ A_t(s_1, 0, 1) &= \mathcal{A}_3 := \mathcal{P}_1 \times \mathcal{P}_2 \times \{(\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})\} \cup \mathcal{P}_1 \times \{f_0^2\} \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}), (\begin{smallmatrix} 1 \\ 0 \end{smallmatrix})\}, s_1 > 0, \\ A_t(s_1, s_2, 1) &= \mathcal{A}_4 := \mathcal{P}_1 \times \mathcal{P}_2 \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}), (\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}), (\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})\}, s_1, s_2 > 0. \end{aligned}$$

Finally, the assignment must be decided the latest at $t = \bar{T}$. Thus, in state $s \in S_{\bar{T}}$ with $s_3 = 1$, the sets of feasible actions are

$$\begin{aligned} A_{\bar{T}}(0, 0, 1) &= \mathcal{A}_1 \setminus \{f_0^1\} \times \{f_0^2\} \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})\}, \\ A_{\bar{T}}(0, s_2, 1) &= \mathcal{A}_2 \setminus \{f_0^1\} \times \mathcal{P}_2 \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})\}, s_2 > 0, \\ A_{\bar{T}}(s_1, 0, 1) &= \mathcal{A}_3 \setminus \mathcal{P}_1 \times \{f_0^2\} \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})\}, s_1 > 0, \\ A_{\bar{T}}(s_1, s_2, 1) &= \mathcal{A}_4 \setminus \mathcal{P}_1 \times \mathcal{P}_2 \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})\}, s_1, s_2 > 0. \end{aligned}$$

(iv) The transition law in $t = T, T - 1, \dots, 1$ for $s \in S_t, s' \in S_{t-1}, a \in A_t(s)$ is

$$\mathbf{p}_t(s, a, s') = \begin{cases} \lambda_t \cdot p_t^i \cdot q_t^i(a_i) & \begin{aligned} s'_1 &= s_1 - \delta_1(i) + d \cdot a_3(1), \\ s'_2 &= s_2 - \delta_2(i) + d \cdot a_3(2), \\ s'_3 &= s_3 - a_3(1) - a_3(2), \\ & i \in \{1, 2\}, \end{aligned} \\ 1 - \sum_{i=1}^2 \lambda_t \cdot p_t^i \cdot q_t^i(a_i) & \begin{aligned} s'_1 &= s_1 + d \cdot a_3(1), \\ s'_2 &= s_2 + d \cdot a_3(2), \\ s'_3 &= s_3 - a_3(1) - a_3(2), \end{aligned} \\ 0 & \text{otherwise.} \end{cases}$$

Here, λ_t denotes the combined probability for an arrival in period t . The thinning probability for flight $i = 1, 2$ is given by p_t^i and $q_t^i(a_i)$ denotes the related purchase probability given price $a_i \geq 0$. $\delta_i(j)$ is the indicator function with $\delta_i(j) = 1$, if $i = j$ and otherwise $\delta_i(j) = 0$.

(v) The reward function for $t = T, T - 1, \dots, 1$ is

$$r_t(s, a) = \sum_{i=1}^2 \lambda_t \cdot p_t^i \cdot q_t^i(a_i) \cdot a_i - k_d^i \cdot a_3(i) \quad , s \in S_t, a \in A_t(s) ,$$

where k_d^i denotes the additional cost of flying the larger airplane on flight i , i.e. the cost for d additional seats on flight i . For at least one $i \in \{1, 2\}$, $k_d^i = 0$ (see discussion below).

(vi) The terminal reward function is $V_0 \equiv 0$. We do not consider overbooking and do not incur any terminal cost.

Defining $\tilde{s}_1 := s_1 + d \cdot a_3(1)$, $\tilde{s}_2 := s_2 + d \cdot a_3(2)$, and $\tilde{s}_3 := s_3 - a_3(1) - a_3(2)$, the optimality equation at $t = T, T - 1, \dots, 1$ for all states $s = (s_1, s_2, s_3) \in S_t$ can then be written as

$$V_t(s) = \max_{a \in A_t(s)} \left\{ \sum_{i=1}^2 \lambda_t \cdot p_t^i \cdot q_t^i(a_i) [a_i + V_{t-1}(\tilde{s}_1 - \delta_1(i), \tilde{s}_2 - \delta_2(i), \tilde{s}_3)] \right. \\ \left. + \left[1 - \sum_{i=1}^2 \lambda_t \cdot p_t^i \cdot q_t^i(a_i) \right] V_{t-1}(\tilde{s}_1, \tilde{s}_2, \tilde{s}_3) - \sum_{i=1}^2 k_d^i \cdot a_3(i) \right\} . \quad (6.1)$$

Usually, in RM models, costs are ignored. Variable costs are considered marginal and thus need not be considered. Fixed costs are considered sunk as the service needs to be provided in any case. In the airline business, a scheduled flight takes place even with very few bookings. Hence, every additional seat sold yields a contribution toward the fixed flight costs while it perishes if not sold. The same is true e.g. in fashion retailing. Once an order has been placed and has been produced fixed costs are sunk.

The fixed costs themselves are just a constant in the objective function to maximize profits. Hence, maximizing revenue approximately maximizes profits.

In our model, we do need to consider one cost component. As in regular RM models, we do not consider variable or fixed costs in general. However, we do need to consider the change in fixed costs by changing the assignment and other swapping costs incurred, e.g. additional crew costs.

We assume the plane with the smaller capacity is cheaper to fly for a fixed flight and there are just enough planes to serve all legs considered. If the updated total demand forecast including the bookings-in-hand at the time of the assignment decision can be satisfied using the plane with the smaller capacity for both flights, one of the legs still needs to be served with the larger aircraft. In that case, we simply take the most cost efficient allocation. Let us assume fixed cost K_1 for that fleet assignment. If planes are swapped, we incur assignment cost $K_2 \geq K_1$. Apart from lost demand (spill), the expected gain in revenue needs to offset the difference in cost $K_2 - K_1$ for the change to be justified. Thus, this difference in fixed costs needs to be considered in our model. We do so in the one-stage reward function r_t , where k_d^i is subtracted of the total value when an assignment decision is made. As we only consider the difference to the most cost efficient allocation, $k_d^i = 0$ for flight i that is assigned the larger airplane in that assignment. For an allocation of the additional capacity, i.e. the larger airplane, to the other flight i' , we incur cost $k_d^{i'} = K_2 - K_1 \geq 0$.

Assignment costs need not be considered in the terminal reward function as we assume $\bar{T} > 0$ to be the latest time in the booking period when the assignment must be finalized.

Note that the consideration of changes in costs allows to consider not only flights that have the same destination ($k_d^i = 0, i = 1, 2$), but also flights serving different markets.

To simplify notation, we define operators and sets before we continue with structural results:

For all $t = T, T - 1, \dots, 1$, $s = (s_1, s_2, s_3) \in S_t$, $a = (a_1, a_2, a_3) \in A_t(s)$, $i = 1, 2$ let

$$\begin{aligned}
Y(s, a) &:= (s_1 + a_3(1)d, s_2 + a_3(2)d, s_3 - a_3(1) - a_3(2)) , \\
Y^i(s, a) &:= \begin{cases} Y((s_1 - \delta_1(i), s_2 - \delta_2(i), s_3), a) & s_i + a_3(i)d > 0 , \\ Y(s, a) & s_i + a_3(i)d = 0 , \end{cases} \\
A_t(s|a_1, a_2) &:= \{a_3 \in \left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\} : (a_1, a_2, a_3) \in A_t(s)\} , \\
A_t(s|a_3) &:= \{(a_1, a_2) \in \mathcal{P}_1 \times \mathcal{P}_2 : (a_1, a_2, a_3) \in A_t(s)\} , \\
A_t^1(s) &:= \{a_1 \in \mathcal{P}_1 : (a_1, a_2, (0, 0)') \in A_t(s), a_2 \in \mathcal{P}_2\} , \\
A_t^2(s) &:= \{a_2 \in \mathcal{P}_2 : (a_1, a_2, (0, 0)') \in A_t(s), a_1 \in \mathcal{P}_1\} ,
\end{aligned}$$

$$\mathfrak{D}_t := \{(s, a) : s \in S_t, a \in A_t(s)\} ,$$

and, for $s_i > 0$,

$$\Delta_{s_i} V_t(s) := V_t(s_1, s_2, s_3) - V_t(s_1 - \delta_1(i), s_2 - \delta_2(i), s_3) .$$

Further, for all $v : S_{t-1} \rightarrow \mathbb{R}$ and $a \in A_t(s)$ let

$$\begin{aligned} Q_t v(s, a) &:= \sum_{i=1}^2 \lambda_t p_t^i q_t^i(a_i) [a_i + v(Y^i(s, a))] \\ &\quad + \left(1 - \sum_{i=1}^2 \lambda_t p_t^i q_t^i(a_i)\right) v(Y(s, a)) - \sum_{i=1}^2 k_d^i a_3(i) , \end{aligned}$$

$$L_t v(s|a_1, a_2) := \max_{a' \in A_t(s|a_1, a_2)} \left\{ Q_t v(s, (a_1, a_2, a')) \right\} ,$$

$$L_t v(s|a_3) := \max_{a' \in A_t(s|a_3)} \left\{ Q_t v(s, (a'_1, a'_2, a_3)) \right\} .$$

We now derive structures of the model that are exploited later in an algorithm to find an optimal policy and in the heuristic approaches.

Theorem 6.1.1 (Decomposition of the value function).

For the MDP with the value function defined in (6.1),

(i) at time $t = T, \dots, 1$, in all states $s = (s_1, s_2, 0) \in S_t$, the process decomposes into two independent single-leg processes with the value functions

$$V_t^i(s_i) := \max_{a_i \in \mathcal{P}_i} \left\{ \lambda_t p_t^i q_t^i(a_i) [a_i + V_{t-1}^i(s_i - 1)] + [1 - \lambda_t p_t^i q_t^i(a_i)] V_{t-1}^i(s_i) \right\}$$

for $s_i > 0$ and $i \in \{1, 2\}$. The boundary conditions without overbooking are $V_t^i(0) = V_{t-1}^i(0)$ and $V_0^i \equiv 0$.

The value function of the original process is then the sum of the two single-leg value functions:

$$V_t(s_1, s_2, 0) = V_t^1(s_1) + V_t^2(s_2) .$$

(ii) Further, for $t = T, \dots, \bar{T}$, $s = (s_1, s_2, 1) \in S_t$, and $a_3 \in \left\{ \binom{1}{0}, \binom{0}{1} \right\}$,

$$L_t V_{t-1}(s|a_3) = \sum_{i=1}^2 V_t^i(s_i + a_3(i)d) - a_3(i)k_d^i .$$

(iii) For $t = T, \dots, \bar{T} + 1$, $s = (s_1, s_2, 1) \in S_t$,

$$L_t V_{t-1}(s|(0, 0)') = \sum_{\substack{i=1 \\ s_i \neq 0}}^2 \left(\max_{a_i \in A_t^i(s)} \lambda_t p_t^i q_t^i(a_i) [a_i - \Delta_{s_i} V_{t-1}(s)] \right) + V_{t-1}(s) .$$

Proof.

- (i) We prove by induction on t . For $t = 0$, let $s^0 = (s_1, s_2, 0) \in S_0$. Since $V_0 \equiv 0$, the assertion holds. Therefore, let $V_t(s) = V_t^1(s_1) + V_t^2(s_2)$ hold for some $t = T - 1, T - 2, \dots, 0$ and $s^t = (s_1, s_2, 0) \in S_t$. Then, in $s = (s_1, s_2, 0) \in S_{t+1}$, since $s_3 = 0$, only actions $a = (a_1, a_2, a_3) \in A_{t+1}(s)$ with $a_3 = (0, 0)'$, and additionally, $a_i = f_0^i$ for $i = 1, 2$ with $s_i = 0$. We get

$$\begin{aligned}
V_{t+1}(s) &= \max_{a \in A_{t+1}(s)} \left\{ \sum_{\substack{i=1 \\ s_i \neq 0}}^2 \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i + V_t(Y^i(s, a))] \right. \\
&\quad \left. + \left[1 - \sum_{\substack{i=1 \\ s_i \neq 0}}^2 \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) \right] V_t(Y(s, a)) \right\} \\
&= \max_{a \in A_{t+1}(s)} \left\{ \sum_{\substack{i=1 \\ s_i \neq 0}}^2 \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i + V_t^1(s_1 - \delta_1(i)) + V_t^2(s_2 - \delta_2(i))] \right. \\
&\quad \left. + \left[1 - \sum_{\substack{i=1 \\ s_i \neq 0}}^2 \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) \right] [V_t^1(s_1) + V_t^2(s_2)] \right\} \\
&= \max_{a \in A_{t+1}(s)} \left\{ \sum_{\substack{i=1 \\ s_i \neq 0}}^2 \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i - \Delta V_t^i(s_i)] + V_t^1(s_1) + V_t^2(s_2) \right\} \\
&= \sum_{\substack{i=1 \\ s_i \neq 0}}^2 \max_{a_i \in A_{t+1}^i(s)} \left\{ \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i - \Delta V_t^i(s_i)] \right\} + V_t^1(s_1) + V_t^2(s_2) \\
&= V_{t+1}^1(s_1) + V_{t+1}^2(s_2) ,
\end{aligned}$$

which completes the proof for (i). (ii) follows straightforward applying (i).

- (iii) For $t = T - 1, T - 2, \dots, \bar{T}$, let $a_3 = (0, 0)'$ and $s = (s_1, s_2, 1) \in S_{t+1}$. Then,

$$\begin{aligned}
&L_{t+1} V_t(s | a_3) \\
&= \max_{(a_1, a_2) \in A_{t+1}(s | a_3)} \left\{ \sum_{\substack{i=1 \\ s_i \neq 0}}^2 \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i - \Delta_{s_i} V_t(s_1, s_2, 1)] \right\} + V_t(s_1, s_2, 1) \\
&= \sum_{\substack{i=1 \\ s_i \neq 0}}^2 \max_{a_i \in A_{t+1}^i(s)} \left\{ \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i - \Delta_{s_i} V_t(s_1, s_2, 1)] \right\} + V_t(s_1, s_2, 1) .
\end{aligned}$$

□

Theorem 6.1.1 (i) allows to use the one-dimensional process with the optimality equation (4.7), adjusted for the thinned probabilities, for each leg separately in states where $s_3 = 0$, i.e. when the option has been executed. Also in states with $s_3 = 1$, if the optimal assignment action $a_3^* \in \{(0, 1)', (1, 0)'\}$ is fixed to execute the option, the one-dimensional processes can be used to calculate the value and the optimal pricing actions

by applying (ii). Additionally, using (iii), when the assignment action $a_3^* = (0, 0)'$ is fixed, i.e. the option is not executed, to derive the optimal prices, instead of maximizing over one two-dimensional set, we can maximize over two one-dimensional action sets. These structures greatly reduce computational effort to find an optimal policy and the expected revenue and allow for parallel computation of the policy for each leg. Theorem 6.1.1 is not only used to develop an efficient algorithm to find an optimal pricing policy, but also to develop a limited lookahead policy heuristic in section 6.3.4.

Theorem 6.1.2 (Monotonicity of the value function).

For the MDP with the value function defined in (6.1),

- (i) for fixed $t = T, T - 1, \dots, 0$, $V_t(s)$ is non-decreasing in s_1 for fixed s_2 and in s_2 for fixed s_1 .
- (ii) $V_t(s_1, s_2, 1) - V_t(s_1, s_2, 0) \geq 0$, for all $t = T, T - 1, \dots, \bar{T}$, and s_1, s_2 with $(s_1, s_2, 1) \in S_t$ and $(s_1, s_2, 0) \in S_t$.
- (iii) $V_t(s)$ is non-decreasing in t .

Proof.

- (i) The proof follows by induction on t . Since $V_0 \equiv 0$, the assertion holds for $t = 0$. Therefore, let (i) hold for some $t = T - 1, T - 2, \dots, 0$. Let $s = (s_1, s_2, s_3) \in S_{t+1}$ with $s_1 \geq 1$, and let $a^* = (a_1^*, a_2^*, a_3^*) \in A_{t+1}(s_1 - 1, s_2, s_3)$ be optimal in $s' = (s_1 - 1, s_2, s_3) \in S_{t+1}$. Then $a^* \in A_{t+1}(s)$, too, and we get

$$\begin{aligned} \Delta_{s_1} V_{t+1}(s) &\geq \sum_{i=1}^2 \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i^*) \Delta_{s_1} V_t(Y^i(s, a^*)) \\ &\quad + \left[1 - \sum_{i=1}^2 \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i^*) \right] \Delta_{s_1} V_t(Y(s, a^*)) \\ &\geq 0, \end{aligned}$$

which completes the proof for s_1 . The monotonicity in s_2 follows analogically.

- (ii) To prove (ii), let $t = T, T - 1, \dots, \bar{T}$. Fix $s' = (s_1, s_2, 1) \in S_t$. Then $s = (s_1, s_2, 0) \in S_t$, too. Further, let $a^* = (a_1^*, a_2^*, (0, 0)') \in A_t(s)$ be optimal in s and suppose $k_d^1 = 0$. Using $a^{**} = (a_1^*, a_2^*, (1, 0)') \in A_t(s')$, and additionally, exploiting the monotonicity of V_{t-1} in s_1 after an assignment, we finally get

$$V_t(s') - V_t(s)$$

$$\begin{aligned}
&\geq \sum_{i=1}^2 \lambda_t p_t^i q_t^i(a_i^*) \left[V_{t-1}(Y^i(s', a^{**})) - V_{t-1}(Y^i(s, a^*)) \right] \\
&\quad + \left[1 - \sum_{i=1}^2 \lambda_t p_t^i q_t^i(a_i^*) \right] \left[V_{t-1}(Y(s', a^{**})) - V_{t-1}(Y(s, a^*)) \right] \\
&\geq 0,
\end{aligned}$$

which confirms (ii) for $k_d^1 = 0$. If $k_d^1 > 0$, then, by assumption, $k_d^2 = 0$ and a similar argument works.

(iii) We split the proof into two parts:

a. First, we assume $s_3 = 0$. We use Theorem 4.6.3 and Theorem 6.1.1 to derive

$$\begin{aligned}
&V_{t+1}(s_1, s_2, 0) - V_t(s_1, s_2, 0) \\
&= V_{t+1}^1(s_1) + V_{t+1}^2(s_2) - V_t^1(s_1) - V_t^2(s_2) \\
&\geq 0.
\end{aligned}$$

b. Now, we consider $s = (s_1, s_2, s_3)$ with $s_3 = 1$ and $t = T - 1, T - 2, \dots, \bar{T}$. Thus, $a = (f_0^1, f_0^2, (0, 0)') \in A_{t+1}(s)$. We get

$$\begin{aligned}
&V_{t+1}(s_1, s_2, 1) - V_t(s_1, s_2, 1) \\
&\geq Q_{t+1} V_t((s_1, s_2, 1), a) - V_t(s_1, s_2, 1) \\
&= 0,
\end{aligned}$$

which completes the proof. \square

Theorem 6.1.2 shows some intuitive attributes of the value function. More available capacity results in higher or equal expected revenue. (ii) further states that the value of the option is not negative. The option might well be worthless though. In the case when there is no demand exceeding the minimal capacity C_1 for either flight, the extra seats do not have a positive value. Also, the more time remains to sell units of capacity, the higher the expected revenue. The monotonicity of the value function is frequently used in subsequent proofs.

Lemma 6.1.3.

For the MDP with the value function defined in (6.1), for $t = T, \dots, \bar{T}$,

$$(i) \quad V_t(s_1, s_2, 1) \geq V_t(s_1 + d, s_2, 0) - k_d^1.$$

$$(ii) \quad V_t(s_1, s_2, 1) \geq V_t(s_1, s_2 + d, 0) - k_d^2.$$

Proof. Let $t = T, \dots, \bar{T}$ and fix $s = (s_1, s_2, 1) \in S_t$. Then there exists an action $a = (a_1, a_2, (1, 0)') \in A_t(s)$ and

$$V_t(s) \geq L_t V_{t-1}(s|(1, 0)')$$

$$= V_t(s_1 + d, s_2, 0) - k_d^1.$$

(ii) can be shown analogically. \square

Lemma 6.1.3 provides lower bounds of the value function in states where $s_3 = 1$, i.e. when the capacity option is available. Applying Theorem 6.1.1, the bounds can be computed using the single-leg value functions V_t^1 and V_t^2 . We use the results in proofs throughout the remainder.

Theorem 6.1.4 (Optimal assignment actions).

For the MDP with the value function defined in (6.1), for any $t = T, \dots, \bar{T} + 1$, in state $s = (s_1, s_2, s_3) \in S_t$, an action $a^* = (a_1^*, a_2^*, a_3^*) \in A_t(s)$ is optimal with

$$(i) \ a_3^* \in \left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}, \text{ if } s_1 > 0.$$

$$(ii) \ a_3^* \in \left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right\}, \text{ if } s_2 > 0.$$

$$(iii) \ a_3^* = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \text{ if } s_1 > 0, s_2 > 0.$$

$$(iv) \ a_3^* \in \left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}, \text{ if } s = (0, 0, 1).$$

Proof.

(i) Let $t = T, \dots, \bar{T} + 1$. Fix $s = (s_1, s_2, 1) \in S_t$ with $s_1 > 0$. Assume $(a_1^*, a_2^*) \in A_t(s|(1, 0)')$ maximizes $L_t V_{t-1}(s|(1, 0)')$. Note that $(a_1^*, a_2^*, (0, 0)') \in A_t(s)$ is also admissible in s . Using Lemma 6.1.3, we then get

$$\begin{aligned} & L_t V_{t-1}(s|(0, 0)') - L_t V_{t-1}(s|(1, 0)') \\ & \geq \sum_{i=1}^2 \lambda_t p_t^i q_t^i(a_i^*) \left[V_{t-1}(Y^i(s, (0, 0)')) - V_{t-1}(Y^i(s, (1, 0)')) \right] \\ & \quad + \left[1 - \sum_{i=1}^2 \lambda_t p_t^i q_t^i(a_i^*) \right] \left[V_{t-1}(Y(s, (0, 0)')) - V_{t-1}(Y(s, (1, 0)')) \right] \\ & \quad + k_d^1 \\ & \geq \sum_{i=1}^2 \lambda_t p_t^i q_t^i(a_i^*) \left[V_{t-1}(Y^i(s, (1, 0)')) - k_d^1 - V_{t-1}(Y^i(s, (1, 0)')) \right] \\ & \quad + \left[1 - \sum_{i=1}^2 \lambda_t p_t^i q_t^i(a_i^*) \right] \left[V_{t-1}(Y(s, (1, 0)')) - k_d^1 - V_{t-1}(Y(s, (1, 0)')) \right] \\ & \quad + k_d^1 \\ & = 0. \end{aligned}$$

In states $s' = (s_1, s_2, 0) \in S_t$ only actions $a \in A_t(s')$ are feasible with the assignment action $a_3 = (0, 0)'$ and the result is trivial.

(ii) can be shown analogically and (iii) follows straightforward from (i) and (ii).

- (iv) In $s = (0, 0, 1) \in S_{\bar{T}}$, suppose $(a_1^*, a_2^*, (1, 0)') \in A_{\bar{T}}(s)$ to be optimal at time $t = \bar{T}$. Then, using Theorem 6.1.1 (ii) and Theorem 4.6.3 (i),

$$\begin{aligned}
& L_{\bar{T}+1}V_{\bar{T}}(s|(1, 0)') \\
&= V_{\bar{T}+1}^1(d) + V_{\bar{T}+1}^2(0) - k_d^1 \\
&\geq V_{\bar{T}}^1(d) + V_{\bar{T}}^2(0) - k_d^1 \\
&= V_{\bar{T}}(s) \\
&= L_{\bar{T}+1}V_{\bar{T}}(s|(0, 0)') .
\end{aligned}$$

Hence, applying $(1, 0)'$ results in a value at least as high as when applying $(0, 0)'$ in s at time $t = \bar{T} + 1$. Repeating the same argument with $\bar{T} + 1$ in place of \bar{T} , the result follows by induction.

A similar argument works if $(a_1^*, a_2^*, (0, 1)') \in A_{\bar{T}}(s)$ is optimal in s at time $t = \bar{T}$. \square

Theorem 6.1.4 reduces the set of actions to maximize over tremendously, which can be exploited to find an optimal policy efficiently. In states $s = (s_1, s_2, s_3)$, where $s_3 = 1$ and $s_1 = 0$ or $s_2 = 0$, one assignment action is at most as good as the others and thus needs not be included in the maximization. For the majority of states, where $s_1, s_2 > 0$, Theorem 6.1.4 can be used to fix the assignment action $a_3^* = (0, 0)'$. Then, in conjunction with Theorem 6.1.1 (iii), the maximization can be separated for each leg, which results in a maximization over two one-dimensional sets of prices instead of having to consider the complete three-dimensional set of feasible actions. Applying (iv), at any time $t = T, \dots, \bar{T} + 1$, given state $s = (0, 0, 1)$, we always assign the option to one of the flights.

Theorem 6.1.5 (Monotonicity of the assignment action at $t = \bar{T}$).

In $s = (s_1, s_2, 1) \in S_{\bar{T}}$,

$$L_{\bar{T}}V_{\bar{T}-1}(s|(1, 0)') - L_{\bar{T}}V_{\bar{T}-1}(s|(0, 1)')$$

is non-increasing in s_1 for fixed s_2 and non-decreasing in s_2 for fixed s_1 .

Proof. Let $s = (s_1, s_2, 1), s' = (s_1 + 1, s_2, 1) \in S_{\bar{T}}$. Note that using Theorem 6.1.1 (ii)

$$\begin{aligned}
& L_{\bar{T}}V_{\bar{T}-1}((s_1 + 1, s_2, 1)|(1, 0)') - L_{\bar{T}}V_{\bar{T}-1}((s_1 + 1, s_2, 1)|(0, 1)') \\
&= L_{\bar{T}}V_{\bar{T}-1}((s_1, s_2, 1)|(1, 0)') - L_{\bar{T}}V_{\bar{T}-1}((s_1, s_2, 1)|(0, 1)') \\
&\quad + V_{\bar{T}}^1(s_1 + 1 + d) - V_{\bar{T}}^1(s_1 + d) - V_{\bar{T}}^1(s_1 + 1) + V_{\bar{T}}^1(s_1)
\end{aligned}$$

Using the concavity of $V_{\bar{T}}^1$ (Theorem 4.6.2), the monotonicity in s_1 follows. A similar argumentation works for s_2 . \square

The results in Theorem 6.1.5 can be exploited to find an optimal policy in $t = \bar{T}$ for states with $s_3 = 1$, i.e. when the capacity option is still available and has to be executed. If in a state $(s_1, s_2, 1) \in S_{\bar{T}}$, an assignment of the additional capacity to flight 2 is optimal, then it is also optimal for $s' = (s'_1, s'_2, 1) \in S_{\bar{T}}$ with $s'_1 \geq s_1$ and $s'_2 \leq s_2$. Similarly, optimal actions assigning the capacity to flight 1 can be found. Figure 6.2 visualizes the structure yielding a switching curve for the assignment action.




Figure 6.2: Example of a switching curve policy for the assignment action at time $t = \bar{T}$ (c.f. Theorem 6.1.5).

Note that although the results in Theorem 6.1.5 hold for any $t = T, \dots, \bar{T}$, we use them only in $t = \bar{T}$ because for $t = T, \dots, \bar{T} + 1$, from Theorem 6.1.4 follows that action $a_3 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ ($a_3 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$) needs not be considered in states $s = (s_1 + 1, s_2, 1)$ (respectively $s = (s_1, s_2 + 1, 1)$).

Next, in order to prove the concavity of the value function, we first define

$$s_1^*(s_2) := \min \left\{ s_1 \in \{0, \dots, C_1\} : V_{\bar{T}}(s_1, s_2, 1) = V_{\bar{T}}^1(s_1) + V_{\bar{T}}^2(s_2 + d) - k_d^2 \right\},$$

$$s_2^*(s_1) := \min \left\{ s_2 \in \{0, \dots, C_1\} : V_{\bar{T}}(s_1, s_2, 1) = V_{\bar{T}}^1(s_1 + d) + V_{\bar{T}}^2(s_2) - k_d^1 \right\},$$

where $\min \emptyset := \infty$.

Additionally, we apply the results of the following two Lemmata:

Lemma 6.1.6.

For the MDP with the value function defined in (6.1), s_1^* is non-decreasing in s_2 and s_2^* is non-decreasing in s_1 .

Proof.

Fix $s = (s_1, s_2, 1)$ with $s_1 < s_1^*(s_2)$. Then $a_3^* = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ is optimal. Applying Theorem 6.1.5 yields that $a_3^* = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ is optimal for $s' = (s_1, s_2 + 1, 1)$ as well.

The monotonicity of s_2^* can be shown analogically. \square

Lemma 6.1.7.

In $t = T, T-1, \dots, \bar{T}$, in $s = (s_1, s_2, 1) \in S_t$ with $t - \bar{T} < s_1 < s_1^*(s_2 - t + \bar{T} - 1)$ or $s_1 \geq s_1^*(s_2) + t - \bar{T} + 1$ for fixed $s_2 > t - \bar{T}$,

$$\Delta_{s_1} V_t(s_1, s_2, 1) - \Delta_{s_1} V_t(s_1, s_2 - 1, 1) = \Delta_{s_2} V_t(s_1, s_2, 1) - \Delta_{s_2} V_t(s_1 - 1, s_2, 1) = 0 .$$

The same holds for $t - \bar{T} < s_2 < s_2^*(s_1 - t + \bar{T} - 1)$ or $s_2 \geq s_2^*(s_1) + t - \bar{T} + 1$ for fixed $s_1 > t - \bar{T}$.

Proof. The proof follows by induction on t . For $t = \bar{T}$ and $s_1 \geq s_1^*(s_2) + 1$ we get

$$\begin{aligned} & \Delta_{s_1} V_{\bar{T}}(s_1, s_2, 1) - \Delta_{s_1} V_{\bar{T}}(s_1, s_2 - 1, 1) \\ &= V_{\bar{T}}(s_1, s_2, 1) - V_{\bar{T}}(s_1 - 1, s_2, 1) - V_{\bar{T}}(s_1, s_2 - 1, 1) + V_{\bar{T}}(s_1 - 1, s_2 - 1, 1) \\ &= V_{\bar{T}}^1(s_1) + V_{\bar{T}}^2(s_2 + d) - k_d^2 - V_{\bar{T}}^1(s_1 - 1) - V_{\bar{T}}^2(s_2 + d) + k_d^2 \\ &\quad - V_{\bar{T}}^1(s_1) - V_{\bar{T}}^2(s_2 - 1 + d) + k_d^2 + V_{\bar{T}}^1(s_1 - 1) + V_{\bar{T}}^2(s_2 - 1 + d) - k_d^2 \\ &= 0 . \end{aligned}$$

Similarly one proves the assertion for $t = \bar{T}$ and $0 < s_1 < s_1^*(s_2 - 1)$. Hence, let the assertion hold for some $t = T-1, T-2, \dots, \bar{T}$. Additionally, for some $j \in \mathbb{N}, 1 < k \leq j$, let $e(k)$ be the j -vector with $e(k) = 1$ and $e(i) = 0$ for all $1 < i \leq j, i \neq k$.

In $t+1$, since $s_1 \geq s_1^*(s_2) + t - \bar{T} + 2 > 0$ and $s_2 > t - \bar{T} + 1 > 0$ an action $a = (a_1, a_2, a_3)$ is optimal with $a_3 = (0, 0)'$ by applying Theorem 6.1.4. Note that additionally $A_{t+1}^i(s) = A_{t+1}^i(s - e(j))$ for all $s \in S_{t+1}$ and $i \neq j$. Then

$$\begin{aligned} & \Delta_{s_1} V_{t+1}(s_1, s_2, 1) - \Delta_{s_1} V_{t+1}(s_1, s_2 - 1, 1) \\ &= \sum_{i=1}^2 \left(\max_{a_i \in A_{t+1}^i(s)} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i - \Delta_{s_i} V_t(s)] \right) + V_t(s) \\ &\quad - \sum_{i=1}^2 \left(\max_{a_i \in A_{t+1}^i(s - e(1))} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i - \Delta_{s_i} V_t(s - e(1))] \right) - V_t(s - e(1)) \\ &\quad - \sum_{i=1}^2 \left(\max_{a_i \in A_{t+1}^i(s - e(2))} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i - \Delta_{s_i} V_t(s - e(2))] \right) - V_t(s - e(2)) \\ &\quad + \sum_{i=1}^2 \left(\max_{a_i \in A_{t+1}^i(s - e(1) - e(2))} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i - \Delta_{s_i} V_t(s - e(1) - e(2))] \right) \\ &\quad + V_t(s - e(1) - e(2)) \\ &= 0 . \end{aligned}$$

For $t - \bar{T} + 1 < s_1 < s_1^*(s_2 - t + \bar{T} - 2)$ and $s_2 > t - \bar{T} + 1$ a similar argumentation works. The proof works analogically for $\Delta_{s_2} V_t(s_1, s_2, 1) - \Delta_{s_2} V_t(s_1 - 1, s_2, 1)$. \square

Hence, on certain domains, the opportunity costs $\Delta_{s_i} V_t$ of one seat on flight $i = 1, 2$ are independent of the number of available seats on the other flight. We use the results to prove the concavity of the value function on these domains in the next Theorem.

Theorem 6.1.8 (Concavity of the value function).

The value function V_t defined in (6.1)

- (i) is concave componentwise in s_1 and s_2 for all $t = T, T - 1, \dots, 0$ and $s = (s_1, s_2, s_3) \in S_t$ with $s_3 = 0$ being fixed.
- (ii) For $t = T, T - 1, \dots, \bar{T}$, in $s = (s_1, s_2, s_3) \in S_t$ with $t - \bar{T} < s_1 < s_1^*(s_2 - t + \bar{T} - 1) - 2$ or $s_1 \geq s_1^*(s_2) + t - \bar{T} + 1$, V_t is concave in s_1 for fixed $s_2 > t - \bar{T}$ and $s_3 = 1$.
- (iii) For $t = T, T - 1, \dots, \bar{T}$, in $s = (s_1, s_2, s_3) \in S_t$ with $t - \bar{T} < s_2 < s_2^*(s_1 - t + \bar{T} - 1) - 2$ or $s_2 \geq s_2^*(s_1) + t - \bar{T} + 1$, V_t is concave in s_2 for fixed $s_1 > t - \bar{T}$ and $s_3 = 1$.

Proof.

- (i) Let $(s_1, s_2, 0) \in S_t$ be fixed. Then, by applying Theorem 6.1.1,

$$V_t(s_1, s_2, 0) = V_t^1(s_1) + V_t^2(s_2) ,$$

where $V_t^1(s_1)$ is concave. Hence, $V_t(s_1, s_2, 0)$ is concave in s_1 .

The concavity in s_2 can be shown analogically.

- (ii) The proof follows by induction on t . For $t = \bar{T}$ the concavity in s_1 follows from the definition of s_1^* and the concavity of $V_{\bar{T}}^1$ (Theorem 4.6.2 in conjunction with Theorem 6.1.1).

Thus, let the assertion hold for some $t = T - 1, T - 2, \dots, \bar{T}$. Again, for some $j \in \mathbb{N}, 1 < k \leq j$, let $e(k)$ denote the j -vector with $e(k) = 1$ and $e(i) = 0$ for all $1 < i \leq j, i \neq k$.

In $t + 1$, let $s = (s_1, s_2, 1), s' = s + e(1), s'' = s + 2e(1) \in S_{t+1}$. Since $t + 1 > \bar{T}$, $s, s', s'' \in S_t$, too. Note that, since $s_1 \geq s_1^*(s_2) + t - \bar{T} + 2 > 0$ or $s_1 > t - \bar{T} + 1 > 0$ and $s_2 > t - \bar{T} + 1 > 0$, actions $a = (\cdot, \cdot, a_3)$ are optimal in s, s', s'' at $t + 1$ with $a_3 = (0, 0)'$ by applying Theorem 6.1.4.

Assume $a_1^* \in A_{t+1}^1(s'')$ maximizes

$$a_1 \rightarrow \lambda_{t+1} p_{t+1}^1 q_{t+1}^1(a_1) [a_1 - \Delta_{s_1} V_t(s'')]$$

and $a_1^{**} \in A_{t+1}^1(s)$ maximizes

$$a_1 \rightarrow \lambda_{t+1} p_{t+1}^1 q_{t+1}^1(a_1)[a_1 - \Delta_{s_1} V_t(s)].$$

Then, $a_1^*, a_1^{**} \in A_{t+1}^1(s')$, too. Making use of Lemma 6.1.7,

$$\begin{aligned} & \Delta_{s_1} V_{t+1}(s_1 + 2, s_2, 1) - \Delta_{s_1} V_{t+1}(s_1 + 1, s_2, 1) \\ &= \sum_{i=1}^2 \left(\max_{a_i \in A_{t+1}^i(s'')} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i)[a_i - \Delta_{s_i} V_t(s'')] \right) + V_t(s'') \\ & \quad - \sum_{i=1}^2 \left(\max_{a_i \in A_{t+1}^i(s')} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i)[a_i - \Delta_{s_i} V_t(s')] \right) - V_t(s') \\ & \quad - \sum_{i=1}^2 \left(\max_{a_i \in A_{t+1}^i(s')} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i)[a_i - \Delta_{s_i} V_t(s')] \right) - V_t(s') \\ & \quad + \sum_{i=1}^2 \left(\max_{a_i \in A_{t+1}^i(s)} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i)[a_i - \Delta_{s_i} V_t(s)] \right) + V_t(s) \\ & \leq \lambda_{t+1} p_{t+1}^1 q_{t+1}^1(a_1^*) \left[\Delta_{s_1} V_t(s') - \Delta_{s_1} V_t(s'') \right] + \Delta_{s_1} V_t(s'') \\ & \quad + \lambda_{t+1} p_{t+1}^1 q_{t+1}^1(a_1^{**}) \left[\Delta_{s_1} V_t(s') - \Delta_{s_1} V_t(s) \right] - \Delta_{s_1} V_t(s') \\ & = \left(1 - \lambda_{t+1} p_{t+1}^1 q_{t+1}^1(a_1^*) \right) \left[\Delta_{s_1} V_t(s'') - \Delta_{s_1} V_t(s') \right] \\ & \quad + \lambda_{t+1} p_{t+1}^1 q_{t+1}^1(a_1^{**}) \left[\Delta_{s_1} V_t(s') - \Delta_{s_1} V_t(s) \right] \\ & \leq 0, \end{aligned}$$

which concludes the proof. (iii) can be shown analogically. \square

Using the concavity of the value function we derive monotone structures of the pricing actions in the remainder.

Theorem 6.1.9 (Monotonicity of the price actions before the final assignment).

In $t = T - 1, T - 2, \dots, \bar{T}$ and $s = (s_1, s_2, 1) \in S_t$ with $s_i > 0, i = 1, 2$, suppose $V_t(s)$ to be concave in s_i for fixed $s_j, j = 1, 2, j \neq i$. Then the largest maximizer $\psi_{t+1}^i(s)$ of the function

$$a_i \mapsto \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i)[a_i - \Delta_{s_i} V_t(s)] \quad (6.2)$$

is non-increasing in s_i .

Proof. The proof follows straightforward from Lemma 4.6.1. \square

Theorem 6.1.9 provides an upper bound of the optimal fare before the assignment has been finalized. In $t = T, T - 1, \dots, \bar{T} + 1$, for state $s \in S_t$ with $s_1, s_2 > 0, s_3 = 1$, from Theorem 6.1.4 follows that an action $a = (a_1, a_2, (0, 0)') \in A_t(s)$ is optimal. The same applies in state $s' = s + e(i) \in S_t$. Then, applying Theorem 6.1.1, the maximizing fare

in either state can be derived by finding the maximizing argument of (6.2), which is non-increasing in s_i on concave intervals of V_{t-1} .

Next, we show that after finalization of the assignment, i.e. in states $s \in S_t$ with $s_3 = 0$, the largest maximizing fare for a flight is non-increasing in the remaining capacity of that flight.

Theorem 6.1.10 (Monotonicity of the price actions after the final assignment).

For a fixed $t = T, T-1, \dots, 1$, $i = 1, 2$, and $s = (s_1, s_2, 0) \in S_t$, the largest fare $\psi_{t+1}^i(s)$ maximizing the function

$$a_i \mapsto \lambda_t p_t^i q_t^i(a_i)[a_i - \Delta V_{t-1}^i(s_i)]$$

is non-increasing in s_i .

Proof. The proof follows straightforward from Theorem 6.1.1 and Theorem 4.6.2. \square

Additionally, in $t = \bar{T}$, only actions $a = (a_1, a_2, a_3)$ with $a_3 \in \{(\frac{1}{0}), (\frac{0}{1})\}$ are admissible. On the intervals $s_1 \geq s_1^*(s_2)$ and $s_1 < s_1^*(s_2) - 1$ (respectively $s_2 \geq s_2^*(s_1)$ and $s_2 < s_2^*(s_1) - 1$), the optimal assignment does not change when the capacity is increased by the definition and monotonicity of s_1^* and s_2^* . Then, applying Theorem 6.1.1 and Theorem 6.1.10, the optimal price of each flight $a_i^* = \psi_t^i(s_i)$ is non-increasing in s_i , $i = 1, 2$, which is stated in the following corollary:

Corollary 6.1.11 (Monotonicity of the price actions at $t = \bar{T}$).

At $t = \bar{T}$, in $s = (s_1, s_2, 1) \in S_t$,

- (i) the largest optimal price $a_1^* = \psi_t^1(s)$ is non-increasing in s_1 on the intervals $[0, \dots, s_1^*(s_2) - 1]$ and $[s_1^*(s_2), \dots, C_1]$ for fixed s_2 .
- (ii) The largest optimal price $a_2^* = \psi_t^2(s)$ is non-increasing in s_2 on the intervals $[0, \dots, s_2^*(s_1) - 1]$ and $[s_2^*(s_1), \dots, C_1]$ for fixed s_1 .

When calculating an optimal pricing policy using backward induction, all states, i.e. all possible remaining capacities, need to be iterated. On the respective domains, the monotonicity structures can be exploited and used as an upper bound when increasing the remaining capacities in each iteration. Then we can additionally use the following lower bound of the optimal actions.

Theorem 6.1.12 (Lower bound of price action).

Assume $t = T-1, T-2, \dots, \bar{T}$ and $s = (s_1, s_2, 1) \in S_t$ with $s_1, s_2 > 0$. Let $\psi_{t+1}^i(s) \in A_{t+1}^i(s)$, $i \in \{1, 2\}$ be the largest price maximizing the function

$$a_i \rightarrow \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i)[a_i - \Delta_{s_i} V_t(s)] .$$

Then, $\min\{\Delta_{s_i} V_t(s), f_0^i\} \leq \psi_{t+1}^i(s) \leq f_0^i$.

Proof. Fix $s = (s_1, s_2, 1) \in S_t$ with $s_1, s_2 > 0$. Then $s \in S_{t+1}$, too. Let $a \in A_{t+1}^i(s) \setminus \{f_0^i\}$ such that $a \leq \Delta_{s_i} V_t(s)$. Then

$$\begin{aligned} & \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a) [a - \Delta_{s_i} V_t(s)] \\ & \leq 0 \\ & = \lambda_{t+1} p_{t+1}^i q_{t+1}^i(f_0^i) [f_0^i - \Delta_{s_i} V_t(s)] , \end{aligned}$$

which implies that a is not the largest maximizer $\psi_{t+1}^i(s)$.

Thus, $\psi_{t+1}^i(s) \geq \min\{\Delta_{s_i} V_t(s), f_0^i\}$. On the other hand, if $a > \Delta_{s_i} V_t(s)$ holds for all $a \in A_{t+1}^i(s)$, the result is trivial. \square

During backward induction, when calculation optimal prices at $t+1$, $\Delta_{s_i} V_t(s)$ is known and can be used as a lower bound reducing the number of actions that need to be considered.

We use the structures found to reduce the complexity and the computational effort to find an optimal policy. A pseudo-code implementation that exploits these structures is presented in Algorithm B.1.1. It is more efficient than the general backward induction that iterates and evaluates every feasible action in each state at every point in time.

6.2 Extension to Sequences

We now extend the results from section 6.1 to consider more legs. While the number of legs considered is arbitrary, we still focus on two types of airplanes with different capacities C_1 and C_2 . We again assume, w.l.o.g., $C_2 > C_1$ and define $d := C_2 - C_1$. The legs considered are divided into two disjunct sets M and N . All legs in one set will be assigned the same aircraft. A swap of equipment now involves several legs, i.e. swapping the larger aircraft from legs in set M to N or vice versa. Let $m := |M|$ and $n := |N|$. We assume the sets to contain the flight indices in increasing order, i.e. $M = \{1, \dots, m\}, N = \{m+1, \dots, m+n\}$. The number of legs in each set is independent of the number of flights in the other set. The costs of assigning the larger capacity to set M are given by k_d^1 , the costs for set N by k_d^2 . As before, we consider only the additional costs compared to the most cost efficient allocation. Hence, $k_d^i = 0$ for at least one $i = 1, 2$. The final assignment has again to be decided the latest at $\bar{T} > 0$.

The model in chapter 6.1 is suitable for but limited to legs between two stations, where an initial change and a swap back can take place, e.g. at two base airports. The extended model can be applied to any two or more legs that are to be considered under DDR regardless of their origins and destinations.

For one or more legs to be considered, two sequences are created that allow a swap such that other legs in the network remain unaffected. The algorithm provided in section 5.4 might be applied. Usually, the two sequences have their origin and destination in common. The start and end of the sequences might even be the same airport. One example is two sequences, each one containing two legs to a destination and back to the base (illustrated in Figure 6.3(a)). Another example is depicted in Figure 6.3(b), where one set contains only one leg connecting two hubs. The second sequence also starts and ends at these hubs, but consists of two legs connecting a third airport.




Figure 6.3: Example of two flight sequences that are eligible for a swap with the same origin and destination.

However, the model is not even limited to sets with connected sequences of legs. We only require that all legs in one set share an equal assignment while other flights of the network remain unaffected by an equipment change. More complex swapping opportunities as depicted in Figure 5.5(b) might be considered, where legs in one set are not linked to each other.

We adhere to the same assumptions and notation as in the last section. Where needed the notation is adapted to the multi-leg case. We define a Markov Decision Process $(T, S, A, \mathbf{p}, r, V_0)$ as follows:

- (i) The planning horizon is $T \in \mathbb{N}$. The time horizon is again divided into time epochs in which at most one customer for any leg arrives. Of course, the larger the number of legs, the smaller each time epoch becomes. The time periods need not be of equal length. Time is counted backwards and $t = T, T - 1, \dots, 0$ denotes the number of periods remaining for ticket sales.
- (ii) The joint state space is

$$\begin{aligned}
 S = & \{0, \dots, C_1\}^m \times \{0, \dots, C_1\}^n \times \{1\} \\
 & \cup \{0, \dots, C_2\}^m \times \{0, \dots, C_1\}^n \times \{0\} \\
 & \cup \{0, \dots, C_1\}^m \times \{0, \dots, C_2\}^n \times \{0\} .
 \end{aligned}$$

For $(s_1, \dots, s_{m+n+1}) \in S$, s_i denotes the (remaining) capacity available on flight $i \in M \cup N$. If the option to assign the larger aircraft is still available $s_{m+n+1} = 1$

and $s_{m+n+1} = 0$, if it has been executed.

The assignment needs to be finalized the latest at $t = \bar{T}$. Hence, $S_t = S$ for $t = T, T-1, \dots, \bar{T}$ and

$$S_t = S \setminus \{0, \dots, C_1\}^m \times \{0, \dots, C_1\}^n \times \{1\}$$

for $t = \bar{T} - 1, \bar{T} - 2, \dots, 0$.

- (iii) The action space is $A = \{(a_1, a_2, \dots, a_m, a_{m+1}, \dots, a_{m+n}, a_{m+n+1}) : a_k \in \mathcal{P}_k : k \in \{1, \dots, m+n\}, a_{m+n+1} \in \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}), (\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}), (\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})\}\}$, where $\mathcal{P}_i = \{f_0^i, f_1^i, \dots, f_{k_i}^i\}$, $k_i \in \mathbb{N}$, $i \in M \cup N$, is the set of allowable prices for flight i including the nullprice f_0^i . The decision to execute the option on set M is $a_{m+n+1}(1) = 1$, i.e. all legs $i \in M$ are assigned the larger aircraft. $a_{m+n+1}(2) = 1$ is the decision to execute the option on legs $i \in N$, while $a_{m+n+1} = (0, 0)'$ denotes the decision not to execute the option.

For $s \in S, a_3 \in \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}), (\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}), (\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})\}$ we define

$$\bar{\mathcal{P}}_k(s, a_3) := \begin{cases} \{f_0^k\} & s_k = 0, k \in M, a_3(1) = 0, \\ \{f_0^k\} & s_k = 0, k \in N, a_3(2) = 0, \\ \mathcal{P}_k & \text{otherwise,} \end{cases}$$

to conveniently define the set of feasible actions $A_t(s) \in A$ in state $s \in S_t$:

For $t = T, T-1, \dots, 1$ and in state $s \in S_t$ with $s_{m+n+1} = 0$ the feasible actions are given by

$$A_t(s) = \bar{\mathcal{P}}_1(s, (\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})) \times \bar{\mathcal{P}}_2(s, (\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})) \times \dots \times \bar{\mathcal{P}}_{m+n}(s, (\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})) \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})\}.$$

For $t = T, T-1, \dots, \bar{T} + 1$ in state $s \in S_t$ with $s_{m+n+1} = 1$, we get

$$\begin{aligned} A_t(s) = & \bar{\mathcal{P}}_1(s, (\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})) \times \bar{\mathcal{P}}_2(s, (\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})) \times \dots \times \bar{\mathcal{P}}_{m+n}(s, (\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})) \times \{(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix})\} \\ & \cup \bar{\mathcal{P}}_1(s, (\begin{smallmatrix} 1 \\ 0 \end{smallmatrix})) \times \bar{\mathcal{P}}_2(s, (\begin{smallmatrix} 1 \\ 0 \end{smallmatrix})) \times \dots \times \bar{\mathcal{P}}_{m+n}(s, (\begin{smallmatrix} 1 \\ 0 \end{smallmatrix})) \times \{(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix})\} \\ & \cup \bar{\mathcal{P}}_1(s, (\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})) \times \bar{\mathcal{P}}_2(s, (\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})) \times \dots \times \bar{\mathcal{P}}_{m+n}(s, (\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})) \times \{(\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})\}. \end{aligned}$$

Finally, in $t = \bar{T}$ and for $s_{m+n+1} = 1$, we have to make a decision on the assignment. Hence, the feasible actions are

$$\begin{aligned} A_t(s) = & \bar{\mathcal{P}}_1(s, (\begin{smallmatrix} 1 \\ 0 \end{smallmatrix})) \times \bar{\mathcal{P}}_2(s, (\begin{smallmatrix} 1 \\ 0 \end{smallmatrix})) \times \dots \times \bar{\mathcal{P}}_{m+n}(s, (\begin{smallmatrix} 1 \\ 0 \end{smallmatrix})) \times \{(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix})\} \\ & \cup \bar{\mathcal{P}}_1(s, (\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})) \times \bar{\mathcal{P}}_2(s, (\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})) \times \dots \times \bar{\mathcal{P}}_{m+n}(s, (\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})) \times \{(\begin{smallmatrix} 0 \\ 1 \end{smallmatrix})\}. \end{aligned}$$

(iv) The transition law in $t = T, T - 1, \dots, 1$ for $s \in S_t, s' \in S_{t-1}, a \in A_t(s)$ is

$$\mathbf{p}_t(s, a, s') = \begin{cases} \lambda_t \cdot p_t^i \cdot q_t^i(a_i) & s'_k = s_k - \delta_k(i) + d \cdot a_{m+n+1}(1), k \in M, \\ & s'_j = s_j - \delta_j(i) + d \cdot a_{m+n+1}(2), j \in N, \\ & s'_{m+n+1} = s_{m+n+1} - a_{m+n+1}(1) - a_{m+n+1}(2), \\ & i \in M \cup N, \\ 1 - \sum_{i \in M \cup N} \lambda_t \cdot p_t^i \cdot q_t^i(a_i) & s'_k = s_k + d \cdot a_{m+n+1}(1), k \in M, \\ & s'_j = s_j + d \cdot a_{m+n+1}(2), j \in N, \\ & s'_{m+n+1} = s_{m+n+1} - a_{m+n+1}(1) - a_{m+n+1}(2), \\ 0 & \text{otherwise.} \end{cases}$$

As before, λ_t denotes the combined probability for an arrival in period t for any flight. The thinning probability for flight $i \in M \cup N$ is given by p_t^i and $q_t^i(a_i)$ denotes the related purchase probability given price $a_i \geq 0$. $\delta_i(j)$ denotes the indicator function.

(v) The reward function for $t = T, T - 1, \dots, 1$ is

$$r_t(s, a) = \sum_{i \in M \cup N} \lambda_t \cdot p_t^i \cdot q_t^i(a_i) \cdot a_i - \sum_{j=1}^2 k_d^j \cdot a_{m+n+1}(j) \quad , s \in S_t, a \in A_t(s) \quad ,$$

where $k_d^i, i = 1, 2$ denote the additional costs of flying the larger airplane on set M or N , respectively. As we only consider the additional costs in regards to the most efficient allocation, for at least one $i \in \{1, 2\}$, $k_d^i = 0$.

(vi) The terminal reward function is $V_0 \equiv 0$. We do not consider overbooking and do not incur any terminal cost.

Next, we adjust the operators and sets to the multi-leg model. For all $t = T, T - 1, \dots, 1$, $s = (s_1, \dots, s_m, s_{m+1}, \dots, s_{m+n+1}) \in S_t$, $a = (a_1, \dots, a_m, a_{m+1}, \dots, a_{m+n+1}) \in A_t(s)$, and $i \in M \cup N$, let

$$Y(s, a) := (s_1 + a_{m+n+1}(1)d, \dots, s_m + a_{m+n+1}(1)d, s_{m+1} + a_{m+n+1}(2)d, \dots, \\ s_{m+n} + a_{m+n+1}(2)d, s_{m+n+1} - a_{m+n+1}(1) - a_{m+n+1}(2)) \quad ,$$

$$Y^i(s, a) := \begin{cases} Y((s_1 - \delta_1(i), \dots, s_m - \delta_m(i), s_{m+1} - \delta_{m+1}(i), \\ \dots, s_{m+n} - \delta_{m+n}(i), s_{m+n+1}), a) & s_i + a_3(1)d > 0, \\ Y((s_1 - \delta_1(i), \dots, s_m - \delta_m(i), s_{m+1} - \delta_{m+1}(i), \\ \dots, s_{m+n} - \delta_{m+n}(i), s_{m+n+1}), a) & s_i + a_3(2)d > 0, \\ Y(s, a) & \text{otherwise,} \end{cases}$$

$$A_t(s|a_1, \dots, a_{m+n}) := \{a_{m+n+1} \in \left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\} : (a_1, \dots, a_{m+n}, a_{m+n+1}) \in A_t(s) \} \quad ,$$

$A_t(s|a_{m+n+1}) := \{(a_1, \dots, a_{m+n}) \in \mathcal{P}_1 \times \dots \times \mathcal{P}_{m+n} : (a_1, \dots, a_{m+n}, a_{m+n+1}) \in A_t(s)\}$,
 $A_t^i(s) := \{a' \in \mathcal{P}_i : (a_1, \dots, a_{i-1}, a', a_{i+1}, \dots, a_{m+n}, (0, 0)') \in A_t(s), a_j \in \mathcal{P}_j, j \in M \cup N, j \neq i\}$,
 and, for $s_i > 0$,

$$\Delta_{s_i} V_t(s) := V_t(s) - V_t(s - e(i)) .$$

With the adapted sets and operators we further get

$$\mathfrak{D}_t := \{(s, a) : s \in S_t, a \in A_t(s)\} ,$$

and, for all functions $v : S_{t-1} \rightarrow \mathbb{R}$, with $a \in A_t(s)$, let

$$Q_t v(s, a) := \sum_{i \in M \cup N} \lambda_t p_t^i q_t^i(a_i) [a_i + v(Y^i(s, a))] \\ + \left(1 - \sum_{i \in M \cup N} \lambda_t p_t^i q_t^i(a_i) \right) v(Y(s, a)) - \sum_{i=1}^2 k_d^i a_{m+n+1}(i) ,$$

$$L_t v(s|a_1, \dots, a_{m+n}) := \max_{a' \in A_t(s|a_1, \dots, a_{m+n})} \left\{ Q_t v(s, (a_1, \dots, a_{m+n}, a')) \right\} ,$$

$$L_t v(s|a_{m+n+1}) := \max_{a' \in A_t(s|a_{m+n+1})} \left\{ Q_t v(s, (a', a_{m+n+1})) \right\} .$$

With these definitions, the optimality equation can simply be written as

$$V_t(s) = \max_{a \in A_t(s)} \left\{ Q_t V_{t-1}(s, a) \right\} . \tag{6.3}$$

We now establish structural results for the multi-leg case. Part of the results in section 6.1 can be extended to the multi-leg case straightforward, while others need to include more aspects. To facilitate a compact presentation of the results, we state those proofs in Appendix A, in which the arguments follow the same ideas as in the preceding section.

Theorem 6.2.1 (Decomposition of the value function).

For the MDP with the value function defined in (6.3),

(i) *at time $t = T, \dots, 1$, in all states $s = (s_1, \dots, s_{m+n}, 0) \in S_t$, the process decomposes into $m + n$ independent single-leg processes with the value functions*

$$V_t^i(s_i) = \max_{a_i \in \mathcal{P}_i} \left\{ \lambda_t p_t^i q_t^i(a_i) [a_i + V_{t-1}^i(s_i - 1)] + [1 - \lambda_t p_t^i q_t^i(a_i)] V_{t-1}^i(s_i) \right\}$$

for $s_i > 0$ and $i \in M \cup N$. The boundary conditions are $V_t^i(0) = V_{t-1}^i(0)$ and $V_0^i \equiv 0$. The value function of the original process is the sum of the single-leg value functions:

$$V_t(s) = \sum_{i \in M \cup N} V_t^i(s_i) .$$

(ii) Further, for $t = T, \dots, \bar{T}$, $s = (s_1, \dots, s_{m+n}, 1) \in S_t$, and $a_{m+n+1} \in \left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}$,

$$\begin{aligned} & L_t V_{t-1}(s | a_{m+n+1}) \\ &= \sum_{j \in M} V_t^j(s_j + a_{m+n+1}(1)d) + \sum_{l \in N} V_t^l(s_l + a_{m+n+1}(2)d) - \sum_{i=1}^2 k_d^i a_{m+n+1}(i). \end{aligned}$$

(iii) For $t = T, \dots, \bar{T} + 1$, $s = (s_1, \dots, s_{m+n}, 1) \in S_t$,

$$\begin{aligned} & L_{t-1} V_t(s | (0, 0)') \\ &= \sum_{\substack{i \in M \cup N \\ s_i \neq 0}} \left(\max_{a_i \in A_t^i(s)} \lambda_t p_t^i q_t^i(a_i) [a_i - \Delta_{s_i} V_{t-1}(s)] \right) + V_{t-1}(s). \end{aligned}$$

Theorem 6.2.1 results from conveying Theorem 6.1.1 to the case with multiple legs. The effect of reducing computational requirements to find an optimal policy is even greater with the increased number of legs considered. Applying (i), the single-leg processes can be used in case the capacity option has been executed, i.e. in states with $s_{m+n+1} = 0$. Similarly, if the optimal assignment action can be predetermined to be $a_{m+n+1}^* \in \{(0, 1)', (1, 0)'\}$ before the maximization, the single-leg processes might be applied using (ii). (iii) will be important later, when we show that in case of remaining capacity $s_i > 0$ for every flight $i \in M \cup N$, an action with $a_{m+n+1} = (0, 0)'$ is always optimal. Then, to determine the optimal prices, instead of maximizing over the complete multidimensional set, we can separate the maximizations.

Theorem 6.2.2 (Monotonicity of the value function).

For the MDP with the value function defined in (6.3),

(i) for fixed $t = T, T - 1, \dots, 0$, $V_t(s)$ is non-decreasing componentwise in s_i for all $i \in M \cup N$.

(ii) With $s = (s_1, \dots, s_{m+n}, 1)$, $s' = (s_1, \dots, s_{m+n}, 0) \in S_t$, $V_t(s) - V_t(s') \geq 0$, for all $t = T, T - 1, \dots, \bar{T}$.

(iii) For fixed $s \in S_t \cap S_{t+1}$, $V_t(s)$ is non-decreasing in t .

(i) and (iii) show that the intuitive results hold that the expected revenue increases the more seats are available for sale and the more time remains to sell them. Theorem 6.2.2 (ii) states that the option price is not negative. The option might well be worthless though. In case we do not expect a demand exceeding the minimal capacity C_1 for any flight, the extra seats do not have a positive value. Then, the more cost efficient fleet assignment is used with no additional assignment costs ($k_d^1 = 0$ or $k_d^2 = 0$, respectively).

Lemma 6.2.3.

For the MDP with the value function defined in (6.3), for $t = T, \dots, \bar{T}$,

- (i) $V_t(s_1, \dots, s_{m+n}, 1) \geq V_t(s_1 + d, \dots, s_m + d, s_{m+1}, \dots, s_{m+n}, 0) - k_d^1$.
- (ii) $V_t(s_1, \dots, s_{m+n}, 1) \geq V_t(s_1, \dots, s_m, s_{m+1} + d, \dots, s_{m+n} + d, 0) - k_d^2$.

Using Lemma 6.2.3, we next develop conditions, when we can omit certain actions when maximizing over the set of feasible actions.

Theorem 6.2.4 (Optimal assignment action).

For the MDP with the value function defined in (6.3), for any $t = T, \dots, \bar{T} + 1$, in state $s \in S_t$, an action $a^* \in A_t(s)$ is optimal with

- (i) $a_{m+n+1}^* \in \left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}$, if $s_i > 0$ for all $i \in M$.
- (ii) $a_{m+n+1}^* \in \left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right\}$, if $s_i > 0$ for all $i \in N$.
- (iii) $a_{m+n+1}^* = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$, if $s_i > 0$ for all $i \in M \cup N$.
- (iv) $a_{m+n+1}^* \in \left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}$, if $s_i = 0$ for all $i \in M \cup N$ and $s_{m+n+1} = 1$.

In case of remaining capacities $s_i > 0$ for every flight $i \in M \cup N$, we can apply (iii) in conjunction with Theorem 6.2.1 (iii) to determine the optimal prices separately for the flights instead of maximizing over the complete multidimensional set. When no capacity remains on any leg, an assignment is always optimal applying (iv). (i) and (ii) state conditions, when one assignment action can be excluded from the maximization because it is at most as good as the other two possibilities.

In contrast to the two-leg model, in the multi-leg case, we cannot always reduce the set of assignment actions to evaluate during backward induction. In states with at least one $s_i = 0, i \in M$, at least one $s_j = 0, j \in N$, and at least one other $s_{i'} > 0, i' \in M$ and at least one $s_{j'} > 0, j' \in N$, all three assignment actions need to be evaluated in the maximization.

Theorem 6.2.5 (Monotonicity of the assignment action).

In $s \in S_t, t = T, T - 1, \dots, \bar{T}$ with $s_{m+n+1} = 1$,

$$L_t V_{t-1}(s|(1, 0)') - L_t V_{t-1}(s|(0, 1)')$$

is componentwise non-increasing in s_i for all $i \in M$ and non-decreasing in s_j for all $j \in N$.

Proof. For some $i \in M$, fix $s, s' := s + e(i) \in S_t, t = T, T - 1, \dots, \bar{T}$ with $s_{m+n+1} = 1$. Using Theorem 6.2.1 (ii)

$$\begin{aligned} & L_t V_{t-1}(s'|(1, 0')) - L_t V_{t-1}(s'|(0, 1')) \\ &= L_t V_{t-1}(s|(1, 0')) - L_t V_{t-1}(s|(0, 1')) \\ & \quad + V_t^i(s_i + 1 + d) - V_t^i(s_i + d) - V_t^i(s_i + 1) + V_t^i(s_i) . \end{aligned}$$

Now, using the concavity of V_t^i (Theorem 4.6.2), the monotonicity in s_i follows. A similar argumentation works for $s_j, j \in N$. \square

In the two-leg case, although the results in Theorem 6.1.5 hold for any $t = T, T - 1, \dots, \bar{T}$, we limited the application to $t = \bar{T}$. At other times, applying Theorem 6.1.4 yields that an action with $a_3 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ ($a_3 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$) needs not be considered in states $s = (s_1 + 1, s_2, 1)$ (respectively $s = (s_1, s_2 + 1, 1)$).

Contrarily, in the multi-leg model, in states with at least one $s_i = 0, i \in M$, at least one $s_j = 0, j \in N$, and at least one other $s_{i'} > 0, i' \in M$ and at least one $s_{j'} > 0, j' \in N$, all three assignment actions need to be considered during backward induction. Then, the monotonicity of the assignment action stated in Theorem 6.2.5 can be exploited to reduce the set of assignment actions that need to be evaluated.

We continue providing upper and lower bounds on the pricing actions, which can be extended straightforward from the two-leg case.

Theorem 6.2.6 (Lower bounds of the price actions).

Assume $t = T - 1, T - 2, \dots, \bar{T}$ and $s \in S_t$ with $s_i > 0$ for all $i \in M \cup N$ and $s_{m+n+1} = 1$. Let $\psi_{t+1}^i(s) \in A_{t+1}^i(s)$ be the largest price maximizing the function

$$a_i \mapsto \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i)[a_i - \Delta_{s_i} V_t(s)] .$$

Then, $\min\{\Delta_{s_i} V_t(s), f_0^i\} \leq \psi_{t+1}^i(s) \leq f_0^i$.

Theorem 6.2.6 together with Theorem 6.2.1 (iii) and 6.2.4 (iii) provides lower bounds on the pricing actions, if $s_i > 0$ for all $i \in M \cup N$. Note that $\Delta_{s_i} V_t(s)$ is known in $t + 1$ and can be exploited during backward induction as a lower bound.

Next, we show that after the finalization of the assignment, i.e. in states $s \in S_t$ with $s_{m+n+1} = 0$, the largest maximizing fare for a flight is non-increasing in the remaining capacity of that flight.

Theorem 6.2.7 (Monotonicity of the price actions after the final assignment).

For a fixed $t = T, T - 1, \dots, 1$, $i \in M \cup N$, and $s \in S_t$ with $s_{m+n+1} = 0$, the largest fare $\psi_{t+1}^i(s)$ maximizing the function

$$a_i \mapsto \lambda_t p_t^i q_t^i(a_i)[a_i - \Delta V_{t-1}^i(s_i)]$$

is non-increasing in s_i .

To show the monotonicity of the prices when the option is still available ($s_{m+n+1} = 1$), analogously to the two-leg case, for $i \in M$ and $j \in N$, we define

$$\begin{aligned} s_i^*(s) &:= \min \left\{ \bar{s}_i \in \{0, \dots, C_1\} : V_{\bar{T}}(s_1, \dots, s_{i-1}, \bar{s}_i, s_{i+1}, \dots, s_{m+n}, 1) \right. \\ &\quad \left. = L_{\bar{T}} V_{\bar{T}-1}((s_1, \dots, s_{i-1}, \bar{s}_i, s_{i+1}, \dots, s_{m+n}, 1)|(0, 1)') \right\}, \\ s_j^*(s) &:= \min \left\{ \bar{s}_j \in \{0, \dots, C_1\} : V_{\bar{T}}(s_1, \dots, s_{j-1}, \bar{s}_j, s_{j+1}, \dots, s_{m+n}, 1) \right. \\ &\quad \left. = L_{\bar{T}} V_{\bar{T}-1}((s_1, \dots, s_{j-1}, \bar{s}_j, s_{j+1}, \dots, s_{m+n}, 1)|(1, 0)') \right\}, \end{aligned}$$

where $\min \emptyset := \infty$.

The monotonicity results for $s_i^*, i \in M \cup N$ need to be extended from the two-leg case to include capacities of flights in the same set.

Lemma 6.2.8.

For the MDP with the value function defined in (6.3),

- (i) $s_i^*, i \in M$ is non-decreasing in $s_j, j \in N$.
- (ii) $s_j^*, j \in N$ is non-decreasing in $s_i, i \in M$.
- (iii) $s_i^*, i \in M$ is non-increasing in $s_j, j \in M, j \neq i$.
- (iv) $s_j^*, j \in N$ is non-increasing in $s_i, i \in N, i \neq j$.

Proof.

- (i) For some $i \in M, j \in N$, fix $s \in S_{\bar{T}}$ with $s_i < s_i^*(s)$ and $s_{m+n+1} = 1$. Then an action a^* with $a_{m+n+1}^* = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ is optimal. In $s' = s + e(j)$, an action a^{**} is optimal with $a_{m+n+1}^{**} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ by applying Theorem 6.2.5.
- (ii) can be shown analogously.
- (iii) For some $i, j \in M, i \neq j$, fix $s \in S_{\bar{T}}$ with $s_i \geq s_i^*(s)$ and $s_{m+n+1} = 1$. Then an action a^* with $a_{m+n+1}^* = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ is optimal. In $s' = s + e(j)$, an action a^{**} is optimal with $a_{m+n+1}^{**} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ by applying Theorem 6.2.5.
- (iv) can be shown analogously. □

Using the monotonicity properties in Lemma 6.2.8 in conjunction with Theorem 6.2.1 and Theorem 6.2.7, we can additionally derive monotonicity structures of the pricing

actions at $t = \bar{T}$. In state $s \in S_{\bar{T}}$ with $s_{m+n+1} = 1$, only actions $a \in A_{\bar{T}}(s)$ are feasible with $a_{m+n+1} \in \left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}$. For some $i \in M$, for $s_i \geq s_i^*(s)$ the optimal assignment does not change when capacity $s_{i'}, i' \in M$ increases or $s_j, j \in N$ decreases. For $s_i < s_i^*(s)$ the optimal assignment does not change when capacity $s_{i'}, i' \in M$ decreases or $s_j, j \in N$ increases. Then, on these domains, the largest optimal price of each flight $a_i^* = \psi_{\bar{T}}^i(s_i)$ is independent of the capacities on other flights and is non-increasing in s_i . Analogical properties hold for flights $j \in N$.

Corollary 6.2.9 (Monotonicity of the price actions at $t = \bar{T}$).

For $s, s' \in S_{\bar{T}}, s_{m+n+1} = s'_{m+n+1} = 1$,

- (i) for some $i \in M$, if $s'_i \geq s_i \geq s_i^*(s)$, $s'_k \geq s_k$ for $k \in M, k \neq i$, and $s'_j \leq s_j$ for $j \in N$, the largest optimal prices

$$\psi_{\bar{T}}^i(s_i) \geq \psi_{\bar{T}}^i(s'_i) .$$

If further $s'_i = s_i$, then $\psi_{\bar{T}}^i(s_i) = \psi_{\bar{T}}^i(s'_i)$.

- (ii) For some $i \in M$, if $s'_i \leq s_i < s_i^*(s)$, $s'_k \leq s_k$ for $k \in M, k \neq i$, and $s'_j \geq s_j$ for $j \in N$, the largest optimal prices

$$\psi_{\bar{T}}^i(s_i) \leq \psi_{\bar{T}}^i(s'_i) .$$

If further $s'_i = s_i$, then $\psi_{\bar{T}}^i(s_i) = \psi_{\bar{T}}^i(s'_i)$.

- (iii) For some $j \in N$, if $s'_j \geq s_j \geq s_j^*(s)$, $s'_k \geq s_k$ for $k \in N, k \neq j$, and $s'_i \leq s_i$ for $i \in M$, the largest optimal prices

$$\psi_{\bar{T}}^j(s_j) \geq \psi_{\bar{T}}^j(s'_j) .$$

If further $s'_j = s_j$, then $\psi_{\bar{T}}^j(s_j) = \psi_{\bar{T}}^j(s'_j)$.

- (iv) For some $j \in N$, if $s'_j \leq s_j < s_j^*(s)$, $s'_k \leq s_k$ for $k \in N, k \neq j$, and $s'_i \geq s_i$ for $i \in M$, the largest optimal prices

$$\psi_{\bar{T}}^j(s_j) \leq \psi_{\bar{T}}^j(s'_j) .$$

If further $s'_j = s_j$, then $\psi_{\bar{T}}^j(s_j) = \psi_{\bar{T}}^j(s'_j)$.

Before an assignment, on domains where V_t is concave in one component, we can exploit the monotonicity of the price actions to find upper bounds for the prices at $t + 1$.

Theorem 6.2.10 (Monotonicity of the price actions before the final assignment).

For $i \in M \cup N, t = T - 1, T - 2, \dots, \bar{T}$, and $s \in S_t$ with $s_i > 0$ and $s_{m+n+1} = 1$,

suppose $V_t(s)$ to be concave in s_i for fixed $s_j > 0, j \in M \cup N, j \neq i$. Then, the largest maximizer $\psi_{t+1}^i(s)$ of the function

$$a_i \mapsto \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i)[a_i - \Delta_{s_i} V_t(s)]$$

is non-increasing in s_i .

However, the concavity does not hold in general on certain intervals as in the two-leg model. Hence, we need to check for concavity in the actual implementation to make use of Theorem 6.2.10.

6.3 Heuristic Approaches

A major drawback of the model formulation in the previous section is the curse of dimensionality. The sizes of the state and action spaces increase rapidly with the number of legs considered. Hence, when more than a few legs are involved in a possible swap and need to be considered simultaneously, computational power might limit the ability to apply the model developed in practice. In this section, we provide heuristics that efficiently compute approximately optimal pricing policies. We start by adapting a common capacity control heuristic to a dynamic pricing setting. The general idea of a limited lookahead policy is introduced in section 6.3.2. We then outline a deterministic approximation commonly used in dynamic pricing problems and extend it to include swapping opportunities. Using the developed deterministic approximation, we present a heuristic that approximates the opportunity cost and applies the concept of a limited lookahead policy.

6.3.1 Period Re-Evaluation of the Fleet Assignment

The most evident approach for a heuristic is to convey the ideas by Berge and Hopferstad (1993) as described in section 5.3 to a dynamic pricing setting. Hence, the single-leg processes developed in section 4.6 are used to find a pricing policy with a given initial fleet assignment. The assignment is re-evaluated at certain planning points during the booking horizon. An adaption of the basic MIP 5.2.1 is used for the optimization. Received bookings at the time constrain feasible solutions of the assignment problem. Other restrictions such as only considering crew-compatible fleets might be imposed as well. Prices are then set by the single-leg processes using the updated capacities. This process is repeated until the assignment needs to be finalized. After the finalization, the single-leg processes are applied to derive a pricing policy for the remaining time until departure. Figure 6.4 depicts the process.




Figure 6.4: Process of Demand-Driven Re-Fleeting Heuristic 1.

Note that simulation studies show that the magnitude of the revenue gain by applying the heuristic is greatly dependent on the time between two planning points when the fleet assignment is re-evaluated. Figure 6.5 depicts the influence for the base-case scenario of the simulation studies in the next section. The full gain can only be exploited when the assignment is revised frequently. Hence, in the remainder, we assume that at the start of every time period, regardless of an arrival of a customer, the fleet assignment is optimized and prices are then set using the updated assignment. The frequent re-optimization might be computational demanding depending on the size of the network considered and the restrictions applied.

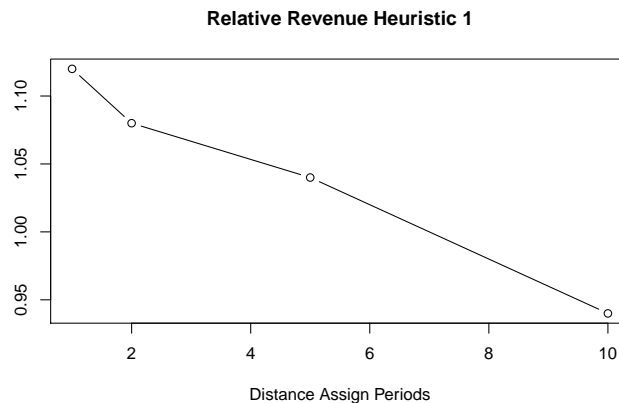


Figure 6.5: Influence of the time distance between re-optimization points. The relative revenue gain by applying Heuristic 1 over a policy without DDR decreases with the distance between two planning points.

The heuristic has been reported to be implemented at several airlines in capacity control settings (e.g. Oba, 2007, Zhao et al., 2007). The applied controls and optimization techniques differ while the general process remains unaffected. Hence, the heuristic can be regarded as industry practice today. However, at many airlines, to varying degrees, the process incorporates manual interaction and must be regarded as decision support rather than an automated procedure to set prices.

6.3.2 Limited Lookahead Policies

In dynamic programming, to reduce the computational effort and memory requirements of finding and storing an optimal policy, a stochastic program might be solved only for the current and possibly several subsequent steps. The value functions of further stages are then approximated by a suitable heuristic. A policy derived in such way is called *limited lookahead policy* (Bertsekas, 1995, chapter 6.3.2).

We speak of a *one-step lookahead policy*, when the value function of the next step is approximated. We do not fix a decision rule for each decision time at the start of the process, but rather choose an action a at the beginning of each time period t only for the current state $i_t \in I_t$. We choose a by replacing the value function V_{t-1} in (3.1) by its approximation \tilde{V}_{t-1} . Hence, an action a is chosen that attains the maximum of

$$\max_{a \in A_t(i_t)} \left\{ r_t(i_t, a) + \sum_{j \in \tilde{I}_{t-1}} \mathbf{p}_t(i_t, a, j) \tilde{V}_{t-1}(j) \right\}, \quad (6.4)$$

where $\tilde{I}_{t-1} := \{j \in I_{t-1} : \mathbf{p}_t(i_t, a, j) > 0, a \in A_t(i_t)\} \subseteq I_{t-1}$ is the subset of states in I_{t-1} that can be reached from i_t . \tilde{V}_{t-1} needs to be approximated only for states $j \in \tilde{I}_{t-1}$.

A *k-step lookahead policy* is derived similarly. Analogously to \tilde{I}_{t-1} , let \tilde{I}_{t-h} , $h = 2, \dots, k$ be the subset of states in I_{t-h} that can be reached from i_t in k steps. Define $\tilde{I}_t := \{i_t\}$. Now, V_{t-k} is approximated by \tilde{V}_{t-k} using a suitable heuristic for all $j \in \tilde{I}_{t-k}$. Then, for $\tau = t, t-1, \dots, t-k+1$, the approximated value function is the unique solution to

$$\tilde{V}_\tau(i) = \max_{a \in A_\tau(i)} \left\{ r_\tau(i, a) + \sum_{j \in \tilde{I}_{\tau-1}} \mathbf{p}_\tau(i, a, j) \tilde{V}_{\tau-1}(j) \right\}, \quad i \in \tilde{I}_\tau. \quad (6.5)$$

At each time $\tau = t, t-1, \dots, t-k+1$, we find an action a_τ maximizing the right-hand side of (6.5) by backward induction replacing (3.1) by (6.5) in Algorithm 3.3.1.

The quality of the policy is obviously greatly dependent on the heuristic used. We develop a suitable heuristic for considering swaps in dynamic pricing problems in the next section. Another influence on the quality is the number of stages approximated. If the approximation of the value function improves in time, the policy approaches the optimal policy the more steps are evaluated before the remaining periods are approximated.

To find a near-optimal policy, more important than the approximation of the actual value function is the quality of the resulting *relative values* (Bertsekas, 1995, p. 267). For a k -step lookahead policy, if for any $j, j' \in \tilde{I}_{t-k}, j \neq j'$, the approximation

$$\tilde{V}_{t-k}(j) - \tilde{V}_{t-k}(j') \approx V_{t-k}(j) - V_{t-k}(j')$$

is reasonably good, so is the resulting policy. This is quite intuitive considering the example of the single-leg model presented in section 4.6. Reformulating (4.7) yields

$$V_{t+1}(c) = \max_{a \in A(c)} \{ \lambda_t q_t(a) [a - \Delta V_t(c)] \} + V_t(c) .$$

If $\Delta V_t(c) = \Delta \tilde{V}_t(c)$, the one-step lookahead policy is actually optimal irrespective of how well the actual values $V_t(c)$ and $V_t(c-1)$ are approximated. We make use of this fact in the next section, where we actually aim to find a good approximation of the opportunity cost $\Delta V_t(c)$.

6.3.3 Dynamic Pricing Linear Program

To develop a heuristic to approximate the opportunity cost in a dynamic pricing problem allowing for swapping aircraft assignments, we start by developing a linear program for the single-leg dynamic pricing problem introduced in section 4.6. To avoid technical complications and facilitate a quick solution by standard optimization techniques, we make the following relaxations to the original problem:

- The capacity is assumed to be a continuous quantity, and accordingly, we also assume demand to be continuous.
- We follow an approach known as *deterministic linear programming* in the RM literature (c.f. e.g. Williamson, 1992). A linear program is used to approximate the value function where all random variables are replaced by their expectations. Hence, we use the expected demand in each period that depends on the offered fare. The approximation based on the deterministic linear programming model should not be mistaken for the solution of a value function using a linear program (c.f. Nickel et al., 2011, chapter 8.3.5).
- Further, we allow a convex combination of prices in each period instead of only a single price chosen from a discrete set. The convex combination can be interpreted as to use each price for a fraction of the period determined by the convex weights.

We adhere to the notation introduced in section 4.6. As decision variables we use the convex weights $\beta_\tau(f)$, $\tau = 1, \dots, t$ corresponding to each allowable fare $f \in \mathcal{P}$. To approximate the value function $\tilde{V}_t(c)$ at time t with remaining capacity c , i.e. state c , we solve the following linear program:

$$\begin{aligned} & \text{(LP 6.3.1)} \\ & \max \sum_{\tau=1}^t \sum_{f \in \mathcal{P}} \beta_\tau(f) \lambda_\tau q_\tau(f) f \end{aligned} \tag{6.6}$$

$$\sum_{\tau=1}^t \sum_{f \in \mathcal{P}} \beta_{\tau}(f) \lambda_{\tau} q_{\tau}(f) \leq c \quad (6.7)$$

$$\sum_{f \in \mathcal{P}} \beta_{\tau}(f) = 1 \quad \text{for all } \tau = 1, \dots, t \quad (6.8)$$

$$\beta_{\tau}(f) \geq 0 \quad \text{for all } \tau = 1, \dots, t, f \in \mathcal{P} \quad (6.9)$$

(6.6) maximizes the total expected revenue using a convex combination of the allowable prices in each time period. (6.7) guarantees that expected sales do not exceed the currently remaining capacity. Note that because $f_0 \in \mathcal{P}$ and $q_{\tau}(f_0) = 0$ for all $\tau = 1, \dots, t$ the constraint can always be fulfilled. (6.8) and (6.9) restrict the decision variables β_{τ} to a convex combination in each period.

By allowing a convex combination of the prices in each period the program is linear in the decision variables. If instead, we had a fare from the discrete set of prices \mathcal{P} as the decision variable in each period, the problem would neither be continuous nor necessarily convex. The linear program presented is similar to that developed in Talluri and Van Ryzin (2004b, chapter 5.2.1.3). However, they use a demand rate as the decision variable by applying the inverse demand function.

Relaxing the problem to continuous decision variables not only facilitates to solve the problem efficiently and quickly by standard means, but also we are able to use the shadow price of restriction (6.7) as an approximation of the opportunity cost $\Delta V_t(c)$ of the original problem in state c .

Similarly, we develop a linear program to approximate the opportunity cost of the dynamic pricing model allowing for Demand-Driven Re-Fleeting developed in section 6.2. We again apply a deterministic linear programming model and replace the random demands by their expectations. As before, we relax the problem to continuous capacities. Consistently, we also assume the additional capacity d to be continuous. We further assume that it can be split between the flight legs in set M and N , respectively. W.l.o.g., let $k_d^1 = 0$ and $k_d^2 \geq 0$. The cost k_d^2 of assigning the additional capacity to set N incurs only for the portion allocated to the set. We denote these seats x_N .

Now, for each leg $i \in M \cup N$, we use a convex combination $\beta_{\tau}^i(f)$ of the allowable prices $f \in \mathcal{P}_i$ as the decision variable in each time period $\tau = 1, \dots, t$. At decision time t in state $s = (s_1, s_2, \dots, s_{m+n}, 1) \in S_t$ we use the linear program

(LP 6.3.2)

$$\max \sum_{\tau=1}^t \sum_{i \in M \cup N} \sum_{f \in \mathcal{P}_i} \beta_{\tau}^i(f) \lambda^{\tau} p_i^{\tau} q_i^{\tau}(f) f - (k_d^2/d) x_N \quad (6.10)$$

$$\sum_{\tau=1}^t \sum_{f \in \mathcal{P}_i} \beta_{\tau}^i(f) \lambda^{\tau} p_i^{\tau} q_i^{\tau}(f) \leq s_i + d - x_N \quad \text{for all } i \in M \quad (6.11)$$

$$\sum_{\tau=1}^t \sum_{f \in \mathcal{P}_i} \beta_{\tau}^i(f) \lambda^{\tau} p_i^{\tau} q_i^{\tau}(f) \leq s_i + x_N \quad \text{for all } i \in N \quad (6.12)$$

$$\sum_{f \in \mathcal{P}_i} \beta_{\tau}^i(f) = 1 \quad \text{for all } i \in M \cup N, \tau = 1, \dots, t \quad (6.13)$$

$$\beta_{\tau}^i(f) \geq 0 \quad \text{for all } i \in M \cup N, f \in \mathcal{P}_i, \quad (6.14)$$

$$\tau = 1, \dots, t$$

$$0 \leq x_N \leq d \quad (6.15)$$

to approximate the value function $V_t(s)$.

The expected total revenue reduced by the cost of allocating x_N to all legs $i \in N$ is maximized in (6.10). Note that the objective does not reflect the expected profit earned as total assignment costs are not considered. Expected future sales are limited to the remaining capacity by (6.11) and (6.12). The current capacity is augmented by the additional (partial) capacity $d - x_N$ and x_N allocated to sets M and N , respectively. Again, the constraint can always be fulfilled because $f_0^i \in \mathcal{P}_i$ and $q_i^{\tau}(f_0^i) = 0$ for all $i \in M \cup N, \tau = 1, \dots, t$. (6.13) and (6.14) restrict the decision variables to a convex combination in each period, while (6.15) constrains the additional capacity to the maximum available additional capacity d .

Note that we aim to approximate the opportunity cost only in states $s \in S_t$ with $s_{m+n+1} = 1$, i.e. when the final assignment has not been decided. To approximate the opportunity cost $\Delta V_t(s)$ of the original problem, we use the shadow prices of (6.11) and (6.12) for the respective flights.

LP 6.3.2 is used in the next section in conjunction with a one-step lookahead policy to derive a dynamic pricing heuristic that allows for swapping assignments of flights. Allocation costs that are incurred when swapping assignments can be included.

6.3.4 Demand-Driven Re-Fleeting Dynamic Pricing Heuristic

We implement a one-step lookahead policy using LP 6.3.2. While the fleet assignment has not been finalized, at the start of each time period, we approximate the opportunity cost of each flight by the shadow prices of (6.11) and (6.12), respectively. We then maximize

$$a_i \rightarrow \lambda_t p_t^i q_t^i(a_i) [a_i - \Delta_{s_i} \tilde{V}_{t-1}(s)]$$

for each flight $i \in M \cup N$ separately according to Theorem 6.1.1 (iii). $\Delta_{s_i} \tilde{V}_{t-1}(s)$ denotes the approximated opportunity cost. At the latest time for the final assignment, or before, if received bookings for any flight reach the capacity of the smaller aircraft, the fleet assignment is finalized. After the finalization, the single-leg pricing policies for

each flight are applied with the updated capacities. Figure 6.6 illustrates the process. A pseudo-code implementation is presented in Algorithm B.2.1.

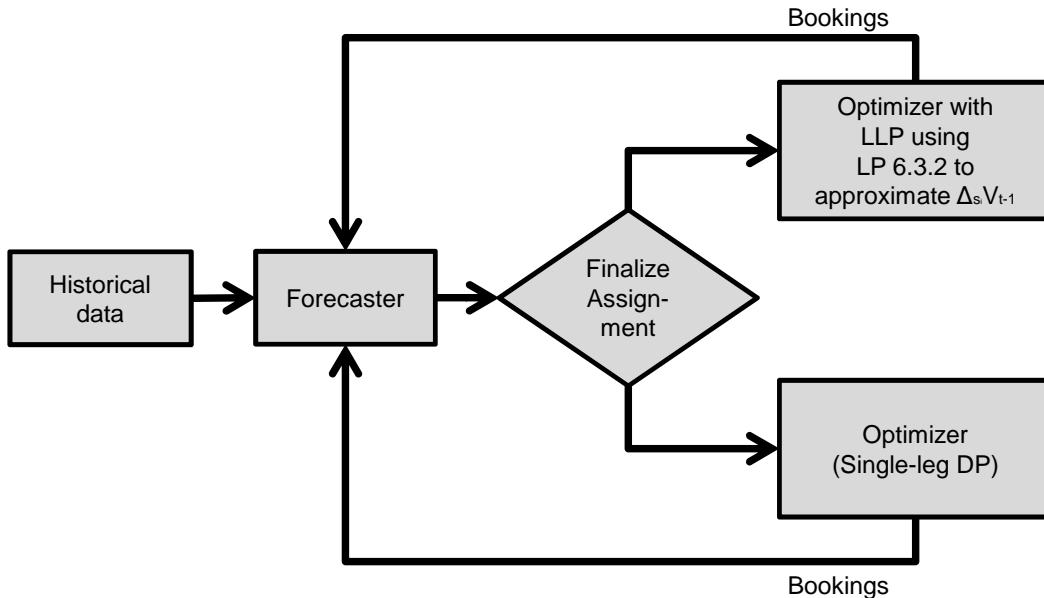


Figure 6.6: Process of Demand-Driven Re-Fleeting Heuristic 2.

6.4 Numerical Studies

In this section, we present simulation studies to analyze the performance of the pricing models and heuristics introduced in the preceding sections. Relevant parameters are varied in different settings to analyze their effects on the performance.

To our knowledge, no simulation study has been published applying Demand-Driven Re-Fleeting in a dynamic pricing environment. Similarities are expected to simulations conducted in capacity control settings, e.g. Berge and Hopperstad (1993), Frank et al. (2006), or Wang and Meng (2008). However, the mentioned publications all discuss heuristic strategies. The heuristics used follow a similar idea as the one developed in section 6.3.1, and hence, comparable results are expected.

We start by briefly describing the simulation environment in the next section. The general setup is outlined in section 6.4.2, before the various scenarios and the corresponding results are presented in detail in section 6.4.3. We summarize the results and conclude with section 6.4.4.

6.4.1 Simulation Environment

The simulation environment is implemented in Java and the simulations were run on a high performance computer cluster. A more realistic setting with larger capacities was sacrificed for a smaller example with shorter run times in order to facilitate a high number of individual simulations in each scenario. A comparison of a realistic large-scale problem and a down-scaled setting yielded comparable results. In each scenario 20,000 simulations were conducted to achieve significant results.

The simulation implements the simple dynamic pricing model without DDR as introduced in section 4.6 (labeled *No DCM* or *No DDR*), the DDR model developed in section 6.2 (labeled *Exact*), and the heuristics described in section 6.3. *Heuristic 1* describes the practice of repeatedly resolving the fleet assignment (see section 6.3.1), while the newly developed heuristic outlined in section 6.3.4 is labeled *Heuristic 2*. To permit a direct comparison, the arrival sequences for each run can either be generated or read from file. Where applicable, we use the same arrival sequences to evaluate different strategies.

The arrivals are generated from an inhomogeneous Bernoulli Process. While any function can be used as an intensity function for the arrival process, we focus on triangular shaped intensity functions. The parameters are described in the next section. The arrival process describes the arrivals combined for all flights. These arrivals are thinned into separate arrival processes for each flight, i.e. each arrival is attributed to a specific flight based on given thinning probabilities. Although the program also allows for dynamically changing thinning probabilities, constant probabilities in time are used for the presented scenarios.

The willingness-to-pay for each flight can be specified separately and any function $q^t : \{T, \dots, 1\} \times \mathcal{P} \rightarrow [0, 1]$ is applicable. For our studies we implement

$$q^t(p) = \min \left\{ 1, e^{\frac{-(p-\beta^t) \ln 2}{(\epsilon^t - 1)\beta^t}} \right\}, \quad (6.16)$$

where β^t is the base price, i.e. the highest price any arriving customer is willing to accept, and ϵ^t is a shape parameter at time t .

(6.16) is also used for the well-known Passenger O&D Simulator (PODS) (Belobaba, 2006). Belobaba (2006) states that $\epsilon = 1.2$ for leisure and $\epsilon = 3$ for business customers. Business and leisure demand can be accounted for by choosing an increasing base price β^t and shape parameter ϵ^t over time. However, we do not distinguish leisure and business demand and use a constant base price and shape parameter. Hence, all demand regardless of the arrival time is valued equally. Accounting for different values of the demand allows to arbitrarily influence the results. High-value demand arriving

after the final assignment results in a low gain by applying DDR. The gain increases when shifting high-value demand further away from departure. The influence of the percentage of the aggregated demand expected before the latest possible assignment, and hence its value in comparison to demand arriving later during the booking process, is analyzed by varying the latest time for the finalization of the assignment.

The flight legs are defined by a willingness-to-pay function and the set of allowed prices including the nullprice. They are grouped into disjunct sets M and N as described in section 6.2. All legs in one set are assigned the same capacity and are thus affected equally by a possible change of equipment. Assignment costs need to be defined for each set. We restrict our examples to two flights, each in one set. To apply DDR strategies, also the latest time for an assignment (LAT) needs to be defined. The assignment needs to be finalized at that time and no further changes are admissible after the LAT.

To get a differentiated picture of how the various parameters influence the value of DDR, we omit estimation errors from the simulations. The presented strategies are based on the true probabilities for arrivals and the acceptance of the offered price. The quality of the forecast greatly affects any efforts of applying RM. Hence, we try to separate the influences of forecasting and of applying DDR. We also follow a conservative approach and do not assume auto-correlation of the arrivals. Auto-correlated arrivals facilitate more precise forecasts the more arrivals have realized. As a result, only the randomness of the arrivals contributes to the value of applying DDR, i.e. the bookings-in-hand reduce the amount of demand that needs to be forecasted and the assignment is based on more realistic data.

Note also that the initial fleet assignment is based on the expectations of the dynamic pricing process. It is thus based on the same data as the pricing, which is usually not the case in practice (see section 4.4 for more details). Just recently, new IT solutions have been developed to help to make data available throughout different departments responsible for the different planning steps and thus to allow to base planning on the same data available at the time (Lufthansa Systems, 2010).

The availability and use of more accurate data in terms of the demand estimation and the granularity is one of the advantages of applying DDR. Thus, both, forecasting errors and a fleet assignment based on less precise data are expected to increase the performance relative to dynamic pricing strategies without DDR. The same holds for auto-correlated arrivals. In our simulations, intentionally, only the realized and thus certain demand contributes to the value of DDR. Including forecasting errors allows to arbitrarily influence results.

Many other scenarios than the ones described below were tested and analyzed. How-

ever, we restrict our outline to relevant scenarios to gain insights into the potential revenue improvement and the sensitivity to the parameters.

6.4.2 General Settings

In different scenarios various parameters are varied from a base-case setting. In each scenario we analyze two flight legs that qualify for an equipment swap. The smaller available capacity is 10 in every scenario. The larger capacity varies between 120% and 180% of the smaller capacity. The flights are intentionally set up equally in terms of the arrival process, willingness-to-pay, possible pricing actions, and capacity costs. For a change of equipment to be profitable, the valued demand for flights in one set has to offset the capacity costs for that set and the valued demand for flights in the other set that exceeds the smaller capacity. Thus, to not arbitrarily influence the outcome of the simulations, the inputs influencing the offset are kept equal for both sets. We do not include assignment costs assuming two legs of approximately the same cost structure.

The set of allowable prices is $\mathcal{P} = \{f_0, 260, 255, 250, \dots, 15, 10\}$ for each flight. The lowest price action is 10, the largest price allowed below the nullprice f_0 is 260. The set contains all prices in between with an increment of 5. The willingness-to-pay is set up according to (6.16). The base price is $\beta^t = 40$ for all time periods t . The influence of the price sensitivity of the customers arriving is analyzed by varying ϵ^t in different scenarios between 1.2 and 3.0. In each scenario, the shape parameter is held constant over time. The purchase probability of an arriving customer is independent of the probability for an arrival.

The intensity function of the arrival process is a scaled triangular function with an added constant in each scenario. The abscissae of the feet are given as a percentage of the length of the booking horizon. The feet are located outside the booking period at -0.3% and 2.9% . The peak is varied in the scenarios from 50% to 100% of the booking period. The constant added to the triangular function is 0.45. The peak probability for an arrival is 0.7. The process is thinned with equal probabilities of 0.5 into arrivals for each flight. The arrival intensity for the base case with a peak at 50% is plotted against the time to departure (TTD) in Figure 6.7.

The latest time for an assignment (LAT) is also given as a percentage of the booking time. The times tested span between full flexibility, when an assignment must be made just before departure, and a finalization of the assignment midway through the booking period.

The length of the booking period is determined by the demand factor, i.e. the expected overall demand relative to the accumulated capacities. The demand factors considered

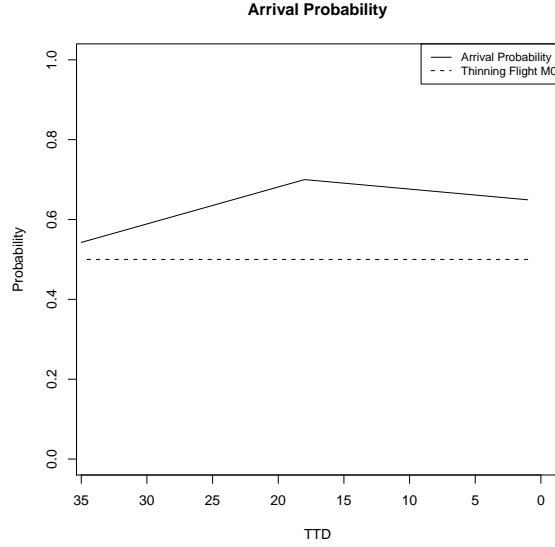


Figure 6.7: Arrival and thinning probabilities in the base-case scenario.

range from 0.7 to 1.3. Given the capacities and the arrival process described above, the horizon is chosen such that the demand factor results as specified.

Table 6.2 provides an overview of the analyzed scenarios. The base-case parameters are highlighted. Note that the coefficient of variation of the arrivals is virtually constant close to 24% when parameters are varied except for the demand factor. Hence, the variation is rather low in regards to values of 20% to 50% in practice for the airline industry (Berge and Hopperstad, 1993, Swan, 2002). In case of the demand factor, the coefficient of variation decreases with higher values because of the longer booking period.

Table 6.2: Scenario Overview

Demand Factor	Capacity Difference	Arrivals Peak	Price Sensitivity (ϵ)	LAT
0.7	20%	50%	1.2	50%
0.8	30%	60%	1.5	60%
0.9	40%	70%	1.8	70%
1.0	50%	80%	2.1	80%
1.1	60%	90%	2.4	90%
1.2	70%	100%	2.7	100%
1.3	80%		3.0	

In an additional scenario with a slightly different setup, the influence of the variation of the arrivals is analyzed. To hold all other parameters constant, especially the demand factor, in this special scenario, the horizon is given and varied between 30 and 100

time periods. The capacities are determined by the length of the booking period, the specified difference of the capacities, and the demand factor. The resulting coefficient of variation for each flight ranges from 22% to 37% which is a realistic range in practice.

6.4.3 Simulation Results

We first present the results of the base-case scenario before varying individual parameters and analyzing their influence on the performance of the various strategies. The base case is set up with a difference of the two available capacities of 20% ($C_1 = 10, C_2 = 12$). We assume an overall demand factor of 1.0. Hence, the number of expected arrivals is in between the two capacities for both flights alike. Retaining a conservative setting, the flights show a low variation of the demand compared to flights observed in practice. The coefficient of variation is 24% for both flights. The peak of the probability for an arrival is located half-way through the booking period. Changes to the assignment are allowed up to 90% of the period. Approximately 90% of the total demand then arrives before the LAT. The price sensitivity is constant with $\epsilon_t = 1.5$. Hence, we assume a reasonably price-sensitive market. The average results are summarized in Table 6.3. The exact model adds an extra revenue of almost 2%. The performance of either heuristic is quite reasonable at around 1.4-1.5%. Even without DDR, the load factor is high at 88%. Applying any DDR strategy increases the load factor. Note that Heuristic 2 achieves the highest load factor, however, also the lowest revenue increase. The exact model yields the highest increase in revenues with a roughly equal load factor as Heuristic 1 suggesting that average ticket prices are the highest when applying the exact model.

Table 6.3: Revenue increase and load factor results.

	No DDR	Heuristic 1	Heuristic 2	Exact
Average Revenue Increase (%)		1.53%	1.42%	1.87%
Average Load Factor (%)	87.5%	89.4%	91.6%	90.2%

Figure 6.8 depicts the 95% confidence intervals of the revenue gains. Because the intervals of the two heuristics overlap, we applied one-sided *Welch's t tests* to verify that the differences of the results are statistical significant. Note that the high number of simulations runs results in tight confidence intervals within a range of about 0.1% in length.

For each scenario, confidence intervals were plotted and statistical tests were run. All results presented below are statistically significant unless they are very close such that results can be regarded as equally good anyhow. To facilitate a compact presentation




Figure 6.8: 95% confidence intervals of the relative revenue increases in the base-case scenario. The results from applying the DDR strategies are plotted relative to the pricing policy without DDR.

of the results, we omit details on confidence intervals and tests in the remainder.

To gain an understanding of the prices set applying the different strategies, in Figure 6.9, Figure 6.10, and Figure 6.11, the strategies are plotted over time for a fixed state with 9 received bookings for each flight. For a comparison, the strategies without DDR are also plotted with the larger and the smaller plane assigned to the respective flight. The prices set applying the exact model stay between the prices set not applying DDR until the very end of the booking process (Figure 6.9). Contrarily, the prices set by Heuristic 1 always coincide with the prices assigning the larger or smaller capacity because the prices are determined by the single-leg processes. This is clearly illustrated in Figure 6.10. Just before the assignment needs to be finalized, the capacities are swapped and the pricing curve switches to the alternative single-leg prices without DDR. Figure 6.11 depicts the prices determined by Heuristic 2. Note that, in contrast to the other strategies, the prices are not bound by the prices of the simple single-leg processes. Figure 6.11(a) shows that, 7 periods to departure, the price for flight $M0$ drops below the single-leg price with the larger plane assigned.

For one exemplary realization, Figure 6.12 depicts the arrivals with their willingness-to-pay and the development of the different pricing strategies in respect to the bookings. The graph highlights the difference of the prices set. Even though the reservation prices of the arrivals in the mid range of the booking period are fairly high, all DDR strategies greatly outperform the traditional pricing strategy without DDR. The final assignment is switched from the initial assignment because of the strong demand for

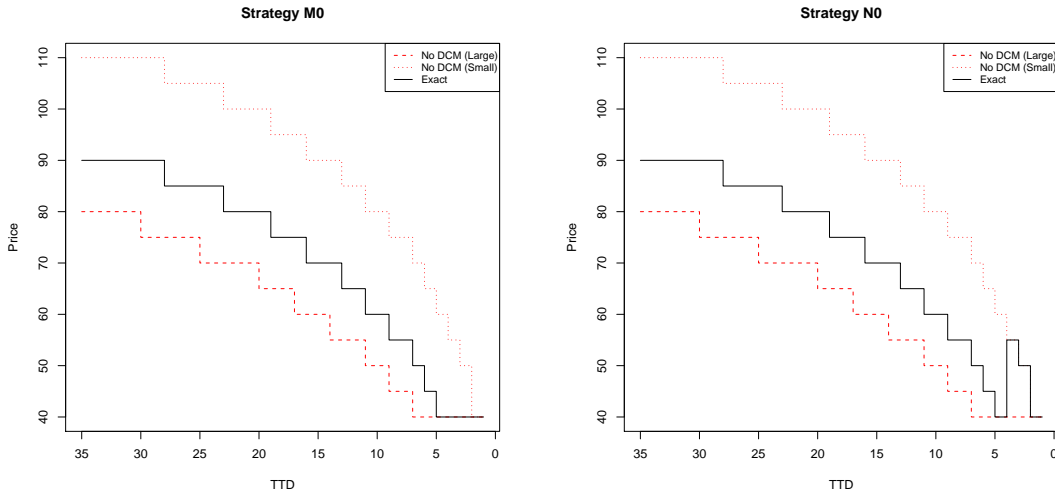
(a) Strategy Flight $M0$ (b) Strategy Flight $N0$

Figure 6.9: Example of a pricing strategy applying the dynamic programming model. The prices are plotted over time for a fixed state with 9 bookings in hand on each flight. Additionally, the single-leg policies are plotted with the larger and smaller capacity assigned to the respective leg.

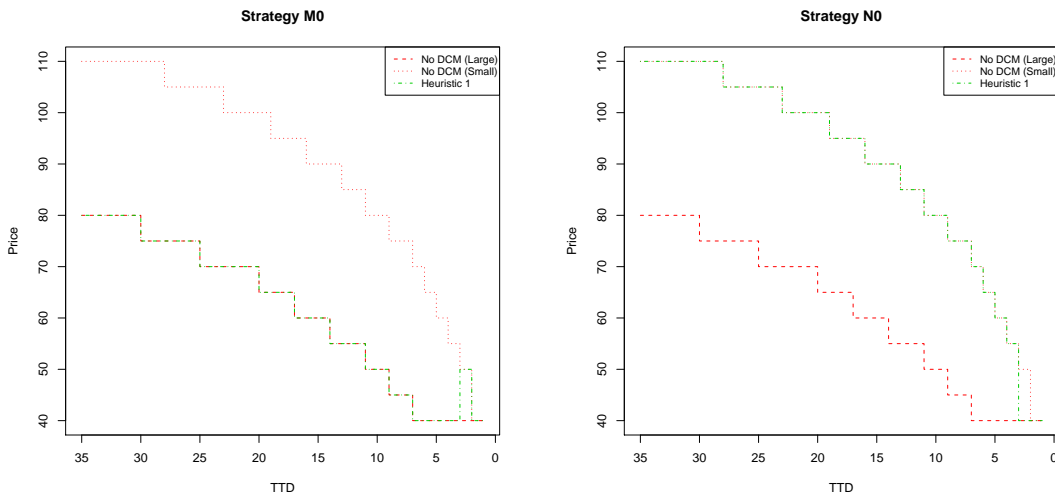
(a) Strategy Flight $M0$ (b) Strategy Flight $N0$

Figure 6.10: Example of a pricing strategy applying Heuristic 1. The prices are plotted over time for a fixed state with 9 bookings in hand on each flight. Additionally, the single-leg policies are plotted with the larger and smaller capacity assigned to the respective leg.

flight $N0$. The customer arriving 10 periods to departure purchases the ticket offered at the second highest price set by the DDR strategies. With the initial assignment only one seat is left at the time of the arrival and the price set applying the traditional


(a) Strategy Flight $M0$ (b) Strategy Flight $N0$

Figure 6.11: Example of a pricing strategy applying Heuristic 2. The prices are plotted over time for a fixed state with 9 bookings in hand on each flight. Additionally, the single-leg policies are plotted with the larger and smaller capacity assigned to the respective leg.

model is too high. The customer is lost.

Demand Factor

To analyze the influence of the magnitude of the demand in relation to the available capacities, the demand factor is varied between 0.7 and 1.3. The demand factor is the total demand for both flights relative to the accumulated capacities. Figure 6.13 depicts the mean relative increase in revenue by applying DDR strategies. A perspicuous result is that the value of applying DDR is the highest for a reasonably good match of demand and capacity. As demand and capacity depart from each other the revenue gain decreases. The revenue gain drops faster for lower demand factors than for higher values. In case of strong demand, when capacities are scarce, assigning the larger capacity to the flight with the higher value of the demand yields the best results. On the opposite, when even the smaller capacity is ample to satisfy demand for either leg, the most cost efficient assignment is preferable.¹ Then, only in very few instances, the demand deviates strongly enough from the expectation such that switching an assignment is profitable and the overall average results converge to those achieved without a DDR strategy.

¹In the presented setting, in case of a very high or low demand factor, the final assignment does not impact results as the flights are set up equally in terms of assignment costs and demand.

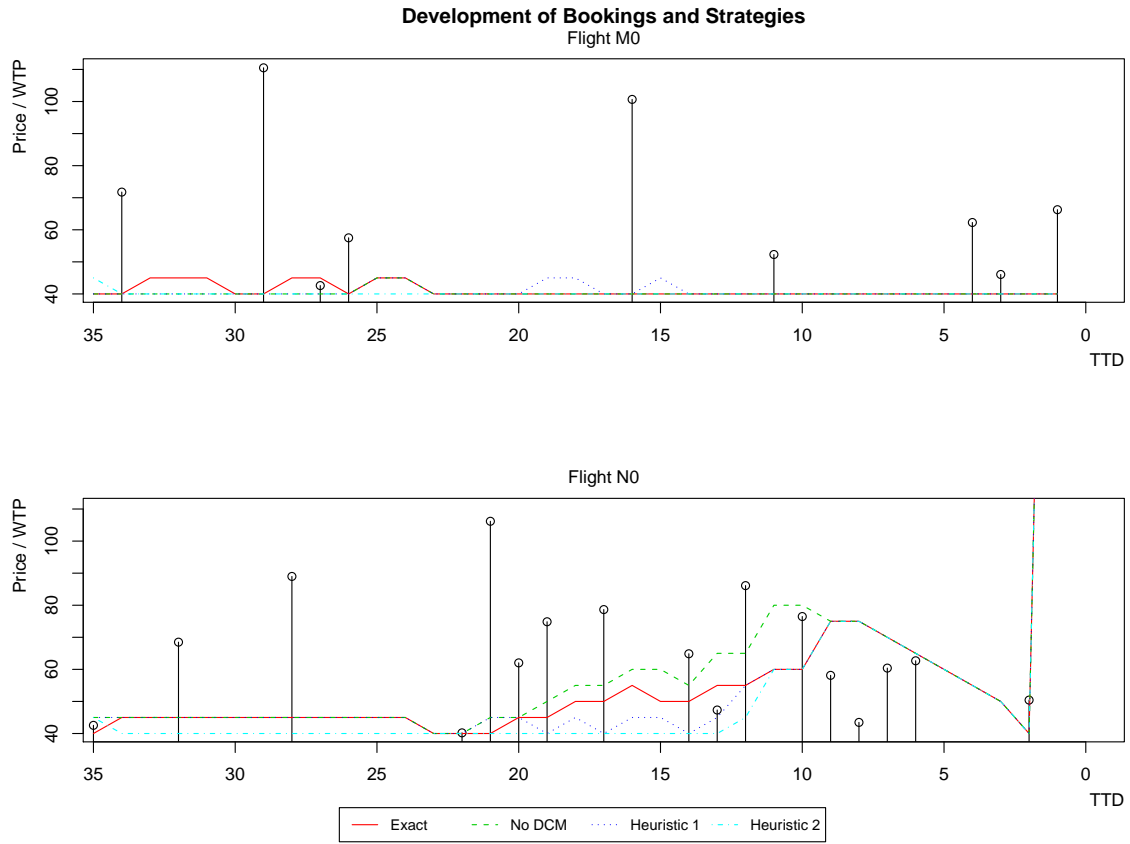


Figure 6.12: Example of a realization of the booking process. For each arrival the respective willingness-to-pay is plotted. The different pricing strategies are plotted for the received bookings at the time.




Figure 6.13: Mean relative revenue increases for varying demand factors. The results from applying the DDR strategies are plotted relative to the pricing policy without DDR.

The difference of the performance of the various strategies increase with the demand factor. While the exact model outperforms both heuristics for settings with a strong demand, the heuristics show a sound performance for a demand factor below one. Also, the heuristics perform approximately equally in situations with low demand and approach the results of the exact model. In case of a stronger demand, Heuristic 1 performs substantially better than Heuristic 2. Note that capacities are typically abundant in relation to demand in the airline industry (Wang and Meng, 2008). However, even in case of the lowest simulated demand factor and the pessimistic simulation setting, applying any DDR strategy yields a substantial gain in revenue of around 0.85% compared to applying simple single-leg dynamic pricing.

Capacity Difference

An increased difference of the two capacities results in a larger revenue gain by applying DDR strategies. The relative increase for various capacity settings is plotted in Figure 6.14. The result is quite intuitive as the expectation of the number of arrivals increases as well because the demand factor is fixed at one. Hence, the smaller capacity is less sufficient to satisfy the expected demand than in settings with a smaller difference of the capacities. Swapping capacities facilitates to assign the larger aircraft to the flight with the stronger demand yielding higher revenues than a fixed assignment.




Figure 6.14: Mean relative revenue increases for varying capacities. The results from applying the DDR strategies are plotted relative to the pricing policy without DDR.

The exact model outperforms both heuristics tested and the improvement increases with the capacity difference. Although Heuristic 2 achieves better results in most

settings than Heuristic 1, the difference is not steady and cannot be regarded as systematic.

Arrivals Peak

Varying the peak of the arrival intensity function does not have a large influence on the performance of any DDR strategy as illustrated in Figure 6.15. Note that the overall portion of demand expected to arrive before the final assignment (LAT) is roughly constant at 88% and the value of the demand arriving before the final assignment is hardly affected by shifting the peak. The exact model performs best and, except for one outlier, Heuristic 1 outperforms Heuristic 2. The results for the different strategies are approximately aligned to each other indicating that the peak of the arrival probability does not have any considerable influence on the revenue gain.





Figure 6.15: Mean relative revenue increases for varying arrival probabilities. The results from applying the DDR strategies are plotted relative to the pricing policy without DDR. The location of the peak of the intensity function is plotted on the abscissa as a percentage of the booking period.

Price Sensitivity

The influence of the price sensitivity of the demand was analyzed by varying the shape parameter ϵ_t from 1.2 to 3.0. According to Belobaba (2006), the former is typical for leisure customers while the later reflects the less sensitive business demand. Figure 6.16(a) shows how willingness-to-pay functions change with different shape parameters.

Note that the shape parameter is assumed constant over the booking period in each scenario. The base price is $\beta_t = 40$ in every scenario.



(a) Exemplary willingness-to-pay functions

(b) Mean relative revenue increase

Figure 6.16: Mean relative revenue increases for varying price sensitivities of the demand. The results from applying the DDR strategies are plotted relative to the pricing policy without DDR.

Figure 6.16(b) depicts the mean results in the different scenarios. The mean revenue gain increases with the price sensitivity of the market. The increase peaks at over 2.5% for a shape parameter of $\epsilon_t = 1.2$ showing the significant potential of applying DDR in price-sensitive markets. In markets with a low sensitivity the gain diminishes. The trend continues down to 0.25% in the scenario with $\epsilon_t = 3.0$.

The advantage of applying the exact model in comparison to the heuristics is statistically significant for scenarios with price-sensitive demand and ranges up to 0.5 points. In these scenarios, Heuristic 1 steadily performs slightly better than Heuristic 2. For large values of the shape parameter in the range of 2.1 to 3.0, the results from the different strategies almost coincide.

Latest Assignment Time

As Figure 6.17 illustrates, the simulation results verify the intuitive expectation that the gain in revenue of applying DDR strategies increases the later during the booking period the assignment needs to be finalized. Note that the percentage of demand expected to arrive before the latest assignment time (LAT) is slightly below the specified LAT, which is given as the percentage of the length of the booking period. Interestingly, even when the assignment needs to be fixed half-way through the booking horizon, DDR strategies still yield an increase of around 1%.

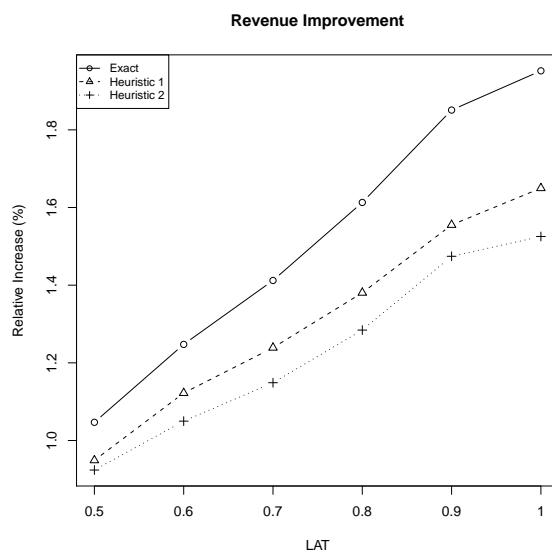


Figure 6.17: Mean relative revenue increases for varying LAT. The results from applying the DDR strategies are plotted relative to the pricing policy without DDR.

The exact model shows the best performance in all settings. The gap between the heuristics and the exact model increases as the LAT approaches the departure time. Heuristic 1 performs continuously just above the results by applying Heuristic 2.

Arrivals Variation

In a last scenario, to test the influence of the variation of the demand, a slightly different setup is used. To keep the important parameters constant, especially the demand factor, the length of the booking period and the capacities are varied. Also, the probability for an arrival in each period is set constantly at $\lambda_t = 0.4$. In Figure 6.18, the revenue increase is plotted against the coefficient of variation (CV) ranging from 20% to 37%. These values are typical for the airline industry.

As expected, a higher variation of the arrivals results in a higher value of applying DDR. The initial assignment is based purely on the expected number of arrivals. Hence, with a higher variation, in more instances, the actual realization differs from the expectation and re-assigning the capacities might yield an increase in revenues. Berge and Hopperstad (1993) and Swan (2002) state that coefficients of variation of up to 50% are typical for the airline industry. Hence, even for a medium variation, the gains are substantial and range up to 2.5%. Both heuristics perform comparably slightly below the exact model.

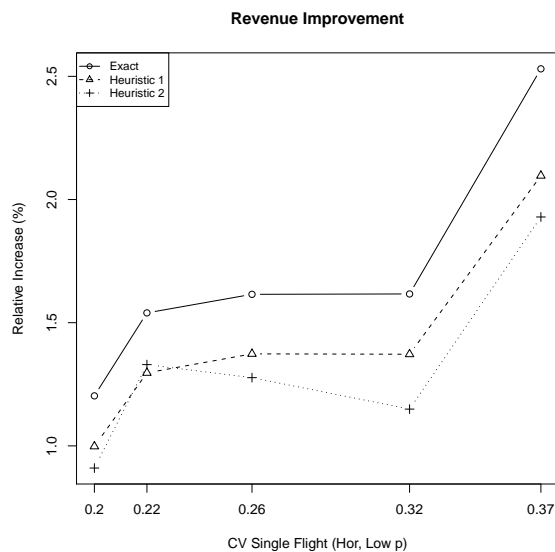


Figure 6.18: Mean relative revenue increases for varying variations of the number of arrivals. The results from applying the DDR strategies are plotted relative to the pricing policy without DDR. The coefficient of variation is plotted on the abscissa.

6.4.4 Summarized Results

Overall, DDR strategies yield a high revenue potential. The simulation setting is conservative as auto-correlation of the arrivals and forecasting errors are excluded. Also, the initial fleet assignment is based on the same data as the dynamic pricing process. In practice, the two planning problems are usually separated and solved based on different data. The fleet assignment is solved before the booking process starts when less precise and less granular data is available. Still, the base case shows an increase in revenue of almost 2% when applying the exact model and around 1.5% when using a heuristic strategy. In light of the transaction volumes typical for airlines, applying the exact model might be worthwhile. In favor of the heuristics is their easy implementation with efficient solvers being widely available. Less space is required to calculate and store the strategies, which is especially important when considering a larger number of flights. Also, other aspects such as robust planning or an improved objective function might be incorporated more easily than in the exact model.

The highest improvement of over 3% shows the scenario with a large difference of the two capacities. Although a difference of 80% is not realistic in an airline setting, other industries, e.g. train or bus transport, might exhibit alternative capacities of such magnitude.

The scenarios with a highly price-sensitive demand show strong improvements while the gain of DDR strategies diminishes when demand is not price-sensitive. Dynamic

pricing in the airline industry is today usually applied in low-price markets. Hence, the demand is highly price-sensitive and DDR yields substantial additional revenue. Demand in other transport industries is usually also quite price-sensitive (e.g. Cole, 2008).

Apart from the scenarios with a highly price-sensitive demand or a large difference of the capacities, the scenarios with a high variation of the arrivals and with a required assignment late during the booking process show a high potential by applying DDR. DDR is worthwhile even when the time of the latest assignment is highly restricted and assignments must be finalized after about half of the demand has realized.

For a low demand factor, the heuristics yield revenue gains comparable to the exact model. In the airline industry, demand below the available capacity is typical (Wang and Meng, 2008). Hence, airlines could achieve good results applying one of the heuristics developed. The performance decreases with a larger gap of the expected arrivals in comparison to the available capacities.

In many settings, Heuristic 1 and Heuristic 2 perform approximately equally. Results by applying Heuristic 1 are more stable and in many cases slightly better than the results by Heuristic 2. Hence, Heuristic 1 is preferable based on the simulation results presented. However, the differences are not steadily considerable and substantial. Before a final recommendation on the heuristics is possible, more tests preferably based on real-world data need to be performed. Depending on the RM processes and systems in place, Heuristic 2 might also be easier to implement into current processes. More advanced systems can incorporate network effects, i.e. primarily mutually dependent demand for legs in a flight network. Often, the network problem is approximated by a decomposition into single-resource problems that are solved separately. Then, the opportunity costs for each leg need to incorporate network effects and are often approximated through (deterministic) linear programs that are very similar to the one used by Heuristic 2. These might be easily adapted to consider possible future equipment swaps. For a detailed introduction on network RM problems we refer the reader to Talluri and Van Ryzin (2004b).

Conclusion

We developed two models that allow to simultaneously optimize dynamic prices and the allocation of two scarce interchangeable capacities. The first model is limited to two flights or products while the second allows an arbitrary number of flights. Different demand and fare structures can be accounted for and additional costs that incur when capacities are swapped can be specified. The two models overcome the limitations of currently available methods that either assume an equal cost structure of the products or do not anticipate possible future changes when optimizing booking controls. Some models further assume equal fare and demand structures. The extended model can be applied to two sequences of flights such that other flights in the network are not affected by an equipment change.

We further presented two heuristics that might be used to find near optimal policies. Especially when the number of flights considered for a swap is large, the dynamic programming models might suffer from the curse of dimensionality. In that case the heuristics become valuable to reduce the computational effort. Additionally, the heuristics might be easier to implement into current systems.

We presented simulations to show the benefits of systematic equipment changes in general and to analyze the influences of different parameters found in practice. The policies determined by applying the two heuristics showed an acceptable performance with results within 1% below the revenue when applying the optimal policies. The dynamic programming models might be applied to fully exploit the benefits of Demand-Driven Re-Fleeting.

The airline industry is the origin of Revenue Management and consistently most publications concentrate on an airline setting. We followed this custom and also focused on the passenger airline case. However, Demand-Driven Re-Fleeting might be beneficial in other industries as well. Other transport segments also face overall limited

capacities that are interchangeable between different products offered at a time. Bus or train companies might change assigned capacities on a short-term basis. Cargo airlines might apply the models to decide whether to fly a dedicated carrier or to use cargo space on passenger flights.

7.1 Summary

We introduced the concept and basic economic theory of price discrimination in general. We then briefly discussed the theory of Markov Decision Processes limited to the scope needed for the models in the subsequent chapters.

The history and evolution of Revenue Management and comprehensive requirements were outlined, before the general process was described and the individual steps were discussed in more detail. Briefly, the traditional concepts in Revenue Management, i.e. static and dynamic capacity control models, and structural results were presented.

A single-leg dynamic pricing model was developed and structures were shown. For a fixed point in time, the value function was shown to be non-decreasing and concave in the capacity remaining for sale. Using these results we further proved the monotonicity of the pricing policy. The optimal price decreases the more capacity is available. We confirmed the intuitive result that the total expected revenue increases with the time left for sales given a fixed capacity. The opportunity costs of selling one seat also increase with the time left. The opportunity costs arise from future sales possibilities that diminish close to departure. They provide a lower bound of the optimal price.

In general the pricing policy is not monotone in time which we showed by a counter-example. We provided a sufficient condition dependent on the purchase probabilities for optimal prices that are monotone in time. In applications like the fashion industry, where customers are willing to pay a premium to acquire the product early during the sales period, e.g. for exclusiveness, the condition is valid. Contrarily, in the airline industry, price-sensitive clients book further in advance and travelers accept to pay a premium price for bookings close to departure.

We introduced the airline planning process and explained the potential benefits of Demand-Driven Re-Fleeting. The general concept and existing frameworks were outlined and shortcomings were discussed. During the fleet assignment the allocation of available aircraft is optimized applying a mixed integer program. We presented an algorithm that finds opportunities for an equipment swap that does not affect the rest of the network given a fleet schedule.

To facilitate the simultaneous optimization of prices and the allocation of capacities,

we developed a two-leg dynamic programming model. During the fleet assignment, two plane types are assigned to the legs and the overall capacity is fixed. Depending on the realization of the bookings, the decision is made which of the two aircraft to operate on which leg. We showed that the problem decomposes into the single-leg processes when the assignment has been finalized or the decision to allocate the capacities is fixed in advance. Further, the optimization of the prices is separable in case the decision is postponed. We proved that postponing the decision is optimal when sales on both flights do not exceed the smaller capacity. Additionally, we provided conditions under which certain assignment options can be excluded from the maximization when calculating an optimal policy. The intuitive results of an increasing value with more capacity remaining and with the time to departure were confirmed. The flexibility to change assignments has a positive value or might be worthless depending on the updated demand forecasts and the bookings-in-hand.

At the latest time, when the allocation has to be decided, e.g. for compliance with union agreements, the optimal assignment decision follows a switching curve. The result facilitates a more efficient calculation of an optimal policy. We specified concave domains of the value function where the optimal prices are monotone in the capacity left for sale.

All of the results were then used to develop an efficient algorithm that determines an optimal policy. The algorithm was presented in pseudo-code. In order to broaden the scope of possible applications, the model was extended to an arbitrary number of flights. The flights need to be separable into two disjunct sets such that all legs in one set share an assigned capacity. While applications of the two-leg model are limited to flights to destinations where a swap can be reversed, e.g. hub airports, the extended model can be applied to any two sequences of flights. Most structures of the two-leg model convey to the extended model. These were presented and differences were discussed.

Because the presented dynamic programming models might suffer from the curse of dimensionality and heuristics enjoy a wide acceptance in the industry, we additionally developed two heuristics. The first heuristic that follows the ideas of Berge and Hopperstad (1993) periodically re-optimizes the fleet assignment throughout the booking horizon. Product prices are optimized without considering possible future swaps of aircraft. The second heuristic builds upon a limited lookahead policy and does anticipate changes of the assigned capacities. A linear program relaxing the integrality of the capacities approximates the opportunity costs of selling one seat. The approximations are then used to optimize the prices.

To assess the overall benefits of Demand-Driven Re-Fleeting and to compare the per-

formance of the various policies derived through the exact models and the heuristics, we simulated a booking process using different scenarios. We followed a conservative approach and did not assume improving forecasts in time. Auto-correlation of arrivals, dependent purchase probabilities, and forecasting errors might further improve the benefits. We found that applying Demand-Driven Re-Fleeting yields revenue increases of up to 3%. The influence of various parameters was analyzed. In our simulations, benefits peaked for a demand factor such that the assigned capacities slightly exceeded the aggregated demand. A larger gap of the two capacities, a required final assignment close to departure, and a large variation of the arrivals improved revenue gains. The benefits of accounting for possible future changes when optimizing prices increased with the price-sensitivity of the demand.

7.2 Future Research

Before applying the suggested models in practice, an ex-post analysis using real-world data should substantiate the results of our simulations based on an artificial data set. We excluded forecasting errors and auto-correlated arrivals. Further, reservation prices were assumed to be mutually independent. Fleet assignment optimization and Revenue Management both heavily rely on forecasts with a good accuracy. We expect benefits of applying Demand-Driven Re-Fleeting to increase when forecasts improve substantially close to departure. Other influences can only be assessed meaningfully in real-world setups that we lacked the data to simulate. For example, the flights were set up equally to demonstrate benefits without influencing results.

Deciding on two alternative capacity assignments should usually provide sufficient flexibility in practice. However, the scope might be extended to a pool of capacities that are subsequently assigned to products with an equal assignment.

Cancellations and no-shows necessitate airlines to consider overbooking in practice. Hence, future research might extend the provided models to include overbooking and cancellations.

If demand for different products is not mutually independent, network effects need to be considered. Often, network problems are decomposed into separate single-leg problems. Then, opportunity costs incurred throughout the network are approximated by a suitable heuristic. The linear program used by Heuristic 2 is similar to a program commonly used to approximate network opportunity costs. Hence, Heuristic 2 might be easily extended to a network setting.

A

Additional Proofs

A.1 Proof of Theorem 6.2.1

Proof.

- (i) We prove by induction on t . Since $V_0 \equiv 0$ and $V_0^i \equiv 0$ for all $i \in M \cup N$, the assertion holds for $t = 0$. Therefore let $V_t(s_1, \dots, s_{m+n}, 0) = \sum_{i \in M \cup N} V_t^i(s_i)$ hold for some $t = T - 1, T - 2, \dots, 0$. Then, in $s = (s_1, \dots, s_{m+n}, 0) \in S_{t+1}$, since $s_{m+n+1} = 0$, only actions $a \in A_{t+1}(s)$ with $a_{m+n+1} = (0, 0)'$ and, additionally, $a_i = f_0^i$ for $i \in M \cup N$ with $s_i = 0$. We then get

$$\begin{aligned} V_{t+1}(s) &= \max_{a \in A_{t+1}(s)} \left\{ \sum_{\substack{i \in M \cup N \\ s_i \neq 0}} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i + V_t(Y^i(s, a))] \right. \\ &\quad \left. + \left[1 - \sum_{\substack{i \in M \cup N \\ s_i \neq 0}} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) \right] V_t(Y(s, a)) \right\} \\ &= \max_{a \in A_{t+1}(s)} \left\{ \sum_{\substack{i \in M \cup N \\ s_i \neq 0}} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) \left[a_i + \sum_{j \in M \cup N} V_t^j(s_j - \delta_j(i)) \right] \right. \\ &\quad \left. + \left[1 - \sum_{\substack{i \in M \cup N \\ s_i \neq 0}} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) \right] \left[\sum_{j \in M \cup N} V_t^j(s_j) \right] \right\} \end{aligned}$$

$$\begin{aligned}
&= \max_{a \in A_{t+1}(s)} \left\{ \sum_{\substack{i \in M \cup N \\ s_i \neq 0}} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) \left[a_i + \underbrace{\sum_{j \in M \cup N} V_t^j(s_j - \delta_j(i)) - \sum_{j \in M \cup N} V_t^j(s_j)}_{=0 \text{ for } i \neq j} \right] \right. \\
&\quad \left. + \sum_{j \in M \cup N} V_t^j(s_j) \right\} \\
&= \max_{a \in A_{t+1}(s)} \left\{ \sum_{\substack{i \in M \cup N \\ s_i \neq 0}} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i - \Delta_{s_i} V_t^i(s_i)] + \sum_{i \in M \cup N} V_t^i(s_i) \right\} \\
&= \sum_{\substack{i \in M \cup N \\ s_i \neq 0}} \max_{a_i \in A_{t+1}^i(s)} \left\{ \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i - \Delta_{s_i} V_t^i(s_i)] + V_t^i(s_i) \right\} \\
&= \sum_{i \in M \cup N} V_{t+1}^i(s_i)
\end{aligned}$$

The proof for (ii) follows straightforward applying (i).

(iii) For $t = T - 1, T - 2, \dots, \bar{T}$, let $s = (s_1, \dots, s_{m+n}, 1) \in S_{t+1}$. Then,

$$\begin{aligned}
&L_{t+1} V_t(s | (0, 0)') \\
&= \max_{a \in A_{t+1}(s | (0, 0)')} \left\{ \sum_{\substack{i \in M \cup N \\ s_i \neq 0}} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i - \Delta_{s_i} V_t(s)] \right\} + V_t(s) \\
&= \sum_{\substack{i \in M \cup N \\ s_i \neq 0}} \max_{a_i \in A_{t+1}^i(s)} \left\{ \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i) [a_i - \Delta_{s_i} V_t(s)] \right\} + V_t(s),
\end{aligned}$$

which completes the proof. \square

A.2 Proof of Theorem 6.2.2

Proof.

(i) The proof follows by induction on t . Since $V_0 \equiv 0$, the assertion holds for $t = 0$.

Therefore, let (i) hold for some $t = T - 1, T - 2, \dots, 0$.

Fix $s, s' = s - e(k) \in S_{t+1}$ with $s_k > 0$ for some $k \in M \cup N$. Let $a^* \in A_{t+1}(s')$ be optimal in s' . Then $a^* \in A_{t+1}(s)$, too. We get

$$\Delta_{s_k} V_{t+1}(s) \geq \sum_{i \in M \cup N} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i^*) [V_t(Y^i(s, a^*)) - V_t(Y^i(s', a^*))]$$

$$\begin{aligned}
& + \left[1 - \sum_{i \in M \cup N} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i^*) \right] [V_t(Y(s, a^*)) - V_t(Y(s', a^*))] \\
= & \sum_{i \in M \cup N} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i^*) [\Delta_{s_k} V_t(Y^i(s, a^*))] \\
& + \left[1 - \sum_{i \in M \cup N} \lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i^*) \right] [\Delta_{s_k} V_t(Y(s, a^*))] \\
\geq & 0,
\end{aligned}$$

which completes the proof.

- (ii) To prove (ii), let $t = T, T-1, \dots, \bar{T}$. Fix $s = (s_1, \dots, s_{m+n}, 1), s' = s - e(m+n+1) \in S_t$. Further, suppose $a^{**} \in A_t(s')$ is optimal in s' . Note that $a_{m+n+1}^{**} = (0, 0)'$. Suppose $k_d^1 = 0$, then by using $a^* = (a_1^{**}, \dots, a_{m+n}^{**}, (1, 0)') \in A_t(s)$, and, additionally, exploiting the monotonicity of V_{t-1} after an assignment, we finally get

$$\begin{aligned}
& V_t(s) - V_t(s') \\
\geq & \sum_{i \in M \cup N} \lambda_t p_t^i q_t^i(a_i^{**}) [V_{t-1}(Y^i(s, a^*)) - V_{t-1}(Y^i(s', a^{**}))] \\
& + \left[1 - \sum_{i \in M \cup N} \lambda_t p_t^i q_t^i(a_i^{**}) \right] [V_{t-1}(Y(s, a^*)) - V_{t-1}(Y(s', a^{**}))] \\
\geq & 0,
\end{aligned}$$

which verifies (ii) for $k_d^1 = 0$. If $k_d^1 > 0$, then, by assumption, $k_d^2 = 0$ and a similar argument works using $a^* = (a_1^{**}, \dots, a_{m+n}^{**}, (0, 1)') \in A_t(s)$.

- (iii) We split the proof into two parts:

- a. First, fix $s \in S_t \cap S_{t+1}$ and suppose $s_{m+n+1} = 0$. Using Theorem 4.6.3 and 6.2.2.

$$\begin{aligned}
& V_{t+1}(s) - V_t(s) \\
= & \sum_{i \in M \cup N} V_{t+1}^i(s_i) - \sum_{j \in M \cup N} V_t^j(s_j) \\
= & \sum_{i \in M \cup N} \Delta_t V_{t+1}^i(s_i) \\
\geq & 0.
\end{aligned}$$

- b. Now, consider $s \in S_t \cap S_{t+1}$ with $s_{m+n+1} = 1$. Note that this implies $t = T-1, T-2, \dots, \bar{T}$. Then $a = (f_0^1, \dots, f_0^{m+n}, (0, 0)') \in A_{t+1}(s)$. We get

$$\begin{aligned}
& V_{t+1}(s) - V_t(s) \\
\geq & Q_{t+1} V_t(s, a) - V_t(s) \\
= & 0.
\end{aligned}$$

□

A.3 Proof of Lemma 6.2.3

Proof. Let $t = T, \dots, \bar{T}$ and fix $s \in S_t$ with $s_{m+n+1} = 1$. Then there exists an action $a \in A_t(s)$ with $a_{m+n+1} = (1, 0)'$ and

$$\begin{aligned} V_t(s) &\geq L_t V_{t-1}(s|(1, 0)') \\ &= V_t(s_1 + d, \dots, s_m + d, s_{m+1}, \dots, s_{m+n}, 0) - k_d^1. \end{aligned}$$

(ii) can be shown analogously. □

A.4 Proof of Theorem 6.2.4

Proof.

- (i) Let $t = T, \dots, \bar{T} + 1$. Fix $s \in S_t$ with $s_i > 0$ for all $i \in M$ and $s_{m+n+1} = 1$. Assume $a^* = (a_1^*, \dots, a_{m+n}^*) \in A_t(s|(1, 0)')$ maximizes $L_t V_{t-1}(s|(1, 0)')$. Note that $(a_1^*, \dots, a_{m+n}^*, (0, 0)') \in A_t(s)$ is also admissible in s . Using Lemma 6.2.3, we then get

$$\begin{aligned} &L_t V_{t-1}(s|(0, 0)') - L_t V_{t-1}(s|(1, 0)') \\ &\geq \sum_{i \in M \cup N} \lambda_t p_t^i q_t^i(a_i^*) \left[V_{t-1}(Y^i(s|(0, 0)')) - V_{t-1}(Y^i(s|(1, 0)')) \right] \\ &\quad + \left[1 - \sum_{i \in M \cup N} \lambda_t p_t^i q_t^i(a_i^*) \right] \left[V_{t-1}(Y(s|(0, 0)')) - V_{t-1}(Y(s|(1, 0)')) \right] \\ &\quad + k_d^1 \\ &\geq \sum_{i \in M \cup N} \lambda_t p_t^i q_t^i(a_i^*) \left[V_{t-1}(Y^i(s|(1, 0)')) - k_d^1 - V_{t-1}(Y^i(s|(1, 0)')) \right] \\ &\quad + \left[1 - \sum_{i \in M \cup N} \lambda_t p_t^i q_t^i(a_i^*) \right] \left[V_{t-1}(Y(s|(1, 0)')) - k_d^1 - V_{t-1}(Y(s|(1, 0)')) \right] \\ &\quad + k_d^1 \\ &= 0. \end{aligned}$$

In states $s' \in S_t$ with $s_{m+n+1} = 0$ only actions $a \in A_t(s')$ are feasible with the assignment action $a_{m+n+1} = (0, 0)'$ and the result is trivial.

(ii) can be shown analogously and (iii) follows straightforward from (i) and (ii).

- (iv) In $s = (0, \dots, 0, 1) \in S_{\bar{T}}$, assume $a^* := (a_1^*, \dots, a_{m+n}^*, (1, 0)') \in A_{\bar{T}}(s)$ to be optimal in s . Then, using Theorem 6.2.1 (ii) and Theorem 6.2.2 (iii),

$$L_{\bar{T}+1} V_{\bar{T}}(s|(1, 0)') = V_{\bar{T}+1}(Y(s, a^*) - k_d^1)$$

$$\begin{aligned}
&\geq V_{\bar{T}}(Y(s, a^*) - k_d^1) \\
&= V_{\bar{T}}(s) \\
&= L_{\bar{T}+1} V_{\bar{T}}(s | (0, 0)') .
\end{aligned}$$

Hence, applying $(1, 0)'$ results in a value at least as high as when applying $(0, 0)'$ in s at time $t = \bar{T} + 1$. Repeating the same argument with $\bar{T} + 1$ in place of \bar{T} , the result follows by induction.

A similar proof works if $(a_1^*, \dots, a_{m+n}^*, (0, 1)') \in A_{\bar{T}}(s)$ is optimal in s at time $t = \bar{T}$. \square

A.5 Proof of Theorem 6.2.6

Proof. Fix $s \in S_t$ with $s_i > 0$ for all $i \in M \cup N$ and $s_{m+n+1} = 1$. Then $s \in S_{t+1}$, too. Let $a \in A_{t+1}^i(s) \setminus \{f_0^i\}$ such that $a \leq \Delta_{s_i} V_t(s)$. Then

$$\begin{aligned}
&\lambda_{t+1} p_{t+1}^i q_{t+1}^i(a_i)(a)[a - \Delta_{s_i} V_t(s)] \\
&\leq 0 \\
&= \lambda_{t+1} p_{t+1}^i q_{t+1}^i(f_0^i)[f_0^i - \Delta_{s_i} V_t(s)] ,
\end{aligned}$$

which implies that a is not the largest maximizer $\psi_{t+1}^i(s)$. Thus, $\psi_{t+1}^i(s) \geq \min\{\Delta_{s_i} V_t(s), f_0^i\}$.

On the other hand, if $a > \Delta_{s_i} V_t(s)$ holds for all $a \in A_{t+1}^i(s)$, the result is trivial. \square

A.6 Proof of Theorem 6.2.7

Proof. The proof follows straightforward from Theorem 6.2.1 and Theorem 4.6.2. \square

A.7 Proof of Theorem 6.2.10

Proof. The proof follows straightforward from Lemma 4.6.1. \square

B

Additional Algorithms

B.1 Two-leg Dynamic Pricing

Algorithm B.1.1 Algorithm to determine an optimal policy for 2 legs subject to swap

Input: Sets of feasible actions $A_t^i(s), i = 1, 2, t \in \{T, \dots, 1\}, s \in S_t$,
single-leg policies and values $a_{i,t}^*, V_t^i, i = 1, 2, t \in \{T, \dots, 1\}$,
arrival probabilities $(\lambda_T, \lambda_{T-1}, \dots, \lambda_1)$,
thinning probabilities $p_t^i, i = 1, 2, t \in \{T, \dots, 1\}$,
purchase probabilities $(q_T, q_{T-1}, \dots, q_1)$,
total capacities C_1, C_2 ,
costs of swap k_d^1, k_d^2 ,
length of booking horizon T ,
latest time for final fleet assignment \bar{T}

Output: Optimal policy $(a_T^*, a_{T-1}^*, \dots, a_1^*)$,
Expected values $(V_T, V_{T-1}, \dots, V_0)$

1. $d = C_2 - C_1$
 2. $t = 0$
 3. (* Use single-leg processes if $s_3 = 0$ (Theorem 6.1.1) *)
 4. **while** $t \leq T$
 5. $s_1 = 0$
 6. **while** $s_1 \leq C_2$
 7. $s_2 = 0$
 8. **while** $s_2 \leq C_1$
 9. $a_t^*(s_1, s_2, 0) = (a_{1,t}^*(s_1), a_{2,t}^*(s_2), (0, 0)')$
-

Algorithm B.1.1 (continued)

-
10. $V_t(s_1, s_2, 0) = V_t^1(s_1) + V_t^2(s_2)$
 11. $s_2 = s_2 + 1$
 12. $s_1 = s_1 + 1$
 13. $s_1 = 0$
 14. **while** $s_1 \leq C_1$
 15. $s_2 = C_1 + 1$
 16. **while** $s_2 \leq C_2$
 17. $a_t^*(s_1, s_2, 0) = (a_{1,t}^*(s_1), a_{2,t}^*(s_2), (0, 0)')$
 18. $V_t(s_1, s_2, 0) = V_t^1(s_1) + V_t^2(s_2)$
 19. $s_2 = s_2 + 1$
 20. $s_1 = s_1 + 1$
 21. $t = t + 1$
 22. (* In $t = \bar{T}$ an assignment is necessary when $s_3 = 1$ *)
 23. $s_1 = 0$
 24. **while** $s_1 \leq C_1$
 25. $s_2 = 0$
 26. **while** $s_2 \leq C_1$
 27. (* Monotonicity assignment actions (Theorem 6.1.5) *)
 28. **if** $s_1 > 0$ **and** $a_t^*(s_1 - 1, s_2, 1)(3) = (0, 1)'$
 29. $a_t^*(s_1, s_2, 1) = (a_{1,t}^*(s_1), a_{2,t}^*(s_2 + d), (0, 1)')$
 30. $V_t(s_1, s_2, 1) = V_t^1(s_1) + V_t^2(s_2 + d) - k_d^2$
 31. **else if** $s_2 > 0$ **and** $a_t^*(s_1, s_2 - 1, 1)(3) = (1, 0)'$
 32. $a_t^*(s_1, s_2, 1) = (a_{1,t}^*(s_1 + d), a_{2,t}^*(s_2), (1, 0)')$
 33. $V_t(s_1, s_2, 1) = V_t^1(s_1 + d) + V_t^2(s_2) - k_d^1$
 34. **else**
 35. **if** $V_t^1(s_1 + d) + V_t^2(s_2) - k_d^1 \geq V_t^1(s_1) + V_t^2(s_2 + d) - k_d^2$
 36. $a_t^*(s_1, s_2, 1) = (a_{1,t}^*(s_1 + d), a_{2,t}^*(s_2), (1, 0)')$
 37. $V_t(s_1, s_2, 1) = V_t^1(s_1 + d) + V_t^2(s_2) - k_d^1$
 38. **else**
 39. $a_t^*(s_1, s_2, 1) = (a_{1,t}^*(s_1), a_{2,t}^*(s_2 + d), (0, 1)')$
 40. $V_t(s_1, s_2, 1) = V_t^1(s_1) + V_t^2(s_2 + d) - k_d^2$
 41. $s_2 = s_2 + 1$
 42. $s_1 = s_1 + 1$
 43. $t = \bar{T} + 1$
 44. **while** $t \leq T$
 45. (* In $s = (0, 0, 1)$ an assignment is optimal (Theorem 6.1.4) *)
-

Algorithm B.1.1 (continued)

```

46.   if  $V_t^1(d) - k_d^1 \geq V_t^2(d) - k_d^2$ 
47.        $a_t^*(0, 0, 1) = (a_{1,t}^*(d), f_0^2, (1, 0)')$ 
48.        $V_t(0, 0, 1) = V_t^1(d) - k_d^1$ 
49.   else
50.        $a_t^*(0, 0, 1) = (f_0^1, a_{2,t}^*(d), (0, 1)')$ 
51.        $V_t(0, 0, 1) = V_t^2(d) - k_d^2$ 
52.   (*  $(0, s_2, 1)$  with  $s_2 > 0$  *)
53.    $s_2 = 1$ 
54.   while  $s_2 \leq C_1$ 
55.        $\tilde{a}^2 = f_0^2$ 
56.        $\tilde{v}^2 = 0$ 
57.       foreach  $f \in A_t^2(0, s_2, 1) \setminus \{f_0^2\}$ 
58.           (* Bounds on price (Theorem 6.1.9 and 6.1.12) *)
59.           if on concave domain of  $V_{t-1}$  in  $s_2$ 
60.                $u = a_t^*(0, s_2 - 1, 1)(2)$ 
61.           else
62.                $u = f_0^2$ 
63.           if  $f < \Delta_{s_2} V_{t-1}(0, s_2, 1)$  or  $f > u$ 
64.               continue
65.                $\tilde{v}_c = \lambda_t p_t^2 q_t^2(f)[f - \Delta_{s_2} V_{t-1}(0, s_2, 1)]$ 
66.               if  $\tilde{v}_c \geq \tilde{v}^2$ 
67.                    $\tilde{v}^2 = \tilde{v}_c$ 
68.                    $\tilde{a}^2 = f$ 
69.           if  $V_t^1(d) + V_t^2(s_2) - k_d^1 \geq \tilde{v}^2 + V_{t-1}(0, s_2, 1)$ 
70.                $a_t^*(0, s_2, 1) = (a_{1,t}^*(d), a_{2,t}^*(s_2), (1, 0)')$ 
71.                $V_t(0, s_2, 1) = V_t^1(d) + V_t^2(s_2) - k_d^1$ 
72.           else
73.                $a_t^*(0, s_2, 1) = (f_0^1, \tilde{a}^2, (0, 0)')$ 
74.                $V_t(0, s_2, 1) = \tilde{v}^2 + V_{t-1}(0, s_2, 1)$ 
75.            $s_2 = s_2 + 1$ 
76.   (*  $(s_1, 0, 1)$  with  $s_1 > 0$  *)
77.    $s_1 = 1$ 
78.   while  $s_1 \leq C_1$ 
79.        $\tilde{a}^1 = f_0^1$ 
80.        $\tilde{v}^1 = 0$ 
81.       foreach  $f \in A_t^1(s_1, 0, 1) \setminus \{f_0^1\}$ 

```

Algorithm B.1.1 (continued)

119. $\tilde{a}^2 = f_0^2$
120. $\tilde{v}^2 = 0$
121. **foreach** $f \in A_t^2(s) \setminus \{f_0^2\}$
122. (* Use lower and upper bound on price (Theorem 6.1.9
and 6.1.12) *)
123. **if** on concave domain of V_{t-1} in s_2
124. $u = a_t^*(s_1, s_2 - 1, 1)(2)$
125. **else**
126. $u = f_0^2$
127. **if** $f < \Delta_{s_2} V_{t-1}(s_1, s_2, 1)$ **or** $f > u$
128. **continue**
129. $\tilde{v}_c = \lambda_t p_t^2 q_t^2(f)[f - \Delta_{s_2} V_{t-1}(s_1, s_2, 1)]$
130. **if** $\tilde{v}_c \geq \tilde{v}^2$
131. $\tilde{v}^2 = \tilde{v}_c$
132. $\tilde{a}^2 = f$
133. $a_t^*(s_1, s_2, 1) = (\tilde{a}^1, \tilde{a}^2, (0, 0)')$
134. $V_t(s_1, s_2, 1) = \tilde{v}^1 + \tilde{v}^2 + V_{t-1}(s_1, s_2, 1)$
135. $s_2 = s_2 + 1$
136. $s_1 = s_1 + 1$
137. $t = t + 1$
-

B.2 Demand-Driven Re-Fleeting Heuristic

Algorithm B.2.1 Heuristic 2

Input: Flight sets M, N ,
 sets of feasible actions $A_t^i(s), i \in M \cup N, t \in \{T, \dots, 1\}, s \in S_t$,
 single-leg policies and values $a_{i,t}^*, V_t^i, i \in M \cup N, t \in \{T, \dots, 1\}$,
 arrival probabilities $(\lambda_T, \lambda_{T-1}, \dots, \lambda_1)$,
 thinning probabilities $p_t^i, i \in M \cup N, t \in \{T, \dots, 1\}$,
 purchase probabilities $(q_T, q_{T-1}, \dots, q_1)$,
 total capacities C_1, C_2 ,
 length of booking horizon T ,
 latest time for final fleet assignment \bar{T}

Output: Bookings received and final fleet assignment

1. $t = T$
 2. $d = C_2 - C_1$
 3. **foreach** $i \in M \cup N$
 4. $s_i = C_1$
 5. $s_{m+n+1} = 1$
 6. **while** $t > \bar{T}$
 7. Solve LP 6.3.2 for $t' = t - 1$ and $s = (s_1, s_2, \dots, s_{m+n+1})$
 8. **foreach** $i \in M$
 9. Set $\Delta_{s_i} \tilde{V}_{t-1}$ to shadow price of (6.11)
 10. **foreach** $i \in N$
 11. Set $\Delta_{s_i} \tilde{V}_{t-1}$ to shadow price of (6.12)
 12. **foreach** $i \in M \cup N$
 13. $a_i^h = \arg \max_{a_i \in A_t^i(s)} \{\lambda_t p_t^i q_t^i(a_i) [a_i - \Delta_{s_i} \tilde{V}_{t-1}]\}$
 14. Publish offer $a^h = (a_1^h, a_2^h, \dots, a_{m+n}^h)$
 15. **if** customer arrives for flight $i \in M \cup N$ and accepts offer a_i^h
 16. $s_i = s_i - 1$
 17. **if** $s_i = 0$
 18. (* Initial capacity has been sold out for one flight *)
-

Algorithm B.2.1 (continued)

19. $t = t - 1$
 20. **break**
 21. $t = t - 1$
 22. **if** $\sum_{i \in M} V_t^i(s_i + d) + \sum_{i \in N} V_t^i(s_i) - k_d^1 \geq \sum_{i \in M} V_t^i(s_i) + \sum_{i \in N} V_t^i(s_i + d) - k_d^2$
 23. **foreach** $i \in M$
 24. $s_i = s_i + d$
 25. **else**
 26. **foreach** $i \in N$
 27. $s_i = s_i + d$
 28. $s_{m+n+1} = 0$
 29. Continue with single-leg processes (Algorithm 4.6.1)
-

Bibliography

- Abara, J. (1989). Applying Integer Linear Programming to the Fleet Assignment Problem. *Interfaces*, 19(4):20–28.
- Ahuja, R. K., Magnanti, T. L., and Orlin, J. B. (1993). *Network flows : theory, algorithms, and applications*. Prentice Hall, Upper Saddle River.
- airliners.de (2011). Germanwings rückt näher an Lufthansa. <http://www.airliners.de/management/marketing/germanwings-rueckt-naeher-an-lufthansa/25654> (06.12.2011).
- Armstrong, M. and Vickers, J. (2001). Competitive price discrimination. *The RAND Journal of Economics*, 32(4):579–605.
- Barnhart, C., Boland, N. L., Clarke, L. W., Johnson, E. L., Nemhauser, G. L., and Shenoi, R. G. (1998). Flight string models for aircraft fleetting and routing. *Transportation Science*, 32(3):208–220.
- Barnhart, C., Farahat, A., and Lohatepanont, M. (2009). Airline fleet assignment with enhanced revenue modeling. *Operations Research*, 57:231–244.
- Barnhart, C., Kniker, T. S., and Lohatepanont, M. (2002a). Itinerary-based airline fleet assignment. *Transportation Science*, 36(2):199–217.
- Barnhart, C., Lu, F., and Shenoi, R. (2002b). Integrated airline schedule planning. In Yu, G., editor, *Operations research in the airline industry*, International series in operations research and management science ; 9, pages 384–403. Kluwer Academic, Boston, 4. print. edition.

- Barz, C. (2007). *Risk-averse capacity control in revenue management*. PhD thesis, Berlin.
- Barz, C. and Waldmann, K. (2007). Risk-sensitive capacity control in revenue management. *Mathematical Methods of Operations Research*, 65:565–579.
- Bellman, R. (1957). A markovian decision process. Technical report, RAND Corporation.
- Belobaba, P. (2006). Introduction to PODS Passenger Choice Model. Lecture notes, Massachusetts Institute of Technology, Dept. of Aeronautics and Astronautics.
- Belobaba, P. (2009a). Airline operating costs and measures of productivity. In Belobaba, P., Odoni, A., and Barnhart, C., editors, *The global airline industry*. Wiley, Chichester.
- Belobaba, P. (2009b). The airline planning process. In Belobaba, P., Odoni, A., and Barnhart, C., editors, *The global airline industry*. Wiley, Chichester.
- Belobaba, P. P. (1987a). *Air travel demand and airline seat inventory management*. PhD thesis.
- Belobaba, P. P. (1987b). Airline yield management. an overview of seat inventory control. *Transportation Science*, 21(2):63.
- Belobaba, P. P. (1989). Application of a probabilistic decision model to airline seat inventory control. *Operations Research*, 37(2):183–197.
- Belobaba, P. P. and Weatherford, L. R. (1996). Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations. *Decision Sciences*, 27(2):343–363.
- Berge, M. E. and Hopperstad, C. A. (1993). Demand driven dispatch: A method for dynamic aircraft capacity assignment, models and algorithms. *Operations Research*, 41(1):153 – 168.
- Bertsekas, D. (1995). *Dynamic programming and optimal control*, volume 1. Athena Scientific, Belmont, Mass.
- Besbes, O. and Zeevi, A. (2006). Blind nonparametric revenue management: Asymptotic optimality of a joint learning and pricing method. Working paper, Decision, Risk, and Operations Division, Columbia University, New York.
- Binggeli, U. and Pompeo, L. (2005). The battle for Europe’s low-fare flyers. *The McKinsey Quarterly*, (August).

- Bitran, G. and Caldentey, R. (2003). An overview of pricing models for revenue management. *Manufacturing Service Operations Management*, 5(3):203–229.
- Bitran, G., Caldentey, R., and Mondschein, S. (1998). Coordinating clearance markdown sales of seasonal products in retail chains. *Operations Research*, 46(5):609–624.
- Bitran, G., Caldentey, R., and Vialz, R. (2005). Pricing policies for perishable products with demand substitution. Working paper.
- Bitran, G. and Wadhwa, H. (1996). A methodology for demand learning with an application to the optimal pricing of seasonal products. Working paper, Sloan School of Management, Massachusetts Institute of Technology, Cambridge.
- Bitran, G. R. and Mondschein, S. V. (1997). Periodic pricing of seasonal products in retailing. *Management Science*, 43(1):64–79.
- Bélangier, N., Desaulniers, G., Soumis, F., and Desrosiers, J. (2006). Periodic airline fleet assignment with time windows, spacing constraints, and time dependent revenues. *European Journal of Operational Research*, 175(3):1754 – 1766.
- Büning, Naeve, Trenkler, and Waldmann (2000). *Mathematik für Ökonomen im Hauptstudium*. Internationale Standardlehrbücher der Wirtschafts- und Sozialwissenschaften. Oldenbourg, München.
- Bodily, S. E. and Weatherford, L. R. (1995). Perishable-asset revenue management: Generic and multiple-price yield management with diversion. *Omega*, 23(2):173 – 185.
- Bouakiz, M. and Kebir, Y. (1995). Target-level criterion in Markov decision processes. *Journal of Optimization Theory and Applications*, 86(1):1–15.
- Boyd, E. A. and Bilegan, I. C. (2003). Revenue management and e-commerce. *Management Science*, 49(10):1363–1386.
- Brumelle, S. and Walczak, D. (2003). Dynamic airline revenue management with multiple semi-markov demand. *Operations Research*, 51(1):137.
- Brumelle, S. L. and McGill, J. I. (1993). Airline seat allocation with multiple nested fare classes. *Operations Research*, 41(1):127.
- Brumelle, S. L., McGill, J. I., Oum, T. H., Sawaki, K., and Tretheway, M. W. (1990). Allocation of airline seats between stochastically dependent demands. *Transportation Science*, 24(3):183.

- Brynjolfsson, E. and Smith, M. D. (2000). Frictionless Commerce? A Comparison of Internet and Conventional Retailers. *Management Science*, 46(4):563–585.
- Chatwin, R. E. (2000). Optimal dynamic pricing of perishable products with stochastic demand and a finite set of prices. *European Journal of Operational Research*, 125(1):149 – 174.
- Chi, Z. (1995). *Airline yield management in a dynamic network environment*. PhD thesis.
- Clark, N. (2009). Downturn opens doors for low-cost airlines. *The New York Times* (24. October).
- Clarke, L. W., Hane, C. A., Johnson, E. L., and Nemhauser, G. L. (1996). Maintenance and crew considerations in fleet assignment. *Transportation Science*, 30(3):249–260.
- Cole, S. (2008). *Applied transport economics : Policy, Management and Decision Making*. Kogan Page, London, 3. ed., repr. edition.
- Cooper, W. L., Homem-de Mello, T., and Kleywegt, A. J. (2006). Models of the spiral-down effect in revenue management. *Operations Research*, 54:968–987.
- Corts, K. S. (1998). Third-degree price discrimination in oligopoly: All-out competition and strategic commitment. *The RAND Journal of Economics*, 29(2):306–323.
- Cote, J.-P. (2010). The future of revenue management with a different perspective: Customer choice modeling, dependent demand forecasting. Conference presentation, IATA Commercial Strategy Symposium.
- Coy, P. (2000). The power of smart pricing: Companies are fine-tuning their price strategies—and it’s paying off. *Business Week* (10. April), pages 160–164.
- Cross, R. G. (1997). *Revenue Management: Hardcore Tactics for Market Domination*. Überreuter, Wien.
- Cui, G. and Choudhury, P. (2002). Marketplace diversity and cost-effective marketing strategies. *Journal of Consumer Marketing*, 19(1):54 – 73.
- Curry, R. E. (1990). Optimal airline seat allocation with fare classes nested by origins and destinations. *Transportation Science*, 24(3):193.
- Daskin, M. S. and Panayotopoulos, N. D. (1989). A Lagrangian Relaxation Approach to Assigning Aircraft to Routes in Hub and Spoke Networks. *Transportation Science*, 23(2):91–99.

- De Boer, S. (2004). The impact of dynamic capacity management on airline seat inventory control. *Journal of Revenue and Pricing Management*, 2(4):315–330.
- DER SPIEGEL (2003). BahnCard bringt wieder 50 Prozent Rabatt. *DER SPIEGEL* (2. July).
- Desaulniers, G., Desrosiers, J., Dumas, Y., Solomon, M. M., and Soumis, F. (1997). Daily aircraft routing and scheduling. *Management Science*, 43(6):841–855.
- Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271.
- Dunn, G. and Dunning-Mitchell, L. (2011). Low-cost carriers: growth expectations. *Airline Business*, 27(5):28 – 36.
- Dupuit, J. (1849). *De l'influence des peages sur l'utilite des voies de communication*, volume 2 of *Annales des ponts et chaussées*. Paris.
- Elmaghraby, W. and Keskinocak, P. (2003). Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science*, 49(10):1287–1309.
- Fehl, U. (1981). Preisdifferenzierung. In Albers, W. and Zottmann, A., editors, *Organisation bis Sozialhilfe und Sozialhilfegesetz*, volume 6 of *Handwörterbuch der Wirtschaftswissenschaften*, pages 160–172. Fischer.
- Feng, Y. and Gallego, G. (1995). Optimal starting times for end-of-season sales and optimal stopping times for promotional fares. *Management Science*, 41(8):1371–1391.
- Feng, Y. and Gallego, G. (2000). Perishable asset revenue management with markovian time dependent demand intensities. *Management Science*, 46(7):941–956.
- Feng, Y. and Xiao, B. (1999). Maximizing revenues of perishable assets with a risk factor. *Operations Research*, 47(2):337–341.
- Feng, Y. and Xiao, B. (2000a). A continuous-time yield management model with multiple prices and reversible price changes. *Management Science*, 46(5):644–657.
- Feng, Y. and Xiao, B. (2000b). Optimal policies of yield management with multiple predetermined prices. *Operations Research*, 48(2):332–343.
- FOCUS Online (2011). Fernbus soll Konkurrent von Bahn und Flugzeug werden. *FOCUS Online* (3. August).

- Frank, M., Friedemann, M., Mederer, M., and Schroeder, A. (2006). Airline revenue management: A simulation of dynamic capacity management. *Journal of Revenue and Pricing Management*, 5(1):62–71.
- Gallego, G. and Hu, M. (2009). Dynamic pricing of perishable assets under competition. Working paper, IEOR Department, Columbia University, New York.
- Gallego, G., Li, L., and Ratliff, R. (2009). Choice-based EMSR methods for single-leg revenue management with demand dependencies. *Journal of Revenue and Pricing Management*, 8(2/3):207 – 240.
- Gallego, G. and Talebian, M. (2010). Demand learning and dynamic pricing for multi-version products. *Journal of Revenue and Pricing Management*.
- Gallego, G. and Van Ryzin, G. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8):999–1020.
- Gallego, G. and Van Ryzin, G. (1997). A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research*, 45(1):24–41.
- Gao, C., Johnson, E., and Smith, B. (2009). Integrated airline fleet and crew robust planning. *Transportation Science*, 43:2–16.
- Gopalan, R. and Talluri, K. T. (1998). Mathematical models in airline schedule planning: A survey. *Annals of Operations Research*, 76:155–185.
- Grothklags, S. (2006). *Flottenzuweisung in der Flugplanung : Modelle, Komplexität und Lösungsverfahren*. PhD thesis.
- Gu, Z., Johnson, E. L., Nemhauser, G. L., and Yinhu, W. (1994). Some properties of the fleet assignment problem. *Operations Research Letters*, 15(2):59 – 71.
- Hane, C. A., Barnhart, C., Johnson, E. L., Marsten, R. E., Nemhauser, G. L., and Sigismondi, G. (1995). The fleet assignment problem: Solving a large-scale integer program. *Mathematical Programming*, 70:211–232.
- Haouari, M., Aissaoui, N., and Mansour, F. Z. (2009). Network flow-based approaches for integrated aircraft fleet and routing. *European Journal of Operational Research*, 193(2):591 – 599.
- Hinderer, K. (1970). *Foundations of non-stationary dynamic programming with discrete time parameter*. Lecture notes in operations research and mathematical systems. Springer, Berlin.

- Hopperstad, C. (2000). Modeling sell-up in pods. Conference presentation, AGIFORS Revenue Management and Distribution Study Group, New York.
- Hu, Q. and Yue, W. (2008). *Markov Decision Processes With Their Applications*, volume 14 of *Advances in Mechanics and Mathematics*. Springer.
- Isler, K. and Imhof, H. (2008). A game theoretic model for airline revenue management and competitive pricing. *Journal of Revenue and Pricing Management*, 7(4):384–396.
- Ivanov, S. (2006). Management of overbookings in the hotel industry - basic concepts and practical challenges. *Tourism Today*, 6:19–32.
- Jacobs, T., Hunt, E., and Korol, M. (2001). Scheduling and revenue management process integration: Benefits and hurdles. Conference presentation, AGIFORS Reservations and Yield Management Study Group, Bangkok.
- Jacobs, T. L., Smith, B. C., and Johnson, E. L. (2008). Incorporating network flow effects into the airline fleet assignment process. *Transportation Science*, 42(4):514–529.
- Kincaid, W. and Darling, D. (1963). An inventory pricing problem. *Journal of Mathematical Analysis and Applications*, 7(2):183–208.
- Lardner, D. (1850). *Railway Economy*. Harper and Brothers, New York.
- Lautenbacher, C. J. and Stidham, S. J. (1999). The underlying markov decision process in the single-leg airline yield-management problem. *Transportation Science*, 33(2):136.
- Lee, A. O. (1990). *Airline reservations forecasting—probabilistic and statistical models of the booking process*. PhD thesis.
- Lee, T. C. and Hersh, M. (1993). A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science*, 27(3):252.
- Levin, Y., McGill, J., and Nediak, M. (2008). Risk in revenue management and dynamic pricing. *Operations Research*, 56(2):326–343.
- Li, M. Z. and Zhuang, W. (2009). Risk-sensitive dynamic pricing for a single perishable product. *Operations Research Letters*, 37(5):327 – 332.
- Li, M. Z. F. and Oum, T. H. (2002). A note on the single leg, multifare seat allocation problem. *Transportation Science*, 36(3):349.
- Lin, K. Y. (2006). Dynamic pricing with real-time demand learning. *European Journal of Operational Research*, 174(1):522 – 538.

- Lin, K. Y. and Sibdari, S. Y. (2009). Dynamic price competition with discrete customer choices. *European Journal of Operational Research*, 197(3):969–980.
- Littlewood, K. (1972). Forecasting and control of passengers. Conference proceedings, AGIFORS 12th Annual Symposium.
- Littlewood, K. (2005). Forecasting and control of passenger bookings. *Journal of Revenue and Pricing Management*, 4(2):111 – 123.
- Lohatepanont, M. (2002). *Airline fleet assignment and schedule design : integrated models and algorithms*. PhD thesis.
- Lohatepanont, M. and Barnhart, C. (2004). Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment. *Transportation Science*, 38(1):19–32.
- Lovell, C. A. K. (1978). A two-stage model of discriminatory pricing. *The Canadian Journal of Economics*, 11(1):56–68.
- Lufthansa (2011). Lufthansa Investor Relations: Kennzahlen des Lufthansa Konzerns im Überblick. <http://investor-relations.lufthansa.com/de/fakten-zum-unternehmen/kennzahlen/konzern.html> (20.01.11).
- Lufthansa Systems (2010). Lufthansa Systems presents innovative Commercial Platform. Press release.
- Maglaras, C. and Meissner, J. (2006). Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing Service Operations Management*, 8(2):136–148.
- Martins, E. Q. V. and Pascoal, M. M. B. (2003). A new implementation of yen’s ranking loopless paths algorithm. *4OR: A Quarterly Journal of Operations Research*, 1:121–133.
- McGill, J. and Van Ryzin, G. (1999). Revenue management: Research overview and prospects. *Transportation science*, 33(2):233.
- Mishra, S., Ratliff, R., and Vinod, B. (2005). New generation in demand forecasting. Conference presentation, AGIFORS 45th Annual Symposium.
- Nickel, S., Stein, O., and Waldmann, K.-H. (2011). *Operations research*. Springer, Berlin.
- Oba, K. (2007). Supply side revenue management: ANA’s experience with an optimized equipment swap system. Conference presentation, AGIFORS Scheduling and Strategic Planning Study Group, Tokyo.

- Oneworld (2011). Designierte Mitglieder. <http://de.oneworld.com/member-airlines/members-elect/> (06.12.2011).
- Ott, A. E. (1959). Zum Problem der Preisdifferenzierung. *Swiss Journal of Economics and Statistics*, 95(4):465–477.
- Papadakos, N. (2009). Integrated airline scheduling. *Computers and Operations Research*, 36:176–195.
- Pastor, J. (1999). Demand driven dispatch: A stochastic approach. Conference presentation, AGIFORS 39th Annual Symposium, New Orleans.
- Pfeifer, P. E. (1989). The airline discount fare allocation problem. *Decision Sciences*, 20(1):149 – 157.
- Phillips, R. (2005). *Pricing and revenue optimization*. Stanford Business Books. Stanford University Press, Stanford.
- Pigou, A. C. (1920). *The Economics of Welfare*. Macmillan and Co., London.
- Pölt, S. (1999). Back to the roots: new results on leg optimization. Conference presentation, AGIFORS Reservations and Yield Management Study Group, London.
- Pölt, S. (2002). Revenue management tutorial. Conference presentation, AGIFORS Reservations and Yield Management Study Group, Berlin.
- Powell, W. B. (2007). *Approximate dynamic programming : solving the curses of dimensionality*. Wiley series in probability and statistics. Wiley-Interscience, Hoboken.
- Puterman, M. L. (1994). *Markov decision processes*. Wiley, New York.
- Rajan, A., Rakesh, and Steinberg, R. (1992). Dynamic pricing and ordering decisions by a monopolist. *Management Science*, 38(2):240–262.
- Rexing, B., Barnhart, C., Kniker, T., Jarrah, A., and Krishnamurthy, N. (2000). Airline fleet assignment with time windows. *Transportation Science*, 34(1):1–20.
- Robinson, L. W. (1995). Optimal and approximate control policies for airline booking with sequential nonmonotonic fare. *Operations Research*, 43(2):252.
- Rosenberg, M. (2010). The relationship between revenue management and sales. do they have the same targets and objectives? Conference presentation, IATA Commercial Strategy Symposium.
- Rothstein, M. (1971). Airline overbooking: The state of the art. *Journal of Transport Economics and Policy*, 5(1):96–99.

- Ryanair (2011a). Über Ryanair - Ryanair Fleet. <http://www.ryanair.com/de/about/fleet> (5.12.2011).
- Ryanair (2011b). Investor relations: Passenger traffic 2002/2011. <http://www.ryanair.com/de/investor/traffic-figures> (20.01.11).
- Sandhu, R. and Klabjan, D. (2007). Integrated airline fleet and crew-pairing decisions. *Operations Research*, 55(3):439–456.
- Sasse, J. and Beleliev, S. (2010). The Integrated Commercial Platform (ICP) and market sensitive Revenue Management as examples for innovative products. Conference presentation, IATA Commercial Strategy Symposium.
- Schäl, M. (1975). Conditions for optimality in dynamic programming and for the limit of n-stage optimal policies to be optimal. *Probability Theory and Related Fields*, 32:179–196.
- Schubert, C. and Mrusek, K. (2009). Konkurrenz für die Deutsche Bahn. *Frankfurter Allgemeine Zeitung* (30. October).
- Schultes, D. (2008). *Route planning in road networks*. PhD thesis.
- Shugan, S. M. (2002). Editorial: Marketing science, models, monopoly models, and why we need them. *Marketing Science*, 21(3):223–228.
- Smith, B. C., Leimkuhler, J. F., and Darrow, R. M. (1992). Yield Management at American Airlines. *Interfaces*, 22(1):8–31.
- Smith, S. A. and Achabal, D. D. (1998). Clearance pricing and inventory policies for retail chains. *Management Science*, 44(3):285–300.
- Stole, L. A. (2007). Chapter 34: Price discrimination and competition. In Armstrong, M. and Porter, R., editors, *Handbook of Industrial Organization*, volume 3 of *Handbooks in Economics*, pages 2221 – 2299. Elsevier.
- Subramanian, J., Stidham Jr., S., and Lautenbacher, C. J. (1999). Airline yield management with overbooking, cancellations, and no-shows. *Transportation Science*, 33(2):147.
- Subramanian, R., Scheff, Richard P., J., Quillinan, J. D., Wiper, D. S., and Marsten, R. E. (1994). Coldstart: Fleet Assignment at Delta Air Lines. *Interfaces*, 24(1):104–120.
- Swan, W. M. (2002). Airline demand distributions: passenger revenue management and spill. *Transportation Research Part E: Logistics and Transportation Review*, 38(3-4):253 – 263.

- Talluri, K. (1996). Swapping applications in a daily airline fleet assignment. *Transportation Science*, 30(3):237–248.
- Talluri, K. and Martínez de Albéniz, V. (2010). Dynamic price competition with fixed capacities. Working paper.
- Talluri, K. and Van Ryzin, G. (1998). An analysis of bid-price controls for network revenue management. *Management Science*, 44(11):1577–1593.
- Talluri, K. and Van Ryzin, G. (2004a). Revenue management under a general discrete choice model of consumer behavior. *Management Science*, 50(1):15 – 33.
- Talluri, K. and Van Ryzin, G. (2004b). *The theory and practice of revenue management*. International series in operations research and management science ; 68. Kluwer Academic, Dordrecht.
- Ulph, D. and Vulkan, N. (2000). Electronic commerce and competitive first-degree price discrimination. Working paper.
- Ulph, D. and Vulkan, N. (2001). E-commerce, mass customisation and price discrimination. Working paper.
- Van Ryzin, G. and McGill, J. (2000). Revenue management without forecasting of optimization: An adaptive algorithm for determining airline seat protection levels. *Management Science*, 46(6):760–775.
- Van Slyke, R. and Young, Y. (2000). Finite horizon stochastic knapsacks with applications to yield management. *Operations Research*, 48(1):155.
- Varian, H. (2000). Versioning information goods. In *Internet publishing and beyond: the economics of digital information and intellectual property*, A publication of the Harvard Information Infrastructure Project, pages 190 – 201. MIT Press.
- Varian, H. R. (1985). Price discrimination and social welfare. *The American Economic Review*, 75(4):870–875.
- Varian, H. R. (1989). Price discrimination. In Schmalensee, R. and Willig, R., editors, *Handbook of Industrial Organization*, volume 1 of *Handbooks in Economics*, pages 597 – 654. Elsevier.
- Varian, H. R. (2004). Competition and market power. In *The Economics of Information Technology : An Introduction*, Raffaele Mattioli Lectures, pages 1 – 48. Cambridge University Press, Cambridge, 1. publ. edition.

- Vinod, B. (2009). Distribution and revenue management: Origins and value proposition. *Journal of Revenue and Pricing Management*, 8(2/3):117 – 133.
- Walczak, D. (2006). Modeling high demand variance in dynamic programming. *Journal of Revenue and Pricing Management*, 5(2):94 – 101.
- Walczak, D., Mardan, S., and Kallesen, R. (2010). Customer choice, fare adjustments and the marginal expected revenue data transformation: A note on using old yield management techniques in the brave new world of pricing. *Journal of Revenue and Pricing Management*, 9(1/2):94 – 109.
- Wang, X. and Meng, Q. (2008). Continuous-time dynamic network yield management with demand driven dispatch in the airline industry. *Transportation Research Part E: Logistics and Transportation Review*, 44(6):1052 – 1073.
- Wang, X. and Regan, A. (2006). Dynamic yield management when aircraft assignments are subject to swap. *Transportation Research Part B: Methodological*, 40(7):563 – 576.
- Weatherford, L., Kimes, S., and Scott, D. (2001). Forecasting for hotel revenue management: Testing aggregation against disaggregation. *The Cornell Hotel and Restaurant Administration Quarterly*, 42(4):53–64.
- Weatherford, L. and Pölt, S. (2002). Better unconstraining of airline demand data in revenue management systems for improved forecast accuracy and greater revenues. *Journal of Revenue and Pricing Management*, 1(3):234–254.
- Weatherford, L. R. (2004). EMSR versus EMSU: Revenue or utility? *Journal of Revenue and Pricing Management*, 3(3):277 – 284.
- Weatherford, L. R. and Kimes, S. E. (2003). A comparison of forecasting methods for hotel revenue management. *International Journal of Forecasting*, 19(3):401 – 415.
- Weber, K. (1956). Zur frage der preisdifferenzierung bei einproduktunternehmen. *Schweizerische Zeitschrift für Volkswirtschaft und Statistik*, pages 17–30.
- White, D. J. (1993). *Markov decision processes*. Wiley, Chichester.
- Wiggershaus, J. (2008). *Ein wertorientiertes Auftragsannahmeverfahren für das Revenue Management in der Sachgüterindustrie*. PhD thesis, Hamburg.
- Williamson, E. L. (1992). *Airline Network Seat Inventory Control: Methodologies and Revenue Impacts*. PhD thesis.

- Wollmer, R. D. (1992). An airline seat management model for a single leg routine when lower fare classes book first. *Operations Research*, 40(1):26.
- Zabel, E. (1972). Multiperiod monopoly under uncertainty. *Journal of Economic Theory*, 5(3):524 – 536.
- Zeni, R. (2001a). *Improved Forecast Accuracy in Airline Revenue Management by Unconstraining Demand Estimates from Censored Data*. PhD thesis.
- Zeni, R. (2001b). Improving forecast accuracy by unconstraining censored demand data. Conference presentation, AGIFORS Reservations and Yield Management Study Group.
- Zhang, D. and Cooper, W. L. (2005). Revenue management for parallel flights with customer-choice behavior. *Operations Research*, 53(3):415–431.
- Zhao, W., Morley, S., and Srinivasan, B. (2007). Flexible capacity at united airlines. Conference presentation, AGIFORS Scheduling and Strategic Planning Study Group, Tokyo.
- Zhao, W. and Zheng, Y.-S. (2000). Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management Science*, 46(3):375–388.
- Zhechev, V. S. and Todorov, A. (2010). The impact of overbooking on hotels operation management. Working paper.