D I S S E R T A T I O N

Uncertainties and Risks in Airline Revenue Management – Capacity Uncertainty as a Showcase

submitted to the
School of Business and Economics
of Freie Universität Berlin
in Candidacy for the Degree of
doctor rerum politicarum
(Dr. rer. pol.)

by
Dipl.-Volksw. Daniel Kadatz

Berlin, 2017

Referees

1. Prof. Dr. Natalia Kliewer, Freie Universität Berlin
2. Prof. Dr. Catherine Cleophas, RWTH Aachen University

Date of Disputation: October 6, 2017

Publications

- Kadatz, D., Kliewer, N., Cleophas, C. (2017). Anticipating Capacity Updates in Airline Revenue Management: The Value of Information. Submitted to *Manufacturing & Service Operations Management*.
- Kadatz, D., Kliewer, N., Cleophas, C. (2017). Exogenous Capacity Changes in Airline Revenue Management: Quantifying the Value of Information. To be published in *OR 2016 Proceedings*.
- Büsing, C., Kadatz, D., Cleophas, C. (2017). Capacity Uncertainty in Airline Revenue Management: Models, Algorithms and Computations. In second revision at *Transportation Science*.
- Cleophas, C., Kadatz, D., Vock, S. (2017). Resilient Revenue Management: A Literature Survey of Recent Theoretical Advances. To be published in *Journal of Revenue and Pricing Management*.

Acknowledgment

I thank Deutsche Lufthansa AG for access to empirical data and domain expertise.

Declaration of Authorship

Except where reference is made in the text of this dissertation, this dissertation contains no material published elsewhere or extracted in whole or in part from a dissertation presented by me for another degree or diploma. No other person's work has been used without due acknowledgment in the main text of the dissertation. This dissertation has not been submitted for the award of any other degree or diploma in any other tertiary institution.

Signed: ______________________________

Date: ______________________________

Danksagung

Auch wenn die hier vorliegende Dissertation nur meinen Namen trägt sind es doch zahlreiche Menschen, die das Ergebnis entscheidend mitgeprägt haben und denen ich zu großem Dank verpflichtet bin.

Einen herzlichen Dank an Natalia Kliewer. Durch deine frühe Förderung habe ich die Freude an der Wirtschaftsinformatik bereits während meines Studiums für mich entdecken können. Unsere gemeinsamen Besprechungstermine haben mir stets dabei geholfen den richtigen Rahmen meiner Arbeit zu ermitteln.

Ganz herzlich danke ich auch Catherine Cleophas. Du hast vom Anfang bis zum Ende an mich geglaubt. Danke für dein Vertrauen, dass du mir als dein erster Mitarbeiter ausgesprochen hast und für jedes Gespräch, jedes Telefonat und jede Nachricht, die mir geholfen haben den richtigen Weg für diese Arbeit einzuschlagen. Alles andere als selbstverständlich ist auch, dass du immer ein offenes Ohr für mich hattest, selbst dann wenn es um Probleme jenseits der Arbeit ging.

Meinen ehemaligen Kolleginnen und Kollegen der Freien Universität Berlin: Bastian Amberg, Lena Wolbeck, Clemens Wickboldt, Charlotte Köhler, Jan Ehmke, Daniel Fürstenau, Nils Olsen, Josephine Reuer, Maximilian Adam, Armando Sandoval und Felix Becker danke ich ebenfalls für jede Diskussion und jedes Gespräch – auch in der Mensa oder über den Feierabend hinaus. Insbesondere möchte ich aber Lucian Ionescu danken. Das Lehren und Lernen mit dir, war immer eine Freude, egal ob innerhalb und außerhalb der Uni.

Ebenfalls zu Dank verpflichtet bin ich den nun größtenteils ehemaligen Lufthanseaten. Der fachliche Austausch und die gemeinsame Arbeit mit euch, Max Gerlach, Benedikt Zimmermann, Philipp Bartke, Martin Friedemann und Steffen Enz war stets produktiv und oft auch hochgradig unterhaltsam. In besonderem Maße möchte ich mich aber bei Sebastian Vock und Johannes Jörg bedanken. Fast immer kam es mir so vor als säßen wir im selben Büro, auch wenn dies nur selten der Fall war. Euer ehrliches Feedback, vor allem auch zu den ersten Fassungen der Arbeit, haben mir enorm geholfen. Als erster Ansprechpartner in Sachen Forschungskooperation bin ich auch Stefan Pölt sehr dankbar. Deine Erfahrung aus der Praxis und Erläuterungen zu den bereitgestellten Daten haben die Abgrenzung der Arbeit deutlich erleichtert.

Ein Dankeschön auch an meine ehemaligen Kolleginnen und Kollegen der RWTH Aachen: Beatrix Lindner, Janina Knepper, Claudia Schütze, Marc Werner, Magdalena Lang, Michael Bastubbe, Jonas Witt, Sarah Kirchner, Annika Thome, Markus Kruber und Christian Purchert. Ihr habt mich in Aachen in kürzester Zeit geradezu familiär aufgenommen und habt mir gezeigt, wie gut lehrstuhlübergreifendes Arbeiten funktioniert. Besonders danken möchte ich aber Felix Willamowski. Nicht nur deine nächtlichen Korrekturleseeinheiten, sondern auch jedes Gespräch zu jeder erdenklichen Tageszeit, haben mir sehr geholfen.

Großer Dank geht auch an meine Familie und Freunde, allen voran an meine Eltern Petra und Achim Kadatz. Ihr standet in meinem Leben immer hinter mir. Auf euch konnte ich mich jederzeit verlassen. Julian Soick und Martin Brehme, ihr seid von klein auf meine Brüder im Geiste und habt mir immer das Gefühl gegeben, dass ich das schaffe – danke dafür.

Zuletzt möchte ich noch einem ganz besonderen Menschen danken. Es begann mit einem Feierabendbier unter Kollegen, nun bist du meine beste Freundin und große Liebe. Laura Hellmann, danke für deine unermessliche Unterstützung und für jedes Opfer, dass auch du erbracht hast. Ohne dich wäre diese Arbeit niemals möglich gewesen.

Abstract

Existing uncertainties and risks in revenue management have motivated a vast body of research. The work in hand identifies uncertainties and risks in airline revenue management. It highlights mature research areas and points out further potentials, especially the problem of uncertain capacities.

Capacity uncertainty contradicts the common revenue management assumption of fixed capacities. However, airlines have to change aircrafts over time, which can result in an altered number of sellable seats. This capacity uncertainty is almost completely neglected in existing research as well as in practice' planning. This neglect motivates the work's major contribution of systematically anticipating capacity uncertainty in revenue management. Therefore, it introduces a scenario-based optimization model considering unexpected capacity changes in advance. A simulation system with a focus on revenue optimization enables a set of computational studies that analyze the influence of capacity uncertainty on revenue management's performance.

Empirically calibrated computational studies compare possible strategies for coping with capacity uncertainty. Here, several strategies take up the challenge of limited information. Results show the potential of anticipating capacity changes in revenue management: the most beneficial strategy uses full information and increases expected revenue on average by 2.47 percent points. But, even a strategy considering limited information on capacity changes can result in revenue improvements in contrast to not anticipating capacity changes. Additional studies analyze those conditions that render anticipating capacity uncertainty to be most beneficial. A systematic anticipation is especially effective when capacity changes occur frequently and late, when capacities differ strongly and when demand arrives early. The studies also test strategies' robustness when distorting forecasted capacities and demand. Here, the most beneficial strategy performs comparatively robust against tested distortions.

Zusammenfassung

Ein Großteil der Forschung im Revenue Management ist durch bestehende Unsicherheiten und Risiken motiviert. Diese Arbeit identifiziert Unsicherheiten und Risiken im Airline Revenue Management und gibt einen Überblick über bestehende Forschungsarbeiten. Während einige Forschungsfelder bereits sehr ausführlich bearbeitet wurden, spricht die Arbeit auch potentielle Forschungslücken an und widmet sich insbesondere dem Problem der Kapazitätsunsicherheit.

Im Gegensatz zur üblichen Annahme, dass Kapazitäten im Revenue Management als fixe Größe betrachtet werden können, zeichnet sich in der Airline-Praxis ein anderes Bild ab. Tatsächlich müssen Airlines ihre Flugzeuge im Buchungsverlauf häufig wechseln, was zu einer veränderten Kapazität, also der Anzahl verkaufbarer Sitzplätze, führen kann. Diese Kapazitätsunsicherheit wurde jedoch bisher weder in der Theorie noch in der Praxis ausreichend berücksichtigt. Die vorliegende Arbeit setzt sich somit eine systematische Antizipation von Kapazitätsunsicherheit zum Ziel. Hierfür wird ein neues szenario-basiertes Revenue-Management-Optimierungsmodell eingeführt, welches unerwartete Kapazitätswechsel bereits im Voraus berücksichtigt. Ein Simulationssystem mit dem Schwerpunkt der Ertragsoptimierung bildet die Grundlage für verschiedene numerische Studien. Diese wurden auf empirischen Airline-Daten kalibriert und ermöglichen den Einfluss von Kapazitätsunsicherheit auf das Ergebnis im Revenue Management zu untersuchen.

In den Studien werden mögliche Strategien im Umgang mit Kapazitätsunsicherheit erprobt und miteinander verglichen. Hierunter fallen auch einige Strategien, die sich der Herausforderung unvollständiger Informationen stellen. Die Ergebnisse zeigen das Potential für künftige Ertragssteigerungen sofern Kapazitätswechsel im Revenue Management antizipiert werden: Die erfolgreichste Strategie nutzt vollständige Informationen und steigert den erwarteten Ertrag um durchschnittlich 2,47 Prozentpunkte. Doch selbst mit unvollständigen Informationen kann der erwartete Ertrag gesteigert werden – im Gegensatz zu einer Vernachlässigung möglicher Kapazitätswechsel. Wei-tere Studien analysieren die Bedingungen unter denen sich eine Berücksichtigung von Kapazitätsunsicherheit am meisten lohnt. Eine systematische Antizipation ist besonders wirksam, wenn Kapazitätswechsel häufig und spät stattfinden, wenn sich Kapazitäten stark unterscheiden und wenn die Nachfrage früh eintrifft. Ebenfalls wird die Robustheit der Strategien gegenüber Prognosefehlern, in Bezug auf Nachfrage und Kapazität, getestet. Die erfolgreichste Strategie ist vergleichsweise robust gegenüber allen getesteten Prognosefehlern.

Contents

List of Tables

List of Figures

Part I

Uncertainties and Risks in Airline Revenue Management

1 Introduction

"Without risks, no company would be able to achieve anything or make a profit" (Lancaster, 2003, p. 158). This holds for revenue management (RM) and in particular for airline revenue management (ARM). While the success of RM is undisputed over decades, the continuously increasing complexity of RM practice requires for approaches coping with arising uncertainties and risks. Based on a focus on airlines, this work first introduces a classification of major uncertainties and risks in RM in Part I, before Part II and III address the specific problem of capacity uncertainty.

In this chapter, Section 1.1 provides the background of RM by summarizing its emergence in the airline domain. Subsequent, Section 1.2 presents a distinction of uncertainties and risks in a business context, before Section 1.3 motivates this work's research objectives.

1.1 Emergence of Airline Revenue Management

The Airline Deregulation Act of 1978 marks the birth of RM as it allowed airlines, for the first time in history, to choose and offer their own prices (B. C. Smith, Leimkuhler & Darrow, 1992). Since then, RM has evolved to one of the most crucial instrument for making profit. This especially holds in the airline sector, but also other industries realized the benefits of RM systems such as car rental (Geraghty & Johnson, 1997), hotels (Choi & Mattila, 2004), retail (Vinod, 2005), casinos (Hendler & Hendler, 2004), cruise lines (Ladany & Arbel, 1991), advertising (Kimms & Müller-Bungart, 2007) and many more. These industries and ongoing developments are, i. a., reviewed by McGill and van Ryzin (1999) and Chiang, Chen and Xu (2007). A comprehensive overview of mathematical models and methods of RM is provided by Talluri and van Ryzin (2004b).

The classic RM challenge is to avoid *spill* and *spoilage*. Spill is the result of selling too much too early, while spoilage is the result of selling too less. An RM system's task is to find the revenue maximizing optimal balance. Here, demand variation, as a natural circumstance, constitutes the biggest problem prohibiting a simple solution approach.

As one approach to the problem, quantity-based RM establishes *fare classes*, provides them with different prices – called fares – and controls each fare class' availability over the booking horizon. The work by Littlewood (1972) is the first path-breaking approach for airlines to systematically maximizing revenue. The author mathematically formulates an intuitive rule: sell tickets in the cheaper of two fare classes as long as the expected marginal utility exceeds the fare of the more expensive fare class. Belobaba (1987) extends this approach to more fare classes, arriving at the today well known expected marginal seat revenue (EMSR) heuristic. The basic RM model was further extended to network effects (Bertsimas & Popescu, 2003; Williamson, 1992) and to incorporate dependent demand (Talluri, van Ryzin, Karaesmen & Vulcano, 2008; Weatherford & Ratliff, 2010). Also the influence of alliances (Gerlach, 2013; Gerlach, Cleophas & Kliewer, 2013; Hu, Caldentey & Vulcano, 2013; Wright, 2014) and competition (W. L. Cooper, Homem-de Mello & Kleywegt, 2015; Grauberger & Kimms, 2016; Zimmermann, 2014) on ARM has lately been considered.

In contrast, a price-based RM approach such as dynamic pricing[1] proposes to continuously adjust fares over time, so that no static fares are required (Şen, 2013). However, due to technical and organizational obstacles, up to now, dynamic pricing is not common in the airline industry (Isler & D'Souza, 2009; Pölt, 2010). However, as both concepts are such widely different, this work confines itself to a quantity-based RM approach representing the current aviation industry standard (Talluri & van Ryzin, 2004b, p. 33). Nevertheless, in particular cases, this work occasionally mentions selected contributions from the field of dynamic pricing if existing quantity-based RM literature is scarce.

Recent research increasingly focuses on the challenge of planning under uncertainty. Cleophas, Kadatz and Vock (2017) is the latest literature overview on uncertainties in RM. The authors review several contributions, primarily with a focus on ARM, by allocating contributions' approaches to the concept of persistence or flexibility. Following this, the next section discusses different theoretical interpretations of the terms *risk* and *uncertainty* and clarifies their usage for the remainder of this work.

1.2 Uncertainties and Risks

Uncertainty and risk are two terms that are often colloquially used. In theory, some contributions omit differentiating risks and uncertainties in the context of RM by using them as synonyms (e. g. Huang and Chang (2009)). This work sticks to a distinction based on the ideas of Knight (1921), which has later been adopted and specified by Runde (1998).

Historically, Knight (1921) introduced a trichotomy, consisting of the so called *a priori probabilities*, *statistical probabilities* and *estimates*. Runde (1998) labeled the probabilities as different forms of risk and estimates as uncertainty. When adding unawareness (Roy, 2010), a list differentiating risk, uncertainty and unawareness involves the following:

Risk in terms of a priori probabilities is given when an outcome probability can be perfectly quantified in advance. Tossing a coin is a picture book example.

Risk in terms of statistical probabilities can be quantified if several realizations of an outcome can be observed. For example, observing the outcome of multiple dice rolls, although the dice's number of sides is unknown, allows assuming a statistical probability that helps predicting future outcomes.

Uncertainty describes knowledge of possible outcomes combined with a lack of quantitative information to estimate outcomes' probabilites; e. g., assuming the previous dice example but leave off the information on previous rolls' outcome.

Unawareness describes a lack of knowledge about possible disturbances or changes. Playing a card game, without being aware of the existence of a joker card, is an example.

However, risks with a priori known probabilities are rarely or almost never present in a business environment. On the other side, unawareness cannot be observed as decision makers are not aware of its presence. Thus, this work focuses on the most present obstacles in RM, tracing back to risks

[1] Over the last decades, a lot of contributions addressed the topic of dynamic pricing in ARM, such as Gallego and van Ryzin (1994), Gallego and van Ryzin (1997), Talluri and van Ryzin (1998), Bitran and Caldentey (2003), Elmaghraby and Keskinocak (2003) and Maglaras and Meissner (2006).

with statistical probabilities – just *risks* in the following – and uncertainties. As long as none of the previous terms have been used, this work speaks of *disturbances*.

In this context, *ignorance* is a possible response to risk or uncertainty. Here, although decision makers know that outcomes may change driven by a disturbance, they choose to ignore it. Beyond ignorance, there are further concepts responding to the challenge of planning under uncertainty and risk such as robustness, anti-fragility and resilience.

Robust solutions promise to perform well for a specific number of scenarios. These solutions, however, have to accept the costs of robustness, quantified as the gap to the performance achievable by non-robust solutions optimized under ideal conditions (Bertsimas & Sim, 2004). More recently, Gorgeon (2015) proposed anti-fragile information systems, which perform well when facing unforeseen disturbances. Earlier, Taleb (2012) stated the possibility of even benefiting from these disturbances.

The term resilience is used in research areas such as psychology, organization or infrastructure planning (Park, Sharman & Rao, 2015) and also describes what RM systems strive to implement: the ability to *bounce back* after disturbances. Cleophas et al. (2017) provide a literature overview of RM and focus on categorizing different approaches to disturbances in RM into persistent or flexible solutions. Examples of contributions considering a distinction and explanation of the concepts of flexibility, adaptability, persistence and stability are: Marschak and Nelson (1962), Pye (1978), Mandelbaum and Buzacott (1990), Lardeux, Goyons and Robelin (2010), Roy (2010) and Wieland and Marcus Wallenburg (2012). The authors distinction is justified by Carvalho, Duarte and Cruz Machado (2011), who emphasize the two aims of resilience: to persist in the face of a disturbance and to recover desirable system states after a disturbance.

In contrast, this work concentrates on identifying the underlying disturbances as a risk or an uncertainty. For that reason, the following chapters discuss many contributions also included in Cleophas et al. (2017), but with a focus on problems' origin instead on categorizing solution approaches. If a risk is identified, researches' effort is to better estimate its probability or to improve its existing consideration in the RM process. Hereby, they try reducing an already tracked and considered risk. In contrast, if identifying an uncertainty, approaches usually try to transform this uncertainty to a risk by first estimating its probabilities before considering it in the RM process. The following section introduces this work's motivation and its research objective.

1.3 Motivation and Outline

This work has two key drivers: The first is to better understand the existence and hazards of uncertainties and risks in ARM. The second motivation is to highlight the particular hazard of capacity uncertainty, serving as a showcase for uncertainty in ARM.

This work is structured into three parts. Part I addresses uncertainties and risks within and beyond ARM systems. It pursues the target of a structured overview providing a distinction of uncertainties and risks in ARM systems. Part II focuses on transferring capacity uncertainty to a risk in an RM optimization problem. Considering capacity uncertainty in RM still constitutes a research gap. Part III investigates the effects of capacity uncertainty and compares different approaches for its consideration based on various information levels and problem setups.

First, Chapter 2 introduces the foundations of RM. The primarily objective of Part I is, however, to introduce existing disturbances ARM systems have to face and assigning them to an uncertainty or risk. The lack of a structured overview on uncertainties and risks in ARM motivates these efforts. First, Chapter 3 addresses uncertainties and risks within ARM, then Chapter 4 addresses uncertainties and risks beyond ARM. Subsequently, this overview categorizes ARM's disturbances first by their origin in the process and second by their solution approach. The overview also distinguishes if solution approaches try transforming an uncertainty to a risk or attempt reducing the hazard of an existing risk. At the end, Chapter 4 introduces capacity uncertainty in ARM; the lack of contributions treating this problem constitutes a research gap, explicated in Chapter 5. The efforts of this work's following two parts then focus on considering the problem of capacity uncertainty in RM systems.

The research objective of Part II is to transform capacity uncertainty to a risk as a showcase of such a transition. The attempt of considering capacity uncertainty in RM is motivated by observing that capacities can change in practice although assuming them to be fixed. Chapter 6 analyzes existing empirical data that helps quantifying these capacity changes. Identifying the critical characteristics of capacity changes allows Chapter 7 to model and thus to consider capacity uncertainty in an RM optimization problem. The chapter also introduces several approaches as response. Then, Chapter 8 establishes a simulation framework that recreates the problem of capacity uncertainty in an RM system. Concluding, this part theoretically transforms capacity uncertainty to a risk.

As its main objective, Part III quantifies capacity uncertainty's influence on RM. Here, Chapter 9 first validates the simulation system before Chapter 10 tests approaches differently responding to capacity changes based on empirical data. This also includes computational studies analyzing capacity uncertainty characteristics and the problem parameters' impact on results. Furthermore, Chapter 11 considers revenue result's dependency on forecasting accuracy. Summarizing, this part presents the numerical benefits of transforming capacity uncertainty into a risk, concluded by Chapter 12.

2 Revenue Management Foundations

The primary objective of *quantity-based RM* is to maximize expected revenue by controlling the availability of similar, but differently priced, services or products. Its challenge is to avoid both selling too much too early and selling too less. Finding the optimal balance is hard, as inherent demand variation prohibits a simple solution.

Section 1.1 already provided a first sketch of ARM's development and history, its major challenges and mentioned some industries also applying the concept of RM. This chapter addresses RM in more detail. First, Section 2.1 presents and explicates conditions for successfully applying RM to a particular problem or industry. Most industries applying a quantity-based RM approach establish a system similar to the one depict in Section 2.2. Here, an RM process view common in the airline domain serves as a road map for the two following chapters addressing uncertainties and risks in RM. Last, Section 2.4 introduces disturbances in RM as treated in this work.

2.1 Conditions for Applying Revenue Management

As stated in Section 1.1, several industries successfully adapted RM as an instrument for revenue maximization (Chiang et al., 2007). When applying the concept of RM, most industries or business cases have the following conditions in common. These conditions are, e. g., discussed by Talluri and van Ryzin (2004b, pp. 13).

Demand variation is a natural condition in RM, as overall demand results from individual and independent customer decisions. Past and future demand often differ strongly, for example, driven by seasonality or trends. Thus, e. g., customer arrival times and the amount of customers per fare class are uncertain numbers. The better demand can be anticipated, the easier is RM's task. An airline, for instance, will never face a certain future demand; a residual fluctuation will always exist.

Customer heterogeneity allows RM to segment customers based on their preferences. By this segmentation, airlines can charge customers different fares, trying to exploit their customers' willingness to pay. Common assumptions are that business passengers, compared to leisure passengers, book closer to departure and rarely travel on weekends.

Price differentiation must be possible, due to the previously addressed customer segmentation, but customers must also agree or at least accept this disparity. On the other side, customers should not perceive prices as a signal of quality; e. g., offering a cheap price does not reflect a worse transportation's quality. In RM, airlines use *fare classes* allowing them to offer different prices for the same service.

Cost structures usually include proportionally high fixed costs such as capital costs, wages and fuel, with simultaneously low variable costs. This ensures companies applying RM to sell a service or product even for comparatively low prices and still making a profit. An additional customer on an aircraft produces small additional costs to an airline but provides the potential of selling a high-priced ticket.

Inflexible production prohibits companies from spontaneously adjusting their produced or offered capacities as a response to demand variation. In ARM, the capacity of a flight is assumed to be fixed over the booking horizon.

Perishable products or services must be offered, otherwise a company could store them. A seat on a flight can, obviously, never be sold after the flight's realization.

Information systems infrastructure must be capable of processing enormous masses of data that are needed, i. a., when forecasting demand or optimizing revenue. Also the system's accuracy is a necessity, free from errors or defects.

If these conditions pertain, a business can apply RM as a valuable instrument. The next section introduces a process view representing a common RM system.

2.2 Revenue Management Process View

As the following chapters introduce disturbances ordered by their location in the RM process, Figure 2.1 depicts such a traditional RM process, based on Talluri and van Ryzin (2004b, Chapter 1). The major three stages of an RM system are represented by rectangles: forecast demand, optimize revenue and accept reservations.

As the next chapter addresses uncertainties and risks within RM systems and the chapter after next those beyond RM systems, the figure highlights system elements in bright blue. In contrast, input on fares, capacity, demand and market lie beyond an RM system's sphere of influence.

RM systems have to store vast amounts of historical data. In Figure 2.1, barrels represent these data. Among others, RM systems record historical customer requests including diverse information such as request time and requested fare class, the past inventory controlled availability of fare classes and also previous bookings. With the help of these data, methods for predicting future demand combine gathered information and result in a demand forecast. Usually, a forecast's accuracy increases by learning effects that take place when analyzing historical data over time. The stage of *forecasting demand* involves two parts: estimation, the descriptive part and forecasting, the predictive part. In ARM, estimation usually calibrates forecasts' parameters while forecasting uses these values for predicting future demand (Talluri & van Ryzin, 2004b, p. 419).

The *forecast* is a prediction on future demand. Among others, this includes information on demand volume per fare class and expected demand arrival times. Often, the forecasting result assumes demand to be deterministic, although aware of its real stochastic nature.

The forecasted demand serves as input to the stage of *revenue optimization*, together with given data on capacity and fares. *Capacity* is the number of sellable seats on an aircraft assigned to a flight and *fares* are the prices customers have to pay for the service of transportation. For normal, the fare decision is made by airlines' pricing department and the capacity allocation is realized by fleet assignment. These two optimization input values are thus given from beyond an RM system. As an optimization approach, e. g., dynamic programming recursion constitutes an optimal solution (Talluri & van Ryzin, 2004b, p. 42). However, as airlines' optimization problem is highly complex, including an enormous amount of dimensions, computationally more efficient heuristics such as the EMSR-b approach are more common in practice (Talluri & van Ryzin, 2004b, pp. 47). As a result of optimizing revenue, inventory controls are transfered to the reservation stage.

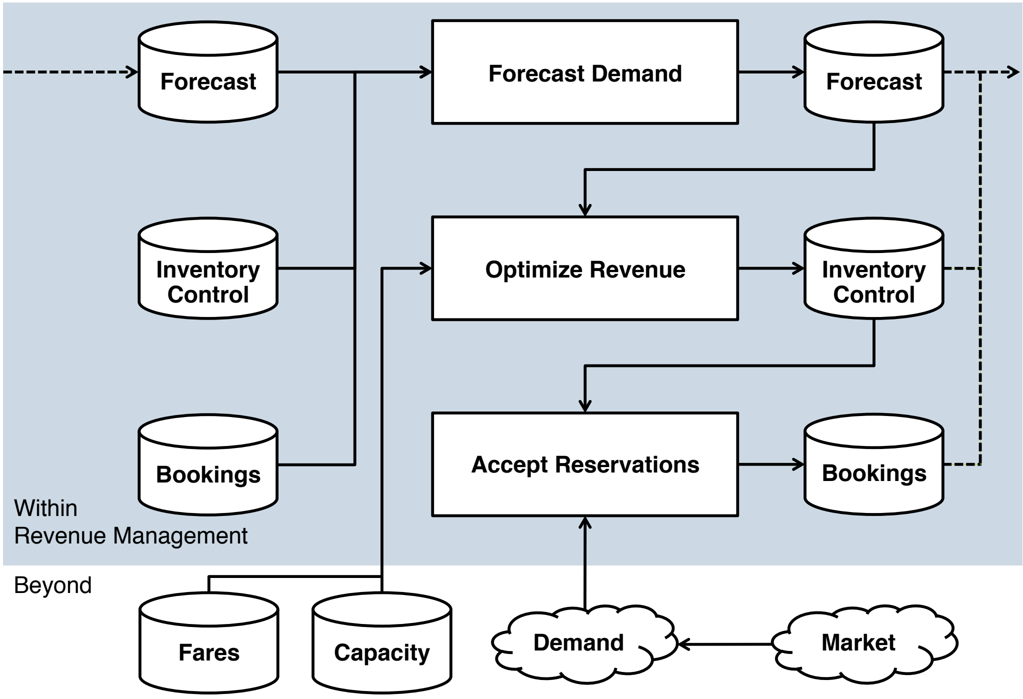

Figure 2.1: Traditional Revenue Management Process

Inventory controls determine the availability of fare classes over time; in the case of quantity-based RM they are also referred to as *capacity controls*. The most common types of controls are: booking limits, protection levels and bid prices (Talluri & van Ryzin, 2004b, pp. 28). While booking limits restrict the sellable capacity per fare class, protection levels state that amount of capacity that will be reserved for future requests per fare class. Both booking limits and protection levels can easily be converted to one another. As their effects are equivalent, only analysts' preferences decide with which to work. In contrast to booking limits and protection levels, bid prices constitute a control mechanism that facilitates considering network effects in RM. A *bid price* determines a value per compartment, such as first, business or economy class in ARM. Compared to virtual fare classes, compartments are aircrafts' physical classes offering a different service. A particular fare class is only available to customers, if its fare exceeds its compartment's bid price.

Independent of the type of control, the RM system transfers optimization's output to the stage of *accepting reservations*. Here, inventory controls decide on accepting or rejecting customer requests. The theoretically optimal controls now face real world's fluctuating and thus uncertain *demand*. Figure 2.1 depicts uncertain values as a cloud. The *market* cloud includes everything that influences airlines' customers, such as economic trends, competition or alliance partners.

Actual demand results in reservations if inventory controls accept customer requests. These reservations are recored as *bookings* in the system. The updated information from every process stage serves as input to the next iteration, indicated by a dotted line. In ARM, such iterations include both, re-optimizing one flight within its booking horizon, as well as re-optimizing the similar flight on a different date. Passing through these iterations constitutes an RM system's feedback loop.

Whenever this work refers to an RM system, it assumes one similar to the one depicted by Figure 2.1. The following section introduces a simplified RM optimization problem.

2.3 Leg-based Deterministic Revenue Management Model

The basic RM problem deals with the question, when to offer which fare class over time. This section provides a simplified problem formulation. The model optimizes revenue on single flight-legs rather than entire networks, considers only a single compartment and assumes demand to be deterministic and independent between fare classes[1]. This section provides the base for Part II, which incorporates and considers capacity uncertainty in RM by extending this section's model.

Thus, the problem is broken down to maximizing revenue from ticket sells and simultaneously keeping down denied boardings. Additionally, ticket sells can neither exceed the number of requesting customers nor the number of seats available on an aircraft. Table 2.1 facilitates reading by providing a summary of this section's mathematical notation.

Type	Symbol	Definition
Variable	$x_{ft} \in \mathbb{N}_{\geq 0}$	Number of tickets to sell in fare class $f \in F$ at time slice $t \in T$
	$e_a \in \{0, 1\}$	Denied boarding indicator of denied boarding $a \in \{1, \ldots, \hat{a}\}$
Set	F	Set of fare classes
	T	Set of time slices
Parameter	$f \in F$	Fare class
	$t \in T$	Time slice, from beginning of booking horizon $\hat{t}$ to departure 0
	$r_f \in \mathbb{Q}_{\geq 0}$	Revenue of fare class $f \in F$, with $r_1 \geq \ldots \geq r_{\lvert F \rvert}$
	$a \in \{1, \ldots, \hat{a}\}$	Denied boarding
	$\hat{a} \in \mathbb{N}_{\geq 0}$	Maximum number of denied boardings
	$k_a \in \mathbb{Q}_{\geq 0}$	Cost of denied boarding $a \in \{1, \ldots, \hat{a}\}$
	$c \in \mathbb{N}_{\geq 0}$	Capacity, an aircraft's number of seats
	$d_{ft} \in \mathbb{N}_{\geq 0}$	Demand in fare class $f \in F$ at time slice $t \in T$

Table 2.1: Definition of Model Symbols Ordered by Type and First Appearance

Let the booking horizon start at time slice $\hat{t} \in \mathbb{N}_{\geq 0}$ and end at flight's departure $t = 0$. Fare class $f \in F$ has a corresponding revenue $r_f \in \mathbb{Q}_{\geq 0}$ that results from selling a ticket. Fare classes are ordered so that $r_1 \geq \ldots \geq r_{\lvert F \rvert}$. The model's maximum number of denied boardings is $\hat{a}$. Parameter $k_a \in \mathbb{Q}_{\geq 0}$ denotes the cost of the ath denied boarding, where $a \in \{1, \ldots, \hat{a}\}$. Decision variable $e_a \in \{0, 1\}$ indicates whether a denied boarding takes place or not. Denied boarding costs increase such as $k_1 \leq \ldots \leq k_{\hat{a}}$. An aircraft's capacity is indicated by $c \in \mathbb{N}_{\geq 0}$. Decision variable $x_{ft} \in \mathbb{N}_{\geq 0}$ states the amount of tickets sold in fare class f at time slice $t \in T$. The expected demand is indicated by $d_{ft} \in \mathbb{N}_{\geq 0}$ at every time slice t for every fare class f.

[1] Those restrictions to a basic RM model are, e. g., discussed by Talluri and van Ryzin (2004b, Chapter 2)

The following integer linear program formulation can be solved to optimality

$$\underset{x_{ft},\, e_a}{\text{maximize}} \sum_{f \in F} r_f \sum_{t \in T} x_{ft} - \sum_{a=1}^{\hat{a}} e_a \cdot k_a \tag{2.3.1}$$

$$\text{s.t.} \quad \sum_{f \in F} \sum_{t \in T} x_{ft} - \sum_{a=1}^{\hat{a}} e_a \leq c \tag{2.3.2}$$

$$x_{ft} \leq d_{ft}, \qquad \forall\, f \in F, t \in T \tag{2.3.3}$$

$$x_{ft} \in \mathbb{N}_{\geq 0}, \qquad \forall\, f \in F, t \in T$$

$$e_a \in \{0, 1\}, \qquad \forall\, a \in \{1, \ldots, \hat{a}\}\,.$$

The problem (2.3.1) is solved by maximizing the revenue from sold tickets, determined by decision variable x_{ft}, multiplied by its related revenue r_f and simultaneously keeping the number of denied boardings $\sum_{a=1}^{\hat{a}} e_a$ small, which is the second decision variable multiplied by its related denied boarding cost k_a. Additionally, the capacity restriction (2.3.2) ensures that the sum of tickets sold $\sum_{t \in T} \sum_{f \in F} x_{ft}$ exceeding an aircraft's capacity c must result in denied boardings e_a. At last, the demand restriction (2.3.3) ensures that sold tickets x_{ft} cannot exceed demand d_{ft} for every fare class f and time slice t.

2.4 Disturbances in Revenue Management

This chapter addressed conditions for applying RM, showed a common RM process view and introduced a simplified problem formulation. On this basis, this section highlights existing disturbances that are categorized either as an uncertainty or a risk. While a disturbance already considered as a risk asks decision makers to reducing the risk's magnitude, an uncertainty constitutes a disturbance that has not been systematically considered and must thus be transformed to a risk. Table 2.2 lists this work's considered disturbances in RM systems based on their location in the process and categorizes them as a risk or uncertainty.

Based on this overview, the following two chapters discuss the addressed disturbances and present appropriate solution approaches for coping with a particular uncertainty or risk: first, uncertainties and risks arising within RM systems, then those beyond RM systems.

Location	Disturbance	Type
within RM	Lacking Booking Data	Uncertainty
	Inaccurate Demand Estimates	Risk
	Inaccurate Demand Model	Uncertainty
	Limited Optimization Objective	Uncertainty
beyond RM	Inherent Demand Variation	Risk
	Inappropriate Static Fare Structure	Uncertainty
	Fixed Capacity Assumption	Uncertainty

Table 2.2: Revenue Management Disturbances' Location and Type

3 Uncertainties and Risks within Revenue Management

This chapter presents an overview on uncertainties and risks that can occur within RM systems. Thereby, it discusses disturbances as listed in Table 2.2 based on the process depict in Figure 2.1. Several contributions addressed in this chapter are also mentioned in Cleophas et al. (2017), but with a focus on uncertainties and risks instead of a focus on persistence and flexibility. First, Section 3.1 discusses the lack of historical booking data. Then, Section 3.2 addresses the problem of inaccurate demand estimates, while Section 3.2 addresses inaccurate demand models. At last, Section 3.4 discusses RM optimization objectives.

As a limitation, this work neglects disturbances that arise from human intervention to the system. Commonly, RM analysts can influence data between the demand forecasting and optimization stage (Talluri & van Ryzin, 2004a, p. 410). Analysts, e. g., increase or decrease expected demand or influence fare class availability based on their expert knowledge. The optimization stage remains typically unaffected by analysts' influences. Although these interventions are expected to be beneficial for an RM system's result, they involve the hazard of unintended outcomes. Thereby, human intervention into the system constitutes an uncertainty, however, it can only be traced back very hardly (Zeni, 2003). Probably for that reason, RM literature on disturbances by RM analysts is very rare.

Likewise, this work neglects disturbances that arise from *communication errors*. These errors' emanate from human nature or from the system: communication errors can arise between multiple humans (e. g., transferring a misinterpreted information), from a human to the system (e. g., feeding the system with flawed input data) or from the system itself (e. g., an erroneous implemented algorithm). The same applies for communication errors as for human interventions: they are hard to determine and theory still lacks for addressing these uncertainties.

On a final note, this and the following chapter neglect disturbances, which are impossible to predict. For example, this applies if an unforeseen event only occurs on rare occasions or a decision maker is not aware of the event's existence. In this context, Section 1.2 spoke of unawareness, which is beyond this work's focus.

3.1 Lacking Booking Data

As depict in Figure 2.1, historical data on forecasted demand, inventory controls and bookings are important inputs for the demand forecasting stage in RM systems. While data on demand forecasts and inventory controls help predicting expected future demand, historical booking data is supposed to be the most crucial information. If it is not available in sufficient quantity and quality, historical booking data cannot be used to forecast demand. This issue particularly complicates RM in new markets or industries (Lennon, 2004).

Under normal circumstances, a decision maker should be aware of lacking booking data on a particular flight or market. However, it is almost impossible to predict and quantify the resulting hazard of lacking booking data on the system's performance. For that reason, lacking booking data constitutes an uncertainty in RM systems. Section 3.1.1 reviews contributions that try to circumvent the uncertainty of lacking booking data by excluding initial demand estimates from the system. In contrast, literature in Section 3.1.2 treats approaches that manually supplement demand forecasts and thus try transforming the uncertainty of lacking booking data into a manageable risk.

3.1.1 Optimization without Initial Demand Estimates

Revenue optimization without initial demand estimates circumvents RM demand estimation's dependency on historical booking data. Approaches from this research area are also relevant when shifts in the market place have rendered existing data irrelevant.

Popovic and Teodorovic (1997) were one of the first, who allowed for missing historical booking data in RM. The missing data is compensated through adaptive inventory controls based on a Bayesian approach.

The solution of Lan, Gao, Ball and Karaesmen (2008) does not require full demand information. Here, however, demand's lower and upper bounds are required, but no information on demand arrivals' timing. Given this limitation, they rate the performance of static and dynamic policies in a single-leg model. This idea is extended by Lan, Ball, Karaesmen, Zhang and Liu (2015), who combine uncertain demand with uncensored knowledge about no-show probabilities for a joint control policy.

A more drastic approach is given by Ball and Queyranne (2009). They present online algorithms that avoid the necessity of booking data in RM models. Thus, they do not require any demand estimates at any given time. The approach also overcomes the traditional assumption of a risk-neutral RM approach, covered in Section 3.4.2.

In the context of dynamic pricing, Besbes and Zeevi (2009) consider risk bounds and near-optimal solution algorithms for two situations: a known parametric demand function with unknown parameter values and an unknown demand function without parametric representation. Extending these results, Wang, Deng and Ye (2014) also differentiate parametric and non-parametric learning. The authors favorably compare the results achievable by dynamic pricing to those achievable by customer-bidding.

While the previously cited references focus on large levels of demand and inventory, den Boer and Zwart (2015) propose an approach with limited active price experimentation. Their approach is suitable for smaller inventory levels and shorter sales horizons.

3.1.2 Manual Supplementation of Demand Forecasts

The most intuitive approach for overcoming the problem of missing historical booking data is manually supplementing the data by human analysts. Zeni (2003) investigates the value of manually improving forecasts. The author formulates a business process, in which RM analysts get feedback on their decisions. Mukhopadhyay, Samaddar and Colville (2007) provide similar suggestions, however, concentrate on forecast quality rather than on achieved revenue. Also Currie and Rowley (2010) highlight the general need to improve forecast accuracy by using RM analysts to be furthermore successful. Weatherford (2015) suggests analysts to add multipliers to existing forecasts for flexibly adapting market information.

Lemke, Riedel and Gabrys (2012) state that forecast data for a similar flight can be manually selected as initial reference for a new flight. In the next periods, the reference is updated. However, manually supplementing is not only helpful when no initial forecast data exist; especially situations with sparse or highly constrained historical booking data may profit from the proposed approaches.

The topic of self-monitoring forecasts[1] – forecasts that regularly check their performance and autonomously adjust themselves – seems to be highly related to this case, however, an initial input of forecasted demand is still needed.

From a more general non-RM perspective, Petropoulos, Fildes and Goodwin (2016) also underpin that manually supplementing automated forecasts can be beneficial when done with deliberation. The authors introduce a new measure for understanding the quality of judgmental adjustments.

3.2 Inaccurate Demand Estimates

Regular quantity-based RM systems feed time series methods with empirical data for predicting the future as best as possible (Talluri & van Ryzin, 2004b, pp. 410). The problem of inaccurate demand estimates is extensively considered in existing RM research. Early on, Weatherford and Belobaba (2002) highlighted the importance of accurate estimates by showing its strong impact on resulting revenue in simulation experiments.

Bartke (2014) analyzes several state-space estimation methods that can be adopted to RM's demand estimation. Maglaras and Eren (2015) suggest to base demand estimation on maximum entropy distributions, while more recently, Gao, Ball and Karaesmen (2016) propose a robust approach using demand's lower and upper bounds instead of probability distributions for estimating demand. Most literature on inaccurate demand estimates focuses on improving one particular aspect. Thus, inaccurate demand estimates represent a risk in RM systems. In contrast to Section 3.3 discussing demand models, this section focuses on demand estimation, primarily calibrating forecast model parameters; Section 2.1 addressed the distinction of estimation and forecasting.

Several aspects of demand must be forecasted to provide the best basis for revenue optimization. As summarized by Cleophas, Frank and Kliewer (2009), historical sales have to be unconstrained to estimate the true demand volume, demand arrival times and customers' choice behavior has to be predicted. Given the focus of recent contributions, this section concentrates on efforts to improve nonparametric demand estimation in Section 3.2.1 and to update demand estimates based on new information in Section 3.2.2. Additionally, Section 3.2.3 addresses techniques for unconstraining demand and Section 3.2.4 discusses contributions for estimating demand arrival times.

3.2.1 Nonparametric Demand Estimates

Modeling and thus accounting for customer choice behavior is the focus of several recent RM research efforts. However, when choice models determine the parameters to be estimated, some assumptions about customers' behavior, for example with regard to customer's utility function, have to enter the estimation procedure.

[1] Contributions treating self-monitoring forecasts are, e. g., Swanson and White (1997), Y. Chen, Yang, Dong and Abraham (2005) and Guidolin and Timmermann (2009).

Employing nonparametric estimation procedures can constitute a solution to the problem of demand estimation. These fit the functional form to the data without being constrained by prior assumptions. For example, Farias, Jagabathula and Shah (2013) estimate a distribution of customers over demand segments that produces the worst-case revenue, being compatible with observations on a given set of availability data. Based on the same dataset, van Ryzin and Vulcano (2014) create maximum likelihood estimates of the choice demand model, based on historical booking and availability data.

Going beyond the knowledge of product availabilities, Azadeh, Hosseinalifam and Savard (2015) provide a nonparametric estimation approach that requires no information about product characteristics, whereas Jagabathula and Vulcano (2015) consider panel data to estimate customers' preference orders.

3.2.2 Updating Demand Estimates and Inventory Controls

In their RM research overview, McGill and van Ryzin (1999) state that "the performance of a given RM system depends, in large part, on the frequency and accuracy of updates". Updating demand forecasts – likewise updating inventory controls as explained in Section 2.2 — throughout the booking horizon allows RM systems to react to new information.

Although online algorithms that update after each booking event would perform best, most systems cannot handle the computational complexity. Frequently repeated offline optimization is used as an approximation. Updated forecasts use newly revealed information exclusively to adjust parameter values.

A widespread approach to real-time RM is to solve the problem via dynamic programming – see Bertsekas (2005) for an overview. At each point in time, request acceptance depends on the current inventory, the expected future demand and past sales. Iteratively updating the dynamic program throughout the booking horizon lets control strategies flexibly consider sudden developments. However, RM systems cannot handle the computational effort resulting from applying exact dynamic programming to realistic problem instances. Therefore, current research primarily focuses on improving computational efficiency. To that end, L. Chen and Homem-de Mello (2010) approximate the multi-stage stochastic programming formulation. More recently, Zou, Ahmed and Sun (2016) propose a stochastic version of the nested decomposition of multi-stage stochastic integer programs.

Jasin and Kumar (2012) study the benefits of re-solving a deterministic linear program by probabilistically implementing the solutions at predetermined times. They provide an upper bound for the expected revenue loss and construct a schedule of re-optimizations to limit this loss. The resulting heuristic is called probabilistic allocation control (PAC). In a follow-up paper, Jasin (2015) shows that frequently re-optimizing PAC, even without re-estimation, reduces the asymptotic revenue impact of uncertain demand.

The work of Jasin (2014) improves the static inventory control model, outperforming some approaches using re-optimization. The proposed self-adjusting heuristic only needs a single optimization stage at the beginning of the booking horizon. Especially in dynamic situations, this provides an computational advantage. Furthermore, implementing the heuristic and performing only a few re-optimizations during the sales period achieves a superior revenue performance.

3.2.3 Demand Unconstraining

A demand model's usefulness depends on the accuracy of its parametrization. Improving the estimation of these demand parameters through enhanced models or methods reduces the risk of inaccurate demand estimates. Unconstraining is one approach for improving these estimates.

As inventory controls constrain bookings, they censor observable demand. To estimate demand parameters from historical bookings, unconstraining techniques are needed. Belobaba and Farkas (1999) illustrate the importance of correctly estimating demand and present different estimation approaches. Zeni (2001) compares the achieved revenue given different unconstraining methods in a simulation study. Weatherford and Pölt (2002) analyze the sensitivity of different methods regarding the share of constrained data. They conclude that in the airline industry, true demand can never be accurately extrapolated.

Queenan, Ferguson, Higbie and Kapoor (2007) formulate a new estimation technique that integrates information about the time when data was constrained. Their numerical results demonstrate efficiency compared to established unconstraining techniques. More recently, Weatherford (2016) gives an overview to the history of unconstraining in RM for nearly the last 40 years. The author compares different unconstraining methods and highlights the relation to particular customer choice models.

3.2.4 Estimation of Demand Arrival Times

Improving the estimation of customer arrival times enhances the underlying probability distribution's fit. This reduces the risk of inaccurate demand estimation. As research on this topic is extensive, only the most influencing contributions are highlighted.

The first attempt to model airline passengers' booking process can be credited to Beckmann and Bobkoski (1958), who highlight the Poisson process' applicability. Still today, modeling bookings through a Poisson process remains industry's standard. Adelson (1966) suggests a compound Poisson process to better estimate total demand.

On empirical airline data, A. O. Lee (1990) tests the fit of a censored Poisson model and shows that it significantly improves parameter value accuracy and revenue. Kimms and Müller-Bungart (2007) propose how to model demand arrival times for simulation studies using a Poisson process. This work refers to Kimms and Müller-Bungart (2007) when modeling stochastic demand generation in Section 8, which is needed for this work's computational studies.

3.3 Inaccurate Demand Model

Neglecting important demand aspects, such as the preference for a particular travel date, constitutes an uncertainty for RM systems' demand models. However, approaches considering and incorporating one of these aspects, resulting in more accurate demand models, help transforming this uncertainty to a risk. In this context, van Ryzin (2005) stresses the importance of understanding and accurately modeling demand, putting the customer-centered consideration of demand in the focus of future research. For airlines, this means especially estimating three aspects of demand: cancellation and no-shows (Section 3.3.1), dependent demand (Section 3.3.2) and strategic customers (Section 3.3.3).

For a general review on customer models in RM, this work suggests the contribution by Shen and Su (2007), which also addresses dependent demand and strategic customers.

3.3.1 Cancellations and No-Shows

Cancellations and no-shows have been targeted early by research in the airline domain. Cancellations indicate customers recalling bookings in advance, while no-shows indicate bookings that are not utilized at the time of departure. Rothstein (1971) provides a first answer in the form of overbooking, i. e., selling an optimally increased virtual capacity. Alstrup, Boas, Madsen and Vidal (1986) discuss overbooking policies for the existence of two customer segments preferring different compartments.

Subramanian, Stidham and Lautenbacher (1999) document one of the first attempts considering cancellations in ARM. The authors model cancellations as a Markov process for a single-leg model with multiple fare classes. The proposed approach explicitly allows for overbooking, its formulation depends on dynamic programing and extends the model proposed by T. Lee and Hersh (1993). The results show that request acceptance depends on the ticket's fare and on the refund at cancellation.

Chiu and Tsao (2004) consider the problem of calculating an appropriate overbooking strategy. The authors propose a genetic algorithm, while explicitly considering demand uncertainty, forecast inaccuracies and competitor strategies.

Aydin, Birbil, Frenk and Noyan (2013) discuss overbooking related to a single-leg RM problem. They propose two novel static models that are risk-based and a dynamic overbooking model. The authors compare them in a computational study with EMSR-based approaches proposed by Lan et al. (2008) and Lan, Ball and Karaesmen (2011). With respect to their test environment, their dynamic model performs best and their static upper bound model outperforms the existing approaches from literature. More recently, Vulcano and Weil (2014) consider a joint optimization of virtual capacities and bid prices.

3.3.2 Customer Choice

The assumption that demand is independent of current alternatives, i. e., that a customer is interested in only a specific combination of product and fare, has a long tradition in RM. However, it is inherently flawed: given relevant information, a rational customer will prefer the cheapest of alternatives or is otherwise indifferent. This leads to so-called *sell-up* or *buy-down*: customers are forced to pay more than the minimal price of the desired product – sell-up – or they pay less than the fare that would best exploit their willingness to pay – buy-down.

Numerous efforts modeling dependent demand are comprehensively reviewed by van Ryzin (2005). Talluri and van Ryzin (2004a) introduce a general discrete-choice model, formulate maximum expected revenue functions and characterize the structure of optimal control policies. Vulcano, van Ryzin and Chaar (2010) quantify improvements that result from using these discrete-choice models in RM in addition to theoretically analyzing a maximum-likelihood approach. The authors' numerical studies show that customer-choice based RM improves results by 1–5% over EMSR.

One of the most hazardous effects of wrongly assumed independent demand is the *spiral-down effect*, analyzed mathematically by W. Cooper, Homem-de Mello and Kleywegt (2006) and numerically by Cleophas (2009). It occurs when an RM system based on a model of independent demand meets customers' dependent choice behavior. Decreasing bookings in high fare classes let forecast algorithms reduce expected demand for those classes, resulting in a structural decline: reduced expected high-fare demand leads to policies that allow for more bookings in cheaper classes, inducing a decline in high fare bookings, and so on. To avoid spiral-down, Gallego, Li and Ratliff (2009) present a choice-based extension of EMSR for a single-leg problem with multiple fare classes, where customer choice follows a multinomial logit model. The authors provide efficient solutions and numerical studies for different orders of customer arrivals.

Fiig, Isler, Hopperstad and Belobaba (2009) develop a method to transform expected demand based on a discrete choice model and fares. Such transformed fares let algorithms designed for independent demand models cope with dependent demand. Meissner and Strauss (2012) model dependent customer choice as a Markov decision process and approximate it with a nonlinear function. Their approach enhances the computational efficiency and derived control policies' accuracy. In Meissner, Strauss and Talluri (2013), the authors extend their approach by formulating a concave approximation for the linear program with customer choice and prove its computational efficiency. More recently, Garrow (2016) gives a comprehensive overview on discrete choice models for air travel demand.

3.3.3 Strategic Customers

Strategic customers can be regarded as a special case of dependent demand, where choice is time-dependent: customers strategically delay their bookings when expecting better offers in the future.

Gönsch, Klein, Neugebauer and Steinhardt (2013) summarize and classify relevant literature on strategic customers in the context of dynamic pricing. Li, Granados and Netessine (2014) estimate the share of strategic customers in U.S. airline transport demand to be 5–19%.

To counteract strategic customers, Jerath, Netessine and Veeraraghavan (2010) propose selling opaque products. The authors focus on last-minute offerings, customers' anticipation of such sales and on the use of intermediaries. They show that intermediaries can support a more efficient use of capacities and induce demand. Gorin, Walczak, Bartke and Friedemann (2012) anticipate and integrate strategic cancellations and re-bookings in network RM.

More recently, Board and Skrzypacz (2016) analyze RM with forward-looking customers and propose a profit-maximizing mechanism based on a sequence of cutoffs. They also account for the customer value as Du, Zhang and Hua (2015), who model strategic customer's with risk preferences, addressing a joint inventory and pricing decision in RM.

3.4 Limited Optimization Objective

Traditionally, RM's optimization objective is maximizing a company's expected short-term revenue, applying for many RM industries and most airlines. The common short-term approach is justified by a company's myopic view: for reasons of monetary liquidity, revenue must by generated as soon as possible. The maximization of expectation in turn is justified by a company's risk-neutral attitude.

However, these two elements do not always have to fit an RM applying company. Here, an inaccurate or limited optimization objective can lead to unintended results and thus constitutes an uncertainty for RM systems. On the one hand, instead of exclusively concentrating on the short-term revenue, some companies also want considering a customer's long-term value, when deciding which customer to accept and reject. This research area and its development is discussed in Section 3.4.1. On the other hand, the assumption that companies always act risk-neutrally is discussed and scrutinized in Section 3.4.2. These two research streams represent the most common alternative optimization objectives in RM. By including alternative objectives to revenue optimization, a company transfers the uncertainty of a limited objective to a risk.

3.4.1 Consideration of Long-Term Customer Value

Considering the long-term value of a customer can be beneficial if a company affords possibly smaller short-term incomes. At the beginning of RM, the basic idea was to sell a ticket only to those customers who are willing to pay the highest prices (e. g. Littlewood (1972)). However, in the late 1990s a progress in non-RM business areas started, where the *customer value* came into focus of companies' long-term goal. Woodruff (1997) gives a general introduction to the customer value and justifies its consideration for generating a competitive advantage. Rejecting a customer who currently is only willing to buy a less expensive product could lead to a general loss of this customer. This poses a threat if the customer otherwise would have planned to buy another product in the future – possibly even for a higher price.

The segment of students, e. g., constitutes a picture book example: during their studies, students' customer values are relatively small, owed to a usually low income, but their long-term value involves a huge potential. Thus, companies can act strategically by offering student fees and subsequently hoping to increase this group's future sales. Additionally, it is necessary to regard the customer hierarchy: some customers have higher influences on buying decisions of other customers. Thus, to win customers over by not rejecting them in short-term is a better long-term decision.

Initially, the majority of customer value related contributions came from the department of marketing.[2] Here, Kumar (2010) provides an extensive overview.

Wirtz, Kimes, Theng and Patterson (2003) discusses problems that could arise when adapting customer relationship management (CRM) techniques to RM. Introducing RM in customer-oriented business or establishing long-term relationship to customers in an already RM applying business can lead to potential customer conflicts. As these contradictions can have a negative impact on companies' long-term success, Wirtz et al. (2003) analyze how these conflicts emerge and propose prevention strategies. Despite these hazards, Noone, Kimes and Renaghan (2003) suggest merging CRM and RM in the domain of hotel RM. Their work is based on the approach of Reinartz and Kumar (2002), suggesting how customer segments should be targeted depending on their lifetime-profitability relationship. Noone et al. (2003) incorporate CRM efforts and RM strategies per customer segment and study the influence of different instruments on the business process, such as lifetime value-based pricing, availability guarantees and ad hoc promotions. Buhl, Klein, Kolb and Landherr (2011) jointly consider CRM and RM in their model – named CR^2M by raising the value in different customer segments. The model's ability to calculate opportunity costs allows a comparison of short- and long-term approaches.

The authors von Martens and Hilbert (2011) divide possible efforts on three levels of management activities: strategic, tactical and operational. Subsequently, with the help of a simulation prototype, they underline that in the long run a joint consideration of CRM and RM will outperform short-term RM control mechanisms – especially if customers' willingness to pay and customer value negatively correlate. In contrast to von Martens and Hilbert (2011), who allow for a general RM statement, Mohaupt and Hilbert (2013) address the domain of cloud computing with a long-term relationship of service provider and customer. They provide a simulation study for evaluating a customer-value based control policy that promises notable revenue benefits.

[2] Related topics from a marketing perspective are: customer value prediction (Verhoef & Donkers, 2001), customer length-strength relation (Reinartz & Kumar, 2002), lifetime value (Reinartz & Kumar, 2003), customer's lifetime value increase (Kumar, Ramani & Bohling, 2004; Venkatesan & Kumar, 2004) and customer equity (Kumar & George, 2007).

In this context, Board and Skrzypacz (2016) distinguish between long-lived and short-lived customers. RM commonly assumes customers being short-lived, leaving the market if they do not buy immediately. Their model includes long-lived customers and also considers customers to act strategically.

The biggest obstacle for practice, however, remains: how are companies able to identify a particular customer and predict its customer value correctly? This is especially difficult to answer when considering that a significant share of buying process' in RM industries happen anonymously. Rygielski, Wang and Yen (2002) are one of the first, who applied data mining techniques on CRM driven questions. Their contribution allows companies identifying valuable customers and predict their future behavior, providing two case studies and explaining the strengths and weaknesses of two particular data mining techniques: CHAID and neuronal networks.

Although technical barriers for collecting and storing huge amounts of necessary data constantly decrease, a thoroughly method for predicting individual customer values still constitutes a core task for jointly considering CRM and RM.

3.4.2 Consideration of Revenue Risk

Considering revenue risk in the optimization stage is the second research field that tries to resolve potential uncertainties of a limited optimization objective. The idea is to explicitly limit the acceptable risk in the mathematical model.

Commonly, an airline constitutes a huge company with a high number of frequent ticket sales. Here, from a pure revenue-driven perspective, the risk-neutral attitude is justified. However, for a company in an industry with rare sales, a risk-averse approach can be useful. Here, the number of profitable sales cannot compensate less profitable sales as it holds for flights in the aviation industry. Event promotion is an example, which can only afford low levels of risk (Koenig & Meissner, 2015). But also airlines can act risk-sensitive, e. g., when risk-averse human analysts override system rules (Isler & Imhof, 2008) or particular RM division goals strive for less fluctuating results, i. e., more planning safety.

In one of the first contributions to risk-sensitive RM, Feng and Xiao (1999) extend the model of Feng and Gallego (1995) by penalizing revenue variation. The resulting objective function accounts for decision makers' risk aversion. The authors present an optimal pricing policy, which more conservatively renders risk aversion. Feng and Xiao (2008) extend this framework and present structural results from an optimal risk-averse policy. Also Huang and Chang (2009) analyse revenue variation: simulation studies show that relaxing the optimality condition in a classic dynamic program reduces revenue variation that results in decreasing average revenues.

Approaches that measure the monetary revenue risk by using the value at risk (VaR) or conditional value at risk (CVaR) become more popular (Gönsch & Hassler, 2014; Koenig & Meissner, 2009, 2010, 2015).[3] In contrast to these absolute risk measures, Lancaster (2003) recommends relative measures. He proposes a revenue-per-available-seat-mile indicator and presents a sensitivity analysis on revenue variation. Mitra and Wang (2005) consider decision makers' risk aversion in the domain of bandwidth provisioning and routing. They suggest revenue deviation rather than tail value at risk (TVaR) as a risk index, which also appears applicable to airlines. Likewise, the approach by Barz and Waldmann (2007) does not explicitly address the airline industry but analyzes the risk-sensitive inventory control for the static and dynamic single-resource problem.

[3] An overview to the concepts of absolute risk measures, such as VaR, CVaR, earnings at risk (EaR) and cash flow at risk (CFaR), is, e. g., provided by Lhabitant and Tinguely (2001).

A first approach to integrating risk-sensitivity in RM's well-known EMSR algorithm can be found in Weatherford (2004). The author proposes an alternative, called expected marginal seat utility (EMSU). The resulting model is not limited to maximizing the expected revenue, but instead maximizes the utility achieved by earning a revenue stream. Also Ball and Queyranne (2009) consider revenue risk but apply online algorithms to the RM problem. However, the authors focus more on establishing a framework that does not rely on demand forecasts, as stated in Section 3.1.1.

Further research directions are considering multiple goals beside revenue maximization as, e. g., seat load factors (Phillips, 2012). Also integrating the influence on an airline's market share for a long-term competitor advantage can be of future interest.

4 Uncertainties and Risks beyond Revenue Management

While the last chapter addressed uncertainties and risks within RM systems, this chapter discusses those arising beyond RM systems. During the tactical planning stage, the pricing department sets fares, whereas fleet assignment determines flight capacities. RM only begins afterwards, with operative planning (Belobaba, 2009, Chapter 6). This chapter discusses such sources of uncertainties from beyond RM. To eliminate both capacity and fare related uncertainties, Feng and Xiao (2006) propose integrating capacity and pricing decisions in RM. The idea's practicability strongly depends on the system's overall performance and on simplifying assumptions, which in turn can cause new uncertainties.

Analogously to the previous, this chapter also discusses several contributions addressed by Cleophas et al. (2017), but with a focus on disturbances' classification into uncertainty and risk. Additionally, the same applies as for Chapter 3: disturbances by human intervention or communication errors are neglected.

At first, Section 4.1 focuses on demand variation's unavoidability and coping with this variation. The following two sections address revenue optimizations' fixed input values that must be obtained from beyond the RM system: static fares in Section 4.2 and static capacities in Section 4.3.

4.1 Inherent Demand Variation

Section 3.3.1 addressed customers that cancel or do not show up. But even when these cancellations and no-shows, the average expected demand or the probability of customers buying a product can be perfectly estimated, inherent demand variation still challenges RM. There is, and will always be, a natural tendency of several factors that influence customers' choice, which cannot be included in airlines' decision making process. Figure 2.1 depicts market influences on individual customers and thus demand as a whole. However, systematically considering market's elements is still at the beginning of its emergence, such as general economy's trends, the influence of airline alliances (Gerlach, 2013) and competition. Several contributions addressing the topic of competition in RM come from a mainly theoretical perspective, such as Belobaba and Wilson (1997), Isler and Imhof (2008), Zimmermann (2014) and W. L. Cooper et al. (2015). These areas still possess a potential for further research.

Inherent demand variation constitutes a risk that can be reduced by offering flexible products (Section 4.1.1), making use of flexible compartments (Section 4.1.2) or by incorporating dynamic capacity management (Section 4.1.3).

4.1.1 Flexible Products

As the name implies, flexible products increase RM flexibility as a response to demand variation. Each flexible product entails multiple alternative manifestations, constituting an allocation of customers to products. The product is specified only after it was sold. For example, booking a ticket for one of three possible destinations while the airline announces each ticket's actual destination not until sales are closed.

Gallego and Phillips (2004) first introduced flexible products in a single-leg, two-class context and developed algorithms for calculating booking limits, addressing additional customers while avoiding cannibalization. Petrick, Gönsch, Steinhardt and Klein (2010) focus on dynamically allocating flexible products to specific resources. The authors identify that, for evaluated methods, forecast quality and revenue gain correlate. Petrick, Steinhardt, Gönsch and Klein (2012) extend the model of Gallego, Iyengar, Phillips and Dubey (2004) to allow for arbitrary notification dates, demonstrating the benefits of flexible products with late notification dates.

Gönsch and Steinhardt (2013) provide a more theoretical view on flexible products by extending the classical dynamic program decomposition. Gönsch, Koch and Steinhardt (2014) provide similar results by adjusting product valuation in the deterministic linear optimization model to capture monetary benefits.

Introducing the concept of opaque products, Fay (2008) considers selling flexible products through an intermediary. The author formulates an analytical model, discusses its assumptions, requirements and managerial implications. Post (2010) extend the concept to variable opaque products for which the customer can select an individual degree of variability.

Continuing this customer-centric perspective, M. Lee, Khelifa, Garrow, Bierlaire and Post (2012) investigate customer preferences for potential manifestations of flexible products. The authors empirically analyze customers' likelihood to exclude alternatives and emphasize the importance of considering preferential customer choice when offering flexible products. Vock (2016) suggests models for considering these preferential customer choices when allocating flexible products. The author's computational studies test several decision heuristics and suggest considering strategic customers. More recently, Vock, Cleophas and Kliewer (2017) analyze the hazard of offering flexible products to strategic customers. Their computational study shows that revenue losses, due to cannibalization and spiral-down effects, can even outweigh the benefits of offering flexible products.

4.1.2 Flexible Compartments

As a response to demand variation, the simplest approach for adjusting capacity is to upgrade passengers when the originally booked compartment is depleted. Alstrup et al. (1986) were the first who formulated such an approach by solving a dynamic overbooking problem with two segments by substitution. Karaesmen and van Ryzin (2004) describe the possibility of multiple substitutable inventory classes.

Shumsky and Zhang (2009) calculate optimal protection limits by backward induction and present heuristics to solve large instances. The authors combine flexible compartments with demand-driven capacity changes; the following section addresses the latter.

4.1.3 Demand-Driven Capacity Changes

The majority of existing research underlines the importance of an optimal fleet assignment on which RM can build. For example, Barnhart, Farahat and Lohatepanont (2009) focus on creating a tractable model of fleet assignment that takes an approximation of revenue maximization into account.[1] However, changes in the market place can worsen the fit of capacity to demand. For that reason, different approaches of dynamically adjusting a fleets capacity to flight legs within the booking horizon emerged in research.

One of the first to propose implementing this concept using aircraft families are Berge and Hopperstad (1993). Their approach is called *demand driven dispatch* (D^3). By reducing the problem, Bish, Suwandechochai and Bish (2004) only take swaps of two aircrafts within one aircraft family into account, calling the approach *demand driven swapping*. Wang and Regan (2002), also study aircraft swaps as an extension of leg-based RM, albeit from a perspective of continuous time. Based on this formulation, Wang and Regan (2006) include costs associated with swaps and provide numerical studies benchmarking the model for different parameterizations.

As an extension to the EMSR-b algorithm, De Boer (2004) considers a dynamic version that adjusts RM policies for effects of capacity adjustments: *dynamic capacity management*. It performs best for a small set of fare classes. The approach by Frank, Friedemann, Mederer and Schroeder (2006) also allows for continuously adjusting capacity. However, the authors focus on a realistic demand model considering demand dependencies between fare classes. Yu, Chen and Zhang (2015) also address dynamic capacity management, but additionally include companies' ability to upgrade customers. As explained in the previous section, Shumsky and Zhang (2009) establish a joint approach of flexible compartments and dynamic capacity management.

4.2 Inappropriate Static Fare Structure

Quantity-based RM systems optimize the availability of discrete fare classes. Although each fare class can represent a set of fares, these fares may not be revenue optimal. A single fare class representing multiple fares adds the challenge of estimating a representative fare for optimization (Weatherford & Belobaba, 2002).

Therefore, an inappropriate static fare structure constitutes an uncertainty for an RM process. The resulting revenue loss cannot be quantified, thus theory strives for a reduction of this uncertainty by integrating the pricing decision in RM (Section 4.2.1) or by shifting the pricing decision to customers (Section 4.2.2).

4.2.1 Integration of Pricing Decision

One way to avoid fare uncertainty is to integrate the pricing decision into RM. By controlling not just fare availability, but the fares themselves, RM systems increase their ability to flexibly adjust to disturbances in the market.

Some first ideas in this direction can be found in Federgruen and Heching (1999). The authors consider the impact of the chosen fares on demand to come, but without focusing on ARM. Their work introduces basic principles applied to the field later on. For example, De Boer (2003) studies

[1] Section 4.3.1 explicitly addresses approaches jointly considering RM and fleet assignment.

the joint pricing and resource allocation problem by presenting structural results and strategies for choosing the optimal fare.

Also Feng and Xiao (2006) analyze integrating pricing and capacity decisions in RM. They propose a control policy that incorporates demand intensity, inventory status and fares. In the same vein, Kocabiyikoglu, Popescu and Stefanescu (2013) evaluate potential benefits when RM and pricing are either coordinated or combined hierarchically. The authors analyze four approaches differing in the degree of coordination and in the stochastic of pricing decisions.

Integrating pricing and RM may be regarded as a step on the way from inventory control to dynamic pricing. Thus, dynamic pricing can be regarded as the ultimate integration of pricing and RM. As pointed out in Şen (2013), dynamic pricing can provide an edge over inventory controls even given well-adjusted fare structures. Recently, Papanastasiou and Savva (2016) consider dynamic pricing in the presence of social learning and strategic consumers, while Dilme and Li (2016) study dynamic pricing problems with fire sales, which are strongly discounted prices. However, while theoretically superior, dynamic pricing is still slow to pervade ARM practice (Isler & D'Souza, 2009; Pölt, 2010).

4.2.2 Customer-Centered Price Setting

Mechanisms such as name-your-own-price (NYOP) and auctions enable customer-driven price as a contrary approach to static fare structures. NYOP leaves the price acceptance decision to the firm, while firms cannot influence fares in auction-based mechanisms.

NYOP mechanisms allow airlines to sell excess inventory at customer-driven prices. The buyer places a bid for a product and the seller decides whether to accept this bid or not. One of the first firms that implemented this mechanism was *Priceline* (Anderson, 2009). Anderson and Wilson (2010) thoroughly review approaches to NYOP, discuss several existing models and formulate implications for future research. Hinz, Hann and Spann (2011) compare different settings in NYOP markets with particular regard to customer satisfaction.

A technical approach to integrate NYOP in existing systems is described by Lochner and Wellman (2004). The authors provide general insights towards customer interaction mechanisms. Wilson and Zhang (2008) show how to optimize this acceptance decision when customers are aware of the probability that their price is accepted and try to maximize individual profit. Wang, Gal Or and Chatterjee (2009) use an analytical model to understand the opportunities and shortcomings of an NYOP selling channel. As benefits, the authors point out improved capacity utilization and reduced demand uncertainty. They state the loss of reliability, when capacity is scarce and fares increase, as a major disadvantage. In the context of hospitality RM, Ampountolas (2016) provides a survey on NYOP mechanisms, also addressing the influence of social media and dynamic pricing.

Auctions present a more traditional approach to customer-driven price setting. Here, customers compete with each other, while in NYOP, they bargain with the seller. An introduction to auctions in RM can be found in Talluri and van Ryzin (2004b, Chapter 6) with further contributions listed by Chiang et al. (2007). More recently, Heo (2016) analyzes group-buying auctions for restaurants in the context of RM. Nevertheless, auctions constitute a research area on its own, beyond the scope of general RM research.

4.3 Fixed Capacity Assumption

RM is in part motivated by the idea of optimally allocating a fixed capacity. However, actual demand can either exceed capacity or fall short. Demand exceeding capacity may be regarded as less severe, as RM inherently reserves capacity for the most valuable customers. Nevertheless, this means that the full revenue potential cannot be realized, resulting in spill. If demand falls short of capacity, some units cannot be sold, and revenue suffers from spoilage. RM research considers two types of capacity-based disturbances: the initial fit between capacity and demand and capacity changes, from beyond RM, that occur over the booking horizon. In both cases, assuming a fixed capacity, constitutes an uncertainty to RM systems that could be transferred to a risk.

The previous Section 4.1.3 already implicitly contradicted the assumption of fixed capacity by addressing the possibility of dynamically adjusting capacities, driven by RM as a response to demand variation. The next section discuses approaches that jointly consider fleet assignment and RM. Section 4.3.2 discusses to anticipate capacity changes that are not induced by RM.

4.3.1 Integration of Fleet Assignment

In the airline industry, allocating capacity to physical products means assigning aircrafts to flights. Intuitively, the aircraft with the biggest capacity are assigned to flights with the most expected requests (De Boer, 2004). Improving fleet assignment increases the fit of capacity and expected demand.

Abara (1989) allocates capacity via a linear integer model allowing for multiple optimization goals, e. g., minimizing costs and maximizing profits. Farkas (1996) accounts for network effects and highlights the importance of correctly estimating spill. Also Barnhart, Kniker and Lohatepanont (2002) concentrate on network effects, presenting a superior fleet assignment problem formulation. Sherali, Bish and Zhu (2006) present an overview of approaches and advancements for fleet assignment. However, the majority of research, such as B. Smith, Ratliff and Jacobs (1998) and Barnhart et al. (2009), still underline the importance of an optimal initial fleet assignment for RM.

4.3.2 Anticipation of Capacity Changes

Among others, technical defects, crew-scheduling problems and weather conditions can cause capacity changes during the booking horizon. In contrast to Section 4.1.3, these changes are, however, not induced by RM as a response to demand variation; they are rather an uncertain input from beyond an RM system.

As a result, a previously optimal RM solution may no longer be valid, when initial capacity and final capacity at departure differ. Thus, the idea of anticipating these capacity changes in RM transforms capacity uncertainty to a risk. To the best of this work's knowledge, only two contributions explicitly anticipate uncertain capacity changes in RM. Wang and Regan (2002) introduce such an anticipation for supporting their framework of repeated aircraft swaps, as discussed in Section 4.1.3. They propose a two-stage approach that divides the time horizon in an interval before and after a potential capacity swap. Wang and Regan (2006) extend this approach by a computational study benchmarking the model for different capacities, demand mixes and markets.

Expert inputs confirmed that capacity changes from beyond RM can constitute an actual threat to airlines, but are not systematically considered in practice yet. The lack of contributions addressing this particular capacity uncertainty motivates this work's Part II and III. The following chapter presents the research gap of capacity uncertainty that serves as a showcase for transforming an RM's uncertainty to a risk.

5 Research Gap

Chapter 3 and 4 showed that uncertainties and risks in RM still motivate research efforts. Today's success of RM's application in airline practice can be mostly attributed to this comprehensive and sustaining progress. The majority of uncertainties and risk in RM has been intensively surveyed and various approaches tried reducing existing risks or transforming uncertainties to manageable risks. The survey gave an insight to some of the most relevant uncertainties and risks in RM and identifies capacity uncertainty as a research gap. The remainder of this work will focus on this particular aspect.

Commonly, RM assumes that capacity is not only limited, but fixed and known in advance (Talluri & van Ryzin, 2004b, p. 14). However, airline practice does not confirm this assumption as capacity changes can occur that are neither induced nor anticipated by RM. The reasons for capacity changes in the airline industry are manifold, such as technical defects, crew planning, special sales, bad weather conditions or strikes. Although it is possible that ARM analysts already intuitively consider these changes when intervening in the system, a systematic consideration is still missing.

Likewise, this particular capacity uncertainty is also almost completely neglected in theory. The lack of sufficient research efforts motivates this work's purpose of understanding capacity changes' influence on ARM results. As its focus, this work addresses systematically anticipating those changes in ARM systems.

The absence of a structural analysis on and a systematic consideration of capacity changes constitutes a research gap. This work's first attempt of closing this gap is a well-suited example for a showcase of transforming an RM uncertainty to a risk. This is addressed by Part II and III.

Distinction to Related Literature

The contributions reviewed in Section 4.1.3 focus on the possibility to adjust aircrafts for better responding to demand variation. This work henceforth refers to the resulting capacity change as *endogenous*. In contrast, capacity changes from beyond an RM systems are referred to as *exogenous*. Endogenous changes are desired and intended by RM, while exogenous capacity changes constitute altered input values to RM systems. If this work in the following speaks of *capacity uncertainty*, it refers to exogenous capacity changes.

Virtual capacities that account for uncertain cancellations, seem to be closely related to the problem of capacity uncertainty. Both insufficient overbooking and exogenous capacity increases may cause spoilage, as seats remain unsold. Both excessive overbooking and exogenous capacity decreases may cause spill, where valuable demand is rejected and denied boardings can occur (Belobaba & Farkas, 1999). However, cancellations are frequently incremental, as individual passengers decide not to travel. Capacity changes can assign an entirely different aircraft type to a flight, resulting in significantly increased or reduced numbers of seats in the included compartments. Therefore, this work's model incorporates this uncertainty in the form of stochastic scenarios, rather than as the dynamic, incremental changes considered by most overbooking research. Here, e. g., Vulcano and Weil (2014) consider a joint optimization of virtual capacities and bid prices. While the approach of virtual capacities could also be promising for coping with capacity uncertainty, the authors primarily refer to overbooking which is demand variation's other side of the coin.

So far, only two contributions explicitly consider exogenous capacity changes in ARM: Wang and Regan (2002) and Wang and Regan (2006). First, Wang and Regan (2002) introduce the concept to support their framework of continuous-time[1] RM, given stochastic demand with repeated aircraft swaps. The authors consider swaps in two stages: the model first maximizes revenue from the swap time up to departure, assuming the new capacity to be fixed. Second, it maximizes revenue from the beginning of booking horizon up to the swap time, taking uncertain future capacity and control policies for the following stage into account. Thus, the authors present an optimal policy for revenue maximization given aircraft swaps and mathematically prove the potential revenue improvement. Wang and Regan (2002) also calculate the best parameterization of penalties for excessive overbooking in their model.

Wang and Regan (2006) extend their previous work by presenting a computational study benchmarking the model for different capacities, demand mixes and markets. However, they also primarily focus on endogenous capacity changes. As both contributions' have a congruent model and their content strongly overlaps, this work henceforth only refers to the published paper of Wang and Regan (2006).

This work in hand also decomposes the booking horizon into periods before and after potential capacity changes, but goes beyond swaps by allowing for more than two potential capacities and more than one change time. In contrast to Wang and Regan (2006), the revenue optimization model is based on a discrete-time RM approach, which is more common in ARM practice (Talluri & van Ryzin, 2004b, Preface).

This work suggests a scenario-based approach for anticipating capacity changes in advance, considering multiple potential capacities. The theoretical background of this work and results of preliminary studies, comparing three control strategies, are also recorded in Büsing, Kadatz and Cleophas (2017). Part III studies the problem in a simulation system fed with input data calibrated on realistic capacity changes and stochastically generated demand. The computational studies feature 13 additional control strategies that allow for analyzing the effects of different information levels, which is very briefly summarized in Kadatz, Kliewer and Cleophas (2017b). The work in hand additionally analyzes the influence of different parameterizations and tests the effects of distorting model's input data. Except the insights from distorting input data, this work's major findings are summarized in Kadatz, Kliewer and Cleophas (2017a).

From now on, this work replaces the term *exogenous capacity changes* with *capacity updates* because an actual change – which is a change of an aircraft – only takes place at departure. An update, however, better expresses the actual condition of an altered value in an RM system. The problem of potential capacity updates is still referred to as *capacity uncertainty*.

For a clear focus, this work is limited in the same ways as Wang and Regan (2006), by making the following assumptions.

Leg-based policy: The approach does not take network effects into account.

Two-stage approach: Optimization is divided into time intervals before and after potential capacity updates occur.

Independent demand: Demand for a fare class is independent of the availability of other fare classes.

[1] Their overall problem framework is based on the work by Liang (1999), who formulated the RM problem as a continuous-time, stochastic dynamic programming model.

No cancellations or no-shows: The model does not account for customer cancellations or no-shows.

Single compartment: Each aircraft only has a single compartment.

In contrast, Table 5.1 shows a distinction of this work in comparison to Wang and Regan (2006).

Model Characteristics	Wang & Regan	This Work
Time	continuous	discrete
Demand	stochastic	deterministic
Potential Capacity Updates	known	known/distorted
Capacity Update Times	1	number of time slices
Number of Capacities	2	unlimited
Max. Number of Capacity Updates	1	unlimited

Table 5.1: This Work's Distinction to Related Literature

Additionally, to the contribution's already addressed differences with regard to time and demand, potential capacity updates are assumed to be known, but this work also analyzes the influence of distorted capacity updates. The model also covers different potential capacity update times: while Wang and Regan (2006) only allow for one potential update at an a priori known time, this work allows capacity updates to occur on every time slice. Also, Wang and Regan (2006) limit the number of potential capacities to swaps, which is two. In contrast, the number of capacities is not limited by this work. While Wang and Regan (2006) only allow for a single capacity update over time, this work's number of capacity updates is unlimited. Concluding, this results in an overall more realistic model including a higher quantity of potential capacity updates.

Altered Revenue Management Process

When anticipating capacity uncertainty in an RM system, the traditional RM process (Figure 2.1) must be adjusted. Figure 5.1 shows necessary modifications, marked in orange.

As obvious, capacity as an uncertain input for optimization is now also depict as a cloud. Additionally to forecasting the expected demand to come, potential capacity updates must be estimated from historical data as well. Finally, this new input can be used to consider capacity updates in a system's optimization stage. While the data on previous capacity updates are already recorded by airlines, forecasting future updates is necessary. However, as this step is similar to forecasting demand and an enormous amount of literature is available on different forecasting and estimation techniques[2], this work assumes forecasting capacity to be performed with sufficient accuracy. Thus, this work's focus is on the altered RM optimization stage when considering capacity uncertainty.

[2]For example, Box, Jenkins, Reinsel and Ljung (2015) provide a comprehensive overview on forecasting techniques.

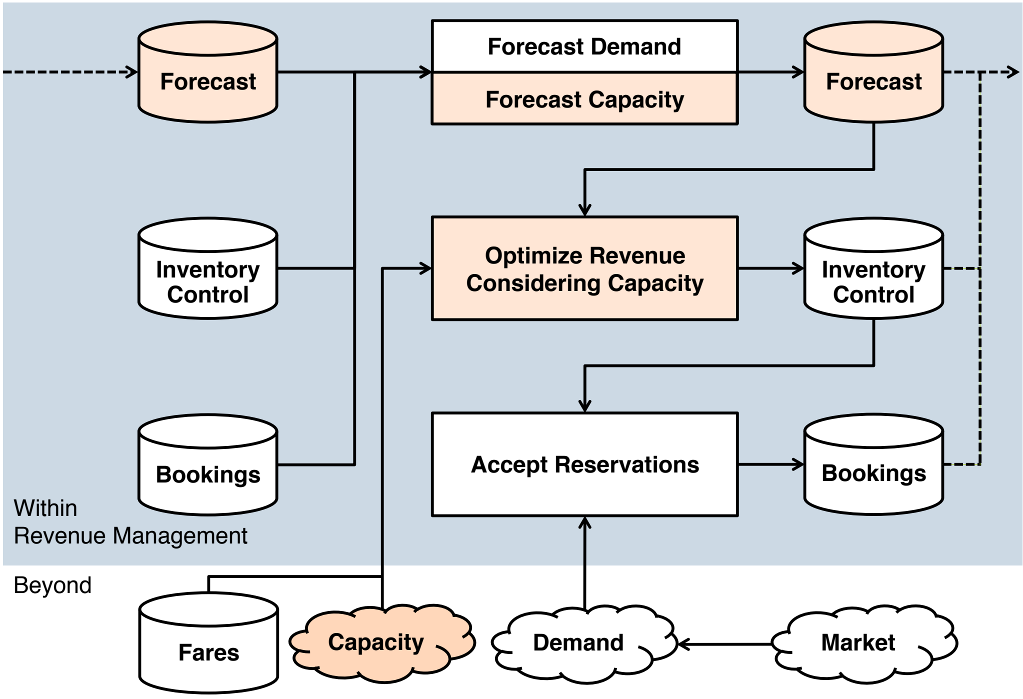

Figure 5.1: Capacity Updates Considering Revenue Management Process

Research Objective and Questions

The presented research gap induces this work's aim of systematically anticipating capacity updates in ARM. For fulfilling this objective, this work's agenda is based on answering the following research questions.

Research Question 1 (RQ1) *How frequently do capacity updates occur in airline practice? Do they affect RM's result?*

While airline practitioners confirm the occurrence of capacity updates, an investigation into airlines' empirical data is necessary (Chapter 6). Analyzing real data allows for evaluating how frequent these updates actually occur in practice. However, even if these updates can be identified, the question still remains whether they actually influence an RM system's revenue performance (Chapter 9 and 10).

Research Question 2 (RQ2) *Which capacity update characteristics are relevant to an RM system? How to model the capacity update characteristics in an RM system?*

Answering this question is necessary for considering capacity updates in RM. Analyzing empirical data on capacity updates (Chapter 6) also helps understanding which updates' characteristics affect an RM system. The identified characteristics are first modeled in the context of revenue optimization (Chapter 7) and finally in an entire RM framework (Chapter 8).

Research Question 3 (RQ3) *What are the benefits of anticipating capacity updates? Are some capacity update characteristics more important to know than others?*

If modeling capacity uncertainty, several strategies allow for anticipating potential capacity updates (Chapter 7). Computational studies test the established strategies and compare their benefits to the fixed capacity assumption and a solution's upper bound (Chapter 10). Here, analyzing the results if only particular capacity update characteristics are known states, which characteristics are more important to know and forecast in advance (Chapter 10).

Research Question 4 (RQ4) *Which factors influence RM's result if capacity updates are anticipated? How does a distorted demand forecast and a distorted forecast on capacity updates affect RM's result?*

The success of strategies that anticipate capacity updates can vary. Thus sensitivity analyses compare strategies' dependency on particular model parameters and state, which of them benefits or harms a strategies' revenue result (Chapter 10). Although this work focuses on an RM system's optimization stage, an always accurate forecast on capacity updates is unlikely. Thus, a study with regard to distorting forecasted demand and forecasted capacity updates helps understanding strategies' robustness on flawed input values (Chapter 11).

Part II

Capacity Uncertainty as a Showcase

6 Analysis of Empirical Data

This chapter includes, but is not limited to, answering research question 1.

RQ1: *How frequently do capacity updates occur in airline practice? Do they affect RM's result?*

At first, the analysis of empirical data concentrates on answering the question's first part. Here, empirical data verifies whether capacity updates actually occur in practice. Then, the frequency of these updates is identified. The research question's second part is partially addressed when analyzing empirical booking data. It is necessary to analyze the actual booking data and compare the time of customer arrivals with actual capacity update times as, RM's result will remain unaffected if no customer request arrives before a capacity update occurs. After an update, the system could simply adjust the seat allocation to the new capacity as if nothing had happened. However, if customers already requested a ticket before the occurrence of an update, revenue results can be affected. Consequently, this chapter analyzes both empirical data on capacity updates as well as data on bookings.

This chapter's analysis also helps answering the first part of research question 2.

RQ2.1: *Which capacity update characteristics are relevant to an RM system?*

For modeling the problem of capacity uncertainty in RM it is necessary to understand the dimensions of a capacity update that can influence the system behavior. The empirical analysis reveals these dimensions and presents their values in practice. The second part of research question 2 is answered in the next chapter, presenting an RM formulation considering capacity updates.

First, Section 6.1 presents flights and capacity updates of Deutsche Lufthansa AG (Lufthansa). The data includes one year of all flights departing in June 2015. Second, Section 6.2 analyses the entire booking data from December 2014 to November 2015. The data has been obtained from two separate systems, which explains the time horizons imperfect match.

6.1 Data on Capacity Updates

The data on flights and capacity updates includes the booking horizon of all flights departing in June 2015. These flights divide into three markets. In this work, they are termed: short-haul (SHORT), medium-haul (MEDIUM) and long-haul (LONG). As a rule of thumbs, flights on SHORT are usually regional flights, on MEDIUM cross-boarder flights and on LONG continent-connecting flights. As a result, flight duration on SHORT is on average shorter than on MEDIUM and obviously much shorter than on LONG.

Preliminary studies showed that capacity updates' magnitude can depend on aircrafts' size. Here, size reflects the number of seats on an aircraft. Motivated by these differences, the work categorizes flights into five aircraft sizes: XS, S, M, L, XL. The aircraft category with the smallest amount of seats, XS, includes all aircrafts with 70 or less seats up to the largest category, XL, with 161–190 seats. Table 6.2 shows the exact categorization.

First, Section 6.1.1 describes empirical flight and capacity data for subsequent analysis. Then, Section 6.1.2 presents key indicators and explains the necessity and the procedure of aggregating capacity updates into clusters.

6.1.1 Data Description

This section first introduces empirical flight data and gives a summary of aggregated key figures. Understanding the flight data is necessary for clearly allocating a capacity update to a particular flight. The following subsection presents empirical data on capacity updates and assigns each update to a flight. Exemplary data tables clarify the appearance of both flight data and data on capacity updates.

Flights

This section analyzes 44,642 flights. The raw data covers data on 48,295 flights; Appendix A explains the need for excluding some of them. The previously explained segmentation into markets assigns 10,300 flights to SHORT (23.07%), 28,438 flights to MEDIUM (63.70%) and 5,904 flights to LONG (13.23%).

Exemplary flight data are depicted in Table 6.1. Every entry consists of a flight number, a departure date, an origin airport and a destination airport. As the example of Table 6.1 shows, a particular flight number is used on different departure dates; each combination of flight number and departure date identifies a unique flight. For example, flight number 0105, flies on January 15, 16 and 17 from AAA to BBB.

Flight Number	Departure Date	Origin Airport	Destination Airport
...	...	...	...
0105	01-15-2015	AAA	BBB
0105	01-16-2015	AAA	BBB
0105	01-17-2015	AAA	BBB
...	...	...	...
0106	01-15-2015	BBB	AAA
...	...	...	...
0107	01-15-2015	AAA	CCC
...	...	...	...

Table 6.1: Exemplary Flights Data

Table 6.2 shows the segmentation into aircraft sizes: the smallest aircraft provides 28 seats and the largest 190 seats. This segmentation identifies 13 groups by combining markets and initial aircraft sizes. For each combination, the median represents the number of seats. In contrast to the mean, the median is more robust against outliers if a distribution is asymmetric (Bulmer, 2012, Chapter 4). Additionally, the median always denotes an aircraft's real number of seats.

Aircraft Size	Number of Seats	SHORT		MEDIUM		LONG	
		Share	Median	Share	Median	Share	Median
XS	28–70	4.59%	60	4.85%	60	–	–
S	71–100	13.92%	90	16.35%	90	–	–
M	101–130	32.94%	114	35.68%	114	9.42%	120
L	131–160	34.40%	144	30.70%	144	31.16%	150
XL	161–190	14.15%	172	12.42%	178	59.42%	178

Table 6.2: Flights' Share and Median Based on Market and Initial Aircraft Size

On SHORT, most initial aircrafts provide 131–160 seats (L), while on MEDIUM, aircrafts provide 101–130 seats (M). The comparatively high share of aircrafts with L or XL on SHORT can be attributed to the market's high demand volume. There are no initial aircrafts with less than 101 seats on LONG. Here, almost 60% of all flights provide 161–190 seats (XL) with a median of 178 seats.

Capacity Updates

Lufthansa automatically records and stores capacity updates, but does not use these data for anticipating future updates. Table 6.3 shows some fictitious capacity updates as an example.

Flight Number	Departure Date	Origin Airport	Destination Airport	Arrival Date	Record Date and Time	Former Number of Seats	New Number of Seats
...	...	...	...	...	...	...	...
0105	01-15-2015	AAA	BBB	01-15-2015	12-12-2014 10:03	90	123
0105	01-15-2015	AAA	BBB	01-15-2015	01-10-2015 12:01	123	96
0105	01-16-2015	AAA	BBB	01-16-2015	01-11-2015 13:07	84	44
...	...	...	...	...	...	...	...
0106	01-15-2015	BBB	AAA	01-15-2015	08-10-2015 15:33	123	96
0107	01-15-2015	AAA	CCC	01-15-2015	01-10-2015 11:05	123	96
...	...	...	...	...	...	...	...

Table 6.3: Exemplary Capacity Updates Data

Each observation includes the same characteristics as the flights data: a flight number, a departure date, an origin airport and a destination airport. Additionally, the data comprises an arrival date, a record date and time, a former number of seats and a new number of seats. In the example of Table 6.3, the first shown capacity update affects the flight with flight number 0105, departing on January 15, 2015 in AAA and arriving in BBB on the same day. The capacity update is recorded to the system on December 12, 2014 at 10:03 a.m. and the former aircraft provides 90 seats while the new allocated aircraft provides 123 seats. The second capacity update affects the same flight (flight number 0105 on January 15, 2015), while the third affects a flight with the same flight number but departing the day after.

From originally 71,489 observations, the following section takes 28,522 of them into account. Appendix A explains, which capacity updates must be neglected from analyses. The 28,522 observations on capacity updates are analyzed per market and divided into 9,202 observations on SHORT (32.27%), 19,222 on MEDIUM (67.39%) and 98 on LONG (0.34%).

Table 6.4 shows empirical frequencies for capacity updates per market. The numbers range from zero to six updates. On SHORT, up to six updates occur, but very rarely (0.02%). On MEDIUM, five updates (0.01%) occur on a particular flight and on LONG maximal four updates (0.02%) occur. The overall chance for at least one capacity update is 65.74% on SHORT, 49.16% on MEDIUM and 1.85% on LONG.

Number of Updates	Probability of SHORT	Probability of MEDIUM	Probability of LONG
0	34.26%	50.84%	98.15%
1	47.42%	34.02%	1.45%
2	14.62%	12.58%	0.31%
3	3.23%	2.28%	0.07%
4	0.39%	0.27%	0.02%
5	0.06%	0.01%	–
6	0.02%	–	–
$\sum$	100.00%	100.00%	100.00%

Table 6.4: Empirical Probabilities for Capacity Updates

Across markets, 51.35% of all flights are affected by at least one capacity update over the booking horizon. If an update takes place, it is most probable to remain the only capacity update, with a probability of 59.83%.

6.1.2 Cluster Analysis

For this work's computational studies, the amount of 28,522 potential capacity updates is reduced by summarizing multiple similar updates. For determining this similarity, the analysis introduces two additional values: capacity update time and capacity update magnitude. The update time determines the difference of the flight's departure and the capacity update date in days before departure (dbd). The magnitude of a capacity update, denoted as a percental change, is positive or negative, depending on whether capacity increases or decreases. Per market and aircraft size, capacity updates' similarity is based on these update times and update magnitudes.

In the example of Table 6.3, flight 0105 with departure at January 15, 2015 is affected by two updates. For the first update, the additionally captured values are update time, 34 days before departure (from 12-12-2014 to 01-15-2015) and update magnitude, +37% (90 to 123 seats). For the second update, the update time is five days before departure (from 01-10-2015 to 01-15-2015) and magnitude −22% (123 to 96 seats), and so on.

At this point, clustering techniques allow for finding the best matching groups of similar observations, called *clusters*. Appendix B provides an overview of clustering techniques and methods tested on the present data. However, this work clusters capacity updates based on input by RM experts. Following their suggestion, updates are split into capacity increases and decreases and based on their update time as it often correlates with particular reasons for an update. E. g., capacity updates between 230–270 days before departure can primarily attributed to a re-optimizing flight plans.

Figure 6.1 illustrates the resulting clusters of aircraft size M capacity updates on MEDIUM. The upper panel shows the entire unallocated observations. The lower panel shows the nine resulting, colored clusters. Here, the previously addressed period from 230–270 days before departure, bears two clusters: cluster 1 involves all updates within this time frame with a positive update magnitude and cluster 5 those with a negative magnitude.

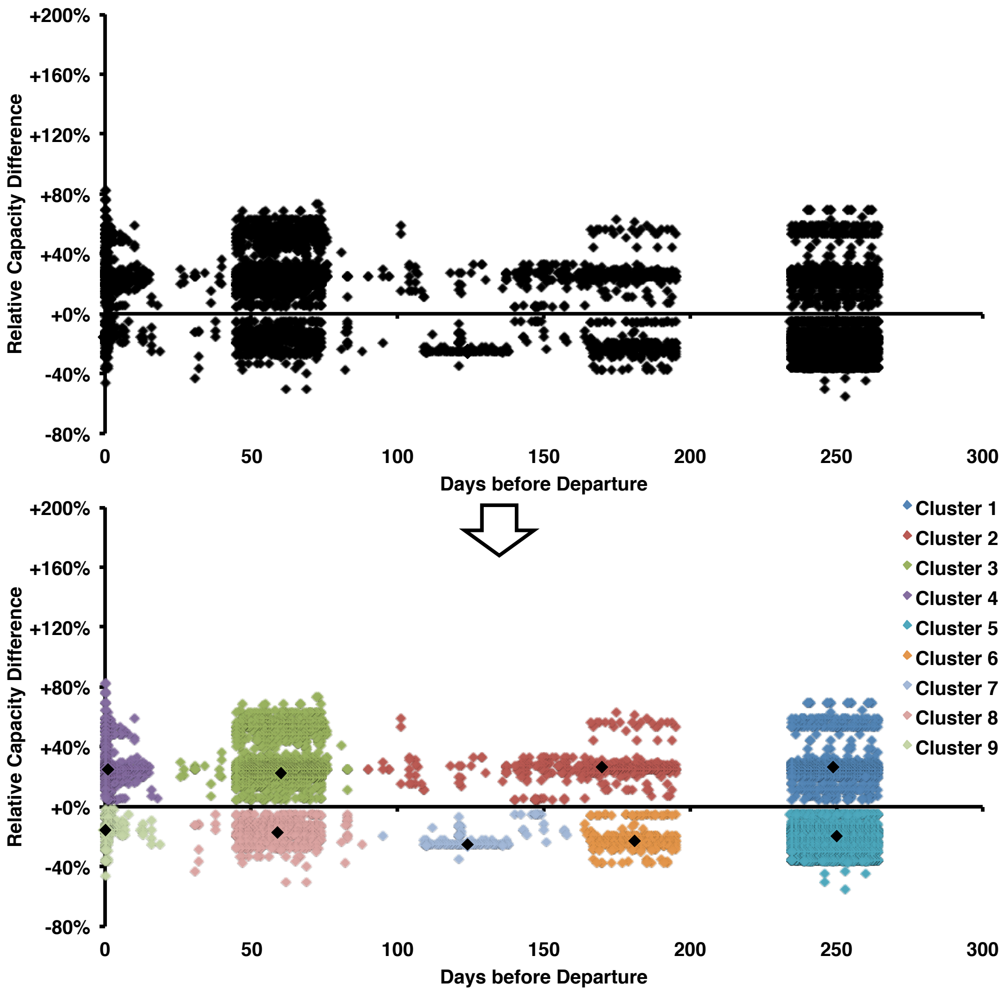

Figure 6.1: Capacity Update Clusters on MEDIUM of aircraft size M

Preliminary studies showed that the given data involves few drastic outliers. Compared to k-means, k-medoids is more robust to outliers (Kaufman & Rousseeuw, 1987; Massart, Kaufman, Rousseeuw & Leroy, 1986), similar to the statistical measure of a median compared to a mean. To prevent the resulting distortion, this work uses k-medoids for calculating clusters' centroids. These centroids are depict in black. In Appendix B, Figures B.1–B.13 show the remaining cluster combinations of aircraft size and market. The centroids serve as an input for calibrating computational studies' capacity scenarios. Table 6.5–6.7 present the resulting cluster centroids per market.

Aircraft Size	Cluster	Cluster Share	Update Magnitude	Update Time
XS	1	73%	+33%	249 dbd
XS	2	14%	+36%	157 dbd
XS	3	13%	+29%	0 dbd
S	1	26%	+25%	250 dbd
S	2	9%	+31%	167 dbd
S	3	8%	+37%	58 dbd
S	4	10%	+33%	0 dbd
S	5	30%	−32%	250 dbd
S	6	4%	−27%	171 dbd
S	7	12%	−38%	60 dbd
S	8	1%	−11%	1 dbd
M	1	15%	+26%	249 dbd
M	2	7%	+28%	174 dbd
M	3	7%	+29%	55 dbd
M	4	13%	+22%	1 dbd
M	5	29%	−24%	250 dbd
M	6	5%	−19%	174 dbd
M	7	19%	−19%	60 dbd
M	8	5%	−14%	1 dbd
L	1	12%	+24%	250 dbd
L	2	3%	+24%	181 dbd
L	3	6%	+20%	91 dbd
L	4	9%	+15%	1 dbd
L	5	18%	−23%	249 dbd
L	6	6%	−23%	182 dbd
L	7	38%	−33%	61 dbd
L	8	8%	−18%	1 dbd
XL	1	19%	−22%	250 dbd
XL	2	2%	−41%	249 dbd
XL	3	2%	−23%	186 dbd
XL	4	43%	−20%	62 dbd
XL	5	19%	−35%	62 dbd
XL	6	15%	−18%	2 dbd

Table 6.5: Cluster Centroids of SHORT as Input for Capacity Update Scenarios

Aircraft Size	Cluster	Cluster Share	Update Magnitude	Update Time
XS	1	78%	+36%	250 dbd
	2	10%	+47%	125 dbd
	3	7%	+27%	1 dbd
	4	1%	−18%	254 dbd
	5	4%	−27%	61 dbd
S	1	22%	+33%	250 dbd
	2	11%	+50%	171 dbd
	3	21%	+27%	61 dbd
	4	8%	+33%	0 dbd
	5	28%	−33%	250 dbd
	6	1%	−38%	182 dbd
	7	1%	−24%	125 dbd
	8	5%	−21%	55 dbd
	9	3%	−21%	0 dbd
M	1	14%	+26%	249 dbd
	2	6%	+26%	170 dbd
	3	28%	+26%	60 dbd
	4	6%	+25%	1 dbd
	5	32%	−20%	250 dbd
	6	3%	−23%	181 dbd
	7	1%	−25%	124 dbd
	8	6%	−18%	59 dbd
	9	4%	−16%	1 dbd
L	1	13%	+24%	250 dbd
	2	6%	+25%	158 dbd
	3	20%	+25%	60 dbd
	4	9%	+19%	1 dbd
	5	27%	−22%	250 dbd
	6	4%	−18%	170 dbd
	7	15%	−22%	60 dbd
	8	5%	−17%	1 dbd
XL	1	31%	−21%	250 dbd
	2	15%	−39%	250 dbd
	3	10%	−18%	154 dbd
	4	18%	−22%	56 dbd
	5	9%	−38%	60 dbd
	6	17%	−18%	3 dbd

Table 6.6: Cluster Centroids of MEDIUM as Input for Capacity Update Scenarios

Aircraft Size	Cluster	Cluster Share	Update Magnitude	Update Time
M	1	62%	+21%	18 dbd
	2	23%	+25%	0 dbd
	3	15%	−24%	48 dbd
L	1	9%	+23%	247 dbd
	2	7%	+21%	2 dbd
	3	51%	−18%	108 dbd
	4	25%	−21%	54 dbd
	5	7%	−18%	4 dbd
XL	1	74%	−16%	99 dbd
	2	17%	−37%	71 dbd
	3	9%	−24%	1 dbd

Table 6.7: Cluster Centroids of LONG as Input for Capacity Update Scenarios

Table 6.8 summarizes the results per market. On SHORT, the probability for a capacity update is highest, with approx. 66%. Also the maximal number of 6 potential capacity updates is highest across markets. Updates occur predominantly late in the booking horizon. Most flights depart on MEDIUM and most updates occur early in the booking horizon, leading more often to an increase than a decrease. On LONG, the probability for a capacity update is very rare. However, in two-thirds of all updates, capacity decreases. Although aircrafts departing on LONG are comparatively large – 150 seats as median – an update only increases or decreases 22 seats on average. As the success of RM depends on requesting customers, Section 6.2 analyses empirical booking data. This section also addresses research question 1; capacity updates are only an issue for RM if customers arrive before potential capacity updates can occur.

Market	Flights Share	Fleet Mix Min./Median/Max.	Average Update Magnitude / Time	Magnitude Ratio Increase/Decrease	Update Max. #/Prob.
SHORT	23%	56/112/179 seats	28 seats/late	45%/55%	6/66%
MEDIUM	64%	44/114/180 seats	28 seats/early	52%/48%	5/49%
LONG	13%	91/150/185 seats	22 seats/mid	34%/66%	4/2%

Table 6.8: Market Characteristics of Empirical Capacity Updates

6.2 Booking Data

This section analyses booking data from December 2014 to November 2015. The data covers net bookings accumulated over time, recorded at the beginning and middle of a month. The given *net bookings* implicitly exclude cancellations and no-shows. Bookings, are categorized into three markets: SHORT, MEDIUM and LONG. The data contains net bookings of 30 fare classes.

Some fare classes have an internal or promotional purpose. As they do not reflect real demand, these fare classes are removed from analysis. Additionally, only considering the economy compartment leads to twelve effective fare classes. For every market, three artificial fare class buckets are created. They are categorized based on their average fare and named fare class 1, 2 and 3. Preliminary studies showed that three fare classes suffice for observing general effects.

As an example, Figure 6.2 illustrates the aggregated results for MEDIUM. The booking curve's peak of fare class 1 and 2 is at 15 days before departure and that of fare class 3 at 30 days before departure. Figure 6.2 implies a common RM assumption: The cheaper a fare class the earlier do customers request on average (Robinson, 1995). The same applies for markets SHORT and LONG, visualized in Appendix C by Figure C.1 and C.3.

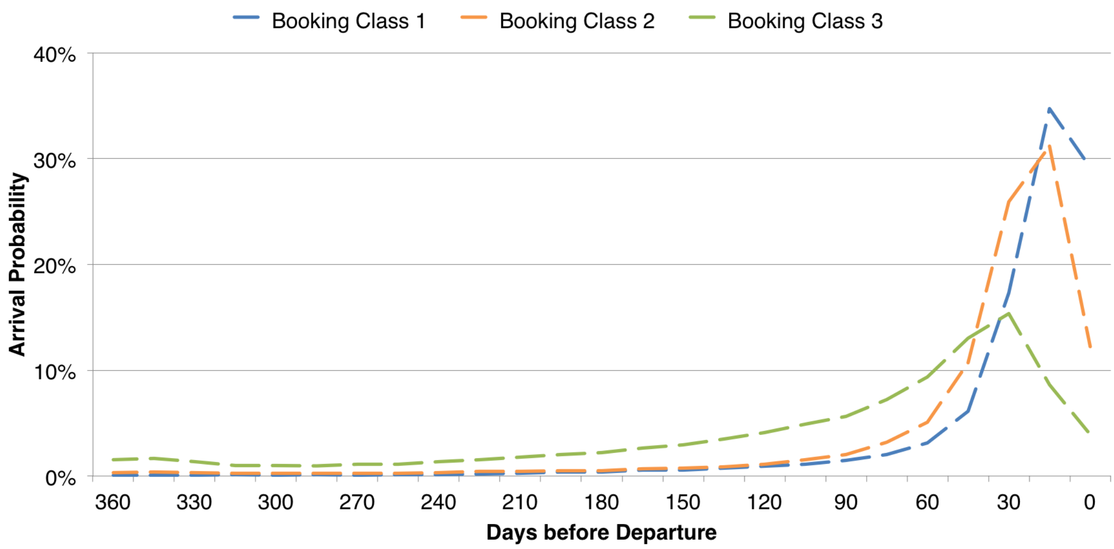

Figure 6.2: Empirical Net Bookings on MEDIUM

This work's computational studies require parameterizing the customer arrival process. It is necessary to realistically represent the customer arrivals over time. The process requests for a parameter determining the arrival probability per market and fare class. Therefore, this work assumes the empirical booking data to be approximately triangular distributed. Request time starts, where the probability for a customer request is larger than 0.5% – neglecting values with a smaller probability. The booking curve's highest point depicts the triangular distribution's mode. On SHORT, customers arrive late, on MEDIUM a bit earlier and on LONG they arrive early.

Table 6.9 shows the resulting triangular distributions' parameter on market level: lower limit, mode and upper limit. They serve as an input for parameterizing the computational studies' customer arrival process in Section 10.1.1.

6.3 Implications of Empirical Data

This chapter shows that capacity updates occur frequently in practice. The very fact that capacity updates affect more than half of all flights (51.35%) and that most updates affect more than 20% of

Market	Triangular Distribution (Lower Limit, Mode, Upper Limit) Fare Class 1	Fare Class 2	Fare Class 3
SHORT	(120, 15, 0)	(135, 15, 0)	(360, 30, 0)
MEDIUM	(180, 15, 0)	(210, 15, 0)	(360, 30, 0)
LONG	(165, 30, 0)	(225, 30, 0)	(360, 45, 0)

Table 6.9: Triangular Distribution of Demand Arrivals per Market and Fare Class

the initial aircraft's seat number justifies this works' expectation. In practice, capacity updates are frequent but strongly depend on a particular market. The data description, e. g., shows that the probability of a flight on SHORT being affected by an update is more than 35 times higher than on LONG. Thus, it could be beneficial in practice to consider capacity updates on some markets, while neglecting them on others.

Additionally, the joint analysis of empirical capacity update and booking data shows that customers frequently book a ticket before capacity updates occur. Thus, this section supposes capacity updates having a negative impact on RM's performance, which encourages systematically anticipating capacity updates in RM optimization.

6.3.1 Derived Dimensions

This chapter also partially answers research question 2. From the empirical data on capacity updates, three particular dimensions appear to be relevant for an RM system. The new capacity, transferred to the system if an update occurs, is obviously important to know. This new capacity defines the number of future sellable seats, when assuming that no further update occurs. The second dimension of a capacity update is its time. The earlier a capacity update occurs, the more time have RM systems to react on new capacities. The third dimension is a capacity update's probability to occur.

These three dimensions of capacity updates – capacity, update time and update probability – build the basis for systematically anticipating capacity updates in RM, as suggested by this work. The mathematical model in the next chapter considers these dimensions and integrates them in the process of finding the optimal number of tickets to sell. Also this work's simulation framework considers capacity based on these three dimensions.

6.3.2 Derived Parameters

For calibrating the simulation model, the computational studies make use of this chapter's analysis on capacity update and booking data. The presented capacity update clusters serve as an input for the computational studies. These clusters represent realistic capacity updates, based on market and aircraft size. Additionally, this chapter's empirical booking data is used for calibrating studies' customer arrival process. Based on the market and an artificial fare class a triangular distribution represents the density function of customer requests. At last, derived triangular distributions determine the non-homogeneous Poisson parameter $\lambda(t)$, which is necessary for generating realistic demand in Section 8.1.1.

7 Modeling Capacity Uncertainty

The empirical analysis in Chapter 6 addresses the first part of research question 2.

RQ2.1: *Which capacity update characteristics are relevant to an RM system?*

In contrast, this chapter addresses the research question's second part. First, it is necessary to model capacity updates in revenue optimization as addressed by this chapter before capacity updates can be modeled in an RM system as explained in the next chapter. Although revenue optimization constitutes this work's focus, other affected elements of an RM system, such as forecasting capacity updates, are subsequently discussed in the following chapter.

This chapter's outline is as follows. First, Section 7.1 introduces a mathematical model considering capacity updates in the leg-based deterministic RM formulation that has been established in Section 2.3. The section explains how characteristics of capacity updates must be considered in a scenario-based RM problem formulation. Section 7.2 introduces different information levels that arise if one or more of these characteristics are unknown. Different information levels represent airlines' knowledge on future capacity updates. Last but not least, Section 7.3 presents multiple possibilities of coping with the previously addressed imperfect knowledge on future capacity updates. Each possibility is a *control strategy*: a unique approach of parameterizing the mathematical model of Section 7.1 based on a particular information level.

7.1 Leg-based Deterministic Model with Capacity Updates

Section 2.3 presented a leg-based deterministic RM model as basis for this work's formulation for anticipating capacity updates. When observing that updates can alter capacity over time, Section 6.3.1 determined three characteristics playing a major role.

Capacity: The new capacity of an update can result in an altered number of seats on a flight leading to a capacity increase or decrease.

Time: The time of a capacity update is the time slice in the booking horizon when an RM system receives the information of an altered capacity. In the following, these times are denoted by days before departure (dbd).

Probability: The probability of a capacity update determines the chance that an update will be observed at a particular time with a particular capacity.

The discrete number of possible combinations allows for building a matrix that includes all potential capacities of aircrafts operating a particular flight times every update time. For that reason, this work decides for a scenario-based formulation representing those updates. Each combination of capacity and time is called a scenario $s \in S$ with a respective probability p^s, where $\sum_{s \in S} p^s = 1$. The update time of a scenario is denoted as t^s and the new capacity as c^s. Striving for computational efficiency, the formulation does not take scenarios with a probability $p^s = 0, \forall\, s \in S$, into account, as they do not provide any additional information. Table 7.1 summarizes the model's symbols; those being new, compared to the base model, are marked in bold.

Type	Symbol	Definition
Variable	$x_{ft} \in \mathbb{N}_{\geq 0}$	Number of tickets to sell in fare class $f \in F$ at time slice $t \in T$
	$\boldsymbol{x^s_{ft}} \in \mathbb{N}_{\geq 0}$	Number of tickets to sell in scenario $s \in S$ in fare class $f \in F$ at time slice $t \in T$
	$\boldsymbol{e^s_a} \in \{0,1\}$	Denied boarding indicator in scenario $s \in S$ of denied boarding $a \in \{1,\ldots,\hat{a}\}$
Set	F	Set of fare classes
	T	Set of time slices
	$\boldsymbol{S}$	Set of capacity update scenarios
Parameter	$f \in F$	Fare class
	$t \in T$	Time slice, from beginning of booking horizon $\hat{t}$ to departure 0
	$r_f \in \mathbb{Q}_{\geq 0}$	Revenue of fare class $f \in F$, with $r_1 \geq \ldots \geq r_{\lvert F \rvert}$
	$a \in \{1,\ldots,\hat{a}\}$	Denied boarding
	$\hat{a} \in \mathbb{N}_{\geq 0}$	Maximum number of denied boardings
	$k_a \in \mathbb{Q}_{\geq 0}$	Cost of denied boarding $a \in \{1,\ldots,\hat{a}\}$
	$d_{ft} \in \mathbb{N}_{\geq 0}$	Demand in fare class $f \in F$ at time slice $t \in T$
	$\boldsymbol{p^s} \in [0,1]$	Probability of scenario $s \in S$, with $\sum_{s \in S} p^s = 1$
	$\boldsymbol{t^s} \in T$	Time of capacity update in scenario $s \in S$
	$\boldsymbol{c^s} \in \mathbb{N}_{\geq 0}$	Capacity in scenario $s \in S$

Table 7.1: Definition of Capacity Uncertainty Considering Model Symbols

Under the assumption of fixed capacity, one decision variable, x_{ft}, suffices for representing the revenue optimizing strategy in the maximization problem (2.3.1). When considering capacity updates, however, it is necessary to introduce and distinguish a global strategy and scenario strategies. All strategies determine the number of desired tickets to sell on a particular time horizon, constituting additional decision variables. However, the *global* strategy, x_{ft}, applies to the time horizon before a capacity update can occur. In contrast, a *scenario* strategy, $x^s_{ft} \in \mathbb{N}_{\geq 0}$, decides which tickets to sell after an update at t^s to a particular capacity c^s occurs. Per scenario, the overall number of tickets to sell is thus given by

$$\sum_{t=\hat{t}}^{t^s+1} x_{ft} + \sum_{t=t^s}^{0} x^s_{ft} \,.$$

As a result, the maximization problem (2.3.1) is extended to (7.1.1). Here, available ticket's decision variables are split into x_{ft}, *before* an update and x^s_{ft}, *after* an update. Also the decision variable that determines denied boardings $e^s_a \in \{0,1\}$ now covers all potential scenarios s.

The capacity restriction (7.1.2) must also be adapted to considering potential capacity updates. Restriction (7.1.2) hence differentiates x_{ft} and x^s_{ft} as well as scenario's possible capacities c^s. Thus, per scenario, every sold ticket exceeding capacity c^s leads to a denied boarding e^s_a.

The previous demand restriction (2.3.3) is split into two. The first restriction (7.1.3) covers the global strategy x_{ft}; the second restriction (7.1.4) covers scenario strategies x^s_{ft}. Thereby, the model still prevents selling tickets to exceed demand, per fare class f and time slice t. The maximization problem is represented by the resulting integer linear program formulation

$$\underset{x_{ft},\, x^s_{ft},\, e^s_a}{\text{maximize}} \sum_{s\in S} p^s \left(\sum_{f\in F} r_f \left(\sum_{t=\hat{t}}^{t^s+1} x_{ft} + \sum_{t=t^s}^{0} x^s_{ft} \right) - \sum_{a=1}^{\hat{a}} \sum_{s\in S} e^s_a \cdot k_a \right) \tag{7.1.1}$$

$$\text{s.t.} \sum_{f\in F} \left(\sum_{t=\hat{t}}^{t^s+1} x_{ft} + \sum_{t=t^s}^{0} x^s_{ft} \right) - \sum_{a=1}^{\hat{a}} e^s_a \leq c^s \qquad \forall\, s\in S \tag{7.1.2}$$

$$x_{ft} \leq d_{ft} \qquad \forall\, f\in F,\, t\in T \tag{7.1.3}$$

$$x^s_{ft} \leq d_{ft} \qquad \forall\, s\in S,\, f\in F,\, t\in T \tag{7.1.4}$$

$$x_{ft},\, x^s_{ft} \in \mathbb{N}_{\geq 0} \qquad \forall\, s\in S,\, f\in F,\, t\in T$$

$$e^s_a \in \{0,1\} \qquad \forall\, s\in S,\, a\in\{1,\ldots,\hat{a}\}\,.$$

This formulation represents a deterministic equivalent of a stochastic problem and can be solved to optimality. If only one scenario is assumed, the scenario-based model collapses to the basic model of Section 2.3, leading to same results. As knowledge on potential capacity update characteristics does not always has to be perfect, the following section introduces different information levels of capacity updates.

7.2 Information Levels of Capacity Updates

The previous section's formulation of a capacity uncertainty considering model allows various possibilities for parameterizing the model's scenario input values c^s, p^s and t^s. This section introduces four different information levels on capacity updates' outcome, times and probabilities. These information levels represent airlines' knowledge on potential capacity updates – full or partial information can be assumed.

Each approach of parameterizing the model is represented by a control strategy. Following this, Section 7.3 presents 16 control strategies, each based on one of the four previously introduced information levels.

As an airline has always knowledge on its own fleet, at least all potential capacities are assumed to be known. This constitutes the lowest information level – termed information level 1. Combining the remaining attributes of capacity updates leads to three additional levels. Table 7.2 summarizes all four information levels. Information level 2a additionally assumes capacity update times to be known; information level 2b assumes update probabilities to be known. Full information on update's outcome, probabilities and times is denoted as information level 3.

Table 7.3 lists four exemplary capacity update scenarios ordered by their probability. Time is given in days before departure. In the first scenario, with probability 0.5, the capacity is updated to 10 seats on the day of departure. The second, third, and fourth scenarios represent updates that might occur six or two days before departure.

The remainder presents minimalistic examples based on these values.

Information Level 1: Capacities
: Information is most limited when only potential capacities are known. For Table 7.3, this translates to only knowing that the final capacity will be five, ten or twelve.

Information Level	Given Information
1	capacities
2a	capacities and update times
2b	capacities and update probabilities
3	capacities, update probabilities and update times

Table 7.2: Information Levels

Scenario	Update Capacity	Update Probability	Update Time
1	10 seats	50%	0 dbd
2	5 seats	25%	6 dbd
3	12 seats	15%	6 dbd
4	5 seats	10%	2 dbd

Table 7.3: Exemplary Capacity Update Scenarios

Information Level 2a: Capacities and Times
In this case, all potential times of capacity updates and all potential resulting capacities are known. For the example in Table 7.3, this means knowing that capacity may be updated to twelve seats, six days before departure, but not knowing the probability of this event.

Information Level 2b: Capacities and Probabilities
For each update, the probability and resulting capacity are known, while no information on the time of updates is available. For the example in Table 7.3, this translates to knowing that capacity will be updated to ten seats with probability 0.5, to five seats with probability 0.35 ($= 0.25 + 0.1$) or to twelve with probability 0.15.

Information Level 3: Capacities, Probabilities and Times
For every potential update, all characteristics are known. Thus, all information on c^s, p^s and t^s $\forall\, s \in S$ can be used for parameterizing optimization. The control strategy employing such full information is denoted as $C^s/P^s/T^s$. Thus, all parameters of every scenario in Table 7.3 serve as input for the optimization model.

Depending on the assumed information level, the capacity updates considering model of Section 7.1 can be differently parameterized. Each approach of parameterizing the model constitutes a control strategy; the following section introduces sixteen possible strategies.

7.3 Control Strategies Depending on Information Levels

Table 7.4 lists 16 control strategies. They rely on parameterizing the scenario-based optimization model (7.1.1)–(7.1.4). This work considers an upper bound strategy relying on perfect foresight (oracle) and one, where the initial capacity is assumed to be the final capacity as a benchmark. All remaining strategies are assigned to one of the four listed levels to indicate the minimally required information.

Information Level	Control Strategy	# of Scenarios	Decision Criteria
Oracle	C^*	1	perfect foresight
Benchmark	$C_{ini}/P_1/T_0$	1	initial capacity
1	$C_{max}/P_1/T_0$	1	maximal capacity
	$C_{min}/P_1/T_0$	1	minimal capacity
	$C(P_{uni})/P_1/T_0$	1	average capacity
	$C^s/P_{uni}/T_0$	$\|C\|$	uniform probabilities, updates at departure
	$C^s/P_{uni}/T_{uni}$	$\|C\| \cdot \|T\|$	uniform probabilities, updates anytime
2a	$C(T_{min})/P_1/T_0$	1	latest update only
	$C(T_{max})/P_1/T_0$	1	earliest update only
	$C^s/P_{uni}/T^s$	$\|S\|$	uniform prob., known capacities and update times
2b	$C(P^s)/P_1/T_0$	1	weighted average capacity
	$C(P_{min})/P_1/T_0$	1	least probable capacity
	$C(P_{max})/P_1/T_0$	1	most probable capacity
	$C^s/P^s/T_0$	$\|C\|$	known probabilities, updates at departure
	$C^s/P^s/T_{uni}$	$\|C\| \cdot \|T\|$	known probabilities, updates anytime
3	$C^s/P^s/T^s$	$\|S\|$	capacities, probabilities and update times

Table 7.4: Control Strategies' Decision Criteria Ordered by Information Level

The upper bound is represented by oracle C^*. The remaining control strategies are all termed according to a common pattern: $C_{...}/P_{...}/T_{...}$, where C represents capacity, P probability and T time. The ancillary terms indicate the respective parameterization: *min* is the minimum and *max* the maximum of a value, *uni* assumes a uniform distribution, *ini* uses the initial value, 0 a value of 0, 1 a value of 1 and s indicates that all scenarios are considered. All control strategies indicated by T_0 assume that the capacity is updated only at departure.

When a strategy has to derive an integer capacity from a decimal value, that value is rounded off to avoid unintentional overbooking. When the information level renders two scenarios indistinguishable, i. e., if scenarios only differ in their update time and the control strategy does not use this information, scenario probabilities are summed-up.

$\boldsymbol{C^*}$ represents an upper bound for revenue maximization under uncertain capacity. It considers the single scenario that results from perfect foresight of the actual final capacity.

$\boldsymbol{C_{ini}/P_1/T_0}$ creates a single scenario. It assumes that the initially assigned capacity will be the final capacity for as long as no capacity update occurs. Whenever the capacity is updated, this strategy assumes that the new capacity will be the final capacity. To this work's best knowledge, this strategy recreates the current industry standard. For Table 7.3, this means using capacity ten for optimization.

$\boldsymbol{C_{max}/P_1/T_0}$ creates a single scenario with assigned probability 1. It assumes that the largest possible capacity will realize. For Table 7.3, this means using a capacity of twelve for optimization.

$\boldsymbol{C_{min}/P_1/T_0}$ creates a single scenario, assuming that the smallest possible capacity will realize. For Table 7.3, this means using a capacity of five for optimization.

$\boldsymbol{C(P_{uni})/P_1/T_0}$ creates a single scenario, which is parameterized with the arithmetic mean of all possible capacities. In Table 7.3, this leads to a capacity of nine $((10+5+12)/3=9)$.

$\boldsymbol{C^s/P_{uni}/T_0}$ creates one scenario per possible capacity and assigns a uniform probability to each. Thus, for the example, capacity five, ten and twelve are all assigned with probability $1/3 = 0.\bar{3}$.

$\boldsymbol{C^s/P_{uni}/T_{uni}}$ creates one scenario per possible capacity and possible time of update. It assigns uniform probabilities to all scenarios. Assuming a time horizon of $|T| = 10$ in Table 7.3, the resulting number of scenarios is $|C| \cdot |T| = 3 \cdot 10 = 30$. Each scenario is equally probable with $1/30 = 0.0\bar{3}$.

$\boldsymbol{C(T_{max})/P_1/T_0}$ uses information on capacities and the potential times of updates. However, without information on scenario probability, only the capacity resulting from the earliest expected update is used to parameterize a single scenario. If more than one capacity can result, the arithmetic average is used. For Table 7.3, the earliest time of a potential update is six days before departure. As this applies for both scenario 2 and 3, the arithmetic mean of the corresponding capacities, $\lceil(5+12)/2\rceil = 8$, is used.

$\boldsymbol{C(T_{min})/P_1/T_0}$ is similar to $C(T_{max}/P_1/T_0)$, but parameterizes the capacity given the latest possible update time. Again, the arithmetic average is used to combine multiple candidates. For Table 7.3, this leads to a capacity of ten, as the corresponding scenario is connected to the latest update, at departure.

$\boldsymbol{C^s/P_{uni}/T^s}$ creates one scenario per capacity and time of update. Without information on scenario probability, it assumes uniformly distributed probabilities. For Table 7.3, this implies that all probabilities are set to $1/4 = 0.25$.

$\boldsymbol{C(P^s)/P_1/T_0}$ considers only a single scenario and parameterizes the expected capacity as the average of all possible capacities weighted by scenario probabilities. If the resulting value is not integer, it is brought down to the next integer. For Table 7.3, this leads to an expected capacity of $\lceil 0.5 \cdot 10 + 0.25 \cdot 5 + 0.15 \cdot 12 + 0.1 \cdot 5\rceil = 8$.

$\boldsymbol{C(P_{min})/P_1/T_0}$ parameterizes the single capacity scenario considered using the least probable capacity. In the example of Table 7.3, the lowest probability is 0.10 in row four, leading to a capacity of five. If there exists more than one scenario with the same minimum probability, the capacity average is used. If the resulting capacity value is not integer, it is brought down to the next integer.

$\boldsymbol{C(P_{max})/P_1/T_0}$ works similarly to $C(P^s)/P_1/T_0$, but parameterizes the single resulting scenario using the most probable capacity. For Table 7.3, this leads to a capacity of ten as 0.5 is the highest probability. If there exist more than one capacity with the same maximum probability, the arithmetic average is used.

$\boldsymbol{C^s/P^s/T_0}$ creates one scenario per potential capacity and assigns each the given probability. For the example in Table 7.3, this implies the probability for capacity ten to be 0.5, for capacity five $0.35(= 0.25 + 0.1)$ and for capacity twelve 0.15.

$\boldsymbol{C^s/P^s/T_{uni}}$ is similar to $C^s/P^s/T_0$ as it employs the full information on potential capacities and their probabilities. However, $C^s/P^s/T_{uni}$ assumes that capacities can be updated at any time with uniform probability. Given a time horizon of $|T| = 10$ for Table 7.3, this leads to a set of $|C| \cdot |T| = 3 \cdot 10 = 30$ scenarios. The probability of each capacity is distributed equally across days before departure; for instance, capacity five is expected to be announced with probability $0.5/10 = 0.05$ on each of the ten days before departure.

$\boldsymbol{C^s/P^s/T^s}$ uses full information to parameterize the scenario-based optimization model. Thus it takes all capacities, update probabilities and times into account.

Expecting different information levels to influence RM's result: more information on update scenarios should enhance anticipating capacity uncertainty. Those control strategies considering only one particular scenario are computationally more efficient in contrast to control strategies considering multiple scenarios. However, strategies that assume a single scenario do not exploit full information available; this could lead to inferior revenue results. Control strategy $C_{ini}/P_1/T_0$ constitutes the current industry standard by not considering potential capacity updates in advance. However, it remains interesting to see, which additional information on potential updates is more valuable: time or probability. It is expected that $C^s/P^s/T^s$ is superior to control strategies with only information on updates' times or probabilities.

Also the differences between information level 1, 2a and 2b are of high interest. The question arises, which information is more important to know for achieving a high revenue: update times, update probabilities or if only knowing the new capacity even suffices? These expectations and questions are addressed in detail by Part III.

Concluding, this chapter partially answers research question 2 by anticipating capacity updates in revenue optimization. A scenario-based formulation takes three dimensions of capacity updates into account. The knowledge on each of these dimensions determines different information levels; the more information available, the more actual parameters serve as input to the capacity update considering model. Finally, this chapter presents an upper bound, a benchmark strategy and one control strategy that takes the full information on capacity updates into account. Additionally, 13 further control strategies are established that cope with imperfect information (level 1–2b). To completely answer research question 2, integrating the capacity update considering model into an RM system is necessary. For this reason, the next chapter presents the simulation framework that allows performing computational studies addressing these questions.

8 Simulation Framework

The previous chapter introduced a scenario-based approach for modeling capacity uncertainty in revenue optimization. Based on this approach, this chapter introduces a simulation framework that supports answering the second part of research question 2.

RQ2.2: *How to model the capacity update characteristics in an RM system?*

As the last chapter established an optimization model including capacity update characteristics, this chapter models these characteristics in an entire system. While this work primarily focuses on an RM system's revenue optimization, this chapter's framework also describes further needed elements of a simplified RM system. Modeling the RM process in a simulation system allows including the realistic assumption of stochastic demand and stochastically occurring capacity updates. Therefore, the simulation system facilitates the computational studies of Part III, addressing research question 3 and 4.

The remainder of this chapter explains the simulation framework based on the process depict in Figure 8.1. Every area shaded in gray refers to a particular section number.

First, Section 8.1 explains simulation's demand generation and the demand forecasting process. Then Section 8.2 explains realizing and forecasting capacity update scenarios. Section 8.3 introduces the fare structure. Subsequent, Section 8.4 introduces the revenue optimization, the calculation of nested booking limits and different re-optimization routines. Section 8.5 presents the simulation system's customer acceptance decision and states result indicators that are needed for comparing simulation results: revenue, seat load factor and denied boardings. Last, Section 8.6 addresses the key points of implementing the simulation system.

8.1 Demand

Section 8.1.1 explains the system's demand generation process. Based on the resulting demand streams in Section 8.1.2, Section 8.1.3 shows forecasting demand, needed as input for revenue optimization.

8.1.1 Demand Generation

Based on the traditional RM model of T. Lee and Hersh (1993), one customer requests only one particular fare class. However, a choice-based customer model could also be implemented, e. g., explained in Talluri and van Ryzin (2004a). In the simulation system, customer cancellations are not taken into account due to the same reason as explained by Vulcano and Weil (2014, p. 9): "*To simplify the analysis, we assume that there are no cancellations during the reservation process. However, our results can be adapted to this case as long as the cancellations are independent among customers, which is a plausible assumption.*"

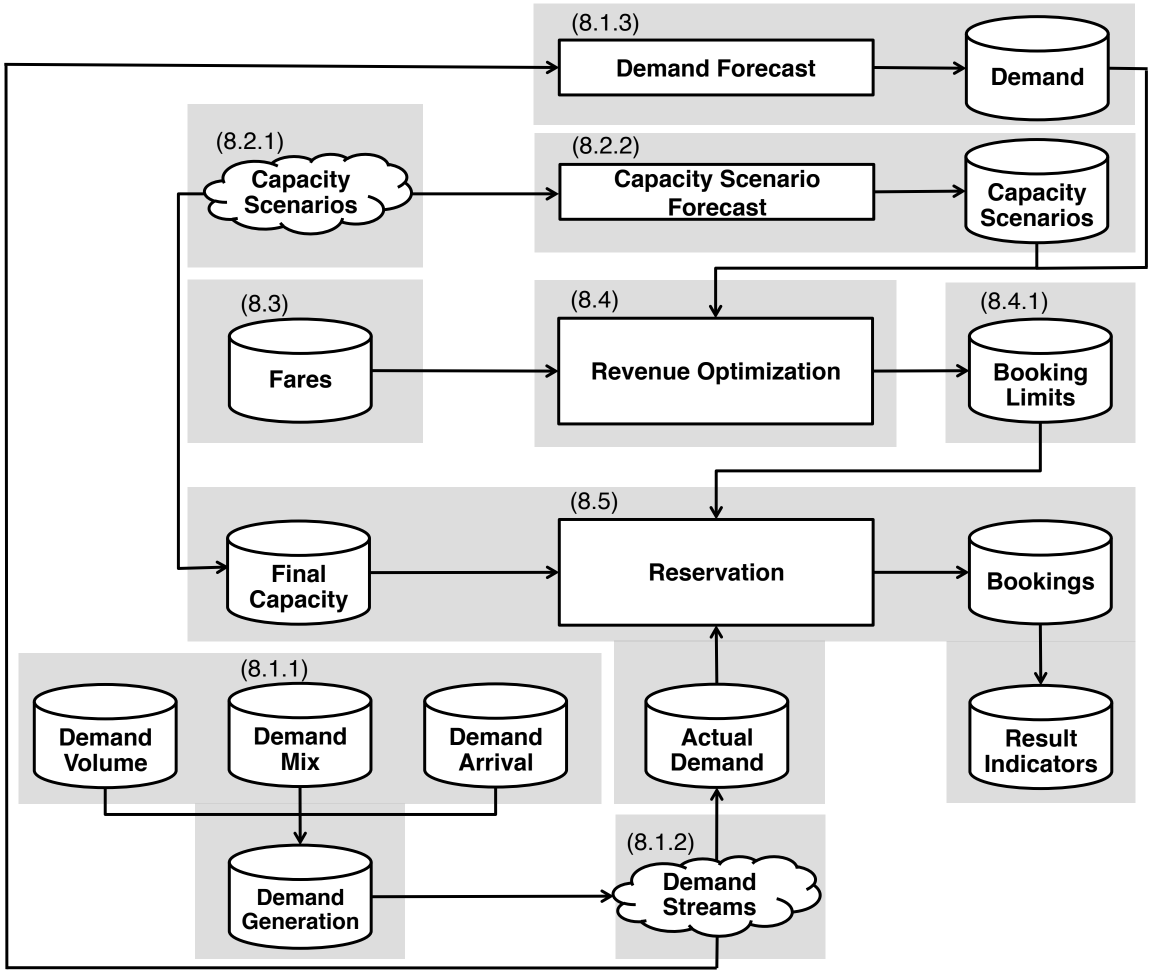

Figure 8.1: Process View of Simulation Framework

Demand arrivals are modeled as recommended by Kimms and Müller Bungart (2007): The simulation's demand generation uses a non-homogeneous Poisson process (NHPP). Assuming a particular customer arrival distribution allows parameterizing the time-dependent Poisson parameter $\lambda(t)$. For every combination of market and initial aircraft size, different demand volumes and demand mixes serve as input.

Demand volumes state the input of overall requesting customers $D \in \mathbb{Q}_{\geq 0}$, as a deterministic factor $v \in \mathbb{Q}_{\geq 0}$, of the initial aircraft's capacity $c^{ini} \in \mathbb{Q}_{\geq 0}$, so

$$D = \lceil v \cdot c^{ini} \rceil \,.$$

The resulting value is rounded to the next integer. A factor $v < 1$ determines that less customers request a ticket than seats are available on the initial aircraft. In contrast, a factor $v > 1$ implies excessive demand.

Demand mixes determine the overall demand's allocation across fare classes. A demand mix $m \in M$ consists of multiple demand mix factors m_f, $\forall\, f \in F$, so that

$$M = \left\{ (m_1, \ldots, m_{|F|}) \middle| m_f \in [0,1]), \sum_{f \in F} m_f = 1 \right\} \,.$$

The overall number of requesting customers D serves as input to the NHPP. Demand generation allows overall demand $\sum_{f \in F} \sum_{t \in T} d_{ft}$ to fluctuate with an a priori determined tolerance $\epsilon \in \mathbb{Q}_{\geq 0}$. This tolerance states the maximal accepted deviation from the initial demand input per fare class D_f in percent, thus

$$D_f \cdot (1 - \epsilon) \leq \sum_{t \in T} d_{ft} \leq D_f \cdot \epsilon, \quad \forall\, f \in F \,.$$

The demand generation hence produces a demand matrix $d_{ft} \in \mathbb{Q}_{\geq 0}$, stating the number of customers to request a ticket per fare class f and time slice t. The NHPP ensures that repetitively generating demand based on the same input values leads to different *demand streams*.

8.1.2 Demand Streams

For every combination $y \in Y$ of initial aircraft size, demand arrival per market, demand volume and demand mix, the simulation system generates multiple demand streams $z \in Z$. In the following, a single demand matrix d_{ft}^{yz} is represented by a combination of input values y and a particular stream z. Each demand matrix d_{ft}^{yz} represents one actual demand in the reservation stage of Section 8.5. Though, the demand forecast of one combination of input values y is based on all of its respective demand streams Z.

8.1.3 Demand Forecast

The optimization model assumes demand to be deterministic. Thus, a deterministic demand forecast is created from all artificial generated demand streams of a particular input combination y. Demand per input combination y sums-up all streams $z \in Z$ and is then divided by the number of streams

$|Z|$ and rounded to the next integer. The resulting demand forecast $\bar{d}^{y}_{ft}$ is calculated for every input combination y, so

$$\bar{d}^{y}_{ft} = \left\lceil \frac{\sum_{z \in Z} d^{yz}_{ft}}{|Z|} \right\rceil , \quad \forall\, y \in Y, \forall\, f \in F, \forall\, t \in T \,.$$

To ensure comparability over all control strategies, they all use the same demand forecast as input. As an exception, upper bound C^*, however, uses a perfect demand foresight on $d^{yz}_{ft}\, \forall\, y \in Y$, ensuring it being a *real* upper bound, which can never be outperformed, thus

$$\bar{d}^{yz}_{ft} = d^{yz}_{ft}, \quad \forall\, y \in Y, \forall z \in Z, \forall\, f \in F, \forall\, t \in T \,.$$

A computational study of this work additionally tests distorting demand forecast's volume and arrival times and analyzes the influence on control strategies' results.

8.2 Capacity

The simulation system requires capacity as input for revenue optimization. While in traditional RM one flight is assigned to one capacity, Section 7.1 demonstrated modeling capacity updates by scenarios. Section 8.2.1 explains the simulation's procedure of selecting capacity updates and Section 8.2.2 states the capacity forecast.

Note that this work considers two potential capacity update patterns. The first is based on the theoretical concept of Wang and Regan (2006), where exactly one capacity update can occur. If capacity is updated once, it is then fixed over time until departure. This pattern is used for the next chapter's validation study.

A more realistic capacity update pattern deals as input for the computational studies of Part III. Here, based on empirical data, multiple capacity updates can occur over time.

8.2.1 Capacity Scenarios

This work's simulation system realizes capacity update scenarios based on their stochastic probabilities. First, an initial aircraft size is assumed, which is capacity c^{ini}. Then, depending on the update pattern and given update probabilities, the simulation system draws a number of capacity updates. If no update takes place, the initial capacity reveals to be the final capacity – scenario 1 in the example of Table 7.3. Otherwise, one or more scenarios are drawn based on their update probability p^s. Thus, the drawn scenario states the update time t^s and new capacity c^s. For the example in Table 7.3, the simulation system could draw scenario 4, updating capacity to five seats two days before departure.

8.2.2 Capacity Scenario Forecast

First, the simulation system assumes perfect foresight on potential capacity updates. Thus, revenue optimization has perfect knowledge on potential capacity update scenarios – information level 3. Here, actual scenario capacities c^s, update times t^s and probabilities p^s serve as input for optimization.

The computational studies also test control strategies with reduced information, as introduced in Section 7.2 – information level 1, 2a and 2b. While scenario capacities are always known to the system, it can lack knowledge on potential update times and scenario probabilities.

However, as the assumption of perfect knowledge on potential update times and scenario probabilities can be doubted in practice, an additional computational study analyses the influence of a distorted capacity forecast.

8.3 Fares

Revenue optimization requires the input of fares $r_f \in \mathbb{Q}_{\geq 0}$ for each fare class f. The fare is the price customers have to pay for a ticket. Fare classes are ordered by revenue, so that r_1 always states the highest fare:

$$r_f > r_{f+1}, \quad \forall\, f \in F\,.$$

8.4 Revenue Optimization

Finally, control strategies parameterize capacity updates for the optimization model of Section 7.1. First, Section 8.4.1 calculates booking limits derived from optimization's result, then Section 8.4.2 introduces different re-optimization patterns.

8.4.1 Booking Limits

The simulation system requires controls to decide, which customer requests to accept in the reservation stage of Section 8.5. This work decides for using booking limits as decision criteria because they are intuitive to read and can, e. g., be easily transferred to protection levels (Talluri & van Ryzin, 2004b, Chapter 2, pp. 28). Every control strategy differently parameterizes the optimization model, which the simulation then solves in order to maximize revenue. As a result, the global number of tickets to sell, x_{ft}, are determined. To calculate booking limits, these contingents are summed up. As the simulation setup considers demand variation, the simulation calculates nested booking limits as they are preferable to partitioned booking limits: they always allow customers of higher fare classes to book a ticket as long as a lower fare class is available. The concept of nested booking limits is explained in Talluri and van Ryzin (2004b, Chapter 2, pp. 28).

Given fare classes ordered by revenue, $r_f > r_{f+1}$, booking limits B_f are calculated as

$$B_f = \sum_{f'=f}^{|F|} \sum_{t=\hat{t}}^{0} x_{f't}, \quad \forall\, f \in F\,.$$

Per control strategy, these booking limits serve as input to the simulation system's reservation stage of Section 8.5.

8.4.2 Re-Optimization

The simulation system contains different re-optimization patterns. All patterns have an initial optimization in common, performed at the beginning of the booking horizon. The re-optimization pattern depends on the capacity update pattern, explained in Section 8.2. The simulation can, but must not, re-optimize revenue if a capacity update occurs.

Next chapter's validation study implements two patterns: one that does not allow re-optimizing and one that does. The simulation system without re-optimization only optimizes revenue once and thus determines booking limits solely before the booking horizon begins. Once a capacity update occurs, the simulation system with re-optimization maximizes revenue anew based on updated capacity and booking information. Thus, the simulation can adjust booking limits exactly once, directly after the occurrence of an update.

This work's computational studies, however, allow multiple re-optimizations over time. Every time capacity is updated, the model also updates information on possible future capacity update scenarios and maximizes revenue anew, potentially leading to adjusted booking limits $B_{f^*} \, \forall \, f \in F$.

8.5 Reservation

The simulation system's reservation stage accepts and rejects customer requests. Here, actual demand and booking limits are compared: As long as booking limits per fare class permit, the simulation accepts customer requests. Throughout the time horizon from $\hat{t}$ to 0, the number of bookings $\tau_f \in \mathbb{Q}_{\geq 0}$ is tracked per fare class f. For each time slice t and fare class f, reservation decides on a customer request $d_{ft} > 0$, based on

$$B_f - \tau_f \begin{cases} > 0, & \text{accept customer} \\ = 0, & \text{reject customer} \end{cases} \qquad \forall \, f \in F \, .$$

At departure, no more customer requests can be accepted. For each simulation run, the tracked bookings enable calculating three fundamental result indicators: *denied boardings*, *revenue* and *seat load factor*. In the computational studies, result indicators are presented as average over all simulation runs.

Denied boardings $\delta \in \mathbb{N}_{\geq 0}$ occur if the overall number of bookings exceeds the final capacity c^*. Denied boardings can never be negative and are given as

$$\delta = \max \left\{ \sum_{f \in F} \tau_f - c^*, \quad 0 \right\} \, .$$

Revenue $\pi \in \mathbb{Q}_{\geq 0}$, of one combination of simulation run and control strategy, is the product of fares and bookings per class, minus denied boarding costs, so that

$$\pi = \sum_{f \in F} r_f \cdot \tau_f - \sum_{a=1}^{\delta} k_a \, .$$

Seat load factor $\sigma \in \mathbb{Q}_{\geq 0}$ is calculated by dividing bookings by the final capacity. This indicator is capped at 100%, as bookings exceeding capacity produce denied boardings:

$$\sigma = \min \left\{ \frac{\sum\limits_{f \in F} \tau_f}{c^*}, \quad 1 \right\} .$$

While revenue states a control strategies success, seat load factors below 1 indicate spoilage resulting from free seats and denied boardings above 0 indicate spill by selling too many tickets.

8.6 Implementation

This chapter presented a simulation framework for considering capacity updates in RM. Based on this framework, a simulation system is build in Java SE 8. The system includes the optimization model presented in Section 7.1, which is solved by IBM ILOG Cplex Optimizer 12.6.1 on an Intel Core i5 2.8 GHz with 16 GB RAM.

For performing computational studies in Part III, this work requires the simulation system's implementation. First, a validation study tests and verifies the simulation system based on artificial input data. Subsequent, several computational studies, calibrated on empirical data of Chapter 6, test possible control strategies for coping with capacity uncertainty and analyze the influence of different problem parameters. Additionally, Part III provides a computational study on the robustness of the best performing control strategies.

Part III

Experiments and Conclusion

9 Validation Study

This chapter presents a validation study based on the developed simulation framework of Chapter 8. The study is motivated by two questions: Does the simulation work free from errors and do capacity updates in the simulation have the same effects as addressed in theory? Thus, this chapter's objective is to validate and verify the simulation system's implementation as well as the credibility of simulation results.

It focuses on two re-optimization patterns – with or without re-optimizing – and two capacity update directions – increase or decrease. The study compares control strategies C^*, $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$ and gives an overview of basic impacts on their performance. First, Section 9.1 states validation study's setup. Then, Section 9.2 and 9.3 formulate various expectations to the simulation system that are subsequently discussed and confirmed by results depicted by multiple visualizations.

Parameter	Value
Booking Horizon	$T = \{30, 29, \dots, 0\}$
Fare Classes	$F = \{1, 2, 3\}$
Fares	$r_f = \{200, 150, 100\}$
Denied Boarding Costs	$k_a = 201 \cdot 1.1^{a-1}$, $\forall\, a = \{1, \dots, \hat{a}\}$
Control Strategies	C^*, $C_{ini}/P_1/T_0$, $C^s/P^s/T^s$
Demand Forecast	perfect $\Rightarrow \bar{d}_{ft} = d_{ft}$, $\forall f \in F$, $\forall\, t \in T$
Capacity Scenario Forecast	perfect
Capacity Updates	one possible update
Re-Optimization	no re-opt. (Section 9.2); re-opt. (Section 9.3)
Simulation Runs	1000 runs per instance
Update Probability	varying from 0%–100% or fixed at 50%
Update Time	varying from 30–0 dbd or fixed at 5 dbd
Capacity	$c^1 = 40$, decrease: $c^2 = 30$; increase: $c^2 = 50$
Demand Volume	decrease: $\sum\limits_{f \in F} \sum\limits_{t \in T} d_{ft} = 48$; increase: $\sum\limits_{f \in F} \sum\limits_{t \in T} d_{ft} = 60$
Demand Mix	$m = (0.25, 0.25, 0.5)$
Demand Arrival	decrease: $U_{(f=1)}(12, 1)$, $U_{(f=2)}(17, 6)$, $U_{(f=3)}(22, 11)$ increase: $U_{(f=1)}(15, 1)$, $U_{(f=2)}(20, 6)$, $U_{(f=3)}(25, 11)$

Table 9.1: Parameterization of Validation Study

9.1 Input Data and Setup of Validation Study

Table 9.1 summarizes validation study's input values discussed in the following sections. The study assumes demand to be deterministic and known in advance. The validation thus allows for a clear view that is unaffected by uncertainty in demand forecast. Neither demand nor capacity calibration is based on the analyzed empirical data of Chapter 6. For every mix of input values, the system performs 1,000 simulation runs and presents the average results over all runs. Statistical convergence was given in all cases.

9.1.1 Demand Arrival

Customer arrival is based on the classic RM assumption that demand of lower fare classes arrives before demand of higher fare classes (Talluri & van Ryzin, 2004a, Chapter 1). However, the validation study allows customer requests of different fare classes to occur at the same time.

The study separately tests capacity increases and capacity decreases. Overall demand is based on the same demand volume factor as explained in the following section. This factor applies both to capacity increases and capacity decreases, so that overall demand and demand arrival times slightly differ.

If capacity decreases can occur, customers of fare class 3 are uniformly distributed between 22 and eleven days before departure, those of fare class 2 between 17 and six days before departure and those of fare class 1 between twelve and one day before departure. As an example, the uniform distributed demand arrival of fare class 1 for a capacity decrease is abbreviated as $U_{(f=1)}(12, 1)$.

If capacity increases can occur, customers of fare class 3 are uniformly distributed between 25 and eleven days before departure, those of fare class 2 between 20 and six days before departure and those of fare class 1 between 15 and one day before departure.

9.1.2 Demand Volume and Mix

The validation study uses the same demand mix as Wang and Regan (2006, Example 1): (0.25, 0.25, 0.5). This indicates, that 25% of all customers request a ticket in fare class 1 and 2 each, while 50% request a ticket in the cheapest fare class 3. Overall demand equals the highest possible capacity, multiplied by factor $v = 1.2$. Preliminary studies showed this factor to enable clearest results.

9.1.3 Capacity Updates and Capacity Scenario Forecast

The booking horizon starts 30 days before departure and includes one potential capacity update. Based on Wang and Regan (2006), the initial capacity is 40. A potential capacity decrease to 30 seats (Wang & Regan, 2006, Example 1) and an increase to 50 seats (Wang & Regan, 2006, Example 3) are separately tested. The simulation assumes a perfect capacity forecast to prevent distorting effects.

9.1.4 Fares and Denied Boarding Costs

The same fare structure as proposed by Wang and Regan (2006) is applied. The respective fare revenues are $r_1 = 200, r_2 = 150$ and $r_3 = 100$. Same as in Wang and Regan (2006) and Vulcano and Weil (2014), the first denied boarding cost is modeled to be higher than the highest fare class price: $k_1 = 201 > 200 = r_1$. Denied boarding costs increase exponentially by a factor of 1.1, so that $k_1 \leq \ldots \leq k_{\hat{a}}$. Thus, the second denied boarding costs $k_2 = k_1 \cdot 1.1 = \lceil 201 \cdot 1.1 \rceil = 221$ and $k_a = 201 \cdot 1.1^{a-1}$ for further denied boardings.

9.1.5 Re-Optimization

If a capacity update occurs, the validation study tests two different responses: no re-optimization in Section 9.2 and re-optimization in Section 9.3.

The optimization model explicitly accounts for re-optimizing by switching from the global to a scenario-based strategy. Thus, without re-optimizing, we expect revenue from $C^s/P^s/T^s$ to be worse compared to the simulation with re-optimizing and probably the same applying to $C_{ini}/P_1/T_0$.

9.2 One Update without Re-Optimization

In this section, one capacity update can occur. However, if capacity is updated, no re-optimization takes place. Thus, booking limits are only set at the beginning of the booking horizon and never adjusted. Section 9.2.1 analyzes potential capacity decreases and Section 9.2.2 potential increases.

9.2.1 Capacity Decrease

With increasing probability for a capacity decrease, we expect the following:

Expectation 1. C^*'s revenue decreases as the average number of sellable seats decreases.

Expectation 2. $C_{ini}/P_1/T_0$'s revenue decreases stronger than C^* due to denied boardings.

Expectation 3. $C^s/P^s/T^s$'s revenue is equal or higher than $C_{ini}/P_1/T_0$'s as $C^s/P^s/T^s$ produces less denied boardings by anticipating potential decreases.

Figure 9.1 illustrates control strategies' average revenue if capacity can decrease from 40 to 30 seats five days before departure.

The ordinate covers the probability for a capacity update. A gray dashed line illustrates the average revenue from C^*; the result is indicated as a factor of its highest value (1.0). Revenue from C^* is highest if no capacity update occurs (0%). If a capacity decrease becomes more probable, revenue from C^* decreases approximately linear (Expectation 1). The black lines represent average revenue of control strategies $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$. Their results are specified in percent of C^*.

If no capacity update can occur (0%), $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$ perform equally to C^*. All control strategies try selling 40 seats as knowing demand in advance. For capacity update probability 0%–30%, $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$ set the same booking limits leading to same revenue results.

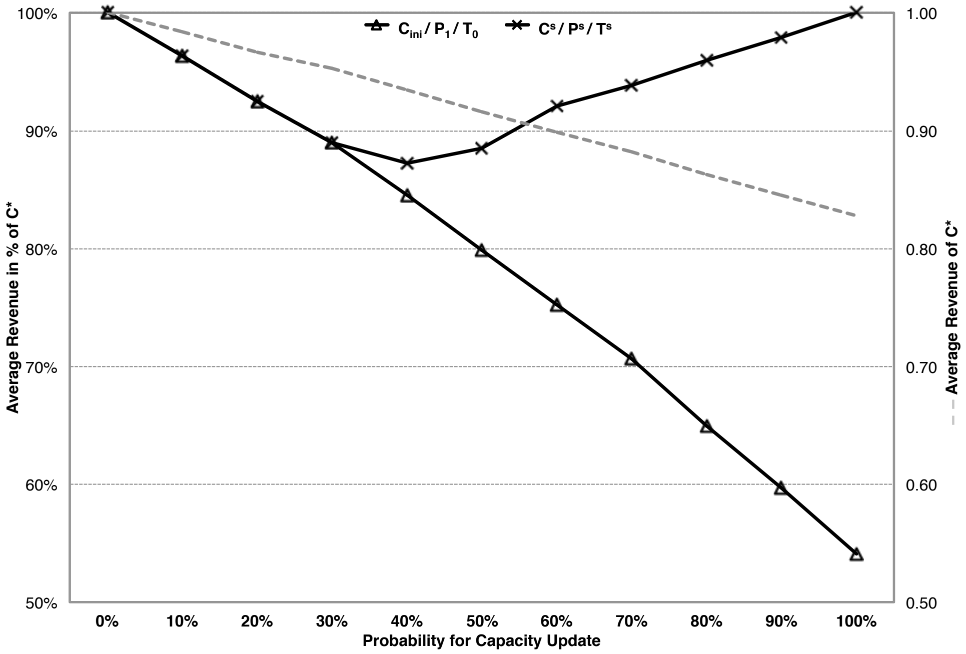

Figure 9.1: One Update without Re-Optimization – Varying Probability of Capacity Update from 40 to 30 for Time 5 dbd

However, if update probability varies between 40% and 100%, $C^s/P^s/T^s$ leads to a revenue higher than $C_{ini}/P_1/T_0$, due to its more conservative booking limits: $C^s/P^s/T^s$ anticipates the potential capacity decrease and offers less tickets than $C_{ini}/P_1/T_0$. In contrast, $C_{ini}/P_1/T_0$ always offers 40 tickets. This leads to an on average higher number of denied boardings (Expectation 2).

The higher the probability for a decrease, the better performs $C^s/P^s/T^s$ (Expectation 3) by anticipating these decreases and the worse performs $C_{ini}/P_1/T_0$. Strategy $C^s/P^s/T^s$ leads to the same revenue result as C^* if a capacity decrease to 30 seats is certain (100%), because here, both strategies try selling 30 seats to the highest valuable customers. Here, $C_{ini}/P_1/T_0$ sells 40 tickets, which always leads to ten denied boardings, resulting in the lowest average revenue of approx. 54%.

With varying update time of a capacity decrease we expect:

Expectation 4. C^*'s revenue is unaffected by varying update times.

Expectation 5. $C_{ini}/P_1/T_0$'s revenue is unaffected by varying update times, but performs worse than C^*.

Expectation 6. $C^s/P^s/T^s$' advantage over $C_{ini}/P_1/T_0$ rises with updates closer to departure.

In Figure 9.2, the update time of a potential capacity decrease from 40 to 30 varies. The update probability is fixed to 50% and the update time varies from 30 to one day before departure.

C^* performs identically (always 1.0), independent of capacity update times (Expectation 4); the upper bound always optimizes revenue based on the final capacity not taking update times into account. From 30 to eight days before departure, $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$ perform equally just below 80% of C^*. Upper bound C^*, however, has an information advantage by knowing final capacity in advance. Both $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$ lead to ten denied boardings if a capacity decrease

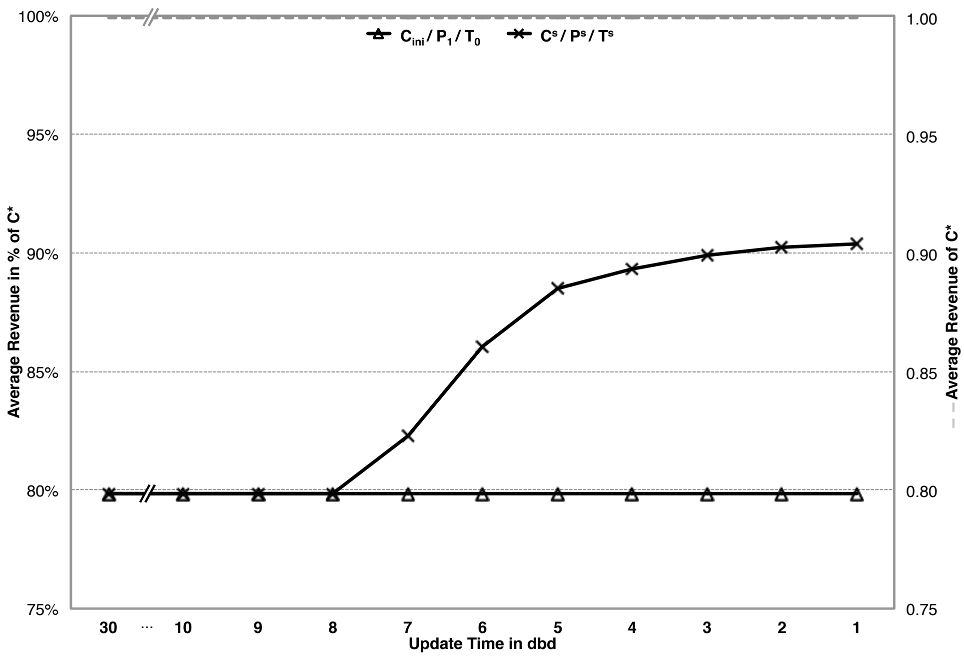

Figure 9.2: One Update without Re-Optimization – Varying Update Time of Capacity Update from 40 to 30 with Probability 50%

occurs. Here, $C_{ini}/P_1/T_0$ remains constant (Expectation 5). However, from update time seven to one day before departure, $C^s/P^s/T^s$ outperforms $C_{ini}/P_1/T_0$ – the closer to departure, the higher the gap (Expectation 6). The later a potential update occurs, the more conservative are $C^s/P^s/T^s$' booking limits.

In general, no bookings from 30 to 23 days before departure could be observed. This is the model's intended behavior as no customers request at that time in any fare class. As explained in Section 9.1.1, the first customers request 22 days before departure.

9.2.2 Capacity Increase

With rising probability for a capacity increase, we expect the following:

Expectation 7. C^*'s revenue increases as the average number of sellable seats increases.

Expectation 8. $C_{ini}/P_1/T_0$'s revenue is unaffected by capacity increases but its performance decreases compared to C^*.

Expectation 9. $C^s/P^s/T^s$' revenue is equal or higher than $C_{ini}/P_1/T_0$'s; the higher the increase probability, the better performs $C^s/P^s/T^s$.

Figure 9.3 illustrates the probability of capacity *increases* from 40 to 50 seats. Here, the same setup of capacity update probabilities and times are applied.

Again, a gray dashed line represents the average revenue from upper bound C^* as a factor of its highest value. The curve rises approximately linear with increasing probability for a capacity increase

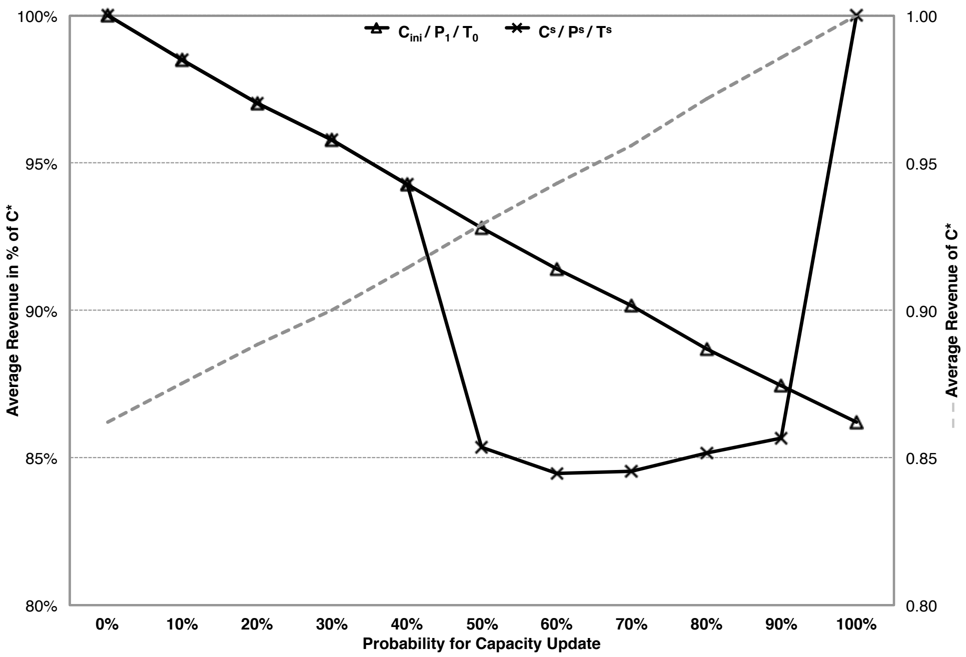

Figure 9.3: One Update without Re-Optimization – Varying Probability of Capacity Update from 40 to 50 for Time 5 dbd

(Expectation 7). This seems logical, as the higher the chance for selling 50 seats the higher C^*'s average revenue.

Control strategies $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$ perform equally for update probability 0%–40%. Both strategies are represented by black lines. $C_{ini}/P_1/T_0$'s performance decreases with an increasing probability for a capacity update. While its average revenue remains constant, from always selling 40 seats, its relative revenue to C^* decreases as C^* increases (Expectation 8).

For capacity update probability 50%–100%, revenue results for $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$ differ strongly. For probability 50%–90%, $C^s/P^s/T^s$ performs clearly worse than $C_{ini}/P_1/T_0$ (contrary to Expectation 9). This effect can be traced back to not re-optimizing if a capacity update occurs. $C^s/P^s/T^s$ sets higher booking limits for the beginning of the booking horizon as it anticipates a potential capacity increase. However, the model bases its success on switching from the global to a scenario-based strategy. This switch is only applicable if a capacity update leads to re-optimizing.

Section 9.2 does not allow for re-optimizing, which explains $C^s/P^s/T^s$' unexpected performance. If the anticipated capacity increase is absent, $C^s/P^s/T^s$' unadjusted booking limits lead to denied boardings. In contrast, if probability for an increase is 100%, $C^s/P^s/T^s$ is as successful as C^*. Then both control strategies optimize revenue for 50 anticipated seats, outperforming $C_{ini}/P_1/T_0$.

With varying update time of a capacity increase, we expect:

Expectation 10. C^*'s revenue is unaffected by varying update times.

Expectation 11. $C_{ini}/P_1/T_0$'s revenue is unaffected by varying update times, but performs worse than C^*.

Expectation 12. $C^s/P^s/T^s$' advantage over $C_{ini}/P_1/T_0$ increases with updates closer to departure.

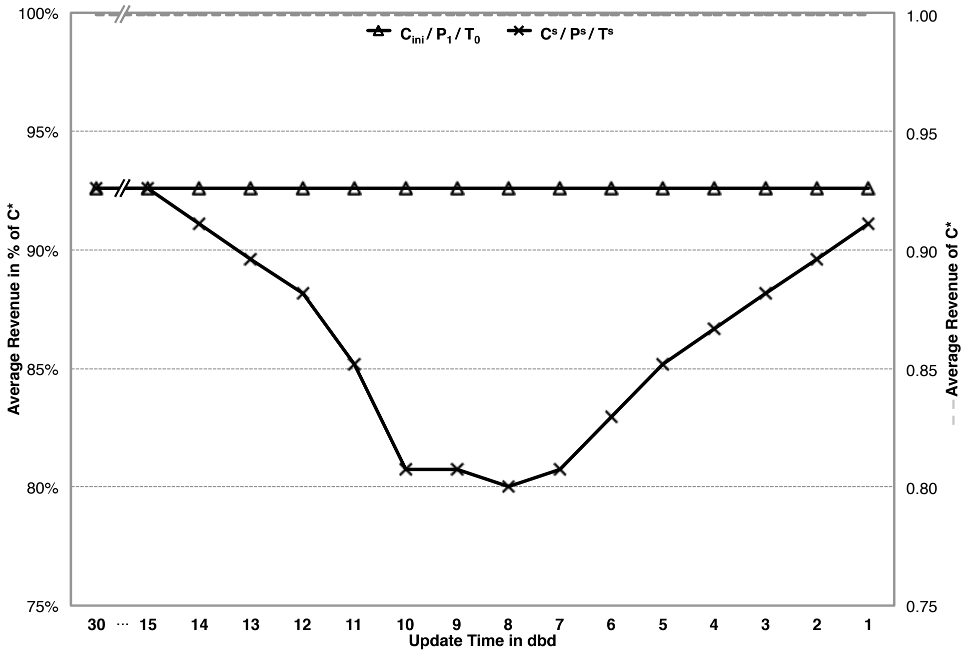

Figure 9.4: One Update without Re-Optimization – Varying Update Time of Capacity Update from 40 to 50 with Probability 50%

Figure 9.4 shows the influence of different potential capacity update times. Here, update probability is held constant at 50% and update time varies from 30 to one day before departure.

As both C^* and $C_{ini}/P_1/T_0$ optimize revenue for one particular capacity and re-optimizing is neglected, their achieved revenue is independent of the potential update time. While C^* performs constantly at factor 1.0 (Expectation 10), $C_{ini}/P_1/T_0$ always sells 40 seats, leading to an average revenue of approx. 93% of C^* (Expectation 11).

From 30 to 15 days before departure, $C^s/P^s/T^s$ and $C_{ini}/P_1/T_0$ perform identically. Here, both control strategies set the same booking limits. However, $C^s/P^s/T^s$ misses the possibility of re-optimizing. From 14 to one day before departure, $C^s/P^s/T^s$' booking limits are higher than that of $C_{ini}/P_1/T_0$. But, if an update occurs, $C^s/P^s/T^s$ cannot adjust its booking limits based on the new capacity information. This leads to a lower revenue performance than $C_{ini}/P_1/T_0$ (contrary to Expectation 12). The downslope of $C^s/P^s/T^s$ proceeds till the update time of ten days before departure. For fare class 3, the last customer requests eleven days before departure. The curve then rises to approx. 92% of C^*, as booking limits almost converge to those of $C_{ini}/P_1/T_0$.

Concluding, Section 9.2 validates expected effects of different capacity update probabilities, times and capacities. In all tests, C^* constitutes an actual upper bound. While Section 9.2.1 illustrates the potential superiority of $C^s/P^s/T^s$ to $C_{ini}/P_1/T_0$, Section 9.2.2 clarifies that the optimization model requires a re-optimization step when capacity is updated. Otherwise, $C^s/P^s/T^s$ fails its optimal revenue result. For that reason, the following section discusses control strategies' performance when the simulation enables re-optimizing.

9.3 One Update with Re-Optimization

This section tests control strategies under the same assumption of Wang and Regan (2006): One potential capacity update can occur leading to a re-optimization. With respect to airline practice, this also constitutes a more realistic setup: Based on RM expert knowledge, airlines consider new capacity information after an update occurred. The updated capacity then serves as input for the re-optimization step. In practice, the date of re-optimizing is usually fixed and scheduled in advance; the simulation, however, re-optimizes immediately after an update.

The same update probabilities, capacities, update times and demand streams of Section 9.2 serve as input. The only difference to the previous section is that after an update the simulation performs a re-optimization. As C^* constitutes the upper bound that never has to re-optimize, this section primarily highlights the performance of $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$.

9.3.1 Capacity Decrease

With increasing probability for a capacity decrease, compared to not re-optimizing, we expect the following:

Expectation 13. C^* is unaffected by re-optimization.

Expectation 14. $C_{ini}/P_1/T_0$'s performance increases by re-optimization.

Expectation 15. $C^s/P^s/T^s$'s performance increases by re-optimization.

Figure 9.5 shows the influence of different update probabilities if capacity can decrease from 40 to 30 seats, five days before departure.

In both figures – Figure 9.1 and 9.5 – C^* performs equally (Expectation 13). The highest revenue is reached if no capacity decrease takes place (0%) and lowest if a decrease is certain (100%). As C^* optimizes with regard to the actual final capacity, a possible re-optimization does not have an influence on its result, which is as expected.

In contrast to C^*, $C_{ini}/P_1/T_0$ as well as $C^s/P^s/T^s$ should be affected by a possible re-optimization step. The downslope of $C_{ini}/P_1/T_0$ with rising probability for a decrease is similar to that of Figure 9.1 without re-optimizing. However, its overall performance is much better (Expectation 14). At worst, $C_{ini}/P_1/T_0$ achieves a revenue of approx. 54% compared to C^* without re-optimizing and approx. 75% with re-optimizing.

The same effect can be observed for $C^s/P^s/T^s$: Its curve shape is similar and the lowest revenue rises from 87% without to 91% with re-optimization (Expectation 15). This is the expected behavior, as both control strategies benefit from re-optimizing.

With varying update time of a capacity decrease, compared to not re-optimizing, we expect the following:

Expectation 16. C^* is unaffected by re-optimization.

Expectation 17. $C_{ini}/P_1/T_0$'s performance varies and increases by re-optimization.

Expectation 18. $C^s/P^s/T^s$'s performance increases by re-optimization and its advantage over $C_{ini}/P_1/T_0$ rises with updates closer to departure.

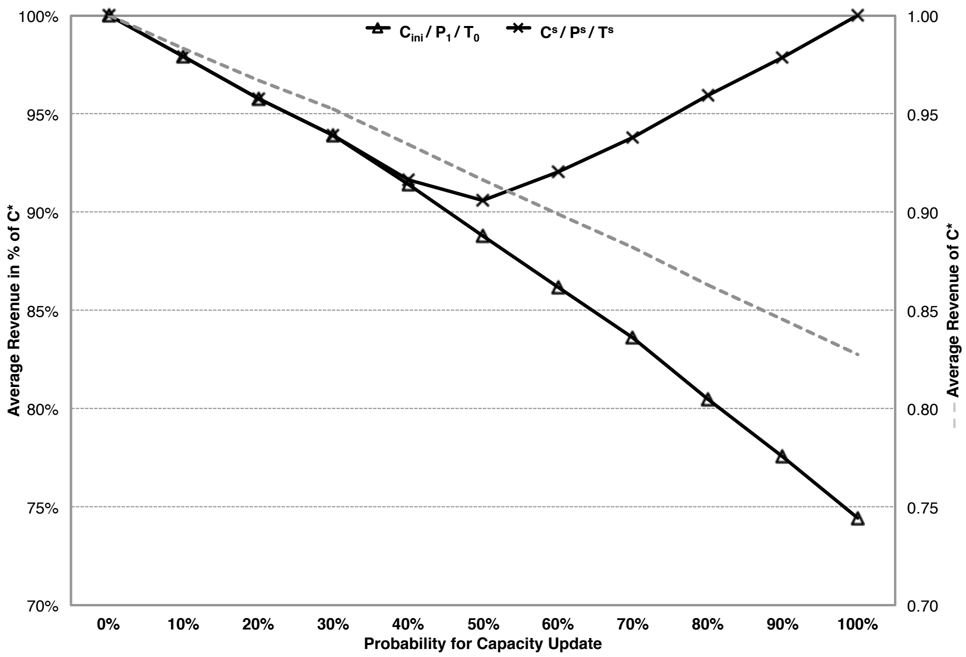

Figure 9.5: One Update with Re-Optimization – Varying Probability of Capacity Update from 40 to 30 for Time 5 dbd

Figure 9.6 illustrates the influence of different update times for one potential update with re-optimization.

Here, as well as in Figure 9.2, C^* achieves a constant revenue as it does not rely on re-optimization. The curves of $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$ have a completely different shape than without re-optimization. The most obvious difference is that both control strategies achieve the same revenue as C^* if the capacity update occurs in the interval from 30 to 19 days before departure. Within this time, both control strategies behave identical, as only few customers of fare class 3 request tickets. The more fare class 3 customer requests arrive, the more probable will they *steal* potential tickets for fare class 2 customers leading to a revenue loss. Thus, for both $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$, revenue decreases for update times closer to departure.

$C_{ini}/P_1/T_0$ does not perform on a constant level, compared to Figure 9.2, as it now benefits from re-optimizing (Expectation 17). Even the lowest revenue from $C_{ini}/P_1/T_0$ with re-optimization (approx. 82%) is higher than the constant revenue without re-optimization (approx. 80%). For $C^s/P^s/T^s$, even the lowest performance with re-optimization is still as successful as the highest without re-optimization – approx. 90% of C^* (Expectation 18). This reveals both control strategies to be superior with re-optimization. Last but not least, Figure 9.6 shows that $C^s/P^s/T^s$ first performs better than $C_{ini}/P_1/T_0$ for seven days before departure, which is one day closer to departure without re-optimization.

9.3.2 Capacity Increase

Within this section, the same input values as in Section 9.3 are applied: five days before departure capacity can be updated from 40 to 50 seats. With increasing probability for a capacity increase,

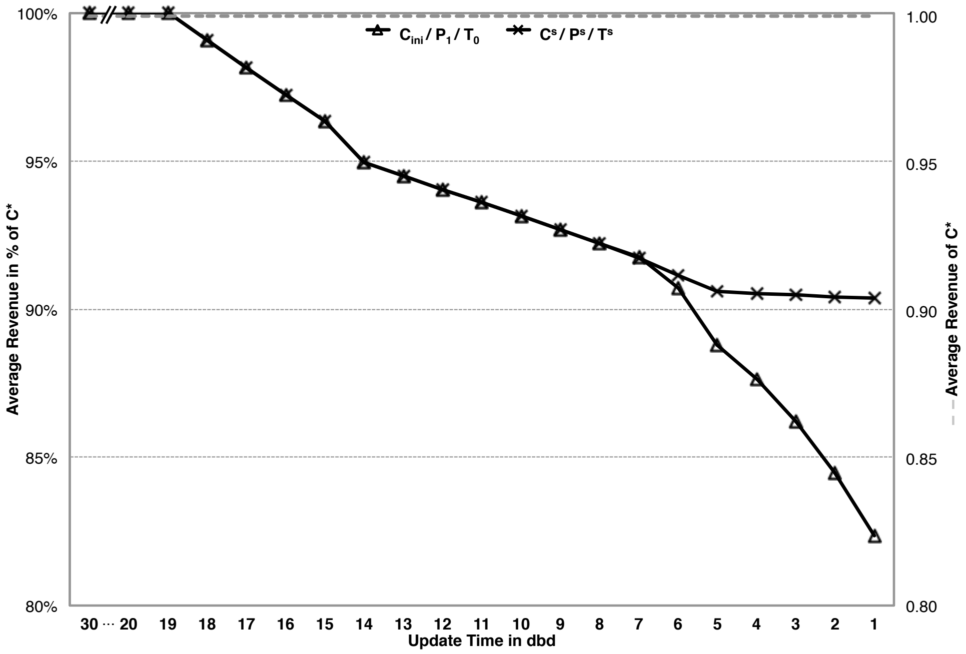

Figure 9.6: One Update with Re-Optimization – Varying Update Time of Capacity Update from 40 to 30 with Probability 50%

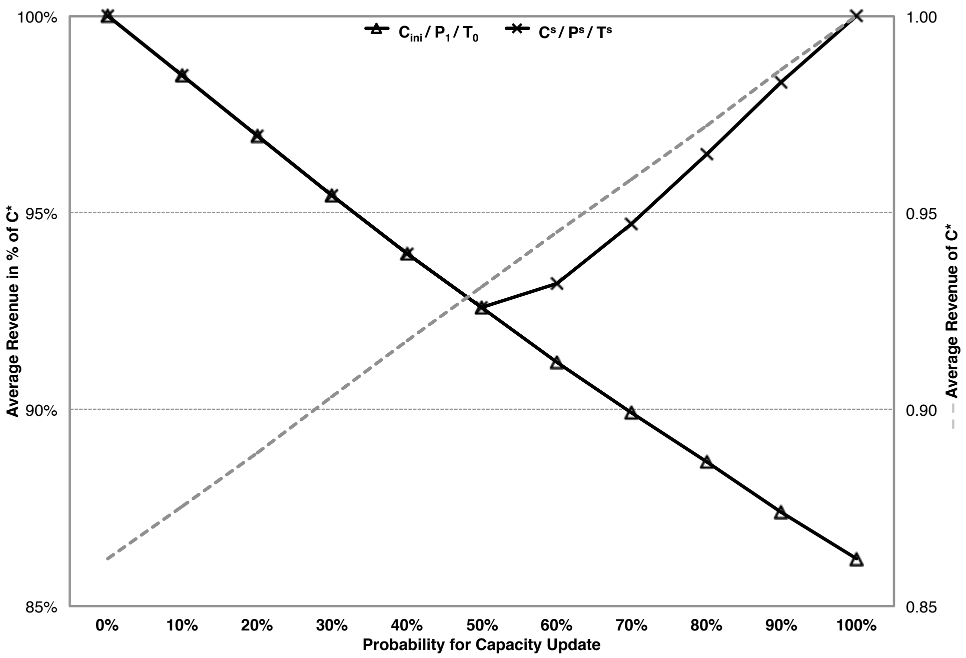

Figure 9.7: One Update with Re-Optimization – Varying Probability of Capacity Update from 40 to 50 for Time 5 dbd

compared to not re-optimizing, we expect the following:

Expectation 19. C^* is unaffected by re-optimization.

Expectation 20. $C_{ini}/P_1/T_0$'s revenue increases by re-optimization.

Expectation 21. $C^s/P^s/T^s$'s revenue increases by re-optimization and is equal or higher than $C_{ini}/P_1/T_0$'s.

Figure 9.7 shows the revenue from C^*, $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$ with respect to different capacity update probabilities.

C^*'s average revenue increases with rising probability for an update, identical to Figure 9.3. As obvious, re-optimizing does not have an influence on the result of C^* (Expectation 19).

However, for $C^s/P^s/T^s$, re-optimizing enables a higher revenue, which increases as probability for an update rises. Without re-optimization, $C^s/P^s/T^s$ always achieves a revenue smaller than approx. 93% of C^* for update probabilities 40%–90%; with re-optimization its revenue always stays above this value (Expectation 21).

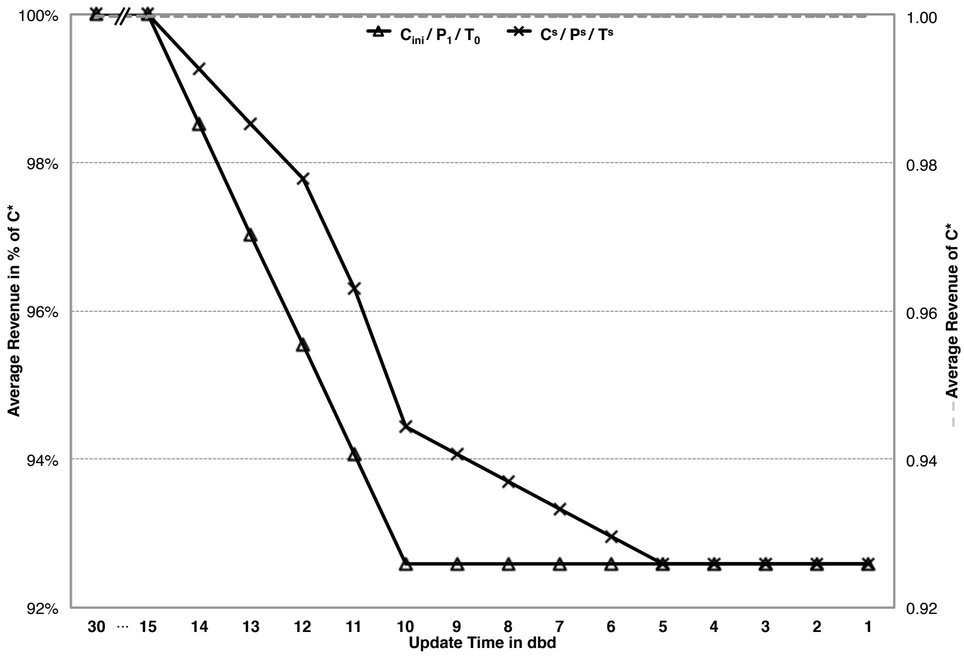

Figure 9.8: One Update with Re-Optimization – Varying Update Time of Capacity Update from 40 to 50 with Probability 50%

In contrast, $C_{ini}/P_1/T_0$ performs equally to the case where no re-optimizing takes place (contrary to Expectation 20). This is surprising, as a benefit from re-optimizing could be expected. However, the update time, five days before departure is too late for $C_{ini}/P_1/T_0$ to utilize the information on ten additional seats as too few customers request. When analyzing the influence of tested update times, this behavior is explained and justified in more detail.

Figure 9.8 illustrates this influence. With varying update time of a capacity increase, compared to not re-optimizing, we expect:

Expectation 22. C^* is unaffected by re-optimization.

Expectation 23. $C_{ini}/P_1/T_0$'s revenue increases by re-optimization.

Expectation 24. $C^s/P^s/T^s$'s revenue increases by re-optimization and is equal or higher than $C_{ini}/P_1/T_0$'s.

As before, C^* is not affected by re-optimization and thus remains on a constant revenue level of 1.0 (Expectation 22).

From 30 to 15 days before departure, $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$ achieve the same revenue as C^* (100%) because early updates allow them to adjust booking limits based on the final capacity. On the interval from 14 to ten days before departure, $C_{ini}/P_1/T_0$'s revenue decreases; from nine to one day before departure it remains constant on the lowest revenue level in relation to C^* (93%). $C_{ini}/P_1/T_0$'s performance is always higher or at least as high as without re-optimization (Expectation 23).

Revenue from $C^s/P^s/T^s$ also decreases from 14 days before departure. However, it reaches its lowest level, five days before departure. This is five days closer to departure than $C_{ini}/P_1/T_0$. From five to one day before departure, $C^s/P^s/T^s$ remains constant on the same revenue level as $C_{ini}/P_1/T_0$ (93%). Thus, $C^s/P^s/T^s$ is only superior to $C_{ini}/P_1/T_0$ on the interval 15 to five days before departure. Earlier updates do not cause a difference, because too few customers arrive that early; later updates do not affect both strategies as only customers of the highest fare class request a ticket, who are allowed to sell a ticket either way. Therefore, $C^s/P^s/T^s$ always performs as good as or better than $C_{ini}/P_1/T_0$ (Expectation 21).

Concluding, Section 9.3 validated expected influences of different capacity update probabilities, times and capacities when enabling re-optimization. The next section summarizes the validation study's most important insights.

9.4 Validation Summary

Section 9.2 validated the expected effects of different capacity update probabilities, times and capacities. However, it also clarified that the optimization model requires a re-optimization when capacity is updated. Otherwise, $C^s/P^s/T^s$ misses its optimal revenue result. For that reason, Section 9.3 discussed the differences compared to Section 9.2, when the simulation enables re-optimizing. The expected benefits could be confirmed.

Both sections validated the simulation's basic properties. Bookings only occurred within the booking horizon; a customer request is only accepted if permitted by a control strategy's booking limits. In general, no unexpected simulation behavior could be observed.

With respect to capacity updates' magnitude, the same effects as in Wang and Regan (2006) arose: capacity increases enable the possibility of higher revenues and capacity decreases can lead to revenue losses. Also can update time and update probability affect the booking limits of a control strategy that anticipates potential capacity updates.

Again, in all tests, C^* constituted an actual upper bound. Under no circumstance could a control strategy achieve a higher revenue than C^*. The upper bound does not rely on re-optimizing due to its information advantage knowing the final capacity in advance. The behavior of $C^s/P^s/T^s$ is as expected: It always performed higher than or at least equally to $C_{ini}/P_1/T_0$. However, the chapter also showed that $C^s/P^s/T^s$ is only beneficial if the simulation actually re-optimizes after an update.

This chapter's study validated and verified the simulation based on artificial demand inputs and artificial capacity update inputs. This clears the way for computational studies, based on empirical airline data previously analyzed. Therefore, Section 10.1 introduces parameterizing the computational studies that follow in Chapter 10 and Chapter 11.

10 Computational Study on Value of Information and Problem Parameters

This chapter presents a computational study on the value of information and on the influence of problem parameters. While the first part addresses the results of different information levels, the second analyzes, which optimization model's problem parameters influence revenue results. The study is based on the simulation framework established in Chapter 8. This chapter's first section, presents the parameterization of the computational study. Here, the results of analyzing empirical data in Chapter 6 are used for calibrating the study's demand and capacity update scenarios. Section 10.2 primarily focuses on answering research question 3.

RQ3: *What are the benefits of anticipating capacity updates? Are some capacity update characteristics more important to know than others?*

To answer this question, the computational study tests 16 control strategies, introduced in Section 7.3. The strategies' results are compared with respect to their underlying information level. Thus, the study reveals the benefits of systematically anticipating capacity updates. Comparing the best performing control strategies per information level allows determining, which capacity update characteristic is more valuable to know: update time or update probability. Finally, Section 10.3 allows answering the first part of research question 4.

RQ4.1: *Which factors influence RM's result if capacity updates are anticipated?*

Here, the computational study's results are analyzed with respect to different problem parameters; multiple visualizations support comparing numerical differences. Section 10.3 addresses those parameters with an influence on the revenue result of the best performing control strategies. These parameters are: demand volume, demand mix, number of capacity updates per flight, final capacity update time as well as update direction and magnitude.

10.1 Input Data and Setup of Computational Study

This section describes parameterizing input values for this chapter's computational study. Altogether, the study tests 2,230,800 combinations of control strategies and instances. Table 10.1 shows the computational study's parameterization. The particular values and their origin are discussed in this section's remainder.

At first, Section 10.1.1 presents demand input and Section 10.1.2 presents simulation's capacity input. Both parameterizations are derived from analyzing the empirical data in Chapter 6. Section 10.1.3 introduces input on fares and denied boardings derived from existing research.

Parameter	Value
Booking Horizon	$T = \{360, 359, \ldots, 0\}$
Fare Classes	$F = \{1, 2, 3\}$
Fares	$r_f = \{200, 150, 100\}$
Denied Boarding Costs	$k_a = 201 \cdot 1.1^{a-1}$
Control Strategies	16 strategies, see Table 7.4
Demand Forecast	$\lceil \sum_{z \in Z} d_{ft}^{yz} \,/\, \lvert Z\rvert \rceil = \bar{d}_{ft}^{y}, \quad \forall\, y \in Y, \forall\, f \in F, \forall\, t \in T$
Demand Streams	$\lvert Z\rvert = 1000$
Market/Aircraft Size	13 combinations, see Table 6.2
Capacity Scenario Forecast	perfect
Capacity Updates	0–6 updates, see Table 6.4
Re-Optimization	immediately after capacity update
Capacity Scenarios	c^s, p^s, t^s, see Table 6.5–6.7
Demand Volumes	$D = \lceil v \cdot c^{ini} \rceil$, with $v = \{0.9, 1.2, 1.6, 1.8\}$
Demand Mixes	$M = \{(m_1, m_2, m_3) \vert m_f \in \{0.25, 0.5\}, m_1 + m_2 + m_3 = 1\}$
Poisson Parameter	$\lambda(t)$, see Table 6.9

Table 10.1: Parameterization of Computational Study

10.1.1 Demand

The study implements independent demand for three fare classes. The following explains parameterizing the demand arrival process, demand volume, demand mix and number of demand streams.

Demand Arrival Process

The demand arrival process is implemented as explained in Section 8.1.1. Section 6.2 reveals empirical booking data to be approximately triangular distributed. Table 6.9 presents distributions' mode and limits for three artificially aggregated fare classes and three markets. In Appendix D, Figures D.1–D.3 visualize the resulting distributions per market. These triangular distributions parameterize the time-dependent Poisson parameter $\lambda(t)$ as input for the NHPP. For example, fare class 3 on SHORT is triangular distributed: $\lambda(t) \sim$ Triangular(360, 30, 0). Here, customer requests can occur from 360 to 0 days before departure and are most probable to occur 30 days before departure.

Demand Volume

The computational study varies demand volume per flight by multiplying the initial capacity by demand factor $v = \{0.9, 1.2, 1.6, 1.8\}$. These values are also analyzed in Vulcano and Weil (2014). Preliminary studies showed these factors to suffice for observing various effects on simulation's result.

Demand Mix

The demand mix denotes the distribution of customer requests across fare classes. The study implements three demand mixes: (0.5, 0.25, 0.25), (0.25, 0.5, 0.25) and (0.25, 0.25, 0.5). Demand mix (0.5, 0.25, 0.25), e. g., says that 50% of overall demand will request fare class 1, while 25% will request fare class 2 and 3 each. Demand mixes (0.5, 0.25, 0.25) and (0.25, 0.25, 0.5) are also evaluated by Wang and Regan (2006).

Demand Streams

Analogously to Wang and Regan (2006) and Vulcano and Weil (2014), the simulation generates $|Z| = 1000$ stochastic demand streams for 13 combinations of initial aircraft size and market, four demand volumes and three demand mixes. This procedure leads to an overall number of 156,000 (=1,000 $\cdot 13 \cdot 4 \cdot 3$) synthetic instances.

10.1.2 Capacity Scenarios

The simulation separately calibrates capacity update scenarios for three markets and five initial aircraft sizes based on empirical data – Table 6.2 summarizes the resulting 13 combinations.

Capacity scenario's input, needed for the computational study, is derived from values shown in Tables 6.5–6.7. Each cluster constitutes a scenario with an update probability (former cluster share), a new capacity (former update magnitude) and an update time. These values determine simulation's capacity update scenarios. The probabilities that no update occurs over the booking horizon are derived from Table 6.4. Furthermore, this chapter's computational study assumes a perfect capacity scenario forecast as explained in Section 8.2.2.

10.1.3 Fares and Denied Boarding Costs

The computational study implements the same fare structure as Wang and Regan (2006): $r_1 = 200, r_2 = 150$ and $r_3 = 100$.

Following Wang and Regan (2006) and Vulcano and Weil (2014), the first denied boarding's cost minimally exceeds the highest fare; here $k_1 = 201 > r_1 = 200$. Denied boarding costs increase exponentially by a factor of 1.1, so that $k_1 < \ldots < k_{\hat{a}}$. Thus, the second denied boarding costs $k_2 = k_1 \cdot 1.1 = \lceil 201 \cdot 1.1 \rceil = 221$ and $k_a = 201 \cdot 1.1^{a-1}$ for further denied boardings.

10.2 Value of Information

This study presents control strategies depending on the required information level as shown in Table 7.4.

Both control strategies assuming uniformly distributed update times, $C^s/P_{uni}/T_{uni}$ and $C^s/P^s/T_{uni}$, exponentially increase the number of scenarios. Such strategies increase average Cplex run time by a factor of approx. 150. For both $C^s/\ldots/T_{uni}$ control strategies, average Cplex run time per strategy is 259.98 ms; for all other strategies it is 1.75 ms. Therefore, the simulation only applies

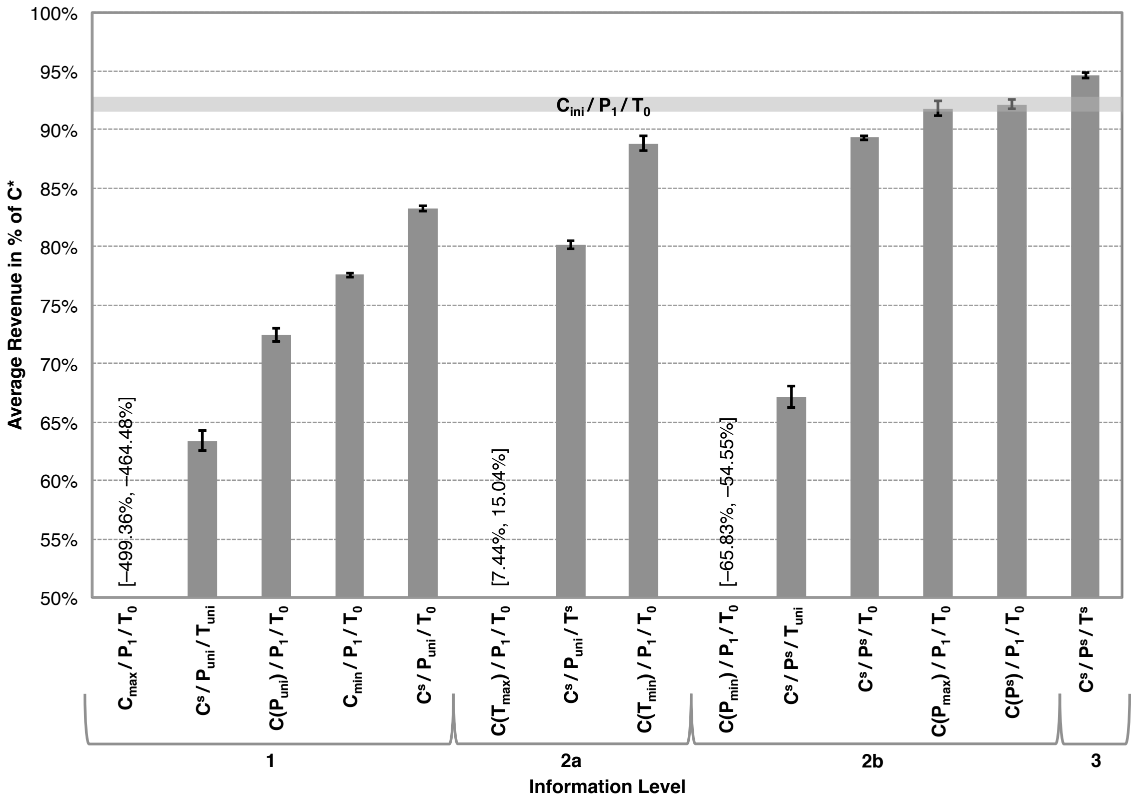

Figure 10.1: Average Revenue per Control Strategy in Percent of C* Across Markets

$C^s/\ldots/T_{uni}$ control strategies to 100 demand streams, whereas it applies all other control strategies to 1,000 demand streams. Statistical convergence was given in all cases.

When evaluating revenue, results state a 95% confidence interval based on the t-distribution. Revenue is always given in percent of the upper bound calculated via C^*. The study aggregates results across markets by calculating an average, weighted by the number of flights stated in Section 6.1.1.

Figure 10.1 shows averaged revenue results across markets. The bars represent revenue in percent of C^*, while error bars illustrate the 95% confidence interval. If a control strategy's revenue is less than 50% of C^*, the figure states the resulting confidence interval instead of a bar. To support comparison, the performance of benchmark strategy $C_{ini}/P_1/T_0$ is indicated by a horizontal band.

For comparing different control strategies with each other, this section introduces two terms:

Strictly dominant: A control strategy strictly dominates another strategy if it returns significantly more revenue on average, so that confidence bands *do not* overlap.

Weakly dominant: A control strategy weakly dominates another strategy if it returns more revenue on average, but the confidence bands overlap.

On average, $C^s/P^s/T^s$ performs best: It strictly dominates all other strategies and generates 2.47 percent points more revenue than $C_{ini}/P_1/T_0$. Revenue from $C^s/P^s/T^s$ lies in a small confidence interval, which implies comparatively stable results; revenue variance from $C_{ini}/P_1/T_0$ exceeds that from $C^s/P^s/T^s$ by more than 500%.

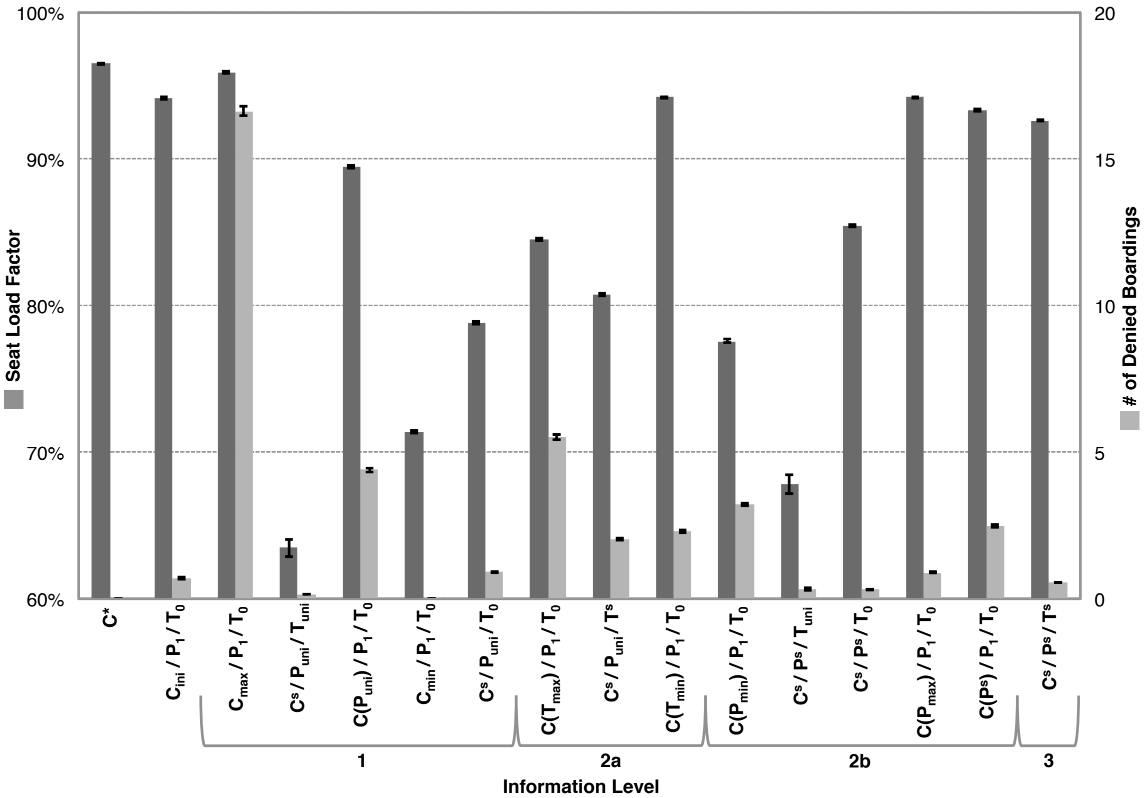

Figure 10.2: Average Seat Load Factor and Average Number of Denied Boardings per Control Strategy Across Markets

Control strategies $C(P^s)/P_1/T_0$ and $C(P_{max})/P_1/T_0$ perform similarly to benchmark strategy $C_{ini}/P_1/T_0$. Both strategies' average revenue lies within the confidence interval of $C_{ini}/P_1/T_0$. Revenue from $C(P_{max})/P_1/T_0$ is 0.38 percent points lower than that from $C_{ini}/P_1/T_0$; revenue variance is also slightly lower. However, $C(P^s)/P_1/T_0$'s average revenue is 0.04 percent points lower than that from $C_{ini}/P_1/T_0$. Furthermore, $C(P^s)/P_1/T_0$ leads to a lower revenue variance, as on average, revenue variance from $C_{ini}/P_1/T_0$ exceeds that from $C(P^s)/P_1/T_0$ by more than 300%.

Control strategy $C(T_{max})/P_1/T_0$ performs worse than 50% of C^*. Furthermore, control strategies $C_{max}/P_1/T_0$ and $C(P_{min})/P_1/T_0$ even lead to negative results.

Figure 10.2 illustrates the average seat load factor and denied boardings across markets. As expected, C^* leads to the highest seat load factor and never induces denied boardings. As the most conservative strategy, $C_{min}/P_1/T_0$ never causes denied boardings either, but results in a much lower seat load factor. Both $C^s/P^s/T^s$ and $C_{ini}/P_1/T_0$ induce comparatively few denied boardings: on average 0.56 and 0.70. However, $C_{ini}/P_1/T_0$ results in a seat load factor that is 1.52 percent points higher than that from $C^s/P^s/T^s$. Concluding, $C_{ini}/P_1/T_0$ sells a marginally larger number of tickets while $C^s/P^s/T^s$ leads to more revenue per ticket sold.

Subsequently, this section's remainder first benchmarks control strategies that require a similar information level. For each level, the study retains the best solution. Subsequently, benchmarking the best performers from each level allows to gauge the value of information.

10.2.1 Information Level 1 – Capacities

Table 10.2 lists revenue results from all control strategies requiring only information on potential capacities, as well as benchmark $C_{ini}/P_1/T_0$. Strategies $C^s/P_{uni}/T_0$ and $C(P_{uni})/P_1/T_0$ are marked in bold, because $C^s/P_{uni}/T_0$ strictly dominates on SHORT and MEDIUM and $C(P_{uni})/P_1/T_0$ strictly dominates on LONG.

Control Strategy	SHORT Mean	SHORT Conf. Int. 95%		MEDIUM Mean	MEDIUM Conf. Int. 95%		LONG Mean	LONG Conf. Int. 95%	
$C_{max}/P_1/T_0$	−1187.96%	[−1226.71%,	−1149.20%]	−332.85%	[−345.80%,	−319.91%]	31.78%	[29.87%,	33.69%]
$C^s/P_{uni}/T_{uni}$	49.01%	[48.06%,	49.96%]	64.29%	[63.41%,	65.17%]	84.19%	[83.75%,	84.63%]
$\boldsymbol{C(P_{uni})/P_1/T_0}$	40.80%	[39.61%,	42.00%]	79.78%	[79.37%,	80.19%]	**92.26%**	[**92.07%**,	**92.44%**]
$C_{min}/P_1/T_0$	79.38%	[79.22%,	79.55%]	76.84%	[76.63%,	77.05%]	77.93%	[77.81%,	78.04%]
$\boldsymbol{C^s/P_{uni}/T_0}$	**84.33%**	[**84.13%**,	**84.53%**]	**82.73%**	[**82.52%**,	**82.95%**]	83.74%	[83.62%,	83.87%]
$C_{ini}/P_1/T_0$	91.39%	[90.75%,	92.02%]	91.26%	[90.51%,	92.02%]	98.03%	[97.79%,	98.27%]

Table 10.2: Information Level 1 – Average Revenue of Control Strategies in % of C*

Of all control strategies, $C_{max}/P_1/T_0$ performs worst, inducing a negative revenue on SHORT and MEDIUM. Here, denied boarding costs exceed revenues from selling tickets. The excessive denied boardings result from always assuming the highest possible capacity to realize. On average, 16 passengers must be compensated for a denial of service at departure. This is very costly, as, for the 16^{th} denied boarding, costs increase to 462% of the highest fare.

Results from $C^s/P_{uni}/T_{uni}$ show that assuming uniformly distributed update probabilities and times is not promising. The performance of $C^s/P_{uni}/T_{uni}$ is inferior to most alternatives. Its lack of success may be attributed mostly to a generally low seat load factor.

Control strategy $C(P_{uni})/P_1/T_0$ provides mixed results. On SHORT, it performs badly, with an average revenue of only 40.80%. On MEDIUM, $C(P_{uni})/P_1/T_0$ performs better but still below a revenue of 80%. At least on LONG, it strictly dominates all other control strategies of information level 1. However, average seat load factor of $C(P_{uni})/P_1/T_0$ is under 90%, which is relatively low and its denied boardings are the third highest of all control strategies.

The most risk-averse strategy, $C_{min}/P_1/T_0$, results in comparatively stable but low revenue across all three markets: from 76.84% to 79.38%. However, it cannot compete with strategies that accept the risk of denied boardings, such as $C^s/P_{uni}/T_0$.

$C^s/P_{uni}/T_0$ performs in comparatively stable in relation to upper bound C^* across all three markets – from 82.73% to 84.33%. When the probability of capacity updates is high, using more information on potential capacities is beneficial. Therefore, $C^s/P_{uni}/T_0$ strictly dominates $C(P_{uni})/P_1/T_0$ on SHORT and MEDIUM. However, on LONG, updates occur rarely, and the computationally faster strategy $C(P_{uni})/P_1/T_0$ strictly dominates $C^s/P_{uni}/T_0$.

In conclusion, when no information on the time or probability of updates is available, $C_{ini}/P_1/T_0$ performs best. In other words, under these circumstances it is best to ignore the possibility of capacity updates until an update actually occurs. Consequently, the study neglects alternative control strategies utilizing only information level 1 in further analyses.

10.2.2 Information Level 2a – Capacities and Times

Table 10.3 shows the revenue resulting from control strategies that utilize only information on potential capacities and update times. Here, $C(T_{min})/P_1/T_0$ strictly dominates $C(T_{max})/P_1/T_0$ and $C^s/P_{uni}/T^s$ across all markets: the lower the potential of capacity updates from SHORT to MEDIUM to LONG, the comparatively better performs $C(T_{min})/P_1/T_0$.

Control Strategy	SHORT Mean	SHORT Conf. Int. 95%		MEDIUM Mean	MEDIUM Conf. Int. 95%		LONG Mean	LONG Conf. Int. 95%	
$C(T_{max})/P_1/T_0$	−82.41%	[−87.52%,	−77.31%]	36.48%	[32.52%,	40.44%]	53.04%	[52.28%,	53.80%]
$C^s/P_{uni}/T^s$	80.30%	[79.96%,	80.65%]	79.48%	[79.08%,	79.88%]	83.11%	[82.95%,	83.28%]
$\boldsymbol{C(T_{min})/P_1/T_0}$	**83.01%**	**[82.30%,**	**83.72%]**	**89.71%**	**[89.02%,**	**90.40%]**	**94.47%**	**[94.19%,**	**94.74%]**
$C_{ini}/P_1/T_0$	91.39%	[90.75%,	92.02%]	91.26%	[90.51%,	92.02%]	98.03%	[97.79%,	98.27%]

Table 10.3: Information Level 2a – Average Revenue of Control Strategies in % of C*

Control strategy $C(T_{max})/P_1/T_0$ performs worst, even with a negative revenue on SHORT. It results in the highest number of average denied boardings of all information level 2a control strategies.

Revenue from $C^s/P_{uni}/T^s$ ranges from 79.48% on MEDIUM to 83.11% on LONG. Across markets, its comparatively low seat load factor – just above 80% – is an indicator for spoilage.

$C(T_{min})/P_1/T_0$'s superior revenue performance is based on the seat load factor achieved, although the average denied boardings slightly exceed those from $C^s/P_{uni}/T^s$, as shown in Figure 10.2.

Comparing $C(T_{min})/P_1/T_0$ with the best performers on information level 1 leads to the following insights: on SHORT, $C^s/P_{uni}/T_0$ strictly dominates $C(T_{min})/P_1/T_0$, while $C(T_{min})/P_1/T_0$ strictly dominates $C^s/P_{uni}/T_0$ on MEDIUM and $C(P_{uni})/P_1/T_0$ on LONG.

By assuming the last announced scenario's capacity to be the final capacity, $C(T_{min})/P_1/T_0$ often imitates $C_{ini}/P_1/T_0$. Strategy $C(T_{min})/P_1/T_0$ only differs if updates can occur at departure, which is the case for aircraft size XS and S on SHORT, for size S on MEDIUM and for size M on LONG – withdrawable from Table 6.5–6.7. In these situations, $C(T_{min})/P_1/T_0$ performs inferior to $C_{ini}/P_1/T_0$.

In conclusion, given only potential capacities and update times, on two out of three markets $C(T_{min})/P_1/T_0$ yields more revenue compared to the best strategies from information level 1. However, benchmark $C_{ini}/P_1/T_0$ remains the most successful strategy. In consequence, the study also neglects $C(T_{min})/P_1/T_0$ in further analyses.

10.2.3 Information Level 2b – Capacities and Probabilities

Of those strategies that utilize only information on potential capacities and their probabilities, $C(P^s)/P_1/T_0$ strictly dominates all alternatives on SHORT, as shown in Table 10.4. However, on MEDIUM and LONG, $C(P_{max})/P_1/T_0$ weakly dominates $C(P^s)/P_1/T_0$ and strictly dominates all other strategies.

Control strategy $C(P_{min})/P_1/T_0$ performs worst on SHORT and LONG. Expecting the capacity with the smallest probability to occur is neither intuitive nor successful. However, by accident, the least

Control Strategy	SHORT Mean	SHORT Conf. Int. 95%		MEDIUM Mean	MEDIUM Conf. Int. 95%		LONG Mean	LONG Conf. Int. 95%	
$C(P_{min})/P_1/T_0$	−509.33%	[−532.46%,	−486.20%]	79.02%	[78.81%,	79.23%]	52.82%	[51.51%,	54.14%]
$C^s/P^s/T_{uni}$	49.73%	[48.75%,	50.72%]	67.33%	[66.40%,	68.26%]	96.49%	[95.82%,	97.16%]
$C^s/P^s/T_0$	88.08%	[87.86%,	88.29%]	88.19%	[87.97%,	88.41%]	96.74%	[96.55%,	96.93%]
$\boldsymbol{C(P_{max})/P_1/T_0}$	89.75%	[89.10%,	90.39%]	**91.26%**	[90.51%,	92.02%]	**98.03%**	[97.79%,	98.27%]
$\boldsymbol{C(P^s)/P_1/T_0}$	**91.48%**	[**91.18%**,	**91.78%**]	91.19%	[90.76%,	91.62%]	97.96%	[97.74%,	98.19%]
$C_{ini}/P_1/T_0$	91.39%	[90.75%,	92.02%]	91.26%	[90.51%,	92.02%]	98.03%	[97.79%,	98.27%]

Table 10.4: Information Level 2b – Average Revenue of Control Strategies in % of C*

probable capacity is often close to the most probable capacity on MEDIUM. Here, $C(P_{min})/P_1/T_0$ performs comparatively good with an average revenue of 79.02%.

$C^s/P^s/T_{uni}$ also performs rather badly across all markets. Its knowledge on capacity probabilities is negatively affected by assuming uniformly distributed update times over the whole booking horizon. In contrast, $C^s/P^s/T_0$ consistently exceeds $C^s/P^s/T_{uni}$'s revenue. Its superiority bases on assuming all capacity updates to occur on departure.

Strategy $C^s/P^s/T_0$ leads to better results, but cannot compete with $C(P_{max})/P_1/T_0$ and $C(P^s)/P_1/T_0$. On MEDIUM and LONG, $C(P_{max})/P_1/T_0$ weakly dominates all other control strategies of information level 2b. It even performs as good as benchmark $C_{ini}/P_1/T_0$ on MEDIUM and LONG.

In comparison to $C_{ini}/P_1/T_0$, $C(P^s)/P_1/T_0$ weakly dominates on SHORT and performs similar to $C(P_{max})/P_1/T_0$ on MEDIUM and LONG.

In conclusion, based on the analyzed data, given no information on update times, strategies that create a single scenario, $C(P_{max})/P_1/T_0$ and $C(P^s)/P_1/T_0$, are computationally efficient and outperform $C^s/P^s/T_{uni}$ and $C^s/P^s/T_0$. Across all markets, strategies that utilize information on update probabilities perform significantly better than those utilizing information on potential update times.

However, $C(P_{max})/P_1/T_0$ frequently imitates $C_{ini}/P_1/T_0$ as in the majority of cases the initial capacity is also the most likely final capacity. The only difference is observable on SHORT and for initial aircraft size XS; here $C(P_{max})/P_1/T_0$ performs inferior to $C_{ini}/P_1/T_0$. For that reason, the study neglects $C(P_{max})/P_1/T_0$ in the further analyses.

10.2.4 Information Level 3 – Capacities, Probabilities and Times

Only $C^s/P^s/T^s$ utilizes all information on the considered characteristics of capacity updates. Table 10.5 compares strategy $C^s/P^s/T^s$ to the resulting revenue from $C(P^s)/P_1/T_0$ that, at least partially, outperforms $C_{ini}/P_1/T_0$.

This section neglects control strategies of information levels 1 and 2a as they perform inferior to $C_{ini}/P_1/T_0$. Across all markets, $C^s/P^s/T^s$ earns more than 90% of upper bound C^* and at least weakly dominates all other strategies, including $C_{ini}/P_1/T_0$.

The difference in revenue gaps from $C^s/P^s/T^s$ to $C_{ini}/P_1/T_0$ can be attributed mainly to two effects. On the one hand, the more probable capacity updates occur, the better performs $C^s/P^s/T^s$ compared to $C_{ini}/P_1/T_0$. On the other hand, the demand arrival time plays an important role. On SHORT, update probabilities are high, but demand arrives comparatively late. On MEDIUM, update

Information Level	Control Strategy	SHORT Mean	SHORT Conf. Int. 95%		MEDIUM Mean	MEDIUM Conf. Int. 95%		LONG Mean	LONG Conf. Int. 95%	
2b	$C(P^s)/P_1/T_0$	91.48%	[91.18%,	91.78%]	91.19%	[90.76%,	91.62%]	97.96%	[97.74%,	98.19%]
3	$\boldsymbol{C^s/P^s/T^s}$	**94.10%**	**[93.85%,**	**94.35%]**	**94.16%**	**[93.88%,**	**94.43%]**	**98.07%**	[97.84%,	98.30%]
Benchmark	$C_{ini}/P_1/T_0$	91.39%	[90.75%,	92.02%]	91.26%	[90.51%,	92.02%]	98.03%	[97.79%,	98.27%]

Table 10.5: Information Level 3 and Best Performers – Average Revenue of Control Strategies in % of C*

probabilities are lower, but demand arrives earlier and therefore $C^s/P^s/T^s$ can benefit more. Here, with 2.90 percent points, $C^s/P^s/T^s$' revenue advantage over $C_{ini}/P_1/T_0$ is highest. LONG involves the earliest demand arrivals, but update probabilities are very low, thus both control strategies display a very similar performance.

10.2.5 Implications of Information Levels

For each information level, the study shows one or more control strategy performing best. Depending on the market, control strategies $C(P_{uni})/P_1/T_0$ and $C^s/P_{uni}/T_0$ perform best if only the information on potential capacities can be utilized (information level 1). Utilizing the additional information on update times (information level 2a) shows $C(T_{min})/P_1/T_0$ to perform best. However, $C_{ini}/P_1/T_0$ strictly dominates both best performers of information level 1 and 2a. Thus, given information only on both capacities and update times, RM should ignore the possibility of capacity updates and stick to the initially announced capacity until an update occurs.

However, at least on SHORT, $C(P^s)/P_1/T_0$ weakly dominates $C_{ini}/P_1/T_0$ (information level 2b). Thus, when the probability of at least one capacity update is high, the results suggest implementing $C(P^s)/P_1/T_0$ when only given information on potential capacities and their probabilities. Thus, comparing the best performing strategies of information levels 2a and 2b reveals that information on the probability of capacity updates is more valuable than information on update times.

Finally, full information on potential capacities, update times and probabilities can improve revenue by more than 2.47 percent points over $C_{ini}/P_1/T_0$. Thus, airlines should strive to utilize capacity data to predict such comprehensive information.

The following section addresses the first part of research question 4 by analyzing the influence of different problem parameters on the revenue result. Here and in the following, only the value of information study's best performing control strategies are considered: $C_{ini}/P_1/T_0$, $C(P^s)/P_1/T_0$ and $C^s/P^s/T^s$.

10.3 Sensitivity Analysis of Problem Parameters

To analyze the effect of problem parameters, this section focuses on the three most successful control strategies: $C_{ini}/P_1/T_0$, $C(P^s)/P_1/T_0$ and $C^s/P^s/T^s$. The following sensitivity analysis considers the demand volume and mix as well as the number, time and magnitude of capacity updates. Figures 10.3–10.7 each depict revenue as a percentage of upper bound C^* on the left y-axis. To indicate the general effect of parameters on achievable revenue, the right y-axis depicts revenue earned by C^* as an index of its highest outcome.

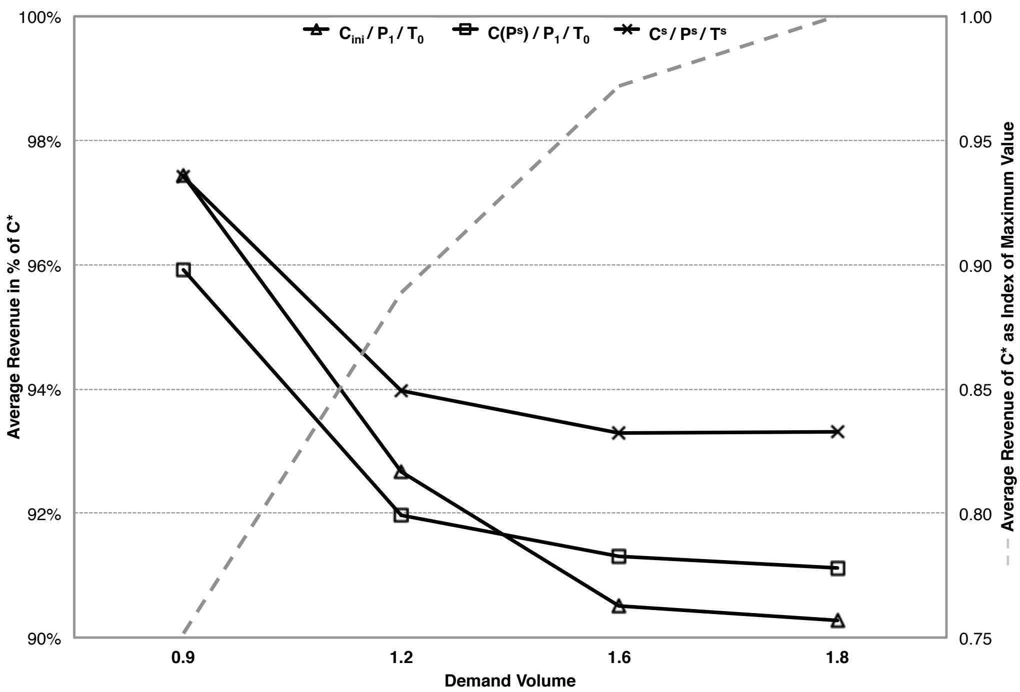

Figure 10.3: Influence of Demand Volume – Average Revenue of Best Performers and C*

10.3.1 Influence of Demand Volume

To compare the control strategies' revenue sensitivity to demand volume, results are aggregated by the demand volume given as the ratio of demand to initial capacity, $v = \{0.9,\ 1.2,\ 1.6,\ 1.8\}$. Figure 10.3 shows the resulting revenue. Note that revenue from C^* increases with demand volume. However, the gap between the compared control strategies and C^* also increases with demand volume.

When capacity exceeds demand ($v = 0.9$), $C_{ini}/P_1/T_0$ and $C^s/P^s/T^s$ perform equally well. Both control strategies accept almost every customer request. Furthermore, when demand is low, even capacity decreases do not cause a high risk of denied boardings.

However, the higher the demand volume, the more clearly outperforms $C^s/P^s/T^s$ benchmark strategy $C_{ini}/P_1/T_0$. This is because $C^s/P^s/T^s$ sells more tickets in more expensive fare classes and $C_{ini}/P_1/T_0$ causes more denied boardings. For demand volumes $v = 1.6$ and $v = 1.8$, $C(P^s)/P_1/T_0$ also outperforms $C_{ini}/P_1/T_0$, but cannot achieve the high performance of $C^s/P^s/T^s$.

Concluding, the benefits of considering uncertain capacity through fully parameterized scenarios increases with demand volume. For high demand volumes and when no information on update times is available, parameterizing capacity through the weighted average of possible capacities – $C(P^s)/P_1/T_0$ – still outperforms the benchmark $C_{ini}/P_1/T_0$.

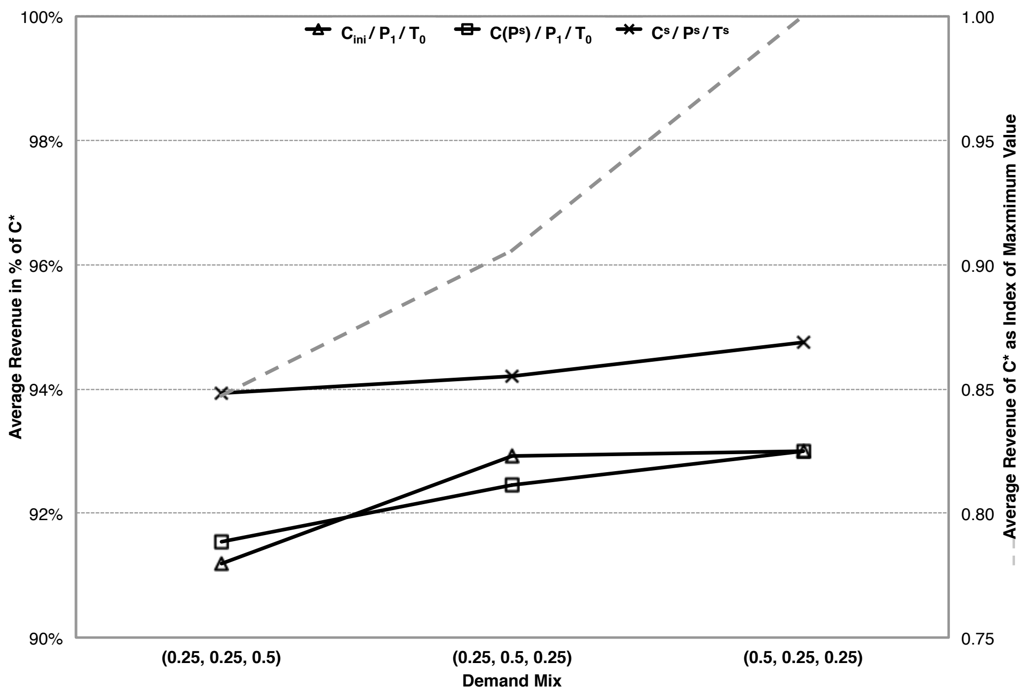

Figure 10.4: Influence of Demand Mix – Average Revenue of Best Performers and C*

10.3.2 Influence of Demand Mix

This section evaluates the effect of three demand mixes: $m = \{(0.25, 0.25, 0.5), (0.25, 0.5, 0.25), (0.5, 0.25, 0.25)\}$. Figure 10.4 shows the aggregated results per demand mix.

While the revenue earned by C^* increases with the value of demand, the gap between C^* and the evaluated control strategies is highest for low value demand (0.25, 0.25, 0.5). This can be explained by the comparatively early arrival of low value demand, resulting in a revenue advantage of $C^s/P^s/T^s$ over $C_{ini}/P_1/T_0$ of 2.74 percent points.

If the number of fare class 2 and 1 customer requests increases, all control strategies' revenue enhance. Across all demand mixes, $C^s/P^s/T^s$ performs best. While $C(P^s)/P_1/T_0$ performs slightly better than $C_{ini}/P_1/T_0$ for low value demand, the reverse applies for demand mix (0.25, 0.5, 0.25). For high value demand, both strategies perform equally.

10.3.3 Influence of Number of Capacity Updates per Flight

The influence of the number of capacity updates is illustrated in Figure 10.5.

Aggregating problem instances based on the number of actual capacity updates per flight leads to two major findings. First, as expected, if no capacity update occurs, $C_{ini}/P_1/T_0$ performs best. Second, if at least one capacity update occurs, $C(P^s)/P_1/T_0$ and $C^s/P^s/T^s$ always outperform $C_{ini}/P_1/T_0$. Here, $C(P^s)/P_1/T_0$ is superior to $C^s/P^s/T^s$ for six capacity updates. From one to five capacity updates, optimizing via $C^s/P^s/T^s$ is the best choice.

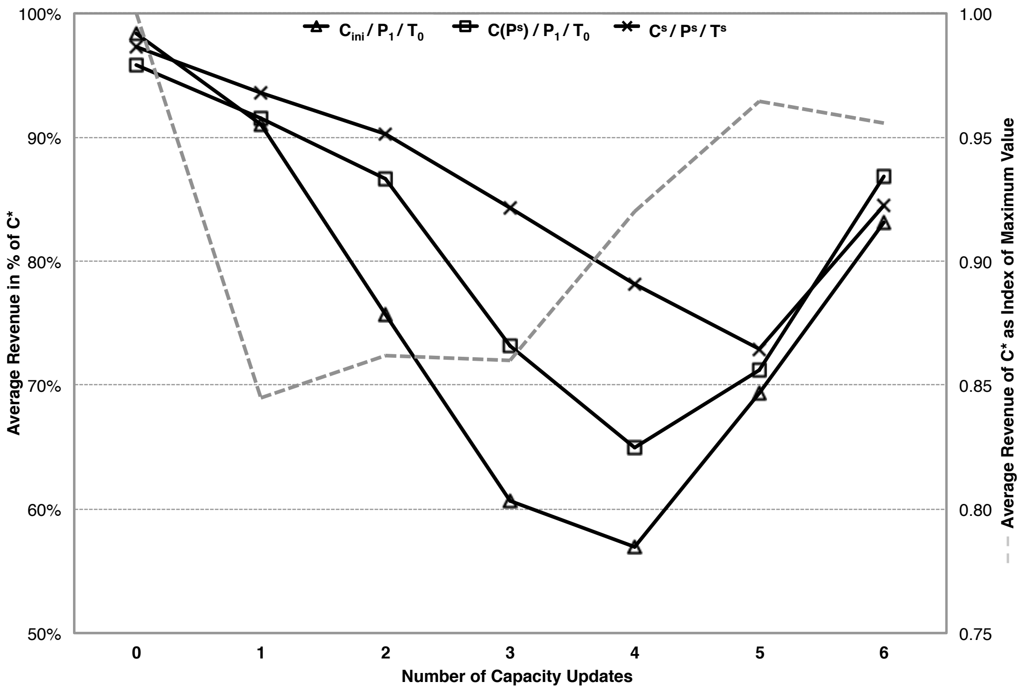

Figure 10.5: Influence of Number of Capacity Updates per Flight – Average Revenue of Best Performers and C*

The figure shows no clear trend for C^*: its revenue decreases from 0 to one update, unsteadily increases from one to five updates and finally decreases from five to six updates. As the analysis of empirical data shows, six updates on a single flight can only occur on SHORT. Here, average capacity is smaller than on MEDIUM and LONG, leading to an overall smaller chance for a high revenue.

This implies other factors having a more direct influence on revenue results than the number of capacity updates per flight. This implication motivates the following sections to address the influence of capacity update times and capacity update magnitudes.

10.3.4 Influence of Final Capacity Update Time

Next, the following analyzes control strategies' sensitivity to the time of capacity updates by aggregating instances based on the time of the final update. Figure 10.6 illustrates the results.

Following input from industry experts, the time horizon splits into three segments as follows. Early updates, which occur from 360 to 201 days before departure, are mostly due to re-optimization by fleet assignment. Late updates, which occur from 25 to 0 days before departure, are primarily due to operational events. Mid-term updates, from 200 to 26 days before departure, cannot be clearly assigned to a single reason.

For late capacity updates, C^* earns more revenue than for mid-term and early updates. This is because, given the empirical calibration of instances, such updates are more likely to increase capacity, increasing revenue opportunities.

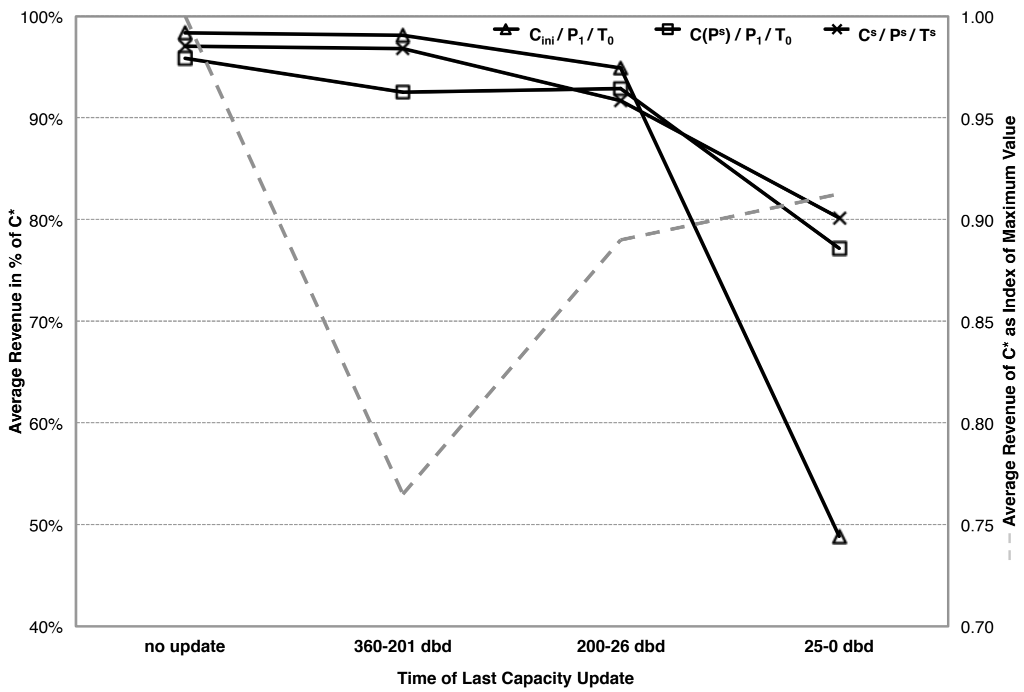

Figure 10.6: Influence of Final Capacity Update Time – Average Revenue of Best Performers and C*

$C^s/P^s/T^s$ most clearly outperforms $C_{ini}/P_1/T_0$ for late updates, by more than 31 percent points. Given early or mid-term updates, $C_{ini}/P_1/T_0$ performs astonishingly well. This can be explained as, in these instances, most high value customer requests arrive after the updates. At that time, the updated policy of $C_{ini}/P_1/T_0$ already accounts for the new capacity.

10.3.5 Influence of Capacity Update Direction and Magnitude

Lastly, Figure 10.7 illustrates the influence of two dimensions of capacity updates: direction and magnitude. The update direction states whether capacity increases or decreases and the update magnitude is defined as *low* when initial and updated capacity differ by less than 20% and as *high* otherwise.

Average revenue from C^* is highest for low capacity increases. This is because high increases occur primarily on smaller aircraft sizes, leading to a smaller revenue opportunity than when low increases affect larger aircrafts. When capacity decreases, the benefit of $C^s/P^s/T^s$ over $C_{ini}/P_1/T_0$ does not increase with the magnitude of updates. This can be explained as low decreases occur mostly late in the booking horizon, while high decreases occur comparatively early.

In this setting, $C^s/P^s/T^s$ does not improve revenue over $C_{ini}/P_1/T_0$ for capacity increases. On the one hand, $C^s/P^s/T^s$' booking limits are more cautious than the myopic approach of $C_{ini}/P_1/T_0$. On the other hand, the simulation's independent demand model excludes cannibalization effects. Thus, a different influence of capacity increases can be expected when customers choose between offered fare classes or when customers even observe prices strategically.

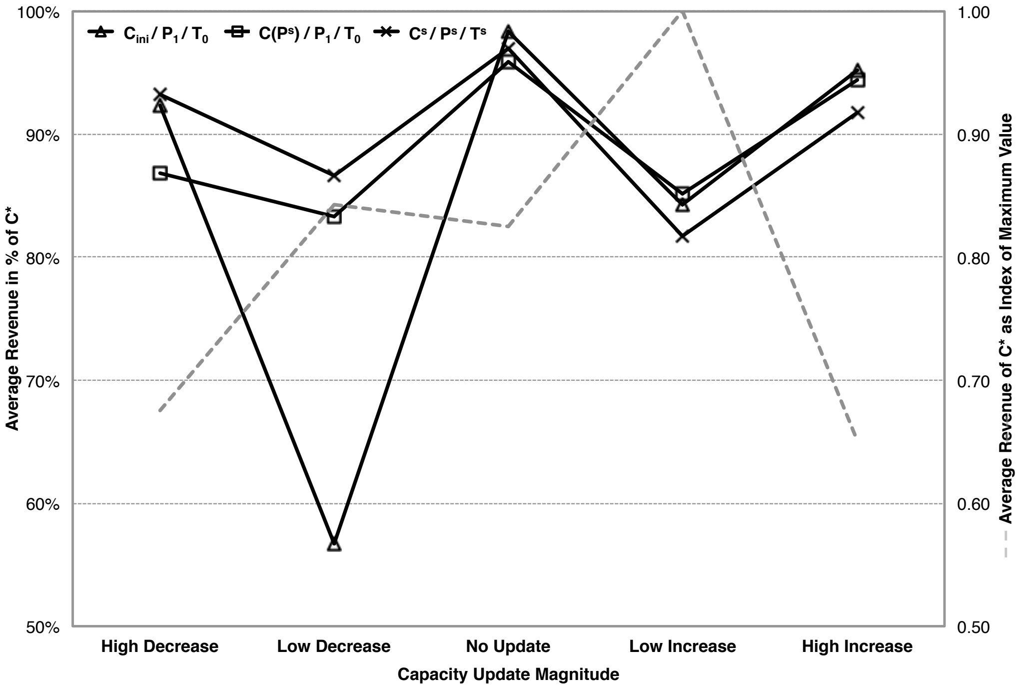

Figure 10.7: Influence of Final Capacity Update Magnitude – Average Revenue of Best Performers and C*

As expected, $C^s/P^s/T^s$ performs best for capacity decreases, as $C(P^s)/P_1/T_0$ and $C_{ini}/P_1/T_0$ produce more denied boardings.

10.3.6 Implications of Influence by Problem Parameters

The computational study on the value of information and problem parameters provides some valuable insights into what influences revenue the most. With respect to demand, the study shows that control strategy $C^s/P^s/T^s$ benefits the most if demand arrives early. The later customers request tickets, the longer can $C_{ini}/P_1/T_0$ postpone the acceptance decision. On a market with demand arriving early, $C^s/P^s/T^s$' revenue advantage increases.

However, when analyzing markets, the study shows few potential of $C^s/P^s/T^s$ on LONG, although customer requests arrive the earliest compared to MEDIUM and short. Here, the crux is capacity updates' probabilities – $C^s/P^s/T^s$ benefits the most, if the probability for an update is high. On LONG, however, the chance that the initial and final capacity differ is less than 2%. Thus, as expected, capacity updates should be considered especially on markets, where updates are frequent.

In contrast, the demand volume and demand mix less drastically influence control strategies' revenue differences. Nevertheless, $C^s/P^s/T^s$ advantage over $C_{ini}/P_1/T_0$ increases with an increasing number of overall demand and if high value demand rarely arrives. On a market, where the sum of customer requests strongly exceeds capacity, $C^s/P^s/T^s$ accepts more customers of higher priced fare classes than any other control strategy. The same applies if the number of high value demand is low. But in overall, the demand structure is not as influential as other problem parameters following.

As expected, if no capacity update occurs, $C_{ini}/P_1/T_0$ performs best. However, if at least a single update occurs, $C^s/P^s/T^s$ approach of anticipating future updates pays off. The exact number of updates on a particular flight, though, does not lead into a clear direction as does the update time.

Analyzing the update time, late updates suggest using $C^s/P^s/T^s$. In contrast, early capacity updates allow $C_{ini}/P_1/T_0$ to earlier adjust its booking limits. But if the chance for late updates exists – due to technical defects or other operational reasons – the study advises using control strategy $C^s/P^s/T^s$.

The study also reveals, that $C^s/P^s/T^s$ bases its success mainly on preventing denied boardings. While capacity increases allow higher revenues from selling more tickets, capacity decreases bear the hazard of even higher denied boarding costs. If increases and decreases can occur, $C^s/P^s/T^s$ sets more cautious booking limits leading to highest revenue results. Here, the magnitude of a capacity update is not as influential as the update time.

11 Computational Study on Robustness

The computational studies in Section 10.2 and 10.3 analyzed control strategies' performance for different information levels and the influence of the model's problem parameters. Both studies used a demand forecast based on actual demand streams (Section 8.1.3) and a perfect forecast on potential capacity scenarios (Section 8.2.2). However, in practice, demand forecasts' accuracy can strongly vary and the same can apply to forecasts on capacity scenarios. Therefore, this chapter's computational study considers inaccurate forecasting by analyzing the influence of distorted demand forecasts and capacity scenario forecasts of varying quality – addressing control strategies' robustness. Thus, this chapter gives an answer to the second part of research question 4.

RQ4.2: *How does a distorted demand forecast and a distorted forecast on capacity updates affect RM's result?*

The following study on robustness uses the same input parameters as previous studies in Section 10.2 and 10.3. However, due to changes concerning the demand forecast and capacity scenario forecast, Section 11.1 states input values differing from previous computational studies. The study on demand forecasts' robustness in Section 11.2 considers three forecasting qualities: at first, a perfect demand forecast is analyzed, then the former forecast is distorted with respect to demand volume and final with respect to demand arrival times. Last but not least, Section 11.3 analyzes a distorted forecast on capacity scenario probabilities and capacity scenario times.

11.1 Input Data Distortion and Setup of Computational Study

This section states the input of different problem parameters, compared to previous computational studies. Here, three major changes take place, concerning the selection of control strategies, the demand forecast and the capacity scenario forecast. Table 11.1 displays these changes, while the following sections introduce them in more detail.

Parameter	Value
Control Strategies	C^*, $C_{ini}/P_1/T_0$, $C(P^s)/P_1/T_0$, $C^s/P^s/T^s$
Demand Forecast	perfect forecast, Section 11.2.1 distorted demand volume, Section 11.2.2 distorted demand arrival time, Section 11.2.3
Capacity Scenario Forecast	distorted capacity scenario probabilities, Section 11.3.1 distorted capacity scenario times, Section 11.3.2

Table 11.1: Differences to Previous Computational Study Parameterization

11.1.1 Control Strategies

As Section 10.2 shows, three control strategies perform best, depending on the information level concerned. The computational studies' results suggest using one of these strategies for achieving good results, close to optimum. Thus, this chapter only considers the addressed strategies – $C_{ini}/P_1/T_0$, $C(P^s)/P_1/T_0$ and $C^s/P^s/T^s$ – as well as upper bound C^*.

11.1.2 Demand Forecast

Section 8.1.3 presented the simulation system's approach of forecasting demand by using the average over all generated demand streams. Changes to this chapter's study on robustness help analyzing the influence of different demand forecasting qualities.

Section 11.2.1 analyzes a *perfect* demand forecast: the tested control strategies receive perfect information on upcoming demand, stating an upper bound on the quality of forecasting demand. This is identical to the previous demand forecast of upper bound C^*. For every demand input combination y and demand stream z, the actual demand, d_{ft}^{yz}, to arrive per fare class f and time slice t, also represents forecasted demand $\bar{d}_{ft}^{yz}$ required by optimization, so

$$\bar{d}_{ft}^{yz} = d_{ft}^{yz}, \quad \forall\, y \in Y, \forall z \in Z, \forall\, f \in F, \forall\, t \in T\,.$$

In contrast, Section 11.2.2 analyzes the influence of distorting a demand forecast based on averaged demand streams. First, the simulation system's forecasted demand volume is systematically distorted by an error term $\phi \in \mathbb{Q}_{\geq 0}$. If the term is positive, demand volume increases; if negative, demand volume decreases:

$$\bar{d}_{ft}^{y} = \left\lceil (1+\phi) \cdot \frac{\sum\limits_{z \in Z} d_{ft}^{yz}}{|Z|} \right\rceil, \quad \forall\, y \in Y, \forall\, f \in F, \forall\, t \in T\,.$$

Section 11.2.3 focuses on systematically distorting demand arrival times and analyzes the resulting influence on control strategies' revenue. The error term $\psi \in \mathbb{N}_{\geq 0}$ determines the shift in arrival times in days before departure. For example, if $\psi = -20$, the simulation system assumes demand for every fare class in every time slice to occur 20 days earlier than the actual forecast. Thus, Section 11.2.3 distorts forecasted demand $\bar{d}_{ft}^{y}$ in the way that

$$\bar{d}_{ft}^{y} = \left\lceil \frac{\sum\limits_{z \in Z} d_{ft'}^{yz}}{|Z|} \right\rceil, \quad \forall\, y \in Y, \forall\, f \in F, \forall\, t \in T,$$

where using t' instead of t ensures forecasted demand to remain in $T = \{\hat{t}, \ldots, 0\}$, the time horizon's boundaries, as

$$t' = \begin{cases} \min\{(t+\psi),\, \hat{t}\}, & \text{if } \psi > 0 \\ t, & \text{if } \psi = 0 \\ \max\{(t+\psi),\, 0\}, & \text{if } \psi < 0\,. \end{cases}$$

11.1.3 Capacity Scenario Forecast

This section introduces distorting capacity scenario forecasts with respect to scenario probabilities and scenario times. The forecast's only uninfluenced parameter is the capacity $\bar{c}^s$ per scenario s. As explained in Section 7.2, an airline's perfect knowledge on its fleet justifies assuming potential capacities c^s to be known in advance. The following equation determines the uninfluenced information on potential capacities:

$$\bar{c}^s = c^s, \quad \forall\, s \in S\,.$$

Section 11.3 systematically distorts forecasted capacity scenario probabilities. An error term $\rho \in \mathbb{Q}_{\geq 0}$ distorts the first scenario's probability p^1, quantifying the likelihood for initial capacity $c^{ini} = c^1$ remaining till departure. The study, thus, analyzes over- and underestimating the probability that no capacity update occurs. Except the first, all remaining scenarios are proportionally adjusted so that $\sum_{s \in S} \bar{p}^s = 1$ still holds. For example, if $\rho = 0.05$, the first scenario's probability p^1 increases by 0.05 and the sum of all remaining scenario probabilities $\sum_{s=2}^{|S|} p^s$ decreases by 0.05. Thus, the forecasted probabilities adjust to

$$\bar{p}^s = \begin{cases} (1+\rho) \cdot p^s, & \text{if } s = 1 \\ \dfrac{p^s}{1+\rho}, & \text{if } s \in \{2, \ldots, |S|\}\,. \end{cases}$$

At last, Section 11.3.1 systematically distorts capacity scenario times. For each scenario s, an error term $\omega \in \mathbb{N}_{\geq 0}$ distorts the actual update time t^s of a scenario s in days before departure. Thereby, the forecasted scenario update time $\bar{t^s}$ serves as simulation input, while still ensuring $\bar{t^s}$ to remain in $T = \{\hat{t}, \ldots, 0\}$, the time horizon's boundaries, so

$$\bar{t^s} = \begin{cases} \min\{(t+\omega),\, \hat{t}\}, & \text{if } \omega > 0 \\ t^s, & \text{if } \omega = 0 \\ \max\{(t+\omega),\, 0\}, & \text{if } \omega < 0 \end{cases} \qquad \forall\, s \in S\,.$$

11.2 Robustness of Demand Forecast

The study on demand forecast's robustness divides into three cases: First, Section 11.2.1 presents results if a perfect demand forecast serves as simulation input. Second, Section 11.2.2 analyzes distorting the forecasted demand volume. Third, Section 11.2.3 analyzes distorting the forecasted demand arrival time.

11.2.1 Perfect Demand Forecast

While previous studies in Chapter 10 assumed a demand forecast based on averaged demand streams, this section utilizes a perfect demand forecast as simulation input. This demand forecast's optimal quality assures knowing the exact future demand in advance, as, so far only applied by upper bound C^*.

Table 11.2 presents the resulting revenue improvements, compared to the results of Chapter 10, where demand is forecasted based on averaged demand streams. Results per market are given as an average increase in percent point and results across markets are marked in bold.

Control Strategy	Revenue Improvement in Percent Points SHORT	MEDIUM	LONG	Across Markets
$C_{ini}/P_1/T_0$	+0.10	+0.05	+1.48	**+0.25**
$C(P^s)/P_1/T_0$	+0.41	+0.01	+1.49	**+0.30**
$C^s/P^s/T^s$	+1.11	+1.12	+1.51	**+1.17**

Table 11.2: Average Revenue Improvement by Perfect Demand Forecast

Revenue improvements across markets are highest for strategy $C^s/P^s/T^s$ with an average surplus of 1.17 percent points. In contrast, $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$ only benefit by 0.25 and 0.30 percent points.

On LONG, all strategies similarly profit with a surplus between 1.48 and 1.51 percent points. Here, capacity updates occur rarely so that control strategies that anticipate potential updates – $C(P^s)/P_1/T_0$ and $C^s/P^s/T^s$ – only benefit slightly more than $C_{ini}/P_1/T_0$.

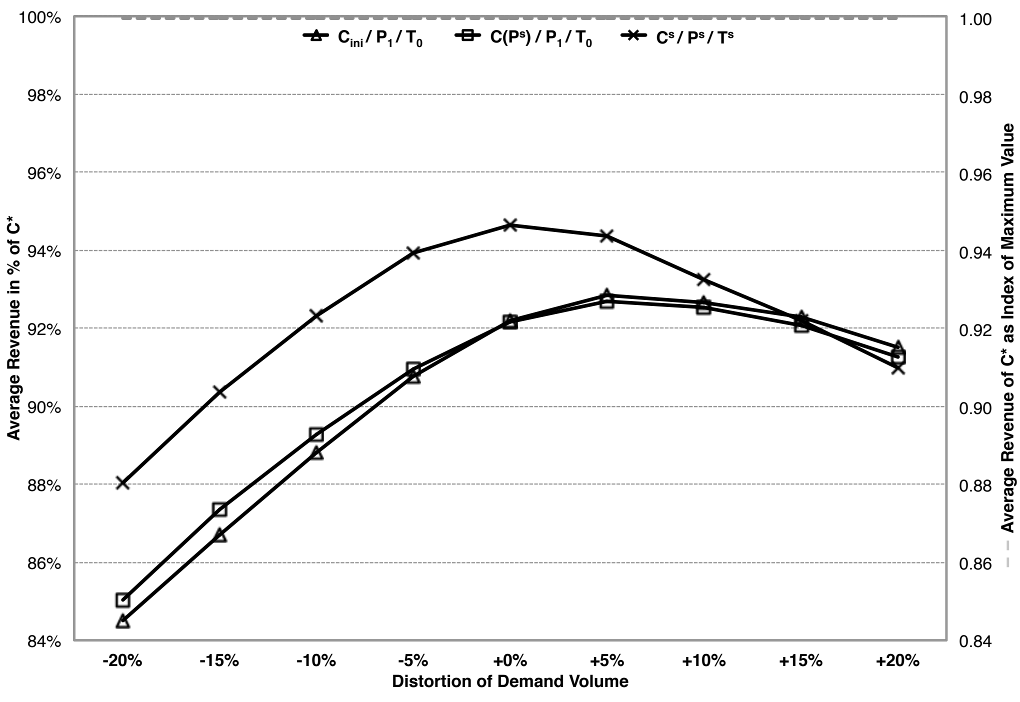

Figure 11.1: Impact of Distortion of Demand Volume Forecast

In contrast, strategies' results on SHORT and MEDIUM strongly differ. On SHORT, $C_{ini}/P_1/T_0$ only benefits by 0.10 percent points, while $C(P^s)/P_1/T_0$ by 0.41 percent points and $C^s/P^s/T^s$ even by 1.11 percent points. On MEDIUM, $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$ profit by 0.05 and 0.01 percent points, whereas $C^s/P^s/T^s$ achieves a surplus of 1.12 percent points. Here, demand arrives earlier than on SHORT, but also updates occur much earlier, so that $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$ cannot benefit that much from a perfect demand forecast.

11.2.2 Inaccurate Demand Volume Forecast

This section analyzes the influence of distorting forecast's demand volume. In contrast to the previous section, where demand is perfectly known, the simulation system's demand forecast is based on an average over all demand streams. Again, this approach represents a good but not perfect demand forecast. Figure 11.1 shows control strategies' revenue changes compared to C^* if demand volume is systematically distorted from -20% to $+20\%$.

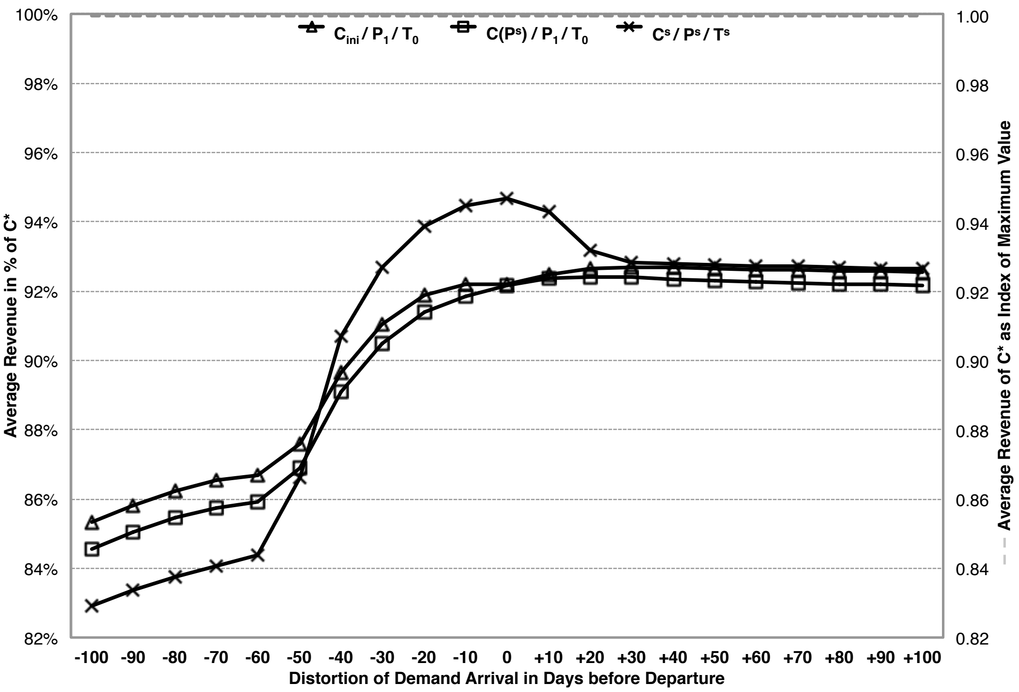

Figure 11.2: Impact of Distortion of Demand Arrival Time Forecast

Upper bound C^* achieves a constant revenue as it is unaffected by distortions. When distorting demand volume from -20% to $+10\%$, control strategy $C^s/P^s/T^s$ outperforms $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$; the less distortion the better its revenue. In contrast, $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$ perform comparatively alike. While $C(P^s)/P_1/T_0$ achieves slightly more revenue when decreasing the forecasted demand volume, $C_{ini}/P_1/T_0$ achieves marginally more revenue when increasing the forecasted demand volume. The results show that both strategies slightly benefit from a small distortion of $+5\%$ as such changes increase both strategies' booking limits simultaneously rarely leading to high denied boarding costs.

Taken as a whole, $C^s/P^s/T^s$ performs better than $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$ if demand volume is underestimated. Here, $C^s/P^s/T^s$ offers less tickets due to lower booking limits retaining its revenue advantage. However, the higher the forecast overestimates demand volume, the less cautious booking limits result from $C^s/P^s/T^s$, leading to a more strongly declining revenue than $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$.

11.2.3 Inaccurate Demand Arrival Time Forecast

While the last section addressed the influence of an inaccurately forecasted demand volume, this section analyzes the influence of inaccurately forecasted demand arrival times. Figure 11.2 shows control strategies' revenue changes if forecasted demand arrival times are distorted from -100 to $+100$ days before departure.

Again, and in the remainder of this chapter, C^* constitutes revenue's upper bound unaffected by distortions. A negative distortion of demand forecast's arrival times causes control strategies to expect customer requests arriving earlier. As long as arrival times are distorted by -40 or less days before departure, $C^s/P^s/T^s$ remains the best performing strategy. But, if arrival times are distorted by -50 or more days before departure, $C_{ini}/P_1/T_0$ achieves the highest revenue.

In contrast, as demand arrival times are shifted closer to departure, $C^s/P^s/T^s$' and $C_{ini}/P_1/T_0$'s revenue results increasingly converge with a remaining marginal advantage of $C^s/P^s/T^s$. As opposed to this, control strategy $C(P^s)/P_1/T_0$ is mostly inferior to at least one of both strategies if forecasted demand arrival times are distorted.

11.2.4 Discussion on Robustness of Demand Forecasts

In conclusion, distorting demand forecasts reveals that $C^s/P^s/T^s$ benefits the most if demand is perfectly forecasted. This is also confirmed by analyzing both, an inaccurate demand forecast with respect to volume as well as arrival times: the higher the accuracy of forecasting future customer requests the more superior performs $C^s/P^s/T^s$.

However, if a demand forecast systematically under- or moderately overestimates demand volumes, $C^s/P^s/T^s$ constitutes the best approach. Also forecasting demand arrivals more than 40 days earlier than they actually arrive does not prevent $C^s/P^s/T^s$ from achieving the best revenue results compared to $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$.

The analyses showed that $C^s/P^s/T^s$'s performance is comparatively robust against systematic errors of demand forecasts. For that reason, the study only suggests using $C_{ini}/P_1/T_0$ if demand volume is highly overestimated or if demand arrivals are falsely expected far too early.

11.3 Robustness of Capacity Scenario Forecast

Previous computational studies of Section 10.2 and 10.3 assumed that the simulation accurately forecasts capacity scenarios. This assumption is now challenged by distorting forecasts on scenario's probabilities and times. First, Section 11.3.1 analyzes inaccurate scenario probabilities, then Section 11.3.2 examines inaccurate scenario times. In any case, the potential capacities to occur on a particular flight remain unaffected. Analogously to the study on demand forecast's robustness, this section tests the same control strategies: $C_{ini}/P_1/T_0$, $C(P^s)/P_1/T_0$ and $C^s/P^s/T^s$.

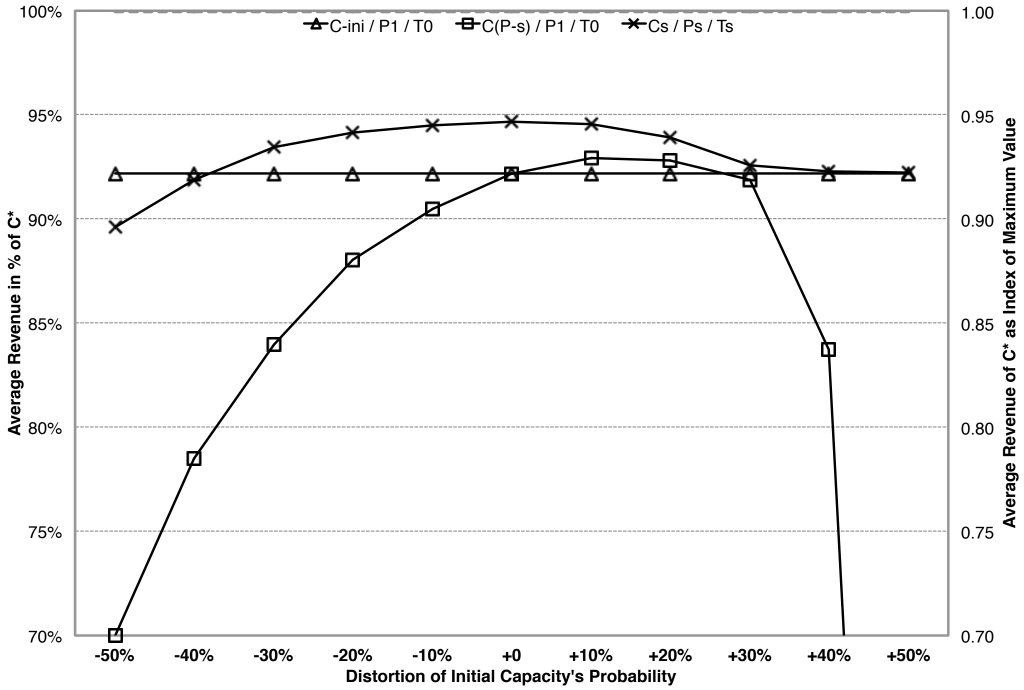

Figure 11.3: Impact of Distortion of Initial Capacity Scenario Probability Forecast

11.3.1 Inaccurate Capacity Scenario Probabilities

When distorting scenario's probabilities, the forecasted probability of the initial capacity's scenario increases or decreases while the remaining scenarios are accordingly adjusted. Figure 11.3 depicts the influence of such distortions on control strategies, ranging from -50% to $+50\%$.

The figure shows $C^s/P^s/T^s$ performing superior to $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$ when distorting scenario probabilities in the range of -30% to $+40\%$. Here, $C^s/P^s/T^s$ achieves the highest revenue if probabilities are undistorted ($+0\%$).

Strategy $C(P^s)/P_1/T_0$ even benefits from slightly increasing initial capacity's forecasted probability by $+10\%$ to $+20\%$. These forecast errors positively influence $C(P^s)/P_1/T_0$'s revenue result by weighting the used capacity closer to the initial capacity. However, $C(P^s)/P_1/T_0$ revenue vastly declines for values higher than $+40\%$. For more clarity, the figure cuts out $C(P^s)/P_1/T_0$'s revenue performance of 12.6% when distorting initial capacity's probability by $+50\%$.

In contrast to $C^s/P^s/T^s$ and $C(P^s)/P_1/T_0$, control strategy $C^{ini}/P_1/T_0$ is unaffected by a distorted scenario probability forecast. $C_{ini}/P_1/T_0$ does not take any scenario information into account and thus, performs on a constant revenue level of approx. 92% of C^*.

11.3.2 Inaccurate Capacity Scenario Times

This section analyzes the influence on control strategy's robustness if capacity scenario times are distorted. It involves systematically distorting scenario update times from -100 to $+100$ days before departure. Figure 11.4 shows control strategies' performance.

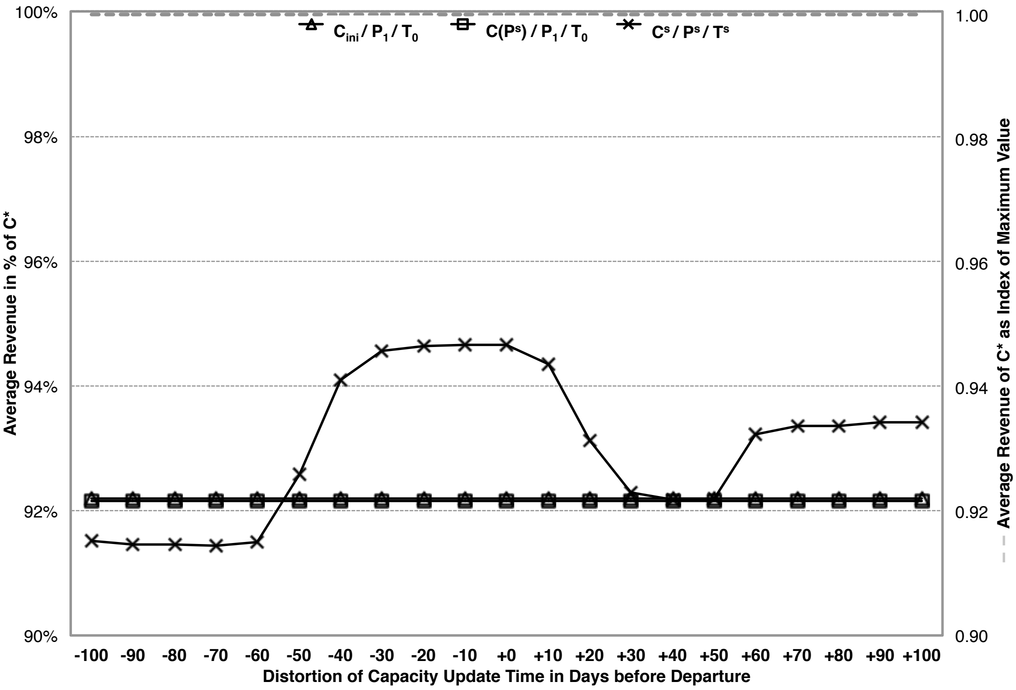

Figure 11.4: Impact of Distortion of Capacity Scenario Time Forecast

As $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$ neglect information on potential scenario times, they obviously remain uninfluenced by those changes. Their revenue performance is identical to the results when analyzing the value of information in Section 10.2. Both strategies – $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$ – achieve an almost equal revenue.

In contrast, $C^s/P^s/T^s$ revenue performance strongly fluctuates as it considers potential capacity scenario times. The revenue is highest if times are undistorted (+0 dbd). For distorting days before departure by +30 to +50 all control strategies perform alike due to various overlapping effects. However, if forecasted scenario times are distorted by −60 days before departure or more, $C^s/P^s/T^s$' performance falls below the revenue of $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$.

11.3.3 Discussion on Robustness of Capacity Scenario Forecasts

The previous sections tested the robustness of capacity scenario forecasts by distorting scenario's probabilities and times. When over- or underestimating the initial capacity's probability by less than 30%, $C^s/P^s/T^s$ achieves the best revenue results. In this way, the control strategy taking all information on capacity scenarios into account is comparatively robust against inaccurate forecasts on capacity scenario probabilities. Results suggest switching to control strategy $C_{ini}/P_1/T_0$ only if forecasted scenario probabilities strongly differ from actual probabilities.

Also with respect to distorted capacity scenario times, $C^s/P^s/T^s$ is rather robust. Only if capacity scenario update times differ by −60 days before departure or more, $C_{ini}/P_1/T_0$ and $C(P^s)/P_1/T_0$ outperform $C^s/P^s/T^s$. However, the study suggests using $C^s/P^s/T^s$ as long as no systematical error assumes capacity scenarios to occur much earlier than they actually appear.

12 Conclusion

The final chapter focuses on capacity uncertainty as addressed in Chapters 5–11. It is structured into three sections: First, Section 12.1 summarizes the work's major findings. Then, Section 12.2 suggests some practical implications for the problem of capacity uncertainty. Last but not least, Section 12.3 discusses limitations and shines a light on future research potentials.

12.1 Summary

This work in hand first presented a general overview of uncertainties and risks in RM with a focus on airlines, before it thoroughly addressed the hazard of capacity uncertainty resulting from unintended capacity updates. The work suggested a scenario-based RM problem formulation that systematically anticipates capacity uncertainty. The established simulation framework allowed for several computational studies analyzing the effects of capacity updates and shows revenue potentials when anticipating future updates with deliberation.

Chapter 1 presented the continuing success of RM in several industries and especially in the airline domain, which is mainly attributed to the vast body of research. Here, most efforts focus on considering existing uncertainties in the RM process or try to reduce remaining risks. As theory lacks an appropriate overview and distinction, this work introduced the concept of uncertainties and risks by transferring its differentiation to the field of RM.

Chapter 2 presented the concept of quantity-based RM with the help of a process view and established a simplified deterministic leg-based model for illustrating the fundamental trade-off between spill and spoilage in RM.

Based on this process, Chapter 3 and 4 presented a literature overview, discussing uncertainties and risks based on their location in the process; both chapters highlighted the historically most successful or recently most promising contributions. Here, Chapter 3 addresses uncertainties and risks within an RM process and Chapter 4 those from beyond.

While both chapters highlighted research opportunities in particular areas, especially the problem of capacity uncertainty is explicated. Capacity uncertainty results from assuming capacity to be fixed over the booking horizon, although it actually can be updated. Only very few contributions, with several limitations and a diverging focus, addressed the problem of capacity uncertainty yet. This gap highly motivated considering capacity uncertainty in ARM as a showcase for transforming an uncertainty to a risk.

Therefore, Chapter 5 distinguished this work's approach from previous contributions and defined research questions that help accomplishing the research objective of systematically anticipating capacity updates in RM.

Part II considered capacity uncertainty in RM by modeling capacity updates in the process and established an RM simulation framework. Based on this framework, Part III provided several computational studies for answering the remaining research questions. The remainder of this section chronologically summarizes the answers to research question 1–4.

RQ1: *How frequently do capacity updates occur in airline practice? Do they affect RM's result?*

Chapter 6 presented Lufthansa's empirical data of one year, ending in 2015, and revealed the occurrence of capacity updates in practice. Analyzing data on flights, capacity updates and bookings showed: across markets, more than 50% of all flights are affected by a capacity update. The majority of updates alters initial capacity by more than 25%. In general, empirical capacity increases and decreases approximately occur equally often. Over the time horizon, capacity updates occur from 270 to 0 days before departure due to various reasons, such as technical defects or reassigning aircrafts.

All these insights justify assuming capacity updates as a threat to RM's performance. Capacity updates' influence on revenue results is analyzed in various studies. First, Chapter 9 tests the influence of capacity updates based on artificial input data, then Chapter 10 provides results based on realistic input data. Both chapters demonstrate that capacity updates can strongly alter RM's result: Under some conditions, revenue decreased by almost 50% for artificial test instances and even led to negative revenues for instances calibrated on empirical data.

Concluding, this work stated that capacity updates can frequently occur in airline practice and involve the hazard of revenue losses.

RQ2: *Which capacity update characteristics are relevant to an RM system? How to model the capacity update characteristics in an RM system?*

Chapter 6 also revealed that three capacity update characteristics are of major interest: the new capacity transferred to the system if an update occurs, the time of an update and its probability that a particular update occurs. Accordingly, Chapter 7 modeled these characteristics in revenue optimization by extending the previously introduced leg-based deterministic RM model of Chapter 2. The extended model considers capacity updates as scenarios, each represented by a particular capacity, time and probability. This formulation allows for anticipating potential updates in advance.

Based on this formulation, Chapter 7 distinguished different information levels that mirror the knowledge on capacity update characteristics: information given solely on potential capacities, on capacities and update times, on capacities and update probabilities or information given on all three characteristics. Altogether, Chapter 7 established 16 control strategies, each based on an information level, including a theoretical upper bound and a benchmark reflecting current airlines' practice. Based on these information levels, each control strategy uniquely parameterizes the optimization model.

Subsequently, Chapter 8 established a simulation framework, including the revenue optimization model, for taking capacity updates into account. The simulation framework, i. a., provides stochastically generated demand and forecasts on potential capacity updates and demand.

Concluding, the work identified three major characteristics of capacity updates: updated capacity, update time and update probability. The established scenario-based formulation considered capacity updates in revenue optimization and the implemented simulation framework modeled these updates in an RM system.

RQ3: *What are the benefits of anticipating capacity updates? Are some capacity update characteristics more important to know than others?*

Chapter 10 presented the results of a computational study, calibrating simulations on empirical data. The simulations' input values are derived from the preprocessed airline data on flights, capacity updates and bookings of Chapter 6. The study compared the established 14 control strategies for coping with capacity uncertainty to an upper bound and a benchmark strategy not anticipating capacity updates.

The study of Chapter 10 highlighted that control strategies having information on the probability of potential capacity updates are more beneficial than control strategies having information on the time of potential capacity updates. Comparing the best performing strategies of these two information levels shows that the value of knowing update probabilities, depending on market conditions, generates 1.55–8.47 percent points more revenue than knowing update times in advance. The best performing strategy with information on potential capacities and update probabilities could, however, only partially outperform the benchmark strategy.

Nevertheless, the study also showed that revenue benefits are the highest if taking complete information on potential capacities, times and probabilities into account. This specific control strategy outperforms the benchmark strategy, depending on market conditions, by 0.04–2.71 percent points – with an average advantage of 2.47 percent points.

Concluding, the more information on potential capacity updates is available, the higher are potential revenue improvements. Here, having information on update probabilities is more valuable than having information on update times. Depending on the specific market, the suggested strategies sometimes more, sometimes less outperformed the strategy not anticipating capacity updates.

RQ4: *Which factors influence RM's result if capacity updates are anticipated? How does a distorted demand forecast and a distorted forecast on capacity updates affect RM's result?*

Chapter 10 provided sensitivity analyses of several problem parameters. Both demand volume and demand mix moderately influence control strategies' revenue: the more high-fare customers request, the higher strategies' resulting revenue. In contrast, the influence of the number of capacity updates per flight and the update magnitude cannot be considered in isolation. The resulting ambiguity can mainly be attributed to the time of the final capacity update. The later the last update occurs, the higher the revenue improvement of control strategies that anticipate potential capacity updates.

Chapter 11 analyzed control strategies' revenue when distorting capacity update forecasts and demand forecasts. The control strategy taking full information on potential capacity updates into account is comparatively robust against distorting forecasted demand volume and demand arrival times. The strategy only falls behind if demand volumes are strongly positively distorted or demand arrival times are assumed to occur much earlier than they actually do.

12.2 Implications for Practice

This work's consideration of capacity uncertainty and its various computational studies imply that ARM practice could strongly benefit from anticipating capacity updates in advance. Although capacity updates are frequent in airline practice, the systems commonly still assume capacities to be fixed over the booking horizon. Comparing several control strategies coping with uncertain capacity demonstrated the existing revenue potential for airlines.

If airlines can forecast future capacity updates with sufficient quality, taking these information into account enables revenue improvements. The studies showed that considering capacity updates in revenue optimization is especially beneficial if customers arrive early and updates occur late. However, if most customers arrive after the occurrence of potential capacity updates, airlines can still neglect capacity updates in RM.

Furthermore, the potential for revenue improvements are highest if capacity updates occur frequently and if updates' former and new capacity strongly differ. In this context, capacity uncertainty probably constitutes a greater hazard for network carriers as for low-cost carriers. Most commonly, network carriers serve heterogeneous markets and flight routes requiring a more diverse fleet. In contrast, low-cost carriers' flight routes are more homogeneous and their pilots are primarily specialized in only a few aircraft types. Thus, low-cost carriers' fleet largely consists of similar aircrafts with a typically similar capacity. A capacity update, hence, will rarely influence revenue results at the same extent as for network carriers.

Also the directions of capacity updates are ex ante important to analyze. Here, potential capacity increases enable some - mostly small - revenue improvements, while considering potential capacity decreases prevents RM systems from possibly vast revenue losses. Although, implementing an independent demand model is expected to stimulate cannibalization. This enables bigger revenue improvements as the most beneficial strategy's advantage mainly relies on picking out those customers, who are willing to pay high prices for a ticket.

As expected, the most beneficial strategy takes full information on potential capacities, update times and update probabilities into account. Nevertheless, airlines still could profit from only considering potential capacities and their update probabilities. This applies if, e. g., capacity update times strongly fluctuate. Then, RM systems inaccurately forecast potential update times, leading to disadvantageous revenue results. In this situation, this work's proposed a single strategy that still enables revenue improvements, compared to not anticipating capacity updates. This strategy averages capacities based on their probability to occur; the averaged capacity serves as input for optimization.

Although this work demonstrated the potential of systematically anticipating capacity updates in RM systems, the computational studies' results only serve as an indicator. The optimization model simplified the real-world optimization problem and empirical data has first been aggregated before used for calibrating simulations. It remains interesting to see if anticipating capacity updates could lead to perhaps even higher benefits in airline practice.

Also, applying the concept of anticipating capacity uncertainty to other RM industries appears to be promising. However, the stated conditions must apply as well: frequent capacity updates, strongly diverging capacities and a demand arriving usually before updates occur. Last but not least, the price of one potential added, or the costs of one potential reduced, capacity is likewise important as the magnitude of capacity updates. For example, adding five capacity units enabling revenue of 10 EUR each is as beneficial as adding ten capacity units for 5 EUR each.

12.3 Limitations and Future Research Potential

Compared to Wang and Regan (2006) this work's approach extended the problem of capacity uncertainty to an overall more realistic framework: potential capacity updates must not perfectly be known, updates can occur at any time and to any new capacity and the number of potential updates is unlimited. However, some limitations still remain.

As most common in airline practice, the established mathematical model for optimizing revenue considering capacity updates is based on a discrete-time RM approach. However, continuous-time approaches would require an appropriate formulation, for example as in Wang and Regan (2006).

Although the simulation system generates demand stochastically, the optimization model assumes demand to be deterministic. In contrast, Wang and Regan (2006) consider stochastic demand. A capacity uncertainty considering model with stochastic demand would be even more realistic.

This work's capacity uncertainty considering model is also based on a two-stage approach, same as the model in Wang and Regan (2006). However, a model accounting for more than one stage with uncertain capacity could be of interest.

The proposed policy refers to a leg-based RM approach. However, optimizing revenue in a network (e. g. van Ryzin and Vulcano (2008)) and simultaneously considering capacity uncertainty is of high interest. Here, forecasted capacity scenarios of different flights are assumed to be linked with each other. However, the assumption of Wang and Regan (2006) that aircrafts are limited to swaps, including only two flights changing their aircrafts, probably does not hold. For example, only a part of an airlines' fleet can be in use or aircrafts could be changed in circles, including three or more aircrafts.

This work assumed demand for fare classes to be independent of other fare classes' availability; however, assuming dependent demand is closer to ARM reality (Talluri et al., 2008; Weatherford & Ratliff, 2010). As already stated, the effects of independent demand could lead to cannibalization effects, enabling additional revenue improvements for this work's suggested strategy. For considering dependent demand, the model could easily adopt the approach of adjusted fare classes (Fiig et al., 2009).

Also, computational studies' fares are comparatively close to each other – factor two from cheapest to most expensive fare. In practice, fares can drastically differ – factor ten and above – enabling even higher revenue improvements.

To prevent overlapping effects, this work ignored cancellations and no-shows. Although, the proposed model could incorporate overbooking (Talluri & van Ryzin, 2004b, Chapter 4), as control strategies can easily use virtual instead of physical capacities. The interplay between virtually increased capacities by overbooking and considering capacity uncertainty appears interesting to analyze.

Furthermore, this work only analyzed a single compartment with three fare classes. Most commonly, compartments are separately controlled, however, buy-ups and buy-downs also occur occasionally from one compartment to another.

This work only focused on capacity uncertainty in the revenue optimization stage. However, RM systems require an actual forecast on capacity updates in advance. Although data on capacity updates' entire history is accessible, considering updates as actual individual scenarios is expected to exceed practice's computational capability. Thus, aggregating similar capacity updates in clusters, as suggested in this work, appears unavoidable.

Last but not least, the computational study only separately analyzed distorting a single attribute of demand forecasts or capacity forecasts. Testing multi-dimensional forecast errors could also be insightful.

References

Abara, J. (1989). Applying integer linear programming to the fleet assignment problem. *Interfaces*, *19*(4), 20–28.

Adelson, R. M. (1966). Compound poisson distributions. *OR*, *17*(1), 73–75.

Alstrup, J., Boas, S., Madsen, O. B. & Vidal, R. V. (1986). Booking policy for flights with two types of passengers. *European Journal of Operational Research*, *27*(3), 274–288.

Ampountolas, A. (2016). *Revenue management, dynamic pricing and social media in the tourism industry: A case study of the name-your-own-price mechanism* (Dissertation). University of Exeter.

Anderson, C. K. (2009). Setting prices on priceline. *Interfaces*, *39*(4), 307–315.

Anderson, C. K. & Wilson, J. G. (2010). Name-your-own price auction mechanisms–Modeling and future implications. *Journal of Revenue and Pricing Management*, *10*(1), 32–39.

Aydin, N., Birbil, S. I., Frenk, J. & Noyan, N. (2013). Single-leg airline revenue management with overbooking. *Transportation Science*, *47*(4).

Azadeh, S. S., Hosseinalifam, M. & Savard, G. (2015). The impact of customer behavior models on revenue management systems. *Computational Management Science*, *12*(1), 99–109.

Ball, M. O. & Queyranne, M. (2009). Toward robust revenue management: Competitive analysis of online booking. *Operations Research*, *57*(4), 950–963.

Barnhart, C., Farahat, A. & Lohatepanont, M. (2009). Airline fleet assignment with enhanced revenue modeling. *Operations Research*, *57*(1), 231–244.

Barnhart, C., Kniker, T. & Lohatepanont, M. (2002). Itinerary-based airline fleet assignment. *Transportation Science*, *36*(2), 199–217.

Bartke, P. (2014). *Demand estimation in airline revenue management* (Dissertation). Freie Universität Berlin.

Barz, C. & Waldmann, K. (2007). Risk-sensitive capacity control in revenue management. *Mathematical Methods of Operations Research*, *65*(3), 565–579.

Beckmann, M. & Bobkoski, F. (1958). Airline demand: An analysis of some frequency distributions. *Naval Research Logistics Quarterly*, *5*(1), 43–51.

Belobaba, P. (1987). *Air travel demand and airline seat inventory management* (Dissertation). Massachusetts Institute of Technology.

Belobaba, P. (2009). The airline planning process. *The Global Airline Industry*, 153–181.

Belobaba, P. & Farkas, A. (1999). Yield management impacts on airline spill estimation. *Transportation Science*, *33*(2), 217.

Belobaba, P. & Wilson, J. (1997). Impacts of yield management in competitive airline markets. *Journal of Air Transport Management*, *3*(1), 3–9.

Berge, M. E. & 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. P. (2005). Dynamic programming and suboptimal control: A survey from ADP to MPC*. *European Journal of Control*, *11*(4), 310–334.

Bertsimas, D. & Popescu, I. (2003). Revenue management in a dynamic network environment. *Transportation Science*, *37*(3), 257–277.

Bertsimas, D. & Sim, M. (2004). The price of robustness. *Operations Research*, *52*(1), 35–53.

Besbes, O. & Zeevi, A. (2009). Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, *57*(6), 1407–1420.

Bish, E. K., Suwandechochai, R. & Bish, D. R. (2004). Strategies for managing the flexible capacity in the airline industry. *Naval Research Logistics (NRL)*, *51*(5), 654–685.

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

Board, S. & Skrzypacz, A. (2016). Revenue management with forward-looking buyers. *Journal of Political Economy*, *124*(4), 1046–1087.

Box, G. E., Jenkins, G. M., Reinsel, G. C. & Ljung, G. M. (2015). *Time series analysis: Forecasting and control.* John Wiley & Sons.

Buhl, H. U., Klein, R., Kolb, J. & Landherr, A. (2011). CR^{2}M – An approach for capacity control considering long-term effects on the value of a customer for the company. *Journal of Management Control*, *22*(2), 187–204.

Bulmer, M. G. (2012). *Principles of statistics.* Courier Corporation.

Büsing, C., Kadatz, D. & Cleophas, C. (2017). *Considering exogenous capacity changes in airline revenue management: Models, algorithms and computations* [Working Paper].

Carvalho, H., Duarte, S. & Cruz Machado, V. (2011). Lean, agile, resilient and green: Divergencies and synergies. *International Journal of Lean Six Sigma*, *2*(2), 151–179.

Chen, L. & Homem-de Mello, T. (2010). Re-solving stochastic programming models for airline revenue management. *Annals of Operations Research*, *177*(1), 91–114.

Chen, Y., Yang, B., Dong, J. & Abraham, A. (2005). Time-series forecasting using flexible neural tree model. *Information Sciences*, *174*(3), 219–235.

Chiang, W.-C., Chen, J. C. & Xu, X. (2007). An overview of research on revenue management: current issues and future research. *International Journal of Revenue Management*, *1*(1), 97–128.

Chiu, C. & Tsao, C. (2004). The optimal airline overbooking strategy under uncertainties. In *Knowledge-Based Intelligent Information and Engineering Systems* (pp. 937–945).

Choi, S. & Mattila, A. S. (2004). Hotel revenue management and its impact on customers' perceptions of fairness. *Journal of Revenue and Pricing Management*, *2*(4), 303–314.

Cleophas, C. (2009). *Simulation-based analysis of forecast performance evaluations for airline revenue management* (Dissertation). University of Paderborn.

Cleophas, C., Frank, M. & Kliewer, N. (2009). Simulation-based key performance indicators for evaluating the quality of airline demand forecasting. *Journal of Revenue and Pricing Management*, *8*(4), 330–342.

Cleophas, C., Kadatz, D. & Vock, S. (2017). Resilient revenue management: A literature survey of recent theoretical advances. *To be published in Journal of Revenue & Pricing Management.*

Cooper, W., Homem-de Mello, T. & Kleywegt, A. (2006). Models of the spiral-down effect in revenue management. *Operations Research*, *54*(5), 968–987.

Cooper, W. L., Homem-de Mello, T. & Kleywegt, A. J. (2015). Learning and pricing with models that do not explicitly incorporate competition. *Operations Research*, *63*(1), 86–103.

Currie, C. & Rowley, I. (2010). Consumer behaviour and sales forecast accuracy: What's going on and how should revenue managers respond? *Journal of Revenue & Pricing Management*, *9*(4), 374–376.

De Boer, S. (2003). *Advances in airline revenue management and pricing* (Dissertation). Massachusetts Institute of Technology.

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.

den Boer, A. V. & Zwart, B. (2015). Dynamic pricing and learning with finite inventories. *Operations research*, *63*(4), 965–978.

Dilme, F. & Li, F. (2016). *Revenue management without commitment: Dynamic pricing and periodic fire sales* [Working Paper].

Du, J., Zhang, J. & Hua, G. (2015). Pricing and inventory management in the presence of strategic customers with risk preference and decreasing value. *International Journal of Production Economics*, *164*, 160–166.

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

Farias, V. F., Jagabathula, S. & Shah, D. (2013). A nonparametric approach to modeling choice with limited data. *Management Science*, *59*(2), 305–322.

Farkas, A. (1996). *The influence of network effects and yield management on Airline Fleet Assignment* (Dissertation). Massachusetts Institute of Technology.

Fay, S. (2008). Selling an opaque product through an intermediary: The case of disguising one's product. *Journal of Retailing*, *84*(1), 59–75.

Federgruen, A. & Heching, A. (1999). Combined pricing and inventory control under uncertainty. *Operations Research*, *47*(3), 454–475.

Feng, Y. & 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. & Xiao, B. (1999). Maximizing revenues of perishable assets with a risk factor. *Operations Research*, *47*(2), 337–341.

Feng, Y. & Xiao, B. (2006). Integration of pricing and capacity allocation for perishable products. *European Journal of Operational Research*, *168*(1), 17–34.

Feng, Y. & Xiao, B. (2008). Technical note – A risk-sensitive model for managing perishable products. *Operations Research*, *56*(5), 1305–1311.

Fiig, T., Isler, K., Hopperstad, C. & Belobaba, P. (2009). Optimization of mixed fare structures: Theory and applications. *Journal of Revenue and Pricing Management*, *9*(1), 152–170.

Frank, M., Friedemann, M., Mederer, M. & Schroeder, A. (2006). Airline revenue management: A simulation of dynamic capacity management. *Journal of Revenue and Pricing Management*, *5*(1), 62–71.

Gallego, G., Iyengar, G., Phillips, R. & Dubey, A. (2004). *Managing flexible products on a network.* Citeseer.

Gallego, G., Li, L. & Ratliff, R. (2009). Choice-based EMSR methods for single-leg revenue management with demand dependencies. *Journal of Revenue and Pricing Management*, *8*(2), 207–240.

Gallego, G. & Phillips, R. (2004). Revenue management of flexible products. *Manufacturing & Service Operations Management*, *6*(4), 321–337.

Gallego, G. & van Ryzin, G. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, *40*(8), 999–1020.

Gallego, G. & van Ryzin, G. (1997). A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research*, *45*(1), 24–41.

Gao, H., Ball, M. O. & Karaesmen, I. Z. (2016). Distribution-free methods for multi-period, single-leg booking control. *Journal of Revenue and Pricing Management*, *15*(6), 425–453.

Garrow, L. A. (2016). *Discrete choice modelling and air travel demand: theory and applications*. Routledge.

Geraghty, M. K. & Johnson, E. (1997). Revenue management saves national car rental. *Interfaces*, *27*(1), 107–127.

Gerlach, M. (2013). *Decentralized code-share revenue management in airline alliances* (Dissertation). Freie Universität Berlin.

Gerlach, M., Cleophas, C. & Kliewer, N. (2013). Airline codeshare alliances. *Business & Information Systems Engineering*, *5*(3), 153–163.

Gönsch, J. & Hassler, M. (2014). Optimizing the conditional value-at-risk in revenue management. *Review of Managerial Science*, *8*(4), 495–521.

Gönsch, J., Klein, R., Neugebauer, M. & Steinhardt, C. (2013). Dynamic pricing with strategic customers. *Journal of Business Economics*, *83*(5), 505–549.

Gönsch, J., Koch, S. & Steinhardt, C. (2014). Revenue management with flexible products: The value of flexibility and its incorporation into DLP-based approaches. *International Journal of Production Economics*, *153*, 280–294.

Gönsch, J. & Steinhardt, C. (2013). Using dynamic programming decomposition for revenue management with opaque products. *BuR-Business Research*, *6*(1), 94–115.

Gorgeon, A. (2015). Anti-fragile information systems. *Thirty Sixth International Conference on Information Systems, Fort Worth.*

Gorin, T., Walczak, D., Bartke, P. & Friedemann, M. (2012). Incorporating cancel and rebook behavior in revenue management optimization. *Journal of Revenue and Pricing Management*, *11*(6), 645–660.

Grauberger, W. & Kimms, A. (2016). Revenue management under horizontal and vertical competition within airline alliances. *Omega*, *59*, 228–237.

Guidolin, M. & Timmermann, A. (2009). Forecasts of US short-term interest rates: A flexible forecast combination approach. *Journal of Econometrics*, *150*(2), 297–311.

Hendler, R. & Hendler, F. (2004). Revenue management in fabulous Las Vegas: Combining customer relationship management and revenue management to maximise profitability. *Journal of Revenue and Pricing Management*, *3*(1), 73–79.

Heo, C. Y. (2016). Exploring group-buying platforms for restaurant revenue management. *International Journal of Hospitality Management*, *52*, 154–159.

Hinz, O., Hann, I. H. & Spann, M. (2011). Price discrimination in E-commerce? An examination of dynamic pricing in name-your-own price markets. *MIS Quarterly-Management Information Systems*, *35*(1), 81.

Hu, X., Caldentey, R. & Vulcano, G. (2013). Revenue sharing in airline alliances. *Management Science*, *59*(5), 1177–1195.

Huang, K. & Chang, K.-C. (2009). A model for airline seat control considering revenue uncertainty and risk. *Journal of Revenue and Pricing Management*, *10*(2), 161–171.

Isler, K. & D'Souza, E. (2009). GDS capabilities, OD control and dynamic pricing. *Journal of Revenue & Pricing Management*, *8*(2), 255–266.

Isler, K. & Imhof, H. (2008). A game theoretic model for airline revenue management and competitive pricing. *Journal of Revenue and Pricing Management*, *7*(4), 384–396.

Jagabathula, S. & Vulcano, G. (2015). A model to estimate individual preferences using panel data. *Available at SSRN*, *2560994*.

Jasin, S. (2014). Reoptimization and self-adjusting price control for network revenue management. *Operations Research*, *62*(5), 1168–1178.

Jasin, S. (2015). Performance of an LP-based control for revenue management with unknown demand parameters. *Operations Research*, *63*(4), 909–915.

Jasin, S. & Kumar, S. (2012). A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research*, *37*(2), 313–345.

Jerath, K., Netessine, S. & Veeraraghavan, S. (2010). Revenue management with strategic customers: Last-minute selling and opaque selling. *Management Science*, *56*(3), 430–448.

Kadatz, D., Kliewer, N. & Cleophas, C. (2017a). *Anticipating capacity updates in airline revenue management: The value of information* [Working Paper].

Kadatz, D., Kliewer, N. & Cleophas, C. (2017b). Exogenous capacity changes in airline revenue management: Quantifying the value of information. *To be published in OR 2016 Proceedings.*

Karaesmen, I. & van Ryzin, G. (2004). Overbooking with substitutable inventory classes. *Operations Research*, *52*(1), 83–104.

Kaufman, L. & Rousseeuw, P. (1987). *Clustering by means of medoids.* Elsevier/North-Holland.

Kaufman, L. & Rousseeuw, P. J. (2009). *Finding groups in data: An introduction to cluster analysis* (Vol. 344). John Wiley & Sons.

Kimms, A. & Müller-Bungart, M. (2007). Revenue management for broadcasting commercials: The channel's problem of selecting and scheduling the advertisements to be aired. *International Journal of Revenue Management*, *1*(1), 28–44.

Kimms, A. & Müller Bungart, M. (2007). Simulation of stochastic demand data streams for network revenue management problems. *OR Spectrum*, *29*(1), 5–20.

Knight, F. H. (1921). Risk, uncertainty and profit. *New York: Hart, Schaffner and Marx.*

Kocabiyikoglu, A., Popescu, I. & Stefanescu, C. (2013). Pricing and revenue management: The value of coordination. *Management Science*, *60*(3), 730–752.

Koenig, M. & Meissner, J. (2009). *Risk minimizing strategies for revenue management problems with target values* [Working Paper].

Koenig, M. & Meissner, J. (2010). *Value-at-risk optimal policies for revenue management problems* [Working Paper].

Koenig, M. & Meissner, J. (2015). Risk management policies for dynamic capacity control. *Computers & Operations Research*.

Kumar, V. (2010). *Customer relationship management*. Wiley Online Library.

Kumar, V. & George, M. (2007). Measuring and maximizing customer equity: A critical analysis. *Journal of the Academy of Marketing Science*, *35*(2), 157–171.

Kumar, V., Ramani, G. & Bohling, T. (2004). Customer lifetime value approaches and best practice applications. *Journal of Interactive Marketing*, *18*(3), 60–72.

Ladany, S. P. & Arbel, A. (1991). Optimal cruise-liner passenger cabin pricing policy. *European Journal of Operational Research*, *55*(2), 136–147.

Lan, Y., Ball, M. O. & Karaesmen, I. Z. (2011). Regret in overbooking and fare-class allocation for single leg. *Manufacturing & Service Operations Management*, *13*(2), 194–208.

Lan, Y., Ball, M. O., Karaesmen, I. Z., Zhang, J. X. & Liu, G. X. (2015). Analysis of seat allocation and overbooking decisions with hybrid information. *European Journal of Operational Research*, *240*(2), 493–504.

Lan, Y., Gao, H., Ball, M. O. & Karaesmen, I. (2008). Revenue management with limited demand information. *Management Science*, *54*(9), 1594–1609.

Lancaster, J. (2003). The financial risk of airline revenue management. *Journal of Revenue and Pricing Management*, *2*(2), 158–165.

Lardeux, B., Goyons, O. & Robelin, C.-A. (2010). Availability calculation based on robust optimization. *Journal of Revenue and Pricing Management*, *9*(4), 313–325.

Lee, A. O. (1990). *Airline reservations forecasting: Probabilistic and statistical models of the booking process.* Technical Report.

Lee, M., Khelifa, A., Garrow, L., Bierlaire, M. & Post, D. (2012). An analysis of destination choice for opaque airline products using multidimensional binary logit models. *Transportation Research Part A: Policy and Practice*, *46*(10), 1641–1653.

Lee, T. & Hersh, M. (1993). A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science*, *27*(3), 252–265.

Lemke, C., Riedel, S. & Gabrys, B. (2012). Evolving forecast combination structures for airline revenue management. *Journal of Revenue and Pricing Management*, *12*(3), 221–234.

Lennon, J. (2004). Revenue management and customer forecasts: A bridge too far for the UK visitor attractions sector? *Journal of Revenue and Pricing Management*, *2*(4), 338–352.

Lhabitant, F.-S. & Tinguely, O. (2001). Financial risk management: An introduction. *Thunderbird International Business Review*, *43*(3), 343–363.

Li, J., Granados, N. & Netessine, S. (2014). Are consumers strategic? Structural estimation from the air-travel industry. *Management Science*, *60*(9), 2114–2137.

Liang, Y. (1999). Solution to the continuous time dynamic yield management model. *Transportation Science*, *33*(1), 117–123.

Littlewood, K. (1972). Forecasting and control of passenger bookings. In *AGIFORS Symposium Proc.* (Vol. 12).

Lochner, K. & Wellman, M. (2004). Rule-based specification of auction mechanisms. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems* (Vol. 2, pp. 818–825). IEEE Computer Society.

MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth berkeley symposium on mathematical statistics and probability* (Vol. 1, pp. 281–297).

Maglaras, C. & Eren, S. (2015). A maximum entropy joint demand estimation and capacity control policy. *Production and Operations Management*, *24*(3), 438–450.

Maglaras, C. & Meissner, J. (2006). Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing & Service Operations Management*, *8*(2), 136–148.

Mandelbaum, M. & Buzacott, J. (1990). Flexibility and decision making. *European Journal of Operational Research*, *44*(1), 17–27.

Mardia, K., Kent, J. & Bibby, J. (1979). Multivariate analysis. *Academic Press Inc London LTD.*

Marschak, T. & Nelson, R. (1962). Flexibility, uncertainty, and economic theory. *Metroeconomica*, *14*(1-2-3), 42–58.

Massart, D. L., Kaufman, L., Rousseeuw, P. J. & Leroy, A. (1986). Least median of squares: A robust method for outlier and model error detection in regression and calibration. *Analytica Chimica Acta*, *187*, 171–179.

McGill, J. & van Ryzin, G. (1999). Revenue management: Research overview and prospects. *Transportation Science*, *33*(2), 233.

Meissner, J. & Strauss, A. (2012). Network revenue management with inventory-sensitive bid prices and customer choice. *European Journal of Operational Research*, *216*(2), 459–468.

Meissner, J., Strauss, A. & Talluri, K. (2013). An enhanced concave program relaxation for choice network revenue management. *Production and Operations Management*, *22*(1), 71–87.

Mitra, D. & Wang, Q. (2005). Stochastic traffic engineering for demand uncertainty and risk-aware network revenue management. *IEEE/ACM Transactions on Networking (TON)*, *13*(2), 221–233.

Mohaupt, M. & Hilbert, A. (2013). Integrated decision support system for establishing a long-term profitable customer base in cloud computing. In *Proceedings of the international multiconference of engineers and computer scientists* (Vol. 1).

Mukhopadhyay, S., Samaddar, S. & Colville, G. (2007). Improving revenue management decision making for airlines by evaluating analyst-adjusted passenger demand forecasts. *Decision Sciences*, *38*(2), 309–327.

Noone, B. M., Kimes, S. E. & Renaghan, L. M. (2003). Integrating customer relationship management and revenue management: A hotel perspective. *Journal of Revenue and Pricing Management*, *2*(1), 7–21.

Papanastasiou, Y. & Savva, N. (2016). Dynamic pricing in the presence of social learning and strategic consumers. *Management Science*.

Park, I., Sharman, R. & Rao, H. R. (2015). Disaster experience and hospital information systems: An examination of perceived information assurance, risk, resilience, and his usefulness. *MIS Quarterly-Management Information Systems Quarterly*, *39*(2), 317–344.

Petrick, A., Gönsch, J., Steinhardt, C. & Klein, R. (2010). Dynamic control mechanisms for revenue management with flexible products. *Computers & Operations Research*, *37*(11), 2027–2039.

Petrick, A., Steinhardt, C., Gönsch, J. & Klein, R. (2012). Using flexible products to cope with demand uncertainty in revenue management. *OR Spectrum*, *34*(1), 215–242.

Petropoulos, F., Fildes, R. & Goodwin, P. (2016). Do 'big losses' in judgmental adjustments to statistical forecasts affect experts' behaviour? *European Journal of Operational Research*, *249*(3), 842–852.

Phillips, R. (2012). Efficient frontiers in revenue management. *Journal of Revenue and Pricing Management*, *11*(4), 371–385.

Pölt, S. (2010). The rise and fall of RM. *Journal of Revenue & Pricing Management*, *10*(1), 23–25.

Popovic, J. & Teodorovic, D. (1997). An adaptive method for generating demand inputs to airline seat inventory control models. *Transportation Research Part B: Methodological*, *31*(2), 159–175.

Post, D. (2010). Variable opaque products in the airline industry: A tool to fill the gaps and increase revenues. *Journal of Revenue and Pricing Management*, *9*(4), 292–299.

Pye, R. (1978). A formal, decision-theoretic approach to flexibility and robustness. *The Journal of the Operational Research Society*, *29*(3), pp.215–227.

Queenan, C. C., Ferguson, M., Higbie, J. & Kapoor, R. (2007). A comparison of unconstraining methods to improve revenue management systems. *Production and Operations Management*, *16*(6), 729–746.

Reinartz, W. J. & Kumar, V. (2002). The mismanagement of customer loyalty. *Harvard business review*, *80*(7), 86–95.

Reinartz, W. J. & Kumar, V. (2003). The impact of customer relationship characteristics on profitable lifetime duration. *Journal of marketing*, *67*(1), 77–99.

Robinson, L. W. (1995). Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes. *Operations Research*, *43*(2), 252–263.

Rothstein, M. (1971). An airline overbooking model. *Transportation Science*, *5*(2), 180–192.

Roy, B. (2010). Robustness in operational research and decision aiding: A multi-faceted issue. *European Journal of Operational Research*, *200*(3), 629–638.

Runde, J. (1998). Clarifying Frank Knight's discussion of the meaning of risk and uncertainty. *Cambridge Journal of Economics*, *22*(5), 539–546.

Rygielski, C., Wang, J.-C. & Yen, D. C. (2002). Data mining techniques for customer relationship management. *Technology in society*, *24*(4), 483–502.

Şen, A. (2013). A comparison of fixed and dynamic pricing policies in revenue management. *Omega*, *41*(3), 586–597.

Shen, Z.-J. M. & Su, X. (2007). Customer behavior modeling in revenue management and auctions: A review and new research opportunities. *Production and Operations Management*, *16*(6), 713–728.

Sherali, H. D., Bish, E. K. & Zhu, X. (2006). Airline fleet assignment concepts, models, and algorithms. *European Journal of Operational Research*, *172*(1), 1–30.

Shumsky, R. A. & Zhang, F. (2009). Dynamic capacity management with substitution. *Operations research*, *57*(3), 671–684.

Smith, B., Ratliff, R. & Jacobs, T. L. (1998). Optimal load factors: Coordinated scheduling, pricing and yield management decisions. In *AGIFORS Symposium Proceedings, Denpesar Bali, Indonesia.*

Smith, B. C., Leimkuhler, J. F. & Darrow, R. M. (1992). Yield Management at American Airlines. *Interfaces*, *22*(1), 8–31.

Steinhaus, H. (1956). Sur la division des corp materiels en parties. *Bull. Acad. Polon. Sci*, *1*, 801–804.

Subramanian, J., Stidham, S. & Lautenbacher, C. (1999). Airline yield management with overbooking, cancellations, and no-shows. *Transportation Science*, *33*(2), 147–167.

Swanson, N. R. & White, H. (1997). Forecasting economic time series using flexible versus fixed specification and linear versus nonlinear econometric models.

International Journal of Forecasting, *13*(4), 439–461.

Taleb, N. N. (2012). *Antifragile: How to live in a world we don't understand* (Vol. 3). Allen Lane London.

Talluri, K. & van Ryzin, G. (1998). An analysis of bid-price controls for network revenue management. *Management Science*, *44*(11-part-1), 1577–1593.

Talluri, K. & van Ryzin, G. (2004a). Revenue management under a general discrete choice model of consumer behavior. *Management Science*, *50*(1), 15–33.

Talluri, K. & van Ryzin, G. (2004b). *The theory and practice of revenue management*. Springer.

Talluri, K., van Ryzin, G., Karaesmen, I. & Vulcano, G. (2008). Revenue management: Models and methods. In *Proceedings of the 40th Conference on Winter Simulation* (pp. 145–156).

van Ryzin, G. (2005). Future of revenue management: Models of demand. *Journal of Revenue & Pricing Management*, *4*(2), 204–210.

van Ryzin, G. & Vulcano, G. (2008). Simulation-based optimization of virtual nesting controls for network revenue management. *Operations Research*, *56*(4), 865–880.

van Ryzin, G. & Vulcano, G. (2014). A market discovery algorithm to estimate a general class of nonparametric choice models. *Management Science*, *61*(2), 281–300.

Venkatesan, R. & Kumar, V. (2004). A customer lifetime value framework for customer selection and resource allocation strategy. *Journal of marketing*, *68*(4), 106–125.

Verhoef, P. C. & Donkers, B. (2001). Predicting customer potential value an application in the insurance industry. *Decision Support Systems*, *32*(2), 189–199.

Vinod, B. (2005). Retail revenue management and the new paradigm of merchandise optimisation. *Journal of Revenue and Pricing Management*, *3*(4), 358–368.

Vock, S. E. (2016). *Customers' preferential choice for manifestations of flexible products* (Dissertation). Freie Universität Berlin.

Vock, S. E., Cleophas, C. & Kliewer, N. (2017). *Opportunistic customer behavior for revenue management with flexible products* [working paper].

von Martens, T. & Hilbert, A. (2011). Customer-value-based revenue management. *Journal of Revenue & Pricing Management*, *10*(1), 87–98.

Vulcano, G., van Ryzin, G. & Chaar, W. (2010). Choice-based revenue management: An empirical study of estimation and optimization. *Manufacturing & Service Operations Management*, *12*(3), 371–392.

Vulcano, G. & Weil, A. (2014). *Joint optimization of virtual capacities and bid-prices for revenue management* [Working Paper].

Wang, T., Gal Or, E. & Chatterjee, R. (2009). The name-your-own-price channel in the travel industry: An analytical exploration. *Management Science*, *55*(6), 968–979.

Wang, X. & Regan, A. (2006). Dynamic yield management when aircraft assignments are subject to swap. *Transportation Research Part B: Methodological*, *40*(7), 563–576.

Wang, X. & Regan, A. C. (2002). *Stochastic and dynamic airline yield management when aircraft assignments are subject to change.* Technical Report.

Wang, Z., Deng, S. & Ye, Y. (2014). Close the gaps: A learning-while-doing algorithm for single-product revenue management problems. *Operations Research*, *62*(2), 318–331.

Weatherford, L. (2016). The history of unconstraining models in revenue management. *Journal of Revenue and Pricing Management*, *15*(3-4), 222–228.

Weatherford, L. & Belobaba, P. (2002). Revenue impacts of fare input and demand forecast accuracy in airline yield management. *Journal of the Operational Research Society*, 811–821.

Weatherford, L. & 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. & Ratliff, R. (2010). Review of revenue management methods with dependent demands. *Journal of Revenue and Pricing Management*, *9*(4), 326–340.

Weatherford, L. R. (2004). Additional paper: EMSR versus EMSU: Revenue or utility? *Journal of Revenue and Pricing Management*, *3*(3), 277–284.

Weatherford, L. R. (2015). Intelligent aggressiveness: Combining forecast multipliers with various unconstraining methods to increase revenue in a global network with four airlines. *Journal of Revenue and Pricing Management*, *14*(2), 84–96.

Wieland, A. & Marcus Wallenburg, C. (2012). Dealing with supply chain risks: Linking risk management practices and strategies to performance. *International Journal of Physical Distribution & Logistics Management*, *42*(10), 887–905.

Williamson, E. L. (1992). *Airline network seat inventory control: Methodologies and revenue impacts* (Dissertation). Massachusetts Institute of Technology, Dept. of Aeronautics & Astronautics.

Wilson, J. & Zhang, G. (2008). Optimal design of a name-your-own price channel. *Journal of Revenue and Pricing Management*, *7*(3), 281–290.

Wirtz, J., Kimes, S. E., Theng, J. H. P. & Patterson, P. (2003). Revenue management: Resolving potential customer conflicts. *Journal of Revenue and Pricing Management*, *2*(3), 216–226.

Woodruff, R. B. (1997). Customer value: The next source for competitive advantage. *Journal of the academy of marketing science*, *25*(2), 139–153.

Wright, C. P. (2014). Decomposing airline alliances: A bid-price approach to revenue management with incomplete information sharing. *Journal of Revenue and Pricing Management*, *13*(3), 164–182.

Yu, Y., Chen, X. & Zhang, F. (2015). Dynamic capacity management with general upgrading. *Operations Research*, *63*(6), 1372–1389.

Zeni, R. (2001). *Improved forecast accuracy in airline revenue management by unconstraining demand estimates from censored data* (Dissertation). State University of New Jersey.

Zeni, R. (2003). The value of analyst interaction with revenue management systems. *Journal of Revenue and Pricing Management*, *2*(1), 37–46.

Zimmermann, B. (2014). *Revenue management with repeated competitive interactions* (Dissertation). Freie Universität Berlin, Germany.

Zou, J., Ahmed, S. & Sun, X. A. (2016). Nested decomposition of multistage stochastic integer programs with binary state variables. *Optimization Online*.

Appendix

A. Report on Capacity Update Data Preparation

Section 6.1.1 prepared empirical flights data as follows. The raw data covers 48,295 flights. Here, data also includes trains and bus transfers to an airport. These transfers are removed from the set, as transfers do not represent actual flights. The same applies to a few flights that are not actually operated by Lufthansa, which were only included for test reasons. Not part of the data, however, are flights operated by an alliance partner, called *code shares*. After removing all unintended data points, 44,642 flights remain.

The presented data on capacity updates in Section 6.1 are prepared in the following way: the raw data on capacity updates spans 71,489 observations; each observation constitutes one data point. The data includes the former and the new aircraft type, which are, however, not taken into account as this work is only interested in the number of seats. In some observations the former and the new number of seats are equal. This can occur if former and new aircraft type differ, but the number of seats remains identical. As those updates do not constitute an actual change in capacity, they are removed from analysis. The original data on capacity updates also involves different compartments. Only those updates affecting the economy compartment are taken into account. The remaining 28,522 observations on capacity updates are analyzed per market.

B. Report on Capacity Update Clusters

Section 6.1.2 aggregated multiple similar capacity updates into single clusters. Kaufman and Rousseeuw (2009) provide a comprehensive overview on clustering techniques. Steinhaus (1956) introduced the well-known clustering approach k-means, posterior termed by MacQueen (1967). Later, Kaufman and Rousseeuw (1987) suggest the k-medoids clustering algorithm. Both approaches, k-means and k-medoids, break up all data points into a number of k clusters and then calculate each clusters' center. Both k-means and k-medoids try to find one point with the smallest distance to a group's data points. In contrast to k-means, where this point can be virtual, k-medoids determines a real data point as center. By that, k-medoids is more robust to outliers (Kaufman & Rousseeuw, 1987; Massart et al., 1986), similar to the statistical measure of a median compared to a mean. Preliminary studies showed that the given data involves few drastic outliers. To prevent the resulting distortion, this work calculates cluster centroids based on a medoids approach.

Although, k-medoids and k-means, require a priori determining the number of clusters, k. There are different rules that allow to determine this number. The rule of thumb by Mardia, Kent and Bibby (1979), e. g., says that $k \approx \sqrt{(n/2)}$. Here, n is the number of observations. For LONG, with $n = 98$ data points, this would lead to $k = 7$, which is still a manageable number of clusters. However, there are $n = 9,201$ data points for SHORT and $n = 19,222$ for MEDIUM, which would result in $k \approx 68$ and $k \approx 98$ clusters. The larger these numbers, the less computationally efficient are the computational studies. Hence, the analysis decides to manually create clusters of capacity updates by magnitude and time based on expert input as mentioned in Section 6.1.2. Figures B.1–B.13 show data points and cluster centroids for every combination of market and aircraft size.

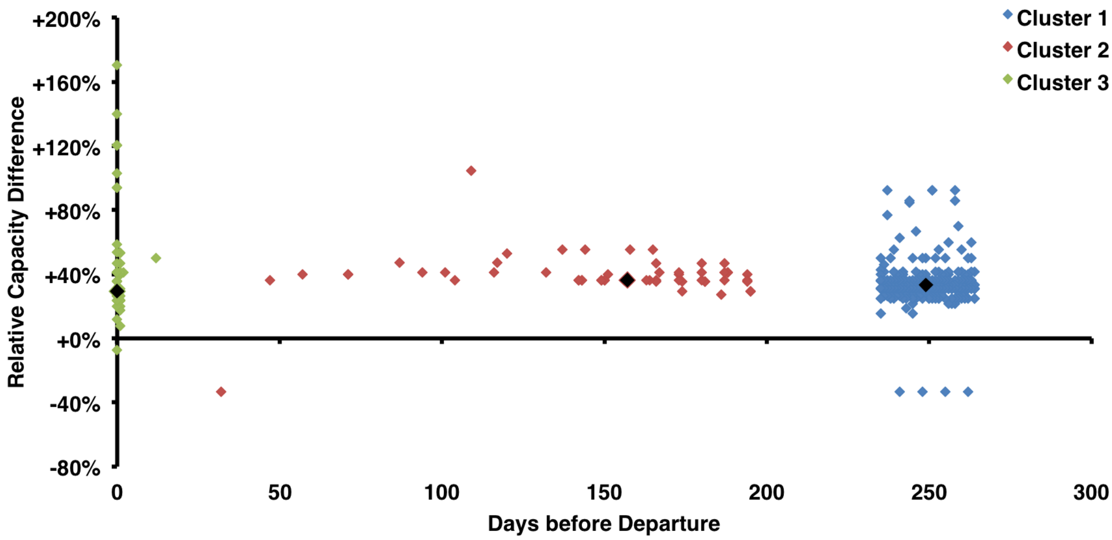

Figure B.1: Clusters at SHORT and Aircraft Size XS

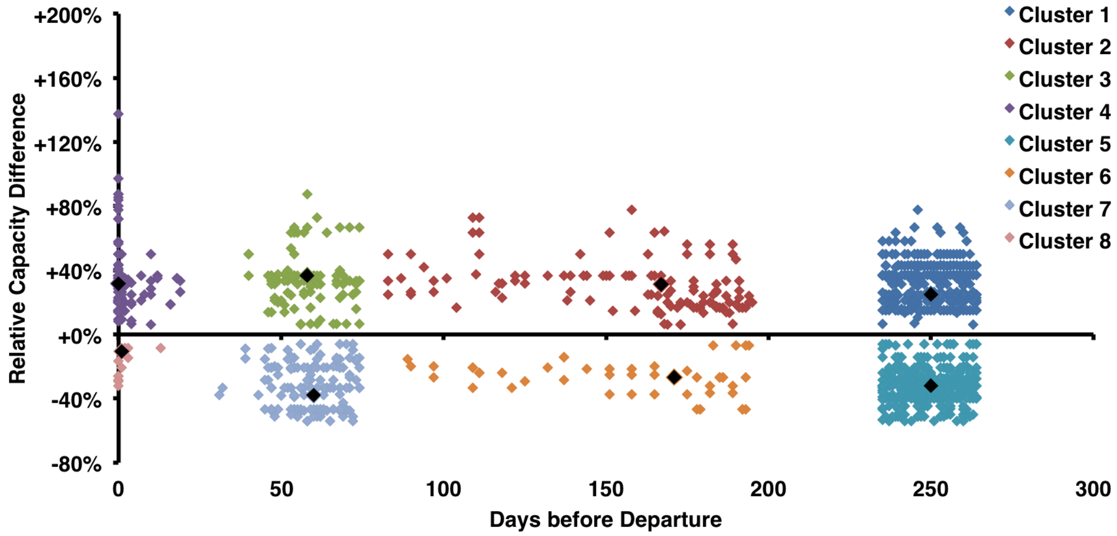

Figure B.2: Clusters at SHORT and Aircraft Size S

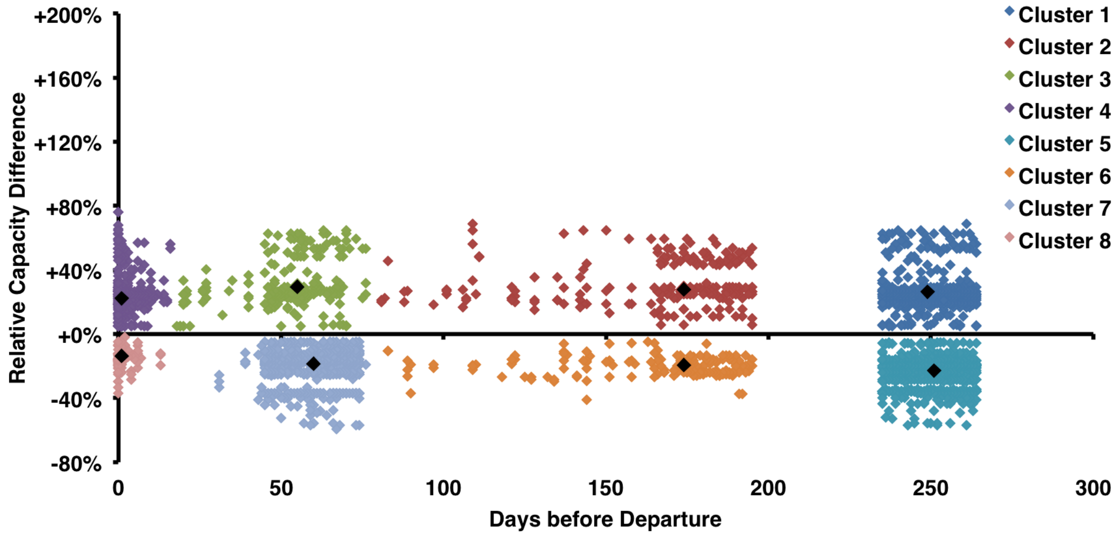

Figure B.3: Clusters at SHORT and Aircraft Size M

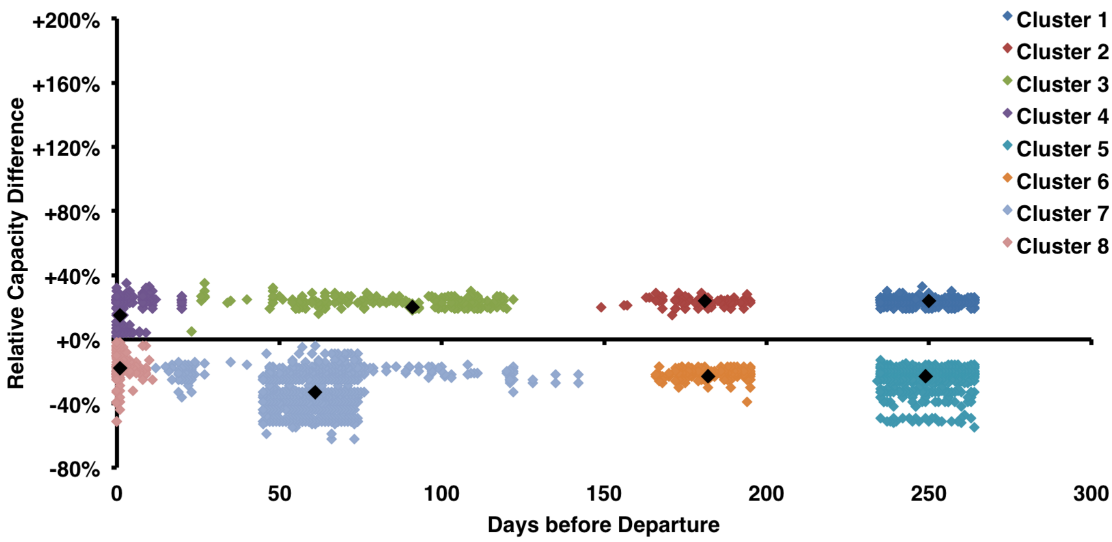

Figure B.4: Clusters at SHORT and Aircraft Size L

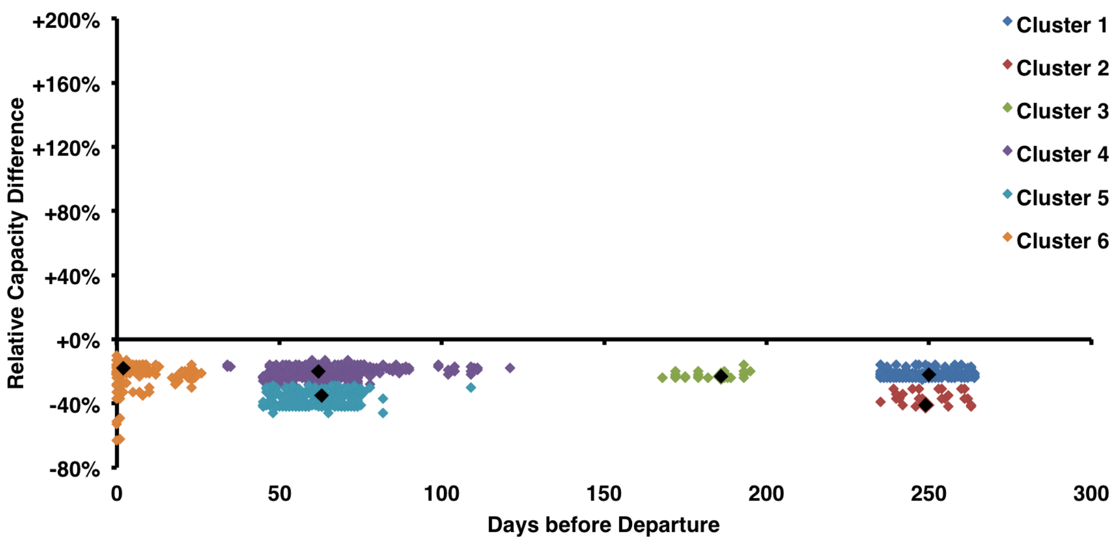

Figure B.5: Clusters at SHORT and Aircraft Size XL

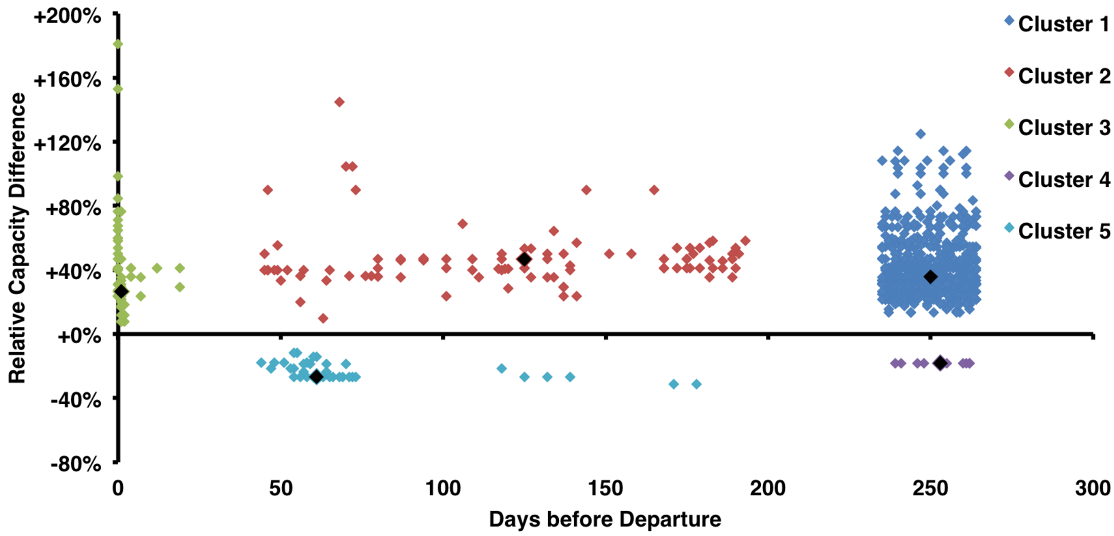

Figure B.6: Clusters at MEDIUM and Aircraft Size XS

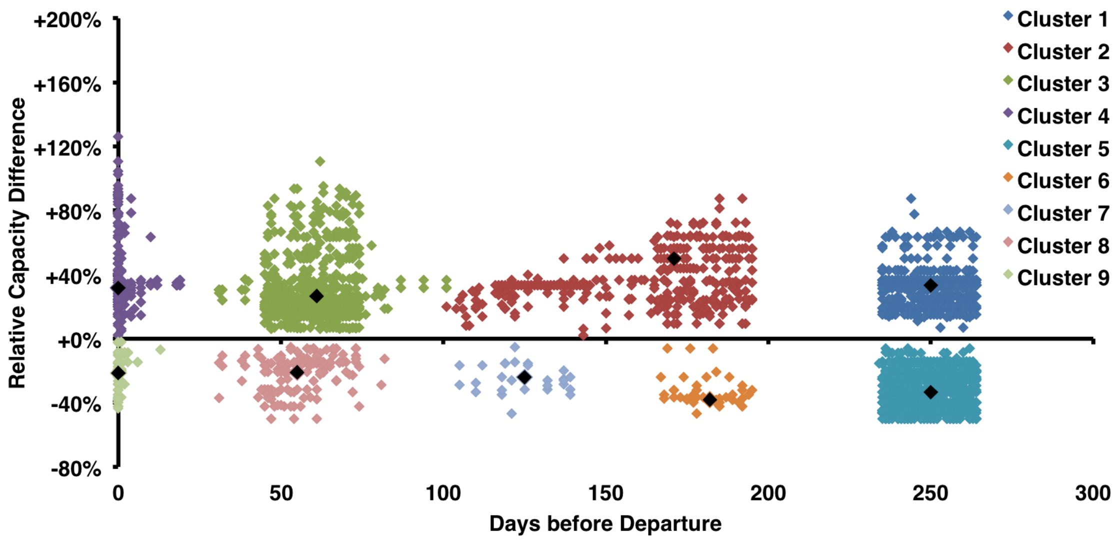

Figure B.7: Clusters at MEDIUM and Aircraft Size S

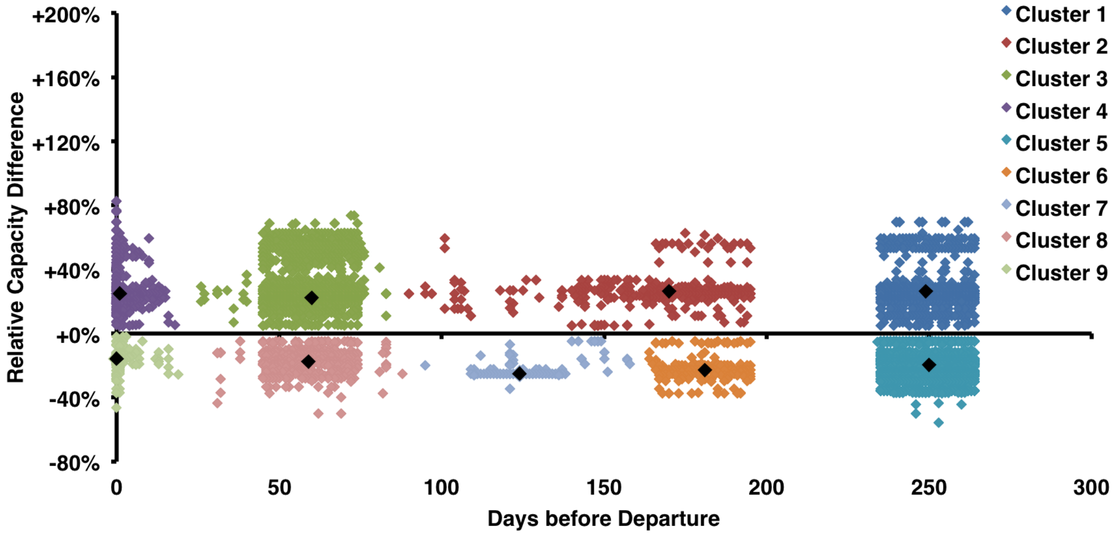

Figure B.8: Clusters at MEDIUM and Aircraft Size M

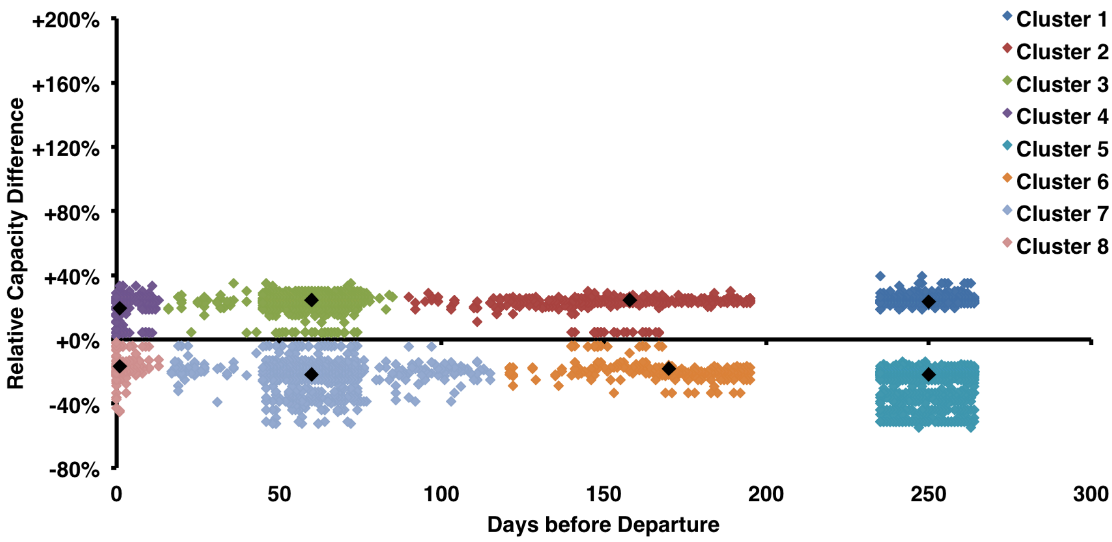

Figure B.9: Clusters at MEDIUM and Aircraft Size L

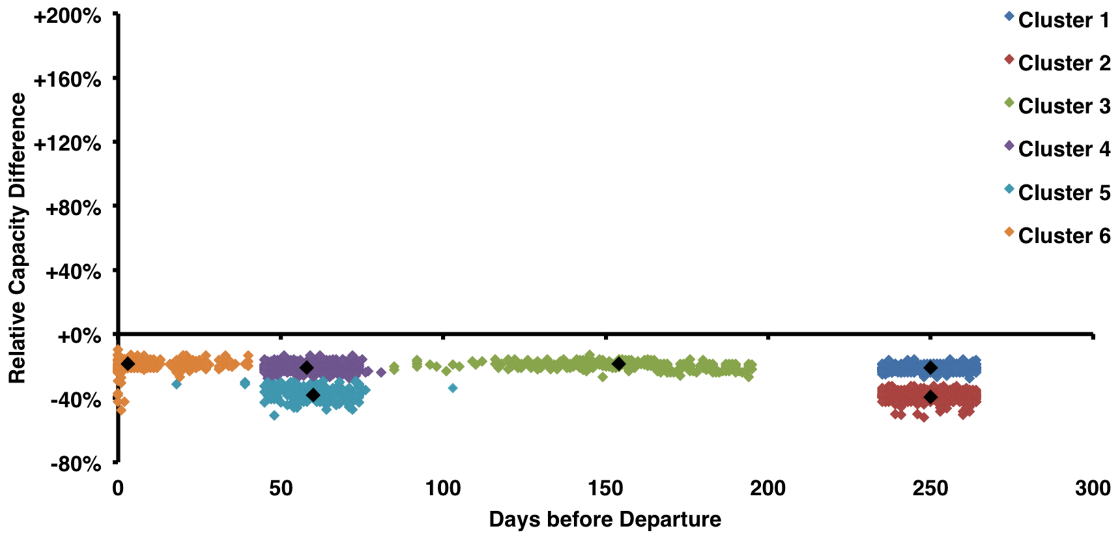

Figure B.10: Clusters at MEDIUM and Aircraft Size XL

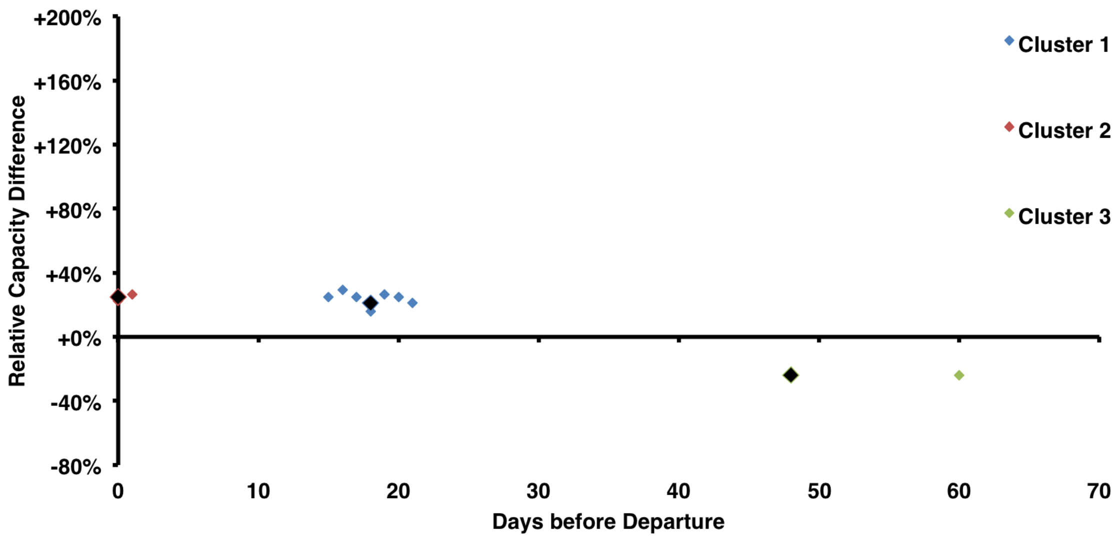

Figure B.11: Clusters at MEDIUM and Aircraft Size M

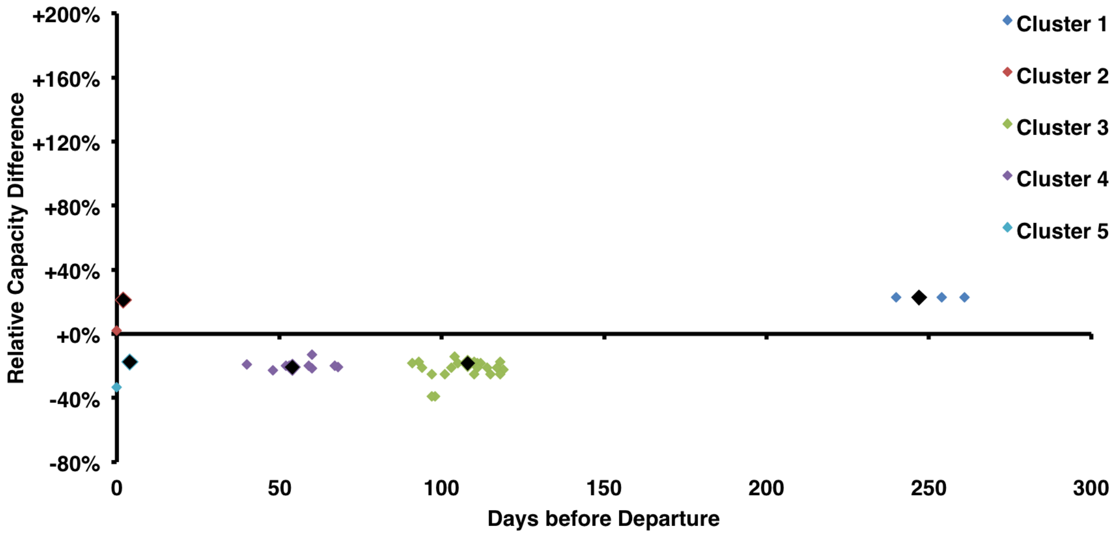

Figure B.12: Clusters at MEDIUM and Aircraft Size L

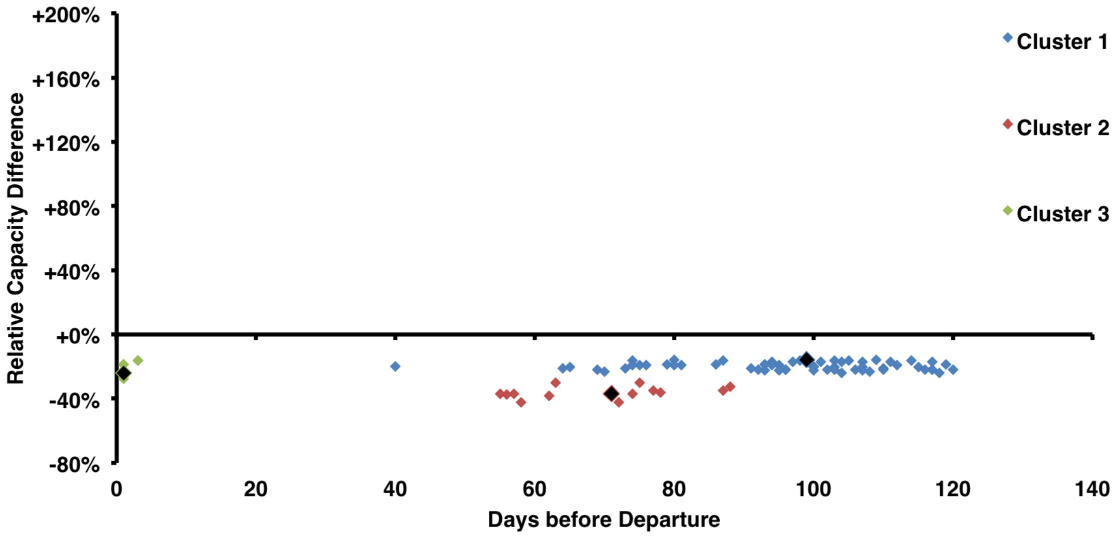

Figure B.13: Clusters at MEDIUM and Aircraft Size XL

C. Empirical Booking Data

Section 6.2 analyzed empirical booking data. Figure C.1–C.3 shows the empirical booking probability over the time horizon of 360 days for three artificially aggregated fare classes. Here and in the appendix' remainder, fare classes are also referred to as *booking classes*.

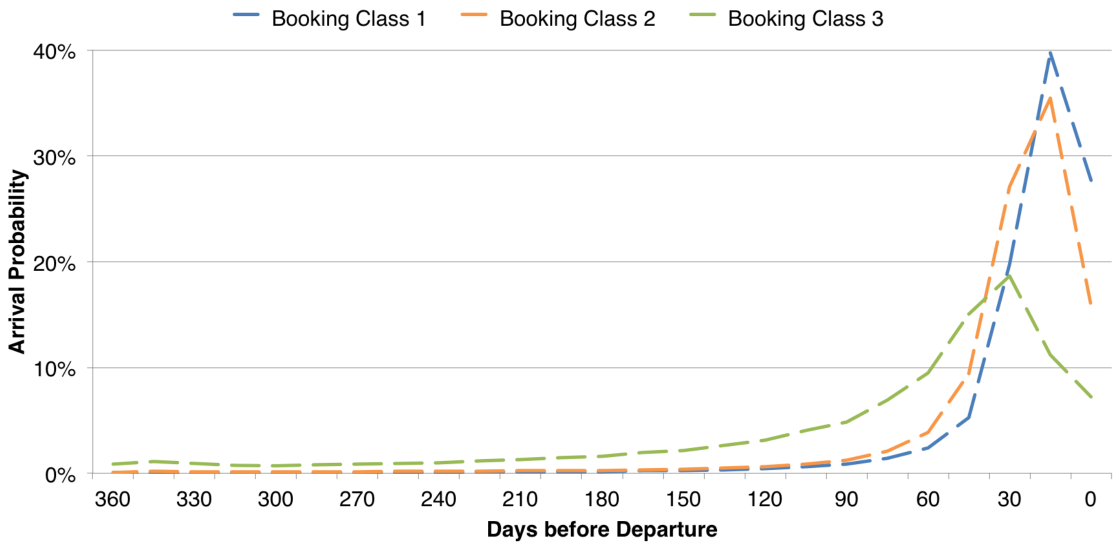

Figure C.1: Empirical Net Bookings on SHORT

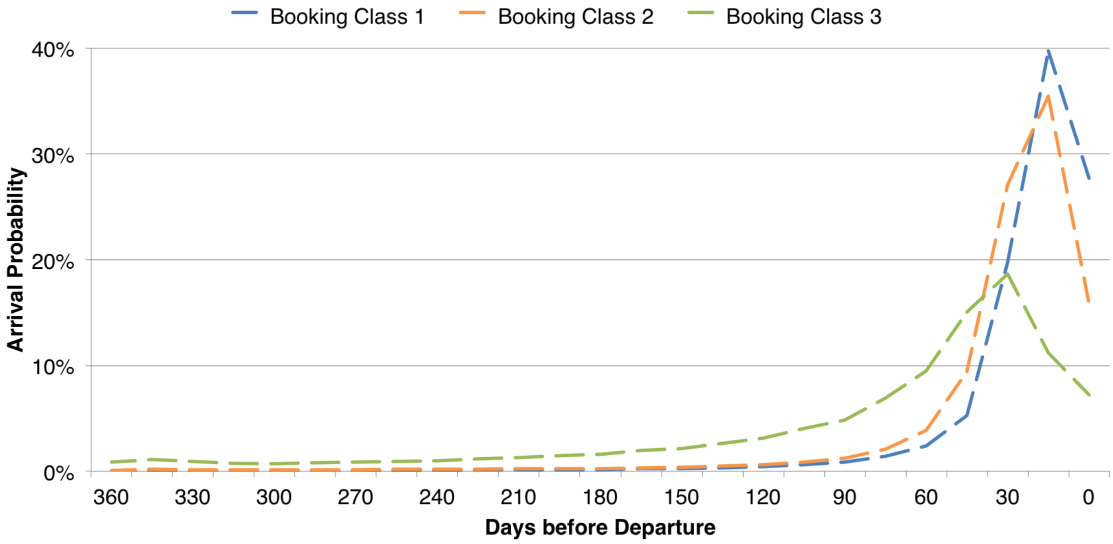

Figure C.2: Empirical Net Bookings on SHORT

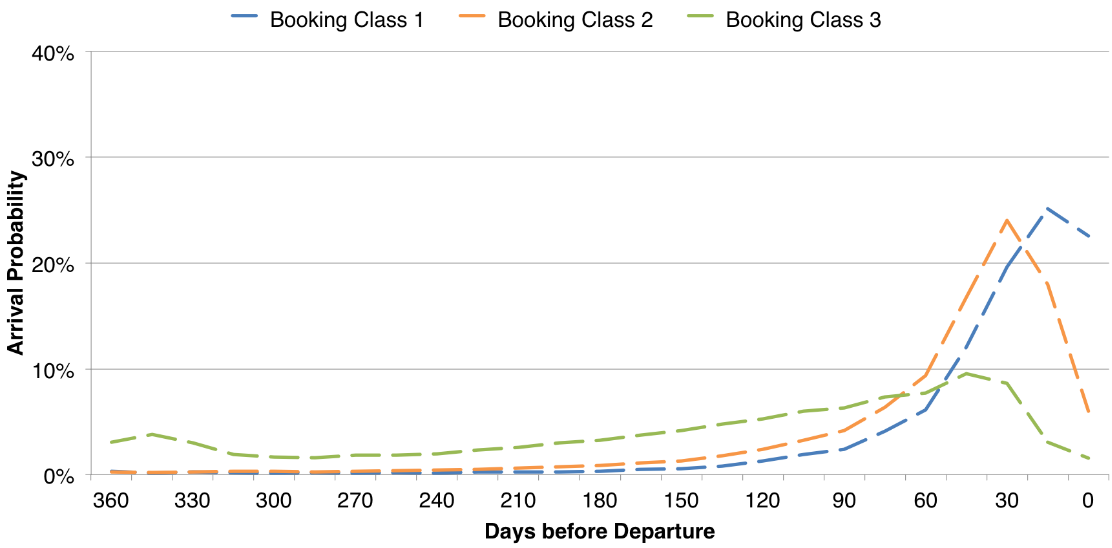

Figure C.3: Empirical Net Bookings on LONG

D. Approximated Demand Arrival

Section 6.2 suggested approximating empirical booking data as triangular distributed demand arrival. Table 6.9 presents the resulting triangular distributions' lower limit, upper limit and mode for three markets and three fare classes. These values serve as input for computational studies' demand, generated as explained in Section 8.1.1. Figure D.1–D.3 visualize the resulting demand arrival probabilities.

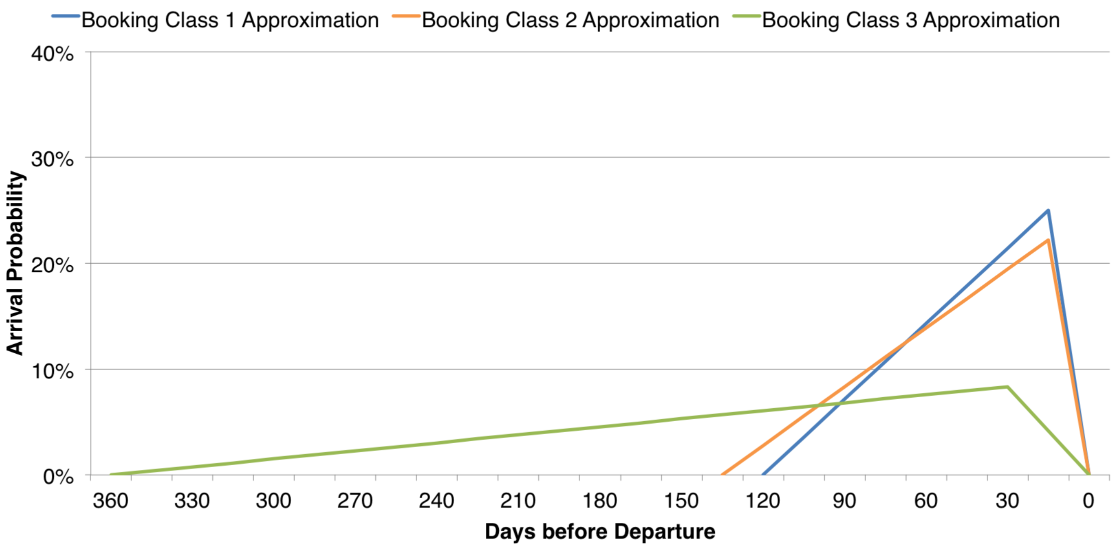

Figure D.1: Customer Arrival Approximation for SHORT Based on Booking Data

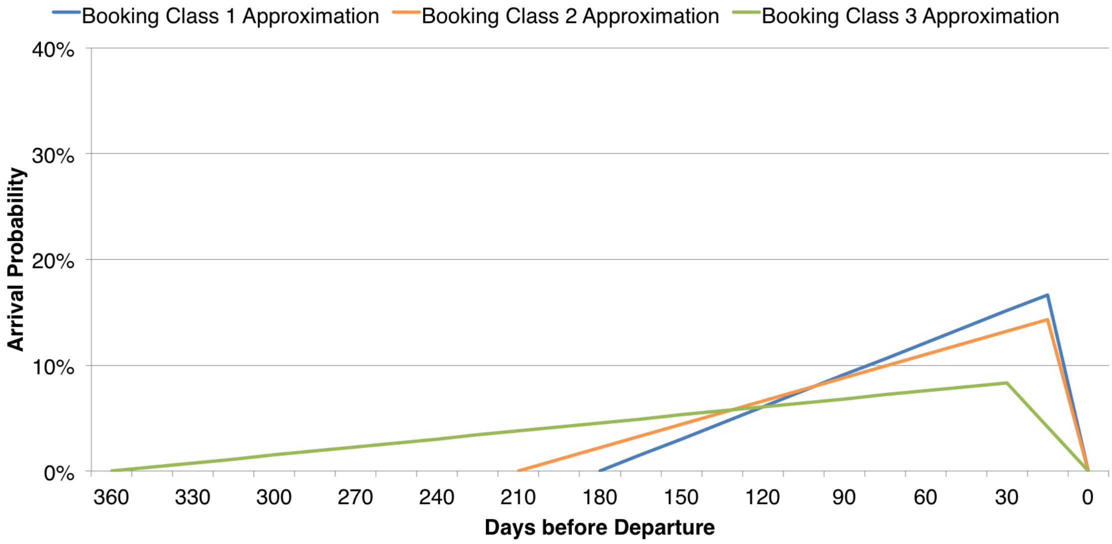

Figure D.2: Customer Arrival Approximation for MEDIUM Based on Booking Data

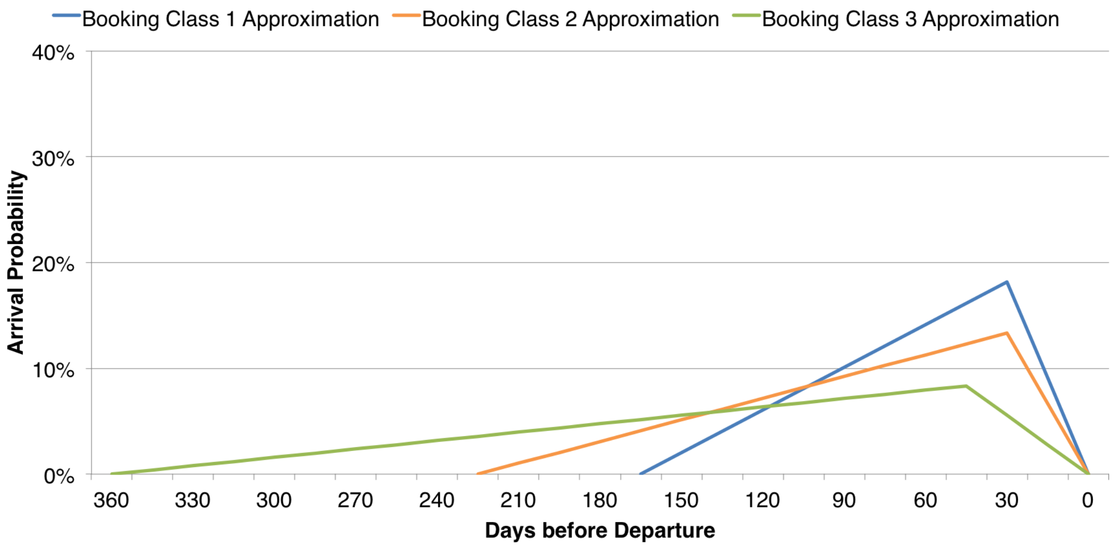

Figure D.3: Customer Arrival Approximation for LONG Based on Booking Data

E. Comparison of Empirical Booking Data and Approximated Demand Arrival

This section mutually draws empirical booking data and approximated demand arrival over time. Figure E.1–E.3 shows comparing booking data and demand arrival on SHORT, Figure E.4–E.6 on MEDIUM and Figure E.4–E.6 on LONG.

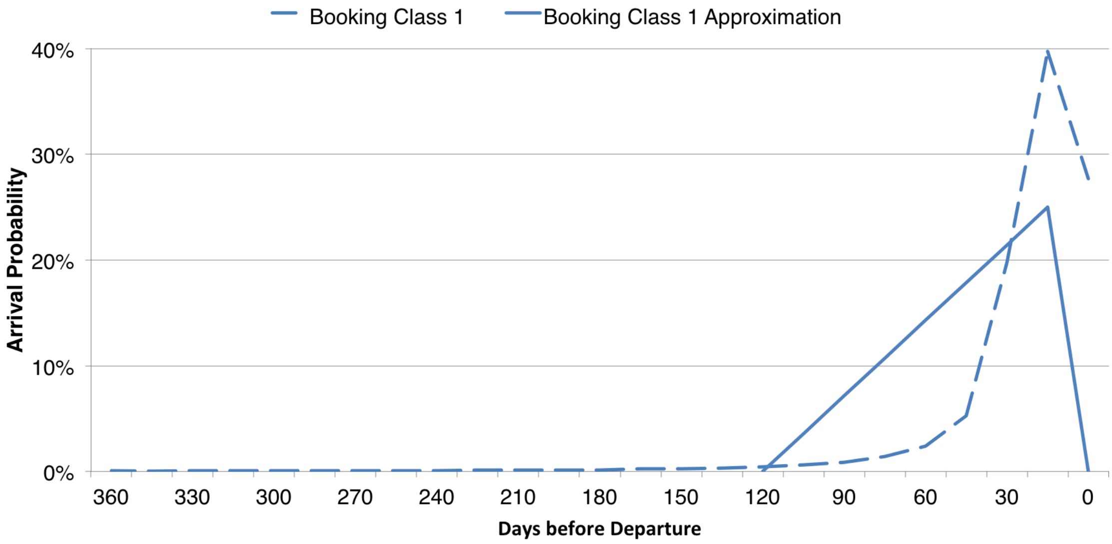

Figure E.1: Fare Class 1 Comparison of Empirical vs. Approximated Customer Arrival Probability on SHORT

Figure E.2: Fare Class 2 Comparison of Empirical vs. Approximated Customer Arrival Probability on SHORT

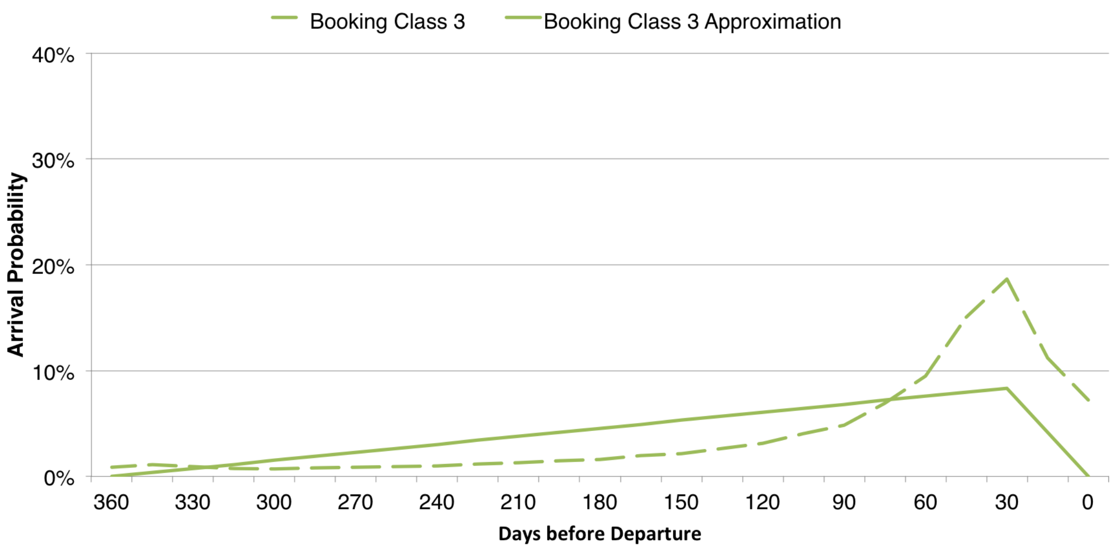

Figure E.3: Fare Class 3 Comparison of Empirical vs. Approximated Customer Arrival Probability on SHORT

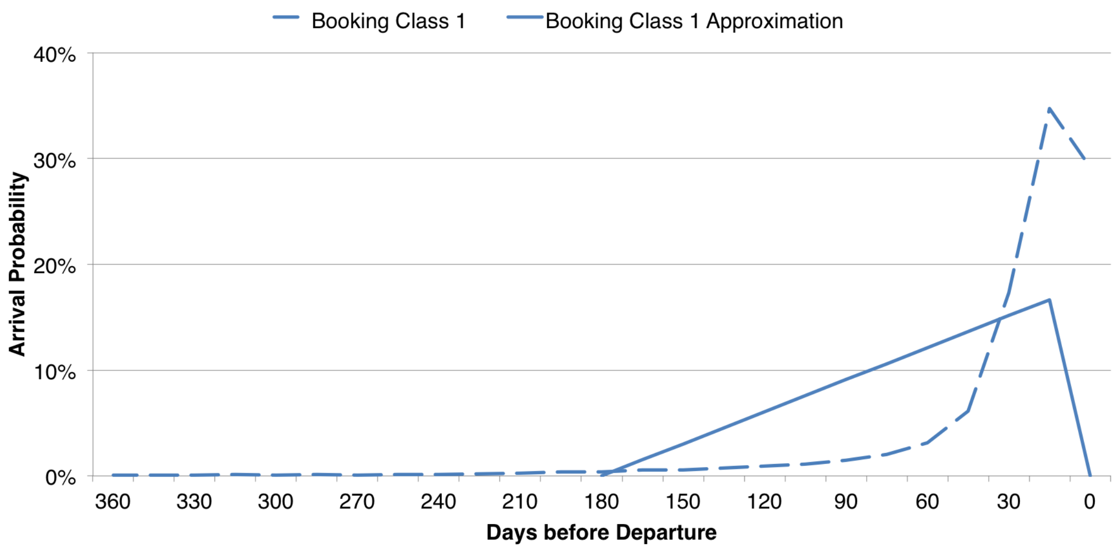

Figure E.4: Fare Class 1 Comparison of Empirical vs. Approximated Customer Arrival Probability on MEDIUM

Figure E.5: Fare Class 2 Comparison of Empirical vs. Approximated Customer Arrival Probability on MEDIUM

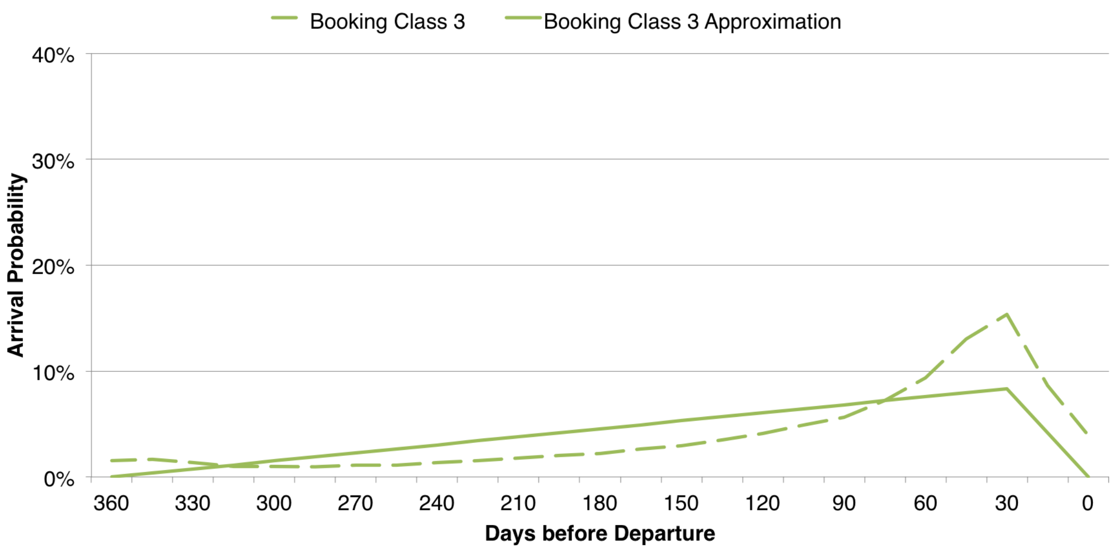

Figure E.6: Fare Class 3 Comparison of Empirical vs. Approximated Customer Arrival Probability on MEDIUM

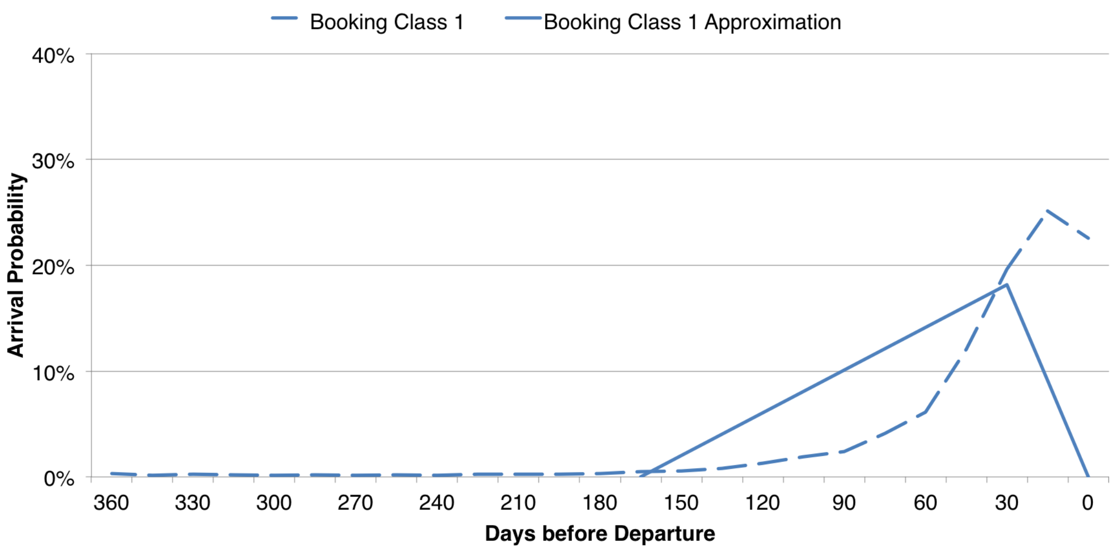

Figure E.7: Fare Class 1 Comparison of Empirical vs. Approximated Customer Arrival Probability on LONG

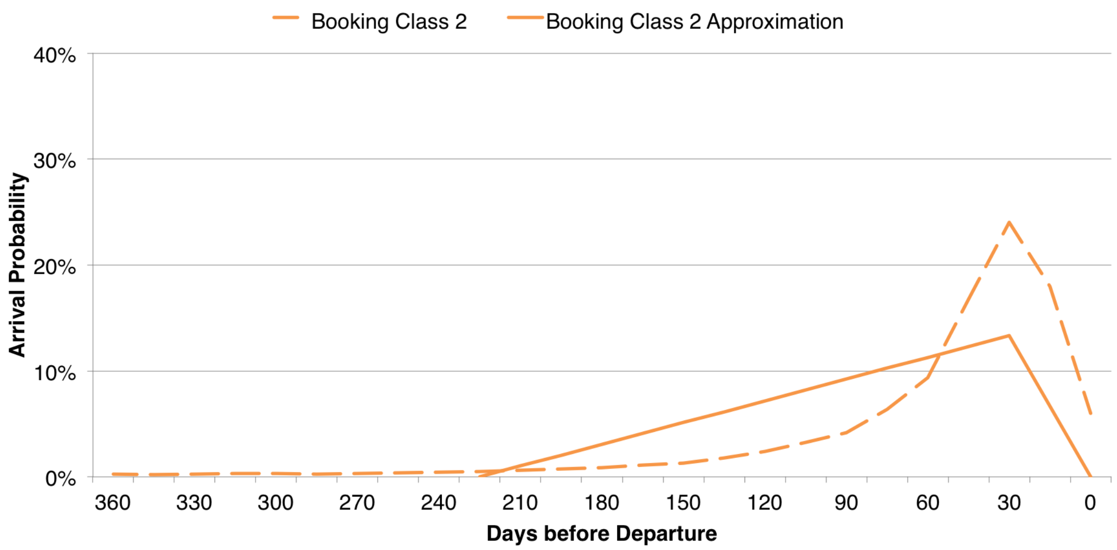

Figure E.8: Fare Class 2 Comparison of Empirical vs. Approximated Customer Arrival Probability on LONG

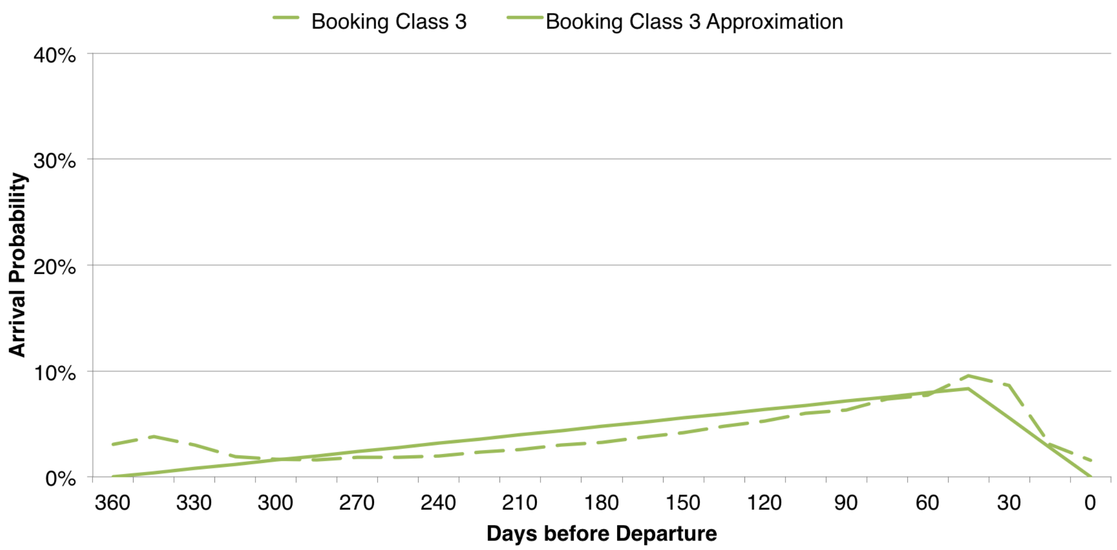

Figure E.9: Fare Class 3 Comparison of Empirical vs. Approximated Customer Arrival Probability on LONG

Acronyms

ARM airline revenue management. 3, 4, 6, 7, 9, 10, 12, 21, 32, 37, 38, 41, 121, 122, 124, 126

CFaR cash flow at risk. 26

CRM customer relationship management. 25, 26

CVaR conditional value at risk. 26

EaR earnings at risk. 26

EMSR expected marginal seat revenue. 4, 11, 22, 27, 31

EMSU expected marginal seat utility. 27

NHPP non-homogeneous Poisson process. 68, 69, 95

NYOP name-your-own-price. 33

RM revenue management. 3–7, 9–12, 14–27, 29–35, 37–42, 45, 49, 53–57, 64, 65, 67, 68, 70, 78, 85, 93, 103, 111, 121–126

TVaR tail value at risk. 26

VaR value at risk. 26