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

Simulation-Based Analysis of Forecast Performance Evaluations for Airline Revenue Management

submitted to
Faculty of Business Administration and Economics
University of Paderborn

by

Dipl.-Wirt.Inf. Catherine Cleophas

Dean of the Faculty of Business Administration and Economics:
Prof. Dr. Peter F. E. Sloane

Referees:
1.) Prof. Dr. Natalia Kliewer
2.) Prof. Dr. Leena Suhl

Paderborn, July 2009

This thesis was created thanks to a cooperation between the International Graduate School Dynamic Intelligent Systems at the University of Paderborn and the German airline Deutsche Lufthansa AG. It includes the consideration of problems occurring in applied revenue management under the aspect of academic research. The goal is to use methodological approaches to airline revenue management, demand forecast and simulation presented in the further text as well as expert knowledge and data available in the industry.

The purpose of this text is the development of a new view of forecast performance, in order to avoid some of the complications connected to evaluation of demand forecasts for revenue management. To enable this, a theoretical concept of decomposing and evaluating forecasts under the laboratory conditions provided by a simulation and using information exclusive to simulation environments is developed. To demonstrate the potential of this concept, the implementation of a simulation environment including a choice-based demand model is documented. Finally, a number of statements about the implications of forecast quality and forecast evaluation is expressed formally and tested using simulation experiments to demonstrate the use of the proposed concept.

Subject classifications: Simulation, Forecasting, Revenue Management, Yield Management, Inventory Control, Pricing, Price-Elasticity, Econometrics

Contents

List of Figures

List of Tables

Part I.

State of the Art and Research Opportunities

In this part, an introduction to the topic of this thesis as well as a background in revenue management and demand forecasting is provided. This includes existing research on overall revenue management, optimization methods, and approaches to demand forecasting. In the course of recapturing the state of the art, current challenges for demand forecasting and research opportunities with regard to forecast evaluation are pointed out.

1. Introduction

Two general approaches to pricing a product exist: Prices may be calculated in order to ensure a *break even* (to cover the costs of production) or to *maximize revenue*. In *airline revenue management*, the latter is achieved by calculating the optimal number of products to sell at a set of prices, given products that create similar production costs. For example, a ticket on a flight from Frankfurt to New York may cost 400 Euro three months before departure, but the price for the same flight under the same conditions may increase to 800 Euro when the ticket is bought three days before departure. The goal is to sell all seats available on the flight at the highest price customers are willing to pay. In the example, this is achieved by differentiating *customer segments* according to the time of booking before departure. Revenue management is often cited to be the art of "selling the right seats to the right customers at the right prices" (American Airlines (1987)).

In order to successfully apply revenue management, knowledge of customers is required. It is impossible to sell the right seats at the right price if there is no information available on the price customers are willing to pay and the conditions under which they are willing to do so. This information is provided by a *demand forecast*. It aspires to predict the amount of customers that will be willing to buy a ticket at a specific price and time.

Based on the forecasted demand, revenue can be maximized by *optimizing* inventory controls. The result should be the maximum of revenue to be earned under the given conditions. Depending on the forecast and the optimization method, these inventory controls maximize revenue per flight or for a complete network of itineraries offered by the airline.

In any case, the process described above indicates that the demand forecast has a decisive influence on the outcome of revenue management. The forecast provides the basis for any optimization and thereby influences the inventory controls. If a forecast is underestimates valuable demand, too many tickets may be sold at a reduced fare, leading to shortages in capacity when customers willing to pay high fares request to book. In the case of overestimation, too many tickets may be reserved for customers expected to pay high fares, leading to unsold tickets and empty seats at the day of departure. The consequence of both errors is falling short of maximum revenue.

Due to their importance, it is necessary to evaluate the *performance* of demand forecast methods. This task is aggravated by two complications. The definition of the term "performance" is ambivalent: Most forecasts are rated by their *accuracy*, their ability to correctly predict the future. Yet in revenue management, experience shows that sometimes a forecast that lacks accuracy can still lead to high revenue - the major indicator of revenue management success. The *revenue* resulting from the implementation of a forecast method may therefore be considered another indicator of a method's performance.

Even the evaluation of forecast accuracy alone is not trivial. The inventory controls that are computed based on the forecast information limit the bookings to be observed. As the historical data created by booking bookings often the the only information available on actual demand, they are frequently used to evaluate the forecast. The forecast's accuracy is estimated based on an indicator it influenced – the risk of *self-fulfilling prophecies* arises.

In this thesis, a *decomposition approach* as a new concept for the analysis of the evaluation of forecasts is presented. For this purpose, a *decomposed view* of revenue management and demand forecasting is introduced. Based on this concept, systematic evaluation processes are developed. A *simulation environment* is documented to demonstrate the implementation of these processes. Finally, a number of hypotheses on the evaluation of demand forecasts are formalized and tested using *simulation experiments*.

1.1. Background and Terminology

Revenue Management plays an important role in the business model of airlines all over the world. Optimizing the seat allocation in order to maximize revenue has a history that goes back to 1950. Before the deregulation of the airline industry introduced competition, revenue management mostly meant *overbooking* settings optimized to avoid empty seats.

For this purpose, forecasts estimate the number of *cancellations* and *no-shows*. Cancellations are returns of booked tickets within the booking horizon; no-shows indicate customers with booked tickets failing to show up at the day of departure. Without overbooking, the number of seats sold is equal to the capacity of the aircraft assigned to a flight. With overbooking, the number of seats sold exceeds capacity not every customer that books a ticket is expected to require a seat. The risk included in the practice of overbooking is the occurrence of *denied boardings*. Denied boardings happen when customers that booked tickets are denied as not enough seats are available due to an overestimation of cancellations and no-shows. They are connected to direct costs as alternative transport or over-night-stays have to be made available by the airline and to indirect costs as customer satisfaction decreases.

The idea of maximizing revenue by offering tickets at different prices gained in importance with the Airline Deregulation Act of 1978. With this act, the American government opened airline markets by removing control from route planning, fares and market entry. This lead to increasing competition and, as a consequence, to competitive pricing.

In order to gain advantages over competition, so-called *early bird offers* were introduced. In this form of *customer segmentation*, as in the example used in the introduction, tickets are offered at reduced fares in the beginning of the booking horizon. As time to departure decreases, fares increase, and tickets become more expensive. The underlying assumption is that those customers that are willing to pay high fares, such as business travelers, request shortly before departure. Customers that book early and at reduced prices are not expected to be willing to pay the regular fare. With the introduction of this concept, demand forecasts predicting the number of customer requests to arrive for specific classes at a given time in the booking horizon gained importance. The decision of how many tickets to sell at reduced fares and how many tickets to reserve for valuable customers is based on such predictions.

In addition to the timing of request within the booking horizon, customer segmentation can also be based on *features* or *restrictions* attached to *booking classes*. Features may include superior physical comfort in the form of seating in the *business compartment* or positive conditions such as a flexible re-booking policies. Restrictions are aimed at rendering tickets sold at reduced fares less attractive, for instance by imposing negative conditions such as *minimum stay* or *weekend stay*.

With increased computing power, the focus shifted from a *flight view* to a *network view* of revenue management. Also referred to as *origin-destination revenue management*, this approach includes the idea that customers do not actually desire to book tickets on isolated flights. Rather, they want to travel from an origin to a destination via a network of *itineraries* offered by one or more airlines. Such itineraries connect two airports and can consist of one or more *legs*. Legs are the edges of the network, described by two vertices (airports) connected by direct flights. Instead of predicting demand to arrive for single flights, a network-based demand forecast predicts requests for itineraries. Instead of maximizing revenue per flight, a network-based optimization for revenue management aims to maximizes revenue for the complete network.

In recent years, the Internet has improved market transparency. Additionally, *no-frills airlines* now offer *restriction-free* classes. These are not differentiated by either features or restrictions, but only by their fare. When customers are able to compare the offers of different airlines easily and are offered classes only differentiated by price, the *independent* or *static* model of demand collapses. Customers no longer request tickets in a specific class. Instead, they may be willing to buy a valuable class but use opportunities for booking at lower fares.

Models now need to include the idea of *correlated* and *flexible* demand. Customer requests depend on which itineraries and classes are made available. With these challenges, the importance of forecast performance (and its evaluation) increases.

1.2. Motivation and Goals

As pointed out for example in Pölt (1998) and Weatherford & Belobaba (2002), the accuracy of the demand forecast has a significant impact on the success of revenue management. The target function of most methods for optimizing revenue includes predicted demand. Depending on the number of customers expected to request tickets at a higher price in the future, current requests are accepted or denied. Examples of optimization methods have been summarized in Weatherford & Bodily (1992), McGill & Ryzin (1999), and Talluri & Van Ryzin (2004b).

Forecast accuracy can have indirect effects on long-term results of revenue management. While the information drawn from the overall results of a revenue management system offers conclusion toward the financial success, it does not provide insights toward the accuracy of the forecast.

While it is important for the performance of an airline's revenue management strategy, measuring *forecast accuracy* is not trivial. Most analyses draw conclusions from comparisons of demand forecasts to historical booking data – the result are *error measurements*. This data has been shaped by the inventory controls in place at a given point of time before departure. These controls the results of revenue optimization techniques. The optimization uses the demand forecast as input. Therefore, evaluating a forecast by comparing it to the bookings that resulted from its application is a biased approach.

Furthermore, while a forecast based on historical data offers information on expected demand, bookings have been *constrained* by limitations in capacity. If demand exceeds capacity, forecasts should exceed bookings. As a consequence, error measurements incorporate *unconstraining* of historical booking data or *constraining* of the forecast. However, the transformation of bookings to demand is a component of the forecast. The transformation of a forecast to compare it to bookings is done using the inventory controls. Both comparisons result in a bias as a part of the forecast method is used for its own evaluation.

Ideally, in order to make statements about the performance of a forecast, one would like to know the actual demand it attempted to predict. As no detailed information about each individual customer's decision processes can be attained, this is not possible in practice. The nearest approximation are customer surveys (*stated* and *revealed preference data*) as documented in Algers & Beser (2001) or *click-streams* hinted at in Nason (2007). Both are not a satisfying solution: While click-streams focus exclusively on web-based points of sale, customer surveys carry the risk of bias – customers may lie consciously or unconsciously.

Accepting that the actual demand may never be known, feasible alternatives have to be explored. One is the comparison of the results gained from the implementation of different forecast methods when all other factors remain the same (*ceteris paribus*). This would mean using first one, then the other method while keeping the same optimization system and in the exact same market environment. Again, this does not seem realistic: As the same seat cannot be sold twice, the same flight event may not be optimized twice. Two different flights either take place in geographically and therefore economically different markets or at different points of time (under circumstances influenced by season and time of day) or are bound to influence one another. Due to this, any implementation of this method in practice is likely to be flawed.

Additionally, the question of whether accurate forecasts are good forecasts is still discussed in the revenue management research community. Simulation experiments have shown that a forecast that is somewhat flawed, depending on the pricing and competition situation given, can lead to higher revenue than what is earned when demand is estimated correctly. For this reason, the further text will make a distinction between forecast performance and accuracy. The latter is regarded as an influential part of the former, but not as its synonym.

The main goal of this thesis is the development and demonstration of a decomposed approach to evaluating forecast performance. To compare the performance of methods, a clear separation between forecast, optimization, and further means of strategic manipulation located in the inventory is necessary. With a decomposition concept, the consequences of changes in any component can be analyzed separately. At the same time, the interpretation of overall system performance is still feasible.

For the demonstration of the decomposed concept, a simulation environment offers many advantages. It provides a stable and fully controllable framework. Two simulation scenarios may confront a system with the exact same customer model. Under such conditions, solely methodical factors can cause differences in bookings and revenue. The applicability of the results of such a simulation can be secured by including a sufficiently complex customer model as well as a realistic network of flights. The results offer insights not just concerning the financial success and the accuracy of different forecasting methods. In addition, conclusions toward the connection of the two and their evaluation may be drawn.

In order to achieve the defined goal and to demonstrate the advantages of a decomposed concept of evaluation in a simulation environment, a number of tasks are defined:

- categorization of existing forecast methods;
- characterization of existing approaches to forecast evaluation;

- conceptualization of a decomposition of revenue management systems and demand forecasts;
- formulation of processes to separately evaluate the components of revenue management systems;
- implementation of a simulation environment to apply the concept;
- formalization of statements on forecast performance evaluation to base simulation experiments on;
- analysis of the results of experiments conducted in the simulation environment.

The next section describes the form in which the approach to these tasks was documented in the following chapters of this thesis. It shortly summarizes the content of each of the three parts.

1.3. Outline

This thesis is divided into three parts. In the first part, an introduction to revenue management and the state of the art regarding forecasts and their evaluation is provided. A research gap is identified and the research opportunities derived from it are listed. In the second part, the solution approach is outlined. For this purpose, the concept of the decomposed evaluation of demand forecasts and the implementation of a simulation environment are documented. In the third part, simulation experiments and the findings resulting from them are summarized. Conclusions toward the evaluation of forecast performance and future research are drawn.

Part One: As the motivation and the findings of this work are based on existing methods of airline revenue management, first an overview of published research in this field is presented in Chapter 2. This chapter starts with a general introduction to the history, the motivation, and the terminology of revenue management. Approaches to maximizing revenue according to linear optimization, dynamic programming, and heuristic methods are presented next. Finally, recent challenges such as the increased overview of the market place provided by Internet search engines and the advent of no-frills airlines are listed.

These developments as well as its general role in maximizing revenue are used to underline the importance of accurate demand forecasting.

Approaches to demand forecasting are summarized in Chapter 3. They are categorized by aspects of demand including overall volume, unconstraining of bookings and flexible demand behavior. The multitude of mathematical methods available to predict the various aspects demand to come becomes apparent in this chapter.

Approaches to quantifying forecast performance are presented in Chapter 4. Furthermore, examples of applied forecast performance measurements are provided.

In Chapter 5, research opportunities with regard to demand forecast performance measurement are pointed out. These research opportunities serve as motivation and provide aims for the further work.

Part Two: A new approach to decomposing and evaluating demand forecasts is presented in Chapter 6. This chapter describes how a simulation system can be used to isolate different components and aspects of the system. These parts can then be evaluated separately. Documented are concepts for analyzing the complete revenue management system, the isolated forecast component and the decomposed aspects of demand volume, unconstraining and behavior.

To implement this concept of decomposition, a simulation environment is documented in Chapter 7. General concepts of simulation control are introduced and the implementation of a supply and demand model is described in detail. In addition, the precise revenue management methods included are explained. Finally, the market implementations prepared for simulation experiments are presented.

Part Three: Based on the simulation environment introduced in Chapter 7, simulation experiments have been conducted to provide examples for the simulation based analysis of forecast performance. The hypotheses that were evaluated empirically are formally listed in Chapter 8 together with the conclusions that can be drawn from the results of the experiments. The aspects under which revenue management in general and demand forecasting in particular were analyzed include the long-term effects of methods, the idea of psychic forecasts used in the simulation, the key performance indicators traditionally used, the consequences of demand uncertainty and first approaches to evaluating price-sensitivity for revenue management.

Finally, Section 9 provides a summary of steps taken and the insights gained. In addition, an outlook to potential further research is offered.

2. Existing Research on Airline Revenue Management

A great body of published research has been devoted to airline revenue management. Literature including overviews of existing research is listed in the first part of this section. Next, some approaches to maximizing revenue given available knowledge of demand are described. Finally, recent challenges that revenue management experts are confronted with are documented.

2.1. Available Overview Literature

A first introduction to revenue management and its importance for business success is provided by Cross (1997). The most important research with regard to demand forecasting for revenue management up to 1999 has been outlined in McGill & Ryzin (1999). More information, especially on general theory, is provided by Pak & Piersma (2002) and Talluri & Van Ryzin (2004b). One of the latest overview articles, Chiang et al. (2007), mostly concentrates the application of revenue management. Weatherford & Bodily (1992), Bitran & Caldentey (2003), and Boyd & Bilegan (2003) focus on other areas of revenue management such as the development of a typology, pricing, and the implications of e-commerce.

Some detail on how to model the restrictions and the objective of revenue management is offered in Wang (1982). An account of the development of a revenue management system at American Airlines is provided by Smith et al. (1992). Revenue management in the broader context of operations research problems in the air transport industry is presented in Barnhart et al. (2003). Methods for approaching the subject with differing degrees of sophistication are presented in Vinod (2006).

2.2. State of the Art of Optimization

The research listed here is concerned both with mathematically optimal solutions and heuristic approaches. Furthermore, most models used in revenue management rely on simplifying assumptions – these are described with the references.

Flight Optimization: In Littlewood (1972), a first approach to optimally allocating availabilities based on demand predicted per flight and booking class is introduced. Its optimality given the condition of static demand was proven by Mayer (1976). Similar topics are considered in Bhatia & Parekh (1973), Richter (1982), Gerchak et al. (1985), Alstrup et al. (1986), Kraft et al. (1986), Pratte (1986), Wollmer (1986a), Gerchak & Parlar (1987), Pfeifer (1989), Wong (1990), Stone & Diamond (1992), Wollmer (1992), M. Li (1997) and M. Z. F. Li & Oum (1998). The expected marginal seat revenue (EMSR) approach to allocating seats, which is widely used in practice, was introduced in P. Belobaba (1987a), P. Belobaba (1987b), and P. Belobaba (1989). The underlying assumption of low-fare demand arriving before high-fare demand is discussed with regard to optimality in Titze & Griesshaber (1983). The concept of nesting booking classes to ensure the availability of valuable classes is outlined in W. Swan (1993c).

In M. Z. F. Li & Oum (2002), optimality conditions of models for flight-based revenue maximization are discussed. One of the first considerations that demand for several classes may be stochastically dependent is presented in Brumelle et al. (1990). Brumelle & McGill (1993) maximizes revenue when demand is dependent on the fare; prices as decision variables are also considered in Weatherford (1997a,b, 2001). The idea that fares may not be monotonically increasing within a fare structure is included in Robinson (1995). A consideration of how fare structures may influence demand and a concept of encouraging sell-up between classes are described in Botimer & Belobaba (1999).

The task of maximizing revenue is modeled as a knapsack problem in Young & van Slyke (1994) and Young & van Slyke (2000). Cancellations and no-shows are included in the model used in Subramanian et al. (1999). Ryzin & McGill (2000) presents an adaptive algorithm aimed at eliminating the strong reliance on demand forecasts. A similar goal is pursued by Ball & Queyranne (2006) via the use of competitive analysis of online algorithms. A situation in which only limited demand information is available is discussed in Lan & Gao (2007). All these papers consider leg-based optimization only.

Flight-based models that consider stochastic uncertainty with no strict assumptions about the timing of demand arrival are often optimized using *dynamic pricing*. This dates back to Kincaid & Darling (1963). For multiple classes and customers that do not exhibit a strict arrival pattern, this is presented in T. C. Lee & Hersh (1993) and Zhao & Zheng (2000). Brumelle & Walczak (1997) includes the possibility of customers arriving in batches, Walczak & Brumelle (2007) builds up on this. A consideration of no-shows and diversion between flights is given in Zhao & Zheng (1998). In Lautenbacher & Stidham (1999), the maximization of revenue with dynamic pricing given dependent demand is discussed. In Cooper & Mello (2002), the limitation of certain stochastic programs to flight-based problems is pointed out.

Network Optimization: While flight-based optimization becomes more sophisticated and relies less on simplifying assumptions, the maximization of revenue over complete networks has come into focus. Models that include several segments being offered on one flight may be regarded as the preparation of this. Ladany & Bedi (1977), Hersh & Ladany (1978), Buhr (1982), Wang (1983), and Simpson (1985) discuss this possibility. Smith & Penn (1988) and Vinod (1995) are also concerned with the optimization of multiple segments.

The first approaches to modeling the passenger flow over a network have been documented in Glover et al. (1982), D'Sylva (1982), Dror et al. (1988), Curry (1990), Vinod & Ratliff (1990); Vinod (1990), Phillips et al. (1991), Wong et al. (1993) and Talluri (1994). As one of the first, Williamson (1992) uses a simulation system to demonstrate the impact of different network and leg-based revenue optimization methods. A linear programming approach to network optimization is outlined in Wollmer (1986b), while a non-linear programming approach is considered in Vinod (1991). In Garcia-Diaz & Kuyumcu (1997), a cutting plane approach is used to solve the network problem. While Williamson (1988) and Vinod (1989) introduce the concept of virtual nesting, Williamson (1992) considers the decision of whether to accept or deny a customer request as the comparison to a bid price.

In Talluri (2001), the idea that passengers may be routed in a way that creates balances high- and low-fare demand over itineraries is introduced. The idea that different itineraries might compete for demand is expanded in Coldren & Koppelman (2005).

Gallego & Ryzin (1997) expanded the dynamic pricing method in order to consider networks rather than legs. In Gallego & Hu (2007), dynamic pricing is applied with special regard to recent challenges such as restriction-free classes and diversion between competing flights.

Cooper & Mello (2003) propose a combination of mathematical programming methods and heuristics to make network revenue management applicable to the practice. A concept to avoid the distinction between forecast and optimization is offered in Chen et al. (2003): Value functions for network revenue management are estimated via statistical learning. Möller et al. (2004) attempts to maximize network revenue via linear programming using stochastic scenarios rather than straight-forward demand forecasts as input. The introduction of a network-based revenue management system in practice is documented in Swift (2002) and Cutshall & Weisbrodt (2006).

In de Boer (2003), among other concepts, the use of simulation to maximize network revenue is proposed; this prepares the way for Bertsimas & de Boer (2005). Ryzin & Vulcano (2006) proposes a simulation-based approach to optimization that extends the work of Bertsimas & de Boer (2005) to include a continuous model of capacity and demand. A concepts for the joint optimization of inventory controls and fleet assignment is introduced in Frank et al. (2006) and evaluated with the help of a simulation system.

2.3. Appraisals of Recent Challenges

Some developments concerning the market place and distribution channels have complicated revenue management for airlines during the last decades. As customers are able to effortlessly compare fares via websites such as opodo.com, travelocity.com, or low-fares.com, they have become more flexible and price-oriented. Considerations of the impact of increased market transparency and the consequences for consumers can be found in Nason (2007). The entry of so-called *no-frills airlines* into the market place helped to further stress the influences of prices and put traditional airlines under pressure. The business model of such airlines is described in some detail in Calder (2006).

General introductions to this situation and the consequences for revenue management can be found in Boyd & Kallesen (2004), Dunleavy & Westermann (2005) and Ratliff & Vinod (2005). In Boyd (2004), special attention is paid to the impact of Internet sales. The rest of this section is devoted to three major challenges: Customers basing their booking decision on which classes are available (*dependent demand*), customers considering *competition* offers, and *strategic customers* delaying their bookings as they hope for reduced fare offers.

Dependent Demand: The consequences of ignoring flexible customers and optimizing revenue under the assumption of static demand are analyzed in Cooper et al. (2006): When customers purchase the lowest available fare, static forecasts become self-fulfilling prophecies. Availabilities assigned to reduced-fare classes are used up by flexible customers, creating bookings – the basis for future forecasts. As more and more demand is predicted for tickets offered at a low fare and the optimization algorithm allocates more availability to low fares, revenue suffers the *spiral-down effect*.

As the perils of not reacting to changes in customer behavior are established, much research has been devoted to new approaches to revenue management. Special attention was directed toward *restriction-free environments*. This term describes the offer of booking classes differentiated only by price. Restrictions such as weekend-stays or minimum-stays are not applied. This strategy is frequently implemented by no-frills airlines and matched by traditional airlines striving for competitiveness as described in Dunleavy & Westermann (2005). Consequences of restriction-free environments are demonstrated using the *Passenger Origin Destination Simulator* (*PODS*, see also C. Hopperstad (2000), Gorin (2000), Gorin (2004) and Reyes (2006) for introductions) in Cusano (2003) and P. Belobaba & Gorin (2004).

In P. Belobaba & Hopperstad (2004), the idea of the spiral-down effect is summarized and a definition of *sell-up* is given: The term describes the probability of a customer buying his ticket for a specific fare, given he would have bought the lowest fare had it been available. The presentation introduces a modified EMSRb algorithm based on sell-up estimates as well as a dynamic program including knowledge on sell-up.

A different view of price-sensitive customers, called *buy-down*, is introduced in Ozdaryal & Saranathan (2004). The term describes the phenomenon of customers who would previously have bought one specific class now buingy a cheaper class, if it is available. As a counter agent, *inventory fences* limiting the availability of cheap classes are proposed.

More comparisons of low-cost environments versus the traditional network models of revenue management can be found in Weber & Thiel (2004). As ways of dealing with price-elasticities the authors suggest neural networks.

A summary of methods for restriction-free environments is provided in Cléaz-Savoyen (2005). The author presents and evaluates two network-optimization methods based on the knowledge of sell-up probabilities: (*Q-Forecasting* as introduced in B. P. Hopperstad C.H. (2004) and *Fare Adjustment* as introduced in Fiig & Isler (2004)). Detailed results for an evaluation of fare adjustment using PODS are also presented in Lua (2007a,b,c) as well as in Kayser (2007).

The consequences of restriction-free environments and a more transparent market place can also be phrased as a change in assumptions. Traditional revenue management regarded demand as *static* or *independent*: Customers buy tickets in one booking class, if that class was not offered, they do not buy at all. The new view includes *dependent* demand: Customers choices depend on the classes offered.

P. Belobaba (1987b) touches on this idea when describing possibilities for incorporation of passenger shifts in the EMSR model. In Pfeifer (1989), decision rules are implemented based on the probability that customers are "shoppers", basing their booking decision on fares. Among the first papers to consider a change in assumptions is also Brumelle et al. (1990). The author expands Littlewood's Rule to include dependent demand.

An approach that is not just limited to the airline industry is documented in Bodily & Weatherford (1995). In this context, buy-down is termed *diversion*. Heuristic decision rules are suggested for multiple nested classes in order to avoid as much diversion as possible. P. Belobaba & Weatherford (1996) also considers diversion and evaluates decision rules to minimize it. In addition, the authors present a new heuristic. Zhao &

Zheng (2001) introduces a dynamic threshold policy in order to limit diversion. The authors differentiate between *static* and *dynamic* policies: Whereas traditional approaches to maximizing revenue such as Littlewood's rule provided static policies, dynamic rules allow for a change customer behavior and according changes in the policy. Furthermore, the model includes the restriction that fare classes, once they are closed, can not be re-opened.

A different view of dependent customer behavior is taken in Gallego & Phillips (2004) and Gallego et al. (2008). Used in order to develop *flexible* or *callable* products, dependent demand is regarded as an opportunity rather than a risk.

Competition: As customers have access to websites summarizing the offers of several airlines, they can comparison-shop much easier. This makes it necessary for airlines to consider competition to a greater degree.

In Fischer & Kamerschen (2003), an analysis of demand aggregated to airport-pairs with regard to the market situation in terms of competition is presented. The authors employed the *Rosse-Panzar test* in order to measure the consequences of competition on airline markets.

The conclusions of this examination are not so much revenue management recommendations as they are general statements about economic implications: The more intense competition is, the lower the average fare. It is stated that with regard to the data used, airline competition is not as perfect as often supposed.

An approach that focuses more on *game theory aspects* of airline revenue management under competition is documented in Netessine & Shumsky (2005). Considered are both *vertical* (different airlines compete on different legs of a multi-leg itinerary) and *horizontal* (different airlines compete on the same leg) competition. Conditions for a *Nash equilibrium* are provided as desirable characteristics of the demand distribution. However, no method of estimating the real consequences on a given market is provided.

A game theoretical view of revenue management is also taken in Gallego et al. (2006). Without special regard for the airline industry, revenue management under competition is regarded as a *sequential* and as a *repeated* game. Again, conditions for a Nash equilibrium are outlined and the advantages of competitors cooperating are pointed out. That condition is very difficult to realize under the legal circumstances of airline revenue management practice.

Another example of the inclusion of competitor information is provided by Walczak (2005). However, while knowledge of the influence of competition on demand is included in a dynamic program according to this presentation, this knowledge is regarded as given. Its estimation is not part of the research.

In Coldren & Koppelman (2005), demand shares for travel along itineraries of competing airlines are predicted. The model includes departure day, brand, and service as factors. Both a multinomial logit model and variations of the nested model are considered in order to analyze the influence of these factors. The price differences between the different airlines' offers are not considered.

In the course of research conducted with PODS, some concepts for including knowledge of competition in the estimation of sell-up rates have been outlined. These and the resulting findings have been presented at several PODS summits. Recent examples can be found inCarrier (2003), Guo (2007), and C. Hopperstad (2007). In addition, research concerned with fare adjustment as for instance Kayser (2007) underlines that the consideration of customer reactions to prices becomes even more important in markets with intense competition. Finally, matching the lowest competitor price is presented as an approach including competition in revenue management that avoids actually forecasting the customer reaction in Lua (2007a,b,c). In all these publications, knowledge of competitor prices is regarded as given.

It can be concluded that much research that considers the effects of competition does so only with regard to its general consequences. Forecasting methods that do consider a wider range of influence factors often offer the possibility of including competition prices or the existence of competition in the model.

In Gallego & Hu (2007), demand models that consider more characteristics of the product as decision factors for customers are examined with special regard to competition. *Dynamic pricing* is presented to include demand forecasts that incorporate such a customer model in the optimization. With consideration of the single-leg model, the authors extend both on Netessine & Shumsky (2005) and Gallego et al. (2006). With regard to the airline industry, cooperation is excluded. An *open-loop Nash equilibrium* is presented. Changes in customer demand as caused by competition are ascribed to the predictions of a customer choice model. A general introduction to dynamic pricing and the opportunities it offers is provided in Westermann (2006).

Strategic Customers: Another phenomenon along with the idea of flexible, informed, and price-sensitive customers is that of *strategic* as opposed to *myopic* customers as described by Talluri & Van Ryzin (2004b). This idea implies that customers learn from their previous booking experiences and use their knowledge for future decisions. For instance, a customer observed an offer for a reduced-fare ticket coming up shortly before departure might delay his next booking, hoping for cheap tickets to be offered late in the booking horizon again. Such strategic behavior needs to be recognized and considered in inventory policies in order to avoid revenue losses.

Such demand behavior is described in Anderson & Wilson (2003); the authors point out the impact of customers implementing strategies to counteract airline revenue management strategies. While not offering a concept on how to quantify or react to strategic customers, the article highlights their possible importance to future revenue management research. Xu & Hopp (2005b) considers strategic behavior in the retail industry as well as in the airline industry and compares empirical observations with regard to changes in price-elasticity. In Cho et al. (2007), *fare track systems* designed to systematically mine airline fare data to enable customers to minimize their expenses are analyzed. Potential benefits of such systems both for customers and airlines are listed.

Levin et al. (2006) describes a dynamic pricing strategy based on the assumption that the market is monopolistic and customers may delay their bookings strategically. Knowledge on customers' strategies is regarded as available. Once more the dire consequences of ignoring strategic customers are pointed out. Another model of dynamic pricing under similar conditions is presented in Su (2007). Wilson et al. (2006) considers a demand model with customers reacting flexibly to offers by buying the lowest available fare and strategically delaying their bookings. An approach to finding optimal booking limits under these conditions is offered.

3. Demand Forecasting for Revenue Management

Forecasting demand is one of the fundamental elements of Revenue Management. Knowledge of the amount and the qualities of demand for seats on flights is crucial for a successful optimization model. Research such as presented in Pölt (1998) shows that forecast accuracy result has a positive impact on revenue. This has also been confirmed in Weatherford & Belobaba (2002). An analysis of business risk presented in Lancaster (2003) shows that much of this risk stems from faulty forecasts.

A succinct characterization of the *demand model* needed for revenue management is offered by van Ryzin (2005):

> [...] it is really the entire system for estimating demand and market response – the data sources, the information technology for collecting and storing data, the various statistical estimations models and algorithms used to process and analyze these data and the infrastructure for deploying model outputs – in short, everything that is required to turn raw data into actionable market information. This is normally called the 'forecasting system' in traditional RM parlance, though forecasting is merely one of its many functions.

Developments such as described in Section 2.3 make forecasting for revenue management even more complex. At the same time, a more automated approach to revenue management, growing data management opportunities, and the resulting need for higher quality forecasts increase the interest in forecast methods (see for example Zaki (2000), Chiang et al. (2007), or Nason (2007)).

In order to provide an overview of forecasting for revenue maximization, a categorization is introduced. Research is listed according to the characteristics of demand that it focuses on. Three characteristics of demand are discussed:

- *Demand volume*: The absolute amount of passengers that request tickets can be split up into the arrival process of demand throughout the booking horizon and the development of overall demand over several departures due to trends and seasonality.
- *Unconstraining*: The transformation of observed sales data to deduce actual demand.
- *Demand behavior*: The reaction of customers requesting tickets to the alternatives offered by the airline and its competitors.

More general information on mathematical methods mentioned as approaches to forecasting are described in the following sections can be found in Armstrong (2001).

3.1. Demand Volume

When within the booking horizon will customers request tickets? How many customers, overall, will demand tickets for a certain flight departure? If information on the amount of demand per product and the timing of its arrival is available, target functions can be formulated to optimally allocate capacity. Research conducted in order to gain estimates for both dimensions will be summarized in this section.

Much early research on forecasting for revenue management has been done to optimize overbooking levels. This is not the key question considered in this thesis; further reading related to overbooking can be found in Taylor (1962); Rothstein & Stone (1967); McGill (1989).

Arrival Process: To determine the optimal availabilities within the booking horizon, three aspects of demand have to be available: A fixed overall demand (regarded as the *state space* according to a terminology derived from A. O. Lee (1990)), the rate of demand arrival as well as the sequence of high- and low-fare demand.

First distributions that describe demand arrival for a single fare class and flight are given in Beckmann & Bobkowski (1958). The resulting description of demand is referred to as *booking curve*. In addition to the positive impact of accepted passenger requests, *cancellations* have a negative impact on this curve.

In order to reduce complexity, it is common to divide the booking horizon up and to then observe bookings as they occurred in time slices. Method that solely considers bookings that have already been accepted on the flight is referred to as *advanced bookings method* according to A. O. Lee (1990). Those concepts that only refer to observed demand on previous departures of the same flight numbers are regarded as *historical booking methods*. It is noted that methods combined from both concepts tend to work best – accordingly, current research uses information both from past flights and already accepted bookings.

Poisson processes as a special case of *Markovian processes* are a common way of modeling demand arrival – one of the first examples is Rothstein (1968). As they assume the state of the system they model to be influenced only by the latest event, the arrival of demand is regarded to be independent of booking controls. Customers request tickets exactly once, if their request is denied they do not return. Recent examples of the application of Poisson processes can be found in Weatherford et al. (1993), Talluri & Van Ryzin (2004a) and Walczak & Brumelle (2007). The arrival rate of the customers is regarded as a parameter to be estimated either via ad-hoc or time-series methods or by analyzing influence factors.

To optimally allocate seats at any time within the booking period, Chen et al. (2003) propose a method of statistical learning that estimates a market value for tickets being purchased at a specific time. The paper extends a model based on a *discrete-time Markov decision problem* in Lautenbacher & Stidham (1999) to the network level. Instead of explicitly forecasting demand, *value functions* for remaining seats are estimated and updated.

A similar topic, demand learning for dynamic pricing, is taken up in Xu & Hopp (2005a). This paper includes the assumption of a *piecewise deterministic and Markovian customer arrival process*, the distribution of which is regarded as known, homogeneous and price independent. It introduces estimates key parameters of the distribution by observing demand as it arrives.

Correlations in demand for different products at the same point of time as well as correlations in demand for the same products at different points of time within the booking horizon are considered in Stefanescu et al. (2004). A *linear mixed effects model* considers a booking time component, weighting matrices for correlations of influences by external shocks, and a normally distributed error term. Data is unconstrained using the estimation maximization algorithm. In order to estimate the parameters of the demand distribution, *maximum likelihood* is applied.

Final Bookings: Overall demand volume can fluctuate over time due to economic trends, the influence of seasons, trade fairs, and holidays, within weekly or daily patterns. As described in Talluri & Van Ryzin (2004b) and Armstrong (2001), a variety of statistical methods can be employed in order to pick up trends and patterns. Two views stand out: Demand may be predicted by considering the data patterns and emulating them with *ad-hoc methods* or it may be predicted by considering possible causal relationships between the booking data and influence factors. The latter approach is also called *associative*.

Much current revenue management research focuses on demand behavior rather than overall demand development. Research considering the macro level of airline demand as needed for airline fleeting, strategy, and general economics stays important. Examples can be found in Andersson (2001), Abed et al. (2001), Bhadra (2003), Battersby (2005) and Cunningham & De Haan (2006). Brons et al. (2002) and Njegovan (2006) consider the influence of prices on overall demand levels.

In Grosche et al. (2007), a rather macroeconomic view of demand estimation is taken, too. Still, the authors do mention the possibility of using their estimates for the optimization of itineraries for which no historical data exists yet. In order to calculate the connection between service-related and geo-economic forces and demand, two *gravity models* are implemented. These assume the influencing factors to be independent.

Even if demand is considered stationary, some fluctuation is likely to occur from one departure to the next. Although focusing on spill estimation for fleet assignment, W. Swan (2002) indicates that a *gamma distribution* of average stationary demand makes sense for revenue management. The paper follows up on research described in W. Swan (1993a,b, 1999).

The idea that overall demand volume may not be exclusively derived the time factor is presented in Sa (1987). The author introduces a model of regression analysis that may be applied to different markets to forecast final bookings. Apart from time and historical bookings, this regression considers socio-demographic factors as well as level of service variables. An earlier overview of regression models for demand estimation can be found in Taneja (1978).

In order to predict demand over a network, Neuling et al. (2004) proposes an analysis of *passenger name records*. This data includes information on each passenger's itinerary, including all booked segments. In addition to a regular demand forecast, a *no-show forecast* is offered.

A theory of time series influenced by a variety of factors is laid out in Armstrong et al. (2004). The authors describe a concept of decomposing the series to represent the causal forces that impact it.

Time series and their reaction to external factors are also the focus of hybrid methods combining traditional statistics and *neural networks*. An example of such an effort is presented in Aburto & Weber (2007), which connects neural networks with an *ARIMA model*. However, the application considers supermarkets rather than airlines.

3.2. Unconstraining

Historical booking data as stored by airlines does not represent actual demand. Customer requests are *constrained* by the amount of tickets offered. Without further information, it is unknown whether more tickets could have been sold, unless offer exceeded demand. A general introduction to this topic is also presented in Pölt (2000). *Unconstraining*, also called *detruncation*, refers to the transformation of bookings to demand.

Nahmias (1994) refers to the demand that is not accepted as bookings as *lost sales*. The paper concentrates on retail rather than the airline industry and describes the application of maximum likelihood estimation, a best linear unbias estimator and an additional new estimator.

In Zeni (2001a) the author extensively presents and compares a number of unconstraining methods. The findings are also summarized in Zeni (2001b). Among the listed concepts are: ignoring the censoring, discarding the censored data, *mean imputation method*, *the booking profile method*, *expectation maximization*, and *projection detruncation*.

The easiest alternative is to ignore the censoring or to discard the censored data. The mean imputation method, a variation of *pickup* unconstraining, and the booking profile method are both ways of estimating demand by extrapolating from the mean of bookings that were not truncated. Expectation maximization uses a normal distribution and maximum likelihood in order to iteratively estimate the parameters influencing demand. Finally, projection detruncation is a variation of expectation maximization that uses an additional parameter to scale the amount of unconstraining applied to the data.

The conclusion of Zeni (2001b) is that anything preferable to ignoring the truncation. While intricate and computationally intense, expectation maximization works best. Such findings are also confirmed by Weatherford (2000) and Weatherford & Pölt (2002).

McGill (1995) introduces a concept of *censored regression* in order to estimate customer bookings. This model extends unconstraining by estimating demand depending on up to nine factors. To that aim, the expectation maximization method is adapted to what the authors call *censored demand estimation maximization*. Other examples of the use of expectation maximization can be found in Talluri & Van Ryzin (2004b), Stefanescu et al. (2004) and Ferguson et al. (2007).

While claiming not to use any forecasting techniques, Ryzin & McGill (2000) solves the problem of unconstraining data by applying *life tables*. This method is taken from *survival analysis* and is implemented to estimate parameters of a survival function. It indicates how many additional requests arrived after a booking class was closed. Tests indicate mixed performance, therefore the implementation is only recommended for small or start-up airlines or when demand is difficult to predict.

A new method of unconstraining is proposed in Ferguson et al. (2007). It uses *double exponential smoothing*, also called *"Holt's Method"*. Two smoothing constants are introduced to calculate the base demand and the trend. An application is described both for monotonously closing booking classes and booking classes being re-opened. Based on simulated customer requests, the new method is compared to an averaging method, estimation maximization and projection detruncation taken from Weatherford & Pölt (2002), as well as the method of life tables described in Ryzin & McGill (2000). The results are generally favorable but do vary with regard to the simulated customer behavior.

In Lan & Gao (2007), another approach to dealing with limited demand information is offered. The authors develop robust inventory controls when only upper and lower boundaries for demand are known. The concept is shown to be effective for single-leg problems. This method as well as its preceding research on *competitive online algorithms* in Ball & Queyranne (2006) would have to be adapted for network models.

While other industries profit from available *turn-down data* (see for instance Zhu (2006)), such data is not yet fully available for airlines. Turn-down data stores customer requests that have been denied in addition to those that were accepted as bookings. If such data was available, the transformation of sales data to get information on demand would become superfluous.

The nearest thing to turn-down data available for airlines are so-called *click-streams*. These record customer behavior as observed on travel websites. They are likely to be biased as customers do not seriously consider all travel-itineraries clicked at. The possible importance of this kind of data for future revenue management forecasting is hinted at in Nason (2007). However, no published research on working with this kind of information in the airline business could be found.

3.3. Demand Behavior

The last sections as well as traditional approaches consider the overall amount of demand to arrive for a distinct product (itinerary and booking class). Changes in customer mentality make other considerations necessary. As laid out for example in van Ryzin (2005), customers that are more informed and flexible require a shift of focus from products to customers. It used to be feasible to merely ask how many units of one product (for instance, business class tickets) would be requested. Now it makes more sense to ask how many business class customers (that might also go for a bargain ticket if it is offered to them) do consider leaving their origin at a certain time for a specific destination. Much of the development that has lead to this kind of flexible customer behavior has been described in Section 2.3

With a new view of demand comes a consideration of dynamic demand behavior. Customers choose between alternatives that are offered by the airline and its competition. It is no longer the absolute number of customers that revenue management needs to consider, but also their *choices* with regard to prices, competition, and other utilities such as schedule time and transfers. As indicated in P. Belobaba (1987b), with such choice behavior, it is possible to recapture rejected customer requests *vertically* (to a different booking class) or *horizontally* (to a different itinerary).

Sell-Up and Buy-Down: Vertical choice behavior leads to *buy-down*, the possibility that customers with a high willingness to pay are offered a low price and accept it. On the other hand, it opens the way to *sell-up*, customers being forced by inventory controls to buy a class that is more valuable than the cheapest class possible. Simulation results with regard to the performance of traditional revenue management given buy-down and sell-up are offered in Cusano (2003) and Ozdaryal & Saranathan (2004).

Under these conditions, the challenge of revenue management becomes to include information on customers' price-elasticity in forecasting models. As will be listed in the following paragraphs, different approaches formulate this using different concepts such as as *willingness to pay*, *Q-forecasting*, or *hybrid demand*.

Willingness to Pay: An appraisal of the influence of prices on customer demand is documented in Castelli et al. (2003). With an *ordinary least squares regression* as well as a *multilevel analysis-based methodology*, the variance of price elasticity on different routes in a network is analyzed. With its focus on overall demand as well as on specific routes for specific airlines, this paper straddles the border of macro- and microeconomics. Weber & Thiel (2004) presents an attempt to estimate customer demand curves based on price. An *artificial neural network* is implemented in order to derive price elasticities from booking data.

First ideas about demand behavior being based on price elasticities date back as early as the 1970s. Notable examples are Lennon (1972) and Jung & Fuji (1976).

Q-Forecasting and Fare Modifiers: Some approaches to forecasting demand under the assumption that customers are not segmented according to restrictions are presented in Cléaz-Savoyen (2005). The findings are accompanied by simulation results as attained through PODS. Building up on research documented in Bohutinsky (1990), P. Belobaba & Weatherford (1996), and Gorin (2000), the thesis concentrates on combinations of *Q-forecasting* and *fare modifiers*.

Q-forecasting models the behavior of customers by calculating the amount of passengers willing to buy the class with the lowest fare - “Q”. The amount of passengers willing to buy the next higher priced class is predicted from *sell-up rates*. These are formulated as *Frat5-rates*: the price-ratio of two classes at which fifty percent of the demand for the lower priced class will be willing to buy the higher-priced class. Frat5-rates are estimated dynamically over the booking horizon via a regression across time-frames.

Fare adjustment using fare modifiers as a concept developed in Fiig & Isler (2004) are introduced to optimize revenue for flights that have both a restricted and an unrestricted fare-structure, thereby catering both to price-sensitive customers and to independent demand. Fare modifiers also rely on sell-up estimates; however, they do not necessarily assume that all lower priced classes are closed (as this is unnecessary in a restricted fare environment).

Hybrid Demand: Building up on Cléaz-Savoyen (2005), Reyes (2006) offers forecasting methods for a combination of restricted and unrestricted classes referred to as *hybrid fare structures*. He describes the challenge of separating *price* and *product-oriented* demand. A similar statement can be found in Boyd & Kallesen (2004), where the two demand segments are called *priceable* and *yieldable*. While a product-oriented customer is only interested in purchasing a ticket with or without specific restrictions, a price-oriented customer will be looking to buy at the lowest available price.

Two understandings of *hybrid forecasting* are introduced in Reyes (2006): The simultaneous deployment of two separate forecasting methods for the two customer segments or the separate forecasting of two fare structures, restricted and unrestricted. The concept is ascribed to P. Belobaba & Hopperstad (2004). Combined are two more methods: Fare adjustment as outlined in Fiig & Isler (2004) and *path categorization*. A similar approach to incorporate market-based demand forecasts in network optimization is also taken in C. Hopperstad (1994).

Path categorization assumes that willingness to pay is related to the amount of transfers a passenger's way over an O&D network includes, as well as to how dominant a market he comes from. Tests conducted with the help of the PODS are cited to show that both fare adjustment and path categorization can significantly improve network revenue if hybrid forecasting is applied. A review of the performance of sell-up algorithms in PODS can be found in Guo (2007).

Customer Choice: While the consideration of prices for customer decisions is gaining importance with the rise of no-frills carriers, other factors contribute to customer choice behavior, too. It is intuitively appealing that when buying tickets, customers should consider the travel time as well as the amount of transfers and connecting times. An example of the examination of travel choices with regard to risk can be found in Theis et al. (2006). Reyes (2006) mentions a connection between willingness to pay and transfers. The paper includes this connection by applying path categorization.

The idea that customer choice depends on the distribution channel and characteristics of the itinerary is outlined in Walczak & Brumelle (2007). This article claims that the decision to buy depends on the price and assumes that the customer's demand function with regard to the price is known. Arrival rates derived from the demand function are included in a Poisson process. *Customer profiles* are included in the arrival process, thereby making customers' demand functions variable. Comprised is also the so-called *market state*, the competitive situation. The article refers to Walczak (2005) at this point (see also Section 2.3).

The quantification of the influence of product characteristics and market state on customer demand is relegated to *consumer choice behavior models*. Among these, *discrete choice models* are of special interest for revenue management - in these, customers have to choose exactly one of several distinct alternatives. The theory of discrete choice models with special regard to revenue management and parameter estimation via maximum likelihood has been described for the first time in Kanafani & Sadoulet (1977) and Ben-Akiva & Lerman (1985). On this theoretical background, a multitude of research has evolved.

In Talluri & van Ryzin (2000), a *multi-nomial logit model of demand* is developed in order to predict customer choice. This model implements *logistic regression* for more than two variables to be included in customers' cost functiosn. The probability of a customer buying a product is calculated as a natural logarithm defined by a linear function. The model parameters (arrival rates as well as choice factors) are estimated via estimation maximization. Given this model, revenue is optimized for the single-leg multiple-fare case with the help of a dynamic program.

The complexity of such programs increases over the available seats as well as over the number of itineraries and classes. This issue is also referred to as the *curse of dimensionality* – see for example Bertsimas & de Boer (2005) for a description as well as an attempt of solution. Therefore, most research considering customer choice models and dynamic programs only consider one leg at a time. A similar model is also applied in Talluri & Van Ryzin (2004a); here, only the prices that are offered are considered as factors.

In Talluri & Van Ryzin (2004b), a *customer choice model* is implemented in order to estimate the connection between prices and customer demand. In this case, customers' cost functions only include fares. After the model parameters are estimated using maximum likelihood, the paper offers optimal policies for an independent view of demand, a *multinomial logit model*, and a model in which customers always purchase the lowest available fare.

Another introduction to choice-based revenue management is provided in Vulcano (2006). *Simulation based optimization* is suggested to build up on forecasts that incorporate this model. Apart from the already mentioned multinomial logit model of choice, alternatives such as *finite-mixture logit*, *Markovian second choice*, and *general random utility* are mentioned. The author also distinguishes between the estimation of choice parameters and that of volume parameters. Forecasting methods for the latter have been presented in Section 3.1. In order to estimate choice parameters, both estimation maximization and maximum likelihood are proposed.

Finally, an examination of price sensitivity with special regard to customers that buy airline tickets online is documented in Garrow et al. (2007). The authors use *stated preference data* in order to estimate a multinomial nested logit model of customer choice behavior. As a consequence of using stated preference data, the article draws special attention to matters of survey design and recruitment. Both ticket prices and sociodemographic factors are included in the analysis in an attempt to explain willingness to pay through other utilities. The use of knowledge of customer behavior is also outlined in Fudenberg & Villas-Boas (2006).

4. Demand Forecast Performance Measurements

As can be seen from the overview presented in the previous sections, much development has taken place with regard to airline revenue management in general and especially demand forecasting. With a shift from product focus to customer focus and a trend toward less restricted products, high quality forecasting has gained significance. The importance of forecast accuracy has been underlined for instance in Weatherford & Belobaba (2002).

Considering the task of finding the best forecast for a given situation, it can also make sense to consider the complete revenue management system. Benchmarks to analyze the efficiency of a set-up of this kind are provided by so-called *revenue opportunity models* (*ROM*). Such models usually attempt to estimate an upper benchmark for the revenue that may be earned in a market by analyzing historical booking data (an example of this is described in Rannou & Melli (2003) with regard to the hotel industry).

Granger & Pesaran (2000) propose the use of *decision theory* in order to evaluate forecast performance. This means evaluating the outcome of decisions based on the forecast in order to make statements about its quality. In the case of predicting demand in order to maximize revenue, the forecast that leads to the highest earnings accordingly must be the best.

General statistics offer several possible *methods* to calculate forecast error measurements. Without regarding the specific challenges of determining future sales by transforming historical bookings, these measurements will be presented in Section 4.1. In the literature presented, the consequences of their application to specific *objects* and *aggregation levels* of comparisons are only hinted at.

For example, the object of comparison for forecast evaluation needs to be agreed on before computing error measurements. Armstrong (2001) suggests analogous data from different geographical areas or control groups as well as *"backcasting"* (predicting the past based on current data). Evaluation methods applied in revenue management tend to use actual booking data (see for example Pölt (1998)). However, they have to chose between different transformations to make booking data comparable to demand forecasts. The consequences of this choice require further examination.

Existing attempts to evaluate any share of the multitude of forecast methods offered can be found in Section 4.2. This includes tests of forecasting concepts conducted in the course of their introduction as well as overview articles comparing existing approaches.

Challenges that arise from the application of the theoretical indicators of forecast accuracy presented in this Chapter to the reality of revenue management are pointed out in Section 5. Regarding these difficulties, research opportunities are identified. Finally, immediate steps that can be taken to create a new approach to forecast evaluation are outlined.

4.1. Theoretical Background

A general introduction to the concept of evaluating forecasts by statistical means is presented in Armstrong (2001). Further literature considering the topic will be summarized in the context of those explanations.

An introductory overview of the challenges and aspects of time series forecasting as well as of the consequences of forecasting errors in economic settings is provided by Makridakis (1986). General guidelines toward the evaluation of forecast for a specific problem is presented in Fildes (1992). The author points out that rather than relying on prior forecasting competitions, researchers should evaluate methods with regard to the data at hand.

The underlying assumption of evaluating forecasts is that the considered alternatives are all methodically valid. This means that the task of choosing the right method for a certain situation has already been tackled. As was shown in the previous chapter, a range of solutions tailored specifically to forecasting for airline revenue management does exist.

However, the question which approach is *best* remains. Note that forecast quality depends on the priorities of the evaluator: It may indicate accuracy, robustness, or a good fit with other revenue management methods. In this chapter, a forecast's ability to predict actual demand, its accuracy, is regarded as its main performance factor.

Armstrong (2001) differentiates between analyzing the *input* and the *output* of forecasting methods. When rating two methods processing the same kind of data and reaching comparable results, a statistical indicator can be computed.

In the case of forecasting for revenue management, similar inputs are accessed by most methods: Historical bookings and in some cases survey data. Forecasting methods for revenue management may be arranged by the input data they use as well as the information they provide on demand. The assumptions used were already mentioned when describing the available methods in Section 3.

Another aspect of error measurement mentioned in Armstrong (2001) may be relevant when testing forecast methods for revenue management: Asymmetrically rated evaluation. If for instance the underestimation of demand is thought to be more harmful to revenue than their overestimation, negative errors in the forecast might be rated more severely. Whether to consider such asymmetries in the initial measurement of forecast errors or to include their analysis in a final interpretation of the evaluation is a matter of context.

In order to evaluate error measures, Armstrong (2001) suggests the comparison of rankings derived from those measurements. However, what measurements should be calculated and compared?

In order to define the forecast error, Fildes (1992) introduces a formal notation. It has been modified to fit the model presented in the remainder of this text.

- Let s be the index of departure days (in the real world) or simulation runs (in a simulation model). In chronological order: $s = 1, ..., N^s$.
- Let t be a point of time in the booking horizon, $t = 0, ..., N^t$ with $t = 0$ indicating the start of the booking horizon and $t = N^t$ indicating its end.
- Let f be the index of flights with $f \in F$.
- Let c be the index of classes with $c \in C$.
- Let $f^{\text{unc}}(f, c, t, s)$ be the forecast of unconstrained demand for flight f in class c between the points of time $t-1$ and t of s.

- Let $b^{\text{unc}}(f,c,t,s)$ be the unconstrained bookings for flight f in class c between points of time $t-1$ and t of s.
- Let $e^{\text{u-u}}_{\circ}(f,c,t,s)$ be the forecast error computed according to method $\circ$ for flight f, class c, and time slice $t-1$ to t in the booking horizon of s. The inputs for the computation of this error are $f^{\text{unc}}(f,c,t,s)$ and $b^{\text{unc}}(f,c,t,s)$.

Errors can be summarized according to two ways: *Series Squared Errors* (*SSE*) and *Series Absolute Percentage Errors* (*SAPE*). Aggregation can happen across a time period considered in one series or across several series given one point of time or the complete period.

Ideally, indicators measuring forecast quality should allow for the criteria given by Diebold & Lopez (1996) as well as for a minimization of error as described by Fildes (1992). To that effect, a combination of error measures can be helpful, describing both the behavior of errors as well as their overall volume.

A preference of *unit-free* measures not affected by the scale of the predictions avoids to unfairly weight different markets with different demand volumes. Biased error measures should also be avoided - an example is the use of the SAPE indicator *mean absolute percent error* (*MAPE*) when considering positive numbers (such as bookings without cancellations). MAPE is calculated by averaging the *absolute percent errors* (*APE*).

- Let $\hat{e}^{\text{u-u}}_{\circ}(s)$ be the series error computed according to method $\circ$ for s.
- Let $e^{\text{u-u}}_{\text{APE}}(f,c,t,s)$ be the APE computed as shown in Definition (4.1).
- Let $\hat{e}^{\text{u-u}}_{\text{MAPE}}(s)$ be the MAPE computed as shown in Definition (4.2).

$$e^{\text{u-u}}_{\text{APE}}(f,c,t,s) := \frac{|f^{\text{unc}}(f,c,t,s) - b^{\text{unc}}(f,c,t,s)|}{b^{\text{unc}}(f,c,t,s)} \cdot 100 \qquad \forall\ f \in F;\ c \in C;\ t = 1,...,N^t - 1;\ s = 1,...,N^s \tag{4.1}$$

$$\hat{e}^{\text{u-u}}_{\text{MAPE}}(s) := \frac{\sum_{f\in F}\sum_{c\in C}\sum_{t=1}^{N^t} e^{\text{u-u}}_{\text{APE}}(f,c,t,s)}{|F|\cdot|C|\cdot N^t} \qquad \forall\ s = 1,...,N^s \tag{4.2}$$

As it calculates the difference between a positive forecast and another positive benchmark, the highest difference possible if demand is underestimated is 100%. Overestimation may be infinitely high. With such measures, an adjustment as the *unbiased absolute percentage error* (*UAPE*) presented in Makridakis (1993) can be in helpful (see Definition (4.3)). This error measurement is summarized in Definition (4.4). The result is also referred to as *symmetrical MAPE* (*SMAPE*) by Tayman & Swanson (1999).

- Let $e^{\text{u-u}}_{\text{UAPE}}(f, c, t, s)$ be the UAPE computed as shown in Definition (4.3).
- Let $\hat{e}^{\text{u-u}}_{\text{SMAPE}}(s)$ be the SMAPE computed for s as shown in Definition (4.4).

$$e^{\text{u-u}}_{\text{UAPE}}(f, c, t, s) := \frac{\left|f^{\text{unc}}(f, c, t, s) - b^{\text{unc}}(f, c, t, s)\right|}{\frac{f^{\text{unc}}(f,c,t,s)+b^{\text{unc}}(f,c,t,s)}{2}} \cdot 100 \qquad \forall\ f \in F;\ c \in V;\ t = 1, ..., N^t - 1;\ s = 1, ..., N^s \tag{4.3}$$

$$\hat{e}^{\text{u-u}}_{\text{SMAPE}}(s) := \frac{\sum_{f\in F}\sum_{c\in C}\sum_{t=1}^{N^t} e^{\text{u-u}}_{\text{UAPE}}(f, c, t, s)}{|F| \cdot |C| \cdot N^t} \qquad \forall\ s = 1, ..., N^s \tag{4.4}$$

Alternatively, in order to compensate for the measure's vulnerability with regard to underestimation, Fildes (1992) suggests the use of the *Median Absolute Percentage Error* (*MdAPE*). The calculation of this measure is shown in Definition (4.5).

- Let $\hat{e}^{\text{u-u}}_{\text{MdAPE}}(s)$ be the MdAPE computed for s as shown in Definition (4.5).
- Let $N = |F| \cdot |C| \cdot N^t$ be the overall number of observations for flights, classes and points of time before departure per s.
- Let $E^{\text{u-u}}_{\text{APE}}(s)$ be the set of all observations of APE for s.
- Let $\arg(E, n)$ be a function returning the nth element of the ordered set E.

$$\hat{e}^{\text{u-u}}_{\text{MdAPE}}(s) := \begin{cases} \frac{\arg\left(E^{\text{u-u}}_{\text{APE}}(s), \frac{N}{2}\right) + \arg\left(E^{\text{u-u}}_{\text{APE}}(s), \frac{N}{2}+1\right)}{2} & N \bmod 2 = 0 \\ \arg\left(E^{\text{u-u}}_{\text{APE}}(s), \frac{N}{2}+1\right) & N \bmod 2 > 0 \end{cases} \qquad \forall\ s = 1, ..., N^s \tag{4.5}$$

Based on the claim that MAPE tends to overstate the error found in skewed forecasts, Tayman & Swanson (1999) reason that a class of measures called *Minimization Estimators* (*M-estimators*) are more valid. The authors list *resistance* and *robustness* as criteria for a good forecast performance estimator. Resistance describes small subsamples having only a limited effect on the results of the evaluation. Robustness implies insensitivity to underlying assumptions e.g. about the distribution. The authors compare the use of MUAPE and M-Estimators. The latter are based on maximum likelihood procedures minimizing objective functions describing the relative deviation of observations

The performance of any forecast depends on the degree in which the data considered *can* be predicted. If forecasts are evaluated between different markets, it makes sense to consider their sensitivity to the degree of *demand uncertainty.* If for example departures of the same flight rather than different markets are compared, each point of data includes the same level of uncertainty. However, when comparing forecasts to actual data the pool of data the researchers can analyze often is limited. In Sullivan et al. (2003), solutions to this problem of *shared data sets* are presented.

A comparison to the naive forecast (*random walk*) provided by the so-called *relative absolute error* (*RAE*) can be advisable to neutralize market effects. The random walk predicts the next value of a series by assuming that it will be the same as the previous value. It is introduced in Definition (4.6).

- Let $\hat{f}\,(f, c, t, s)$ be the naive forecast for class c on flight f between points of time $t-1$ and t of s as shown in (4.6).
- Let $e^{\text{u-u}}_{\text{RAE}}\,(f, c, t, s)$ be the relative absolute error calculated for the forecast generated for demand to arrive for flight f, class c, between points of time $t-1$ and t of run s as shown in Definition (4.7).
- Let $\hat{e}^{\text{u-u}}_{GMRAE}\,(s)$ be the RAE summarized over its geometric mean (*GMRAE*) as presented in Definition (4.8).
- Let $\hat{e}^{\text{u-u}}_{MdRAE}\,(s)$ be the RAE summarized over its median (*MdRAE*) as presented in Definition (4.9).

$$\hat{f}\,(f,c,t,s) := \begin{cases} 0 & s=1 \\ b^{\text{unc}}\,(f,c,t,s-1) & s>1 \end{cases} \qquad (4.6)$$
$$\forall\; f \in F;\;\; c \in C;\;\; t = 1,...,N^t$$

$$e^{\text{u-u}}_{\text{RAE}}\,(f,c,t,s) := \frac{\left|f^{\text{unc}}\,(f,c,t,s) - b^{\text{unc}}\,(f,c,t,s)\right|}{\left|\hat{f}^{\text{unc}}\,(f,c,t,s) - b^{\text{unc}}\,(f,c,t,s)\right|} \qquad (4.7)$$
$$\forall\; f \in F;\;\; c \in C;\;\; t = 1,...,N^t;\;\; s = 1,...,N^s$$

$$\hat{e}^{\text{u-u}}_{\text{GMRAE}}\,(s) := \left(\prod_{f\in F}\prod_{c\in C}\prod_{t=1}^{N^t} e^{\text{u-u}}_{\text{RAE}}\,(f,c,t,s)\right)^{1/\left(|F|+|C|+\left(N^t\right)\right)} \qquad (4.8)$$
$$\forall\; s = 1,...,N^s$$

Note that the geometric mean is costly to calculate for large $|F| + |C| + N^t$.

$$\hat{e}^{\text{u-u}}_{\text{MdRAE}}\,(s) := \begin{cases} \arg\left(E^{\text{u-u}}_{\text{RAE}}\,(s)\,,\,\frac{N+1}{2}\right) & N \bmod 2 = 0 \\ \frac{\arg\left(E^{\text{u-u}}_{\text{RAE}}(s),\frac{N-1}{2}\right)+\arg\left(E^{\text{u-u}}_{\text{RAE}}(s),\frac{N+1}{2}\right)}{2} & N \bmod 2 > 0 \end{cases} \qquad (4.9)$$
$$\forall\; s = 1,...N^s$$

According to Armstrong & Collopy (1992), RAE provides a valid alternative to another measure called Theil's $U2$ as presented in Theil (1966) and discussed and recommended in Bliemel (1973) (see Definition (4.12). In order to make RAE comparable to MAPE, Armstrong & Collopy (1992) suggests the calculation of a *Relative Absolute Percent Error* as shown in Definition (4.10).

- Let $e^{\text{u-u}}_{\text{RAPE}}\,(f,c,t,s)$ be the relative absolute percent error computed as shown in Definition (4.10).
- Let $\hat{e}^{\text{u-u}}_{\text{MRAPE}}\,(s)$ be the mean relative absolute percent error computed for s as shown in Definition (4.11).

- Let $\hat{e}^{\text{u-u}}_{\text{U2}}(s)$ be Theil's U2 as shown in Definition (4.12).

$$e^{\text{u-u}}_{\text{RAPE}}(f,c,t,s) := \frac{\frac{|f^{\text{unc}}(f,c,t,s)-b(f,c,t,s)|}{b(f,c,t,s)}}{\frac{\left|\hat{f}^{\text{unc}}(f,c,t,s)-b(f,c,t,s)\right|}{b(f,c,t,s)}}$$
$$\forall\ f\in F;\ c\in C;\ t=1,...,N^t;\ s=1,...,N^s \tag{4.10}$$

$$\hat{e}^{\text{u-u}}_{\text{MRAPE}}(s) := \frac{\sum_{f\in F}\sum_{c\in C}\sum_{t=1}^{N^t} e^{\text{u-u}}_{\text{RAPE}}(f,c,t,s)}{|F|\cdot|C|\cdot N^t}$$
$$\forall\ s=1,...,N^s \tag{4.11}$$

$$\hat{e}^{\text{u-u}}_{\text{U2}}(s) := \frac{\sqrt{\sum_{f\in F}\sum_{c\in C}\sum_{t=1}^{N^t}\left(f^{\text{unc}}(f,c,t,s)-b(f,c,t,s)\right)^2}}{\sqrt{\sum_{f\in F}\sum_{c\in C}\sum_{t=1}^{N^t} b(f,c,t,s)}}$$
$$\forall\ s=1,...,N^s \tag{4.12}$$

While the relative absolute percent error is bounded between 0 and 1, RAE and U2 have no finite upper boundary. Another estimate of forecast quality based on the comparison to a random walk is *Percent Better*. As shown in Definition (4.15), this measure calculates the percentage of forecasts for which a given method is more accurate than the random walk or naive forecast.

- Let $\delta^{\text{PB}}_{f,c,t,s}$ as shown in Definition (4.13) be an indicator to define whether for one combination of flight f, class c, and time slice $t-1$ to t of the booking horizon of s the difference of the forecast to the bookings is smaller than that of the naive forecast to the bookings.

- Let $n^{\text{u-u}}_{\text{PB}}(s)$ as shown in Definition (4.14) be the number of cases in which the difference between the unconstrained forecast computed the unconstrained bookings is smaller than that between the naive forecast and unconstrained bookings for s.

- Let $\hat{e}^{\text{u-u}}_{\text{PB}}(s)$ as shown in Definition (4.15) be the percent better indicator.

$$\delta^{\text{PB}}_{f,c,t,s} := \begin{cases} 1 & \left| f^{\text{unc}}\left(f,c,t,s\right) - b\left(f,c,t,s\right) \right| < \left| \hat{f}^{\text{unc}}\left(f,c,t,s\right) - b\left(f,c,t,s\right) \right| \\ 0 & \text{else.} \end{cases} \qquad (4.13)$$
$$\forall\ f \in F;\ c \in C;\ t = 1, ..., N^t;\ s = 1, ..., N^s$$

$$n^{\text{u-u}}_{\text{PB}}\left(s\right) := \sum_{f \in F} \sum_{c \in C} \sum_{t=1}^{N^t} \delta^{\text{PB}}_{f,c,t,s} \qquad (4.14)$$
$$\forall\ s = 1, ..., N^s$$

$$\hat{e}^{\text{u-u}}_{\text{PB}}\left(s\right) := \frac{n^{\text{u-u}}_{\text{PB}}\left(s\right)}{|F| \cdot |C| \cdot \left(N^t\right)} \cdot 100 \qquad (4.15)$$
$$\forall\ s = 1, ..., N^s$$

Armstrong & Collopy (1992) emphasize the ease of explanation as an advantage of RAE over U2. However, as described in Makridakis (1993), the use of RAE is not without problems, either, as it can be blown out of proportion due to overly large random values. On the other hand, the use of Theil's U2 can be justified in a scientific context when managerial understanding is not the first priority.

Armstrong (2001) argues against the use of both a *coefficient of determination*, R^2, and *Root Mean Square Error* (*RMSE*) when evaluating time series forecasting methods. R^2, as shown in Definition (4.16), is calculated from the ratio of the sum of squared errors and the total sum of squares of a set of predicted and observed values. While an R^2 of 1 is regarded as an indication of a perfect fit of predicted values to the observed data, this can be misleading. On the one hand, this figure does not consider a systematical bias. On the other hand, with a variation of zero in the data, R^2 can turn out as zero, indicating no correlation, even if the forecast is correct.

- Let $\overline{b}\left(s\right)$ be the average number of bookings per flight and class and slice of time of the booking horizon observed in s.
- Let $R^2\left(s\right)$ as shown in Definition (4.16) be the coefficient of determination for forecasts computed in s.

$$\bar{b}(s) := \frac{\sum_{f\in F}\sum_{c\in C}\sum_{t=1}^{N^t} b(f,c,t,s)}{|F|\cdot|C|\cdot N^t}$$

$$R^2(s) := \frac{\sum_{f\in F}\sum_{c\in C}\sum_{t=1}^{N^t}\left(f^{\text{unc}}(f,c,t,s) - b(f,c,t,s)\right)^2}{\left(b(f,c,t,s) - \bar{b}(s)\right)^2} \tag{4.16}$$

$$\forall\ s = 1,...,N^s$$

The RMSE as shown in Definition (4.17) is calculated as the square root of the *mean squared error* (*MSE*) also referred to as *mean square forecast error* (*MSFE*) or *mean square deviation* (*MSD*). It offers an alternative to the *mean absolute error* (*MAE*) or *mean absolute deviation* (*MAD*). The RMSE is cautioned against in Armstrong & Collopy (1992), Armstrong & Fildes (1995) and Armstrong (2001) as it overstates the impact of large errors. Its interpretative pitfalls are also emphasized in Fildes (1992). While this makes sense in so far as that larger errors can lead to larger economic losses, it mixes the interpretation of results with their calculation. For the same reason, the use of MSE is criticized in Chatfield (1988). It is recommended for use only when the predicted quantities are of a comparable order. In the case of comparing for instance departures of the same flight, this can hold true.

- Let $\hat{e}^{\text{u-u}}_{\text{RMSE}}(s)$ be the RMSE computed as shown in Definition (4.17).

$$e^{\text{u-u}}_{\text{RMSE}}(s) := \sqrt{\frac{\sum_{f\in F}\sum_{c\in C}\sum_{t=1}^{N^t}\left(f^{\text{unc}}(f,c,t,s) - b(f,c,t,s)\right)^2}{|F|\cdot|C|\cdot N^t}} \tag{4.17}$$

$$\forall\ s = 1,...,N^s$$

In Fildes (1992), the calculation of a *Geometric Root MSE Squared* is suggested in order to summarize the indicator. For reasons of computational effort, like U2, RMSE will be averaged arithmetically in the further text.

Scale-dependent measures such as MSE cannot be used to compare forecast performance over diverse series. Measures based on percentage errors such as MAPE are undefined if zero values are predicted (division by zero), and measures based on relative errors such as MRAE can overstate extreme values. Based on these problems, Hyndman & Koehler (2006) introduce a new measure. The *Mean Absolute Scaled Error* (MASE) is calculated as a mean of the RAE, scaling the forecast error to the error of the naive forecast (see Definition (4.7)).

- Let $\hat{e}^{\text{u-u}}_{\text{MASE}}(s)$ be the MASE computed as shown in Definition (4.18).

$$\hat{e}^{\text{u-u}}_{\text{MASE}}(s) := \frac{\sum_{f\in F}\sum_{c\in C}\sum_{t=1}^{N^t}\left|f^{\text{unc}}(f,c,t,s) - b(f,c,t,s)\right|}{|F|\cdot|C|\cdot N^t} \qquad \forall\ s = 1, ..., N^s \tag{4.18}$$

Using the concept for computing MASE, the authors also define the *Root Mean Squared Scaled Error* (*RMSSE*).

- Let $\hat{e}^{\text{u-u}}_{\text{RMSSE}}(s)$ be the RMSSE computed as shown in Definition (4.19).

$$\hat{e}^{\text{u-u}}_{\text{RMSSE}}(s) = \sqrt{\sum_{f\in F}\sum_{c\in C}\sum_{t=1}^{N^t}\frac{\left(f^{\text{unc}}(f,c,t,s) - b(f,c,t,s)\right)^2}{\left(\hat{f}^{\text{unc}}(f,c,t,s) - b(f,c,t,s)\right)^2}} \qquad \forall\ s = 1, ..., N^s \tag{4.19}$$

In order to directly compare the accuracy of two competing forecasts, Diebold & Mariano (1995) suggests the use of an evaluation of a null hypothesis stating a lack of difference. The authors claim that their approach allows for a more diverse range of forecast errors. They formulate the hypothesis as the statement that the population mean of the loss-differential series of the two forecasts is 0. A summary of such tests can be found in Harvey (1997).

With regard to the variety of available forecast measures, some limitation is called for. When considering the idea of forecast measures as such rather than comparing two specified methods, rankings appear to make more sense than straight comparisons. Therefore, tests based on a hypothesis assuming equal forecast quality will be neglected from this point on.

One flaw inherent to scaled quality indicators becomes clear on further consideration. Whether forecasts are scaled to the performance of a naive approach or to the spread of actual values, in the case of small numbers, division by zero is a constant issue. Especially with regard to revenue management, however, this cannot be ignored: whether a forecast predicts 5 or 100 bookings for a situation in which 0 bookings occur can make all the difference.

The only indicators that avoid this trap are those that are not scaled. While for example RMSE does not adjust for scale, it also considers every error as it occurs with the gravity in which it occurs. Due to the lack of scaling, RMSE values cannot be compared over data sets of different proportions. A simulation could allow for the artificial stabilizing of market sizes: Three different methods can be evaluated based on the same customer model and therefore within the same order of magnitude.

4.2. Applied Forecast Performance Evaluation

Much research is concerned with the theory of evaluating forecast accuracy. However, there are also several reports of applied forecast evaluation in the area of revenue management. The topic is split between evaluations based on data from real-life systems and evaluations using a simulation environment. While real-life data necessarily includes the actual degree of complexity, it is also truncated and limited with regard to data collection methods. In contrast to this, simulation-based approaches offer more insight into customer decisions but are limited by implicit model assumptions.

W. M. Swan (1990) represents one of the first reports of forecast evaluations based on a simulation. In this case, the evaluation is focused on the *spill* (rejected customer requests) caused by capacity allocation decisions based on demand forecasts. The results of the analysis performed indicate that it is crucial – unconstraining is referred to as an "estimate of spill" in this study.

Zeni (2001a) describes the evaluation of a range of unconstraining methods using a simulation to constrain data and forecasting methods to unconstrain it again. The constrained data is compared to actual airline data sets. The text has already been summarized in 3.2, the findings are also summarized in Zeni (2001b). The author judges any approach preferrable to that of ignoring the truncation of data. He indicates that expectation maximization methods work best. Such findings are also confirmed by Weatherford (2000) and Weatherford & Pölt (2002).

Another research study concerned with evaluating unconstraining methods is presented in Ferguson et al. (2007). This study actually uses hotel data but relates the methods and tests considered to those applied to airline data in other studies. Estimation maximization methods and a newly proposed double exponential smoothing approach are compared under the assumption of strictly static demand.

Ratliffe (2008) considers the problem of unconstraining under the aspect of customers flexibly choosing from a range of flights offered. The results presented are based on a simulation study. Evaluation is performed by a combination of MAD and MAPE ranking, comparing multi-flight and single-flight estimation maximization methods. The conclusion is that overall demand volume can make a difference for the success of forecasting methods.

In Ryzin & McGill (2000), the success of a revenue management system without a systematic forecast is evaluated using a simulation system. A simulation framework for the evaluation of revenue management strategies is also presented in Abdelghany & Abdelghany (2007) and Abdelghany & Abdelghany (2008). With regard to the success of the introduction of a network-based forecast method, a performance analysis is presented in Rockmann & Alder (2009). In Pölt (1998), some thoughts on forecast methods and their evaluation are presented. P. P. Belobaba (1998) considers the same issue based on PODS simulations. Forecasting approaches to estimating customer behavior, with special regard to sell-up, are evaluated in C. Hopperstad (2007).

In Frank et al. (2008), a number of general principles for the development of simulations to evaluate revenue management systems is provided. Thoughts regarding the calibration of stochastic demand data for such simulations are offered in Kimms & Müller-Bungart (2007).

5. Research Gap and Opportunities

As has been shown, a body of research on the theory as well as some documentation on the actual application of forecast evaluation for revenue management is available. Research opportunities arise as a gap exists in the available literature. While simulation methods have been applied to estimate the success of revenue management strategies, potential of exclusive knowledge on the demand model included in a simulation has not been used extensively. Additionally, few simulations include flexible demand as opposed to demand streams conforming to the static assumptions included in many forecast methods.

Most theoretical evaluation methods focus on the calculation of forecast accuracy. This is usually regarded to be the difference between actual observations and predictions. In demand forecasting for revenue management, observations are indirectly influenced by predictions – bookings only manifest if optimized availabilities allow them. The optimization uses the forecast as input. New error measurements may be computed by not comparing predictions to bookings but instead comparing them to a suitable transformation of demand knowledge as it is available in a simulation environment.

As a systematic connection between observed values and predictions rarely is considered in theory, developments in error measurements over time are neglected. When predictions may influence observations, phenomena such as self-fulfilling prophecies can arise. Signs for such methodical flaws can be found in the indicator development. As a simulation enables the quick modeling of long-term effects, the consequences of systematic developments can also be analyzed.

The revenue consequences of the use of a forecast method are often used as indicators when the use of methods in the real world is examined. However, revenue may also be impacted by the interaction of forecast and optimization. It does not in fact offer information on how correct a forecast is but rather indicates how financially successful its use in a certain environment has proven to be. A decomposition of the revenue management system would offer new opportunities of comparing forecasts without necessarily considering the effect they have on revenue.

Using a transparent demand model in a simulation, traditional evaluation methods may be reconsidered. New approaches to forecasting that attempt to describe demand behavior can be evaluated by comparing parameters drawn from the forecast and the model.

From the research gap described, three further steps can be derived:

Decomposition: When the components of a system can be isolated, processes can be designed to evaluate them separately. This way, the performance of individual parts of the system can be analyzed, while the other parts are kept stable (*ceteris paribus*). Such a decomposition may be applied to a revenue management system or to an approach to forecasting. This way, the accuracy of a forecast and its consequences for revenue can be considered separately.

Simulation Environment: A simulation environment offers the opportunity of implementing a decomposed model of revenue management and of interchanging separate modules. In addition to parts of the revenue management system, the market as presented by a demand model can be influenced in a simulation. The demand model is transparent in such an environment and can be mined for analysis. Furthermore, the cycles of forecasting, optimization and customers booking tickets can be sped up to observe long-term effects.

Concept Application: Based on a decomposed model and knowledge about forecast and forecast evaluation methods, expectations toward the evaluation of forecasts can be stated. Simulation experiments can be designed and conducted to test them. From the results of these simulation experiments, insights toward forecast evaluation may be drawn.

Part II.

Solution Approach - Concept and Implementation

A revenue management system can be described as three processes: Demand is predicted in the forecast component, availabilities to maximize revenue are computed in the optimization component, and seats are allocated according to this optimization and possibly strategic goals in the inventory. Corresponding to the aspects described in Chapter 3, the process of forecasting may be divided up further.

Using a simulation system, the theoretical decomposition of revenue management can be realized under laboratory conditions. The evaluation of methods of revenue management in simulations has already been described in Section 4.2. The concept described in this chapter goes one step further by systematically introducing knowledge of the demand model (as is exclusive to the simulation environment) to the evaluation of forecasts.

As this thesis concentrates on the evaluation of demand forecasts, the decomposed view of optimization and inventory is not described in further detail here. However, it may be realized in a similar fashion.

6. Simulation for Decomposition and Evaluation of RM Systems

First, a concept for the overall evaluation of a revenue management system in a simulation environment is described in the chapter. Next, the forecast component as a whole as well as the aspects of predicting demand volume, unconstraining, and behavior aspects are considered.

6.1. Overall System View

The consideration of the performance of a complete revenue management system follows the ideas of decision theory as outlined in Granger & Pesaran (2000). This approach does not conclude whether any part of the system is a decisive factor in its success.

Traditionally, in order to evaluate a whole system, one looks at whether it does what it is supposed to do: Maximize revenue. When considering mathematical methods, this can be achieved by mathematically proving that given correct forecasts, under certain assumptions, an optimization algorithm determines inventory controls that yield maximal revenue. An example of this approach can be found in Mayer (1976).

However, the knowledge that *given correct forecasts, under certain assumptions*, a revenue management system works optimally is not helpful in practice. The assumptions used in mathematical proofs are often simplified – one example is the "low fare demand before high fare demand" rule described in Section 2.2. Forecast quality cannot be assumed to be perfect or even constant. Results based on those conditions do not necessarily apply to real-world systems.

A more practical view of evaluating a complete revenue management system is to look at its outcome. This may be done by implementing and testing it on a real-world market or under laboratory conditions using artificial demand.

The advantage of a real-world implementation obviously lies in the confrontation of the system with just the amount of complexity it is supposed to handle. At the same time, results from the real-world may be tainted by events or economic trends and therefore may not give clear information on the performance of a new method.

Artificial demand can be designed to include only a desired degree of volatility. Furthermore, the same demand may be used to test two or more methods. However, performance when confronted with simplified model may not allow for conclusions on real-world performance.

The results of a revenue management system are *bookings*. Beyond these bookings, aggregated by departure, fare class, and time of booking, monetary indicators are *yield* (the average fare paid) and *revenue* (the sum of fares paid by all booking customers). While the first goal of revenue management is to maximize revenue, indicators on productivity such as *seat load factors* can be helpful, too.

A simulation, in which detailed knowledge on demand behavior is available, offers the opportunity of generating some additional results. These, as will be described subsequently, pertain information on rejected customers as well as on dependent demand behavior. As each customer can be observed requesting tickets and and being accepted or turned down, the equivalents of *turn-down* or *click-stream data* as described in Section 3.2 can be derived. Therefore, in the simulation, one can easily find out how many requests were rejected and for which reasons. This may happen for two reasons, termed *spill* and *spoilage*.

If customer requests are rejected due to limited capacity, this is called spill. Ideally, the phenomenon should apply to low-fare demand, which is rejected in favor of customers with a higher willingness to pay. If customer requests are rejected in order to reserve seats for more valuable demand that fails to materialize, this is called spoilage. Any rejection can be categorized as one or the other. While spill can rarely be avoided, given that aircrafts do have limited capacity, spoilage is the consequence of imperfect revenue management.

Vertically dependent demand may be observed within a simulation in terms of buy-down and sell-up. As demand is artificially generated to include a cost function as well as maximum constraints, each individual passenger's maximum willingness to pay is known. The demand model may include flexible choice behavior letting each customer minimize costs by choosing the cheapest available booking class within the parameters of product preference. This leads to an effect also observed in practice: Customers often pay less than what they theoretically are willing to. As information on the acceptance of product restrictions and the acceptance of price is available, comparisons between what was booked and what might have been booked are feasible in a simulation. Each booking can be placed in a continuum of up-sell (the passenger was forced to pay more than the lowest price for the requested product) and buy-down (the passenger's maximum willingness to pay was not fully exploited).

Horizontally dependent demand, customers choosing between available itineraries, can be documented as *recapture*. Recapture can be observed when a customer, on not being able to book his first choice due to limited availabilities, books a seat on a different itinerary rather than not booking at all. In a demand model with flexible customers, measuring recapture is not trivial. When many factors play a role in customer choice, it can be difficult to decide what the originally desired and what the alternatively accepted itinerary was. One solution is to define "first choice" as the itinerary a customer would have chosen if the all prices were the same.

Another feature of a simulation is the possibility of repeating processes while keeping individual components (such as the demand model) stable. This way, long-term consequences of the use of methods can be analyzed. In practice, economic trends as well as changes in the market situation (i.e. the entry of competitors) may distort the impact of a new method. In a simulation, artificial demand can be used repeatedly and methods can be tested *ceteris paribus*. This way, for example, the consequences of using historical data as generated using a forecast method can be observed by comparing the results over time. If the bookings, revenue, spill, spoilage, booking behaviors or recapture change, this will not be due to changes in customer behavior.

The scheme of how to compare such a system's results is shown in Figure 6.1. Step (1) is to generate artificial demand for as many periods as the simulation aims at modeling. Some of this data will be used for history building in step (2). In the case depicted, this includes periods 1 to $t-1$. Requests and knowledge about the demand model are used to provide the basis for the forecasts to be evaluated. This is realized by creating inventory controls using an optimization based on forecasts derived from the knowledge of artificial

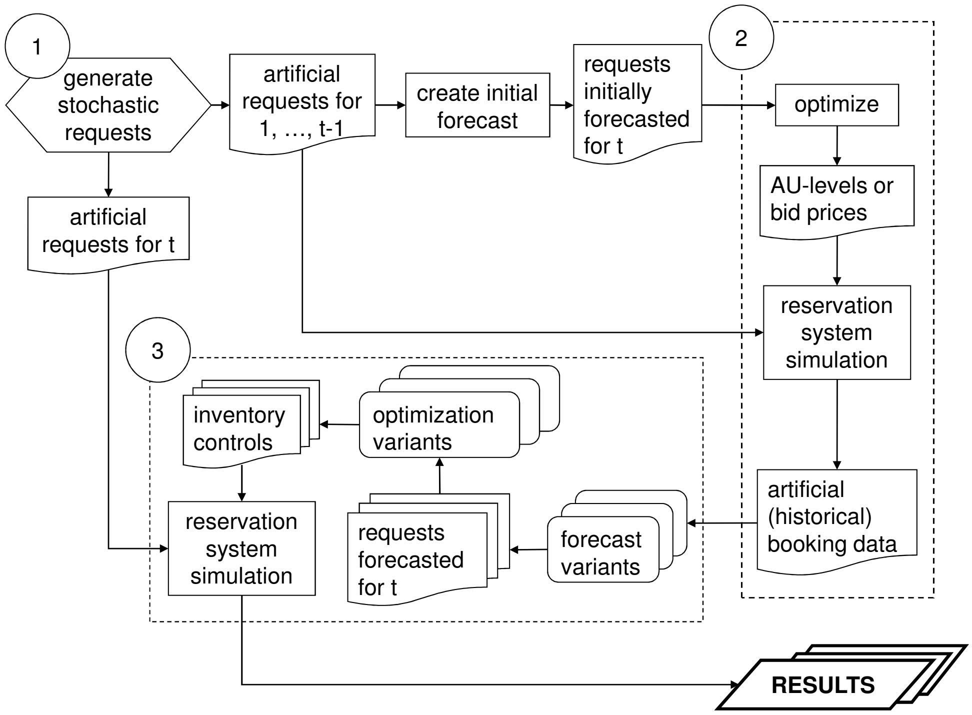

Figure 6.1.: Evaluating the RM System

requests. These inventory controls are then used to turn the requests into booking data, constraining them.

The historical booking data created in step (2) presents the basis for further calculations in step (3). Now different forecast methods come into play. They all share the same data basis – historical bookings from step (2). As the simulation can be repeated using the same demand while employing different methods, each method can be tested. The results of the forecasts are each handed to the same optimization method to generate inventory controls. In a reservation system, each of the different sets of inventory controls is used to channel the requests generated for period t. The results of this process can be compared over the variety of methods evaluated.

In the set-up illustrated, only forecast performance over the course of one period, t, is actually measured. By generating historical booking data for periods 1 to $t - n$ and

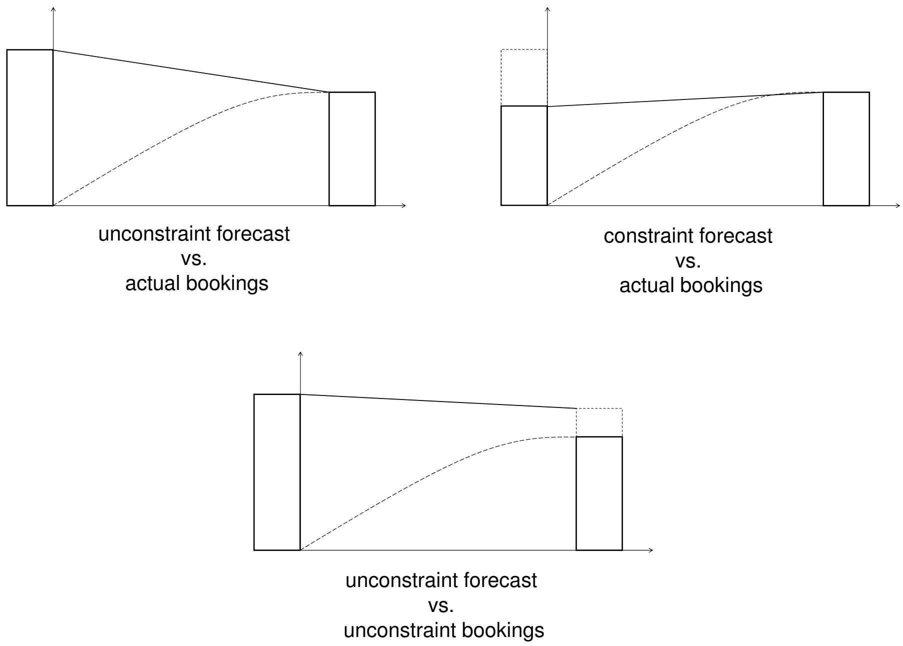

Figure 6.2.: Comparing Forecasts and Bookings

using the compared methods for optimizing availabilities during periods $t - n + 1$ to t, any number n of periods can be included. This way, long-term evaluations are feasible.

6.2. Forecasting Component

The quest for evaluating the performance of the forecast component in a revenue management system is not new. Some research conducted with regard to it has been introduced in Chapter 4. Most attempts at evaluating the quality of a forecast that do not consider the complete system such as described in Section 6.1 use one of three available options presented in Figure 6.2.

The straight-forward approach is to compare predicted demand to the actual bookings. However, these can be quite different quantities as bookings are constrained. Therefore, it seems sensible to transform either the forecast or the bookings in order to make a comparison more meaningful. In order to do this, one either compares the results of the forecast with unconstrained actual bookings, or constrained forecast results to actual bookings.

These comparisons, however, do contain one major bias. The same method applied to unconstrain historical booking data in order to build the forecast is applied to transform actual bookings for the comparison. When constraining the forecast, the inventory controls that were applied according to an optimization based on the forecasts are put into place. Any forecast evaluation that is conducted like this will either include the unconstraining method of the forecast or the inventory controls based on it.

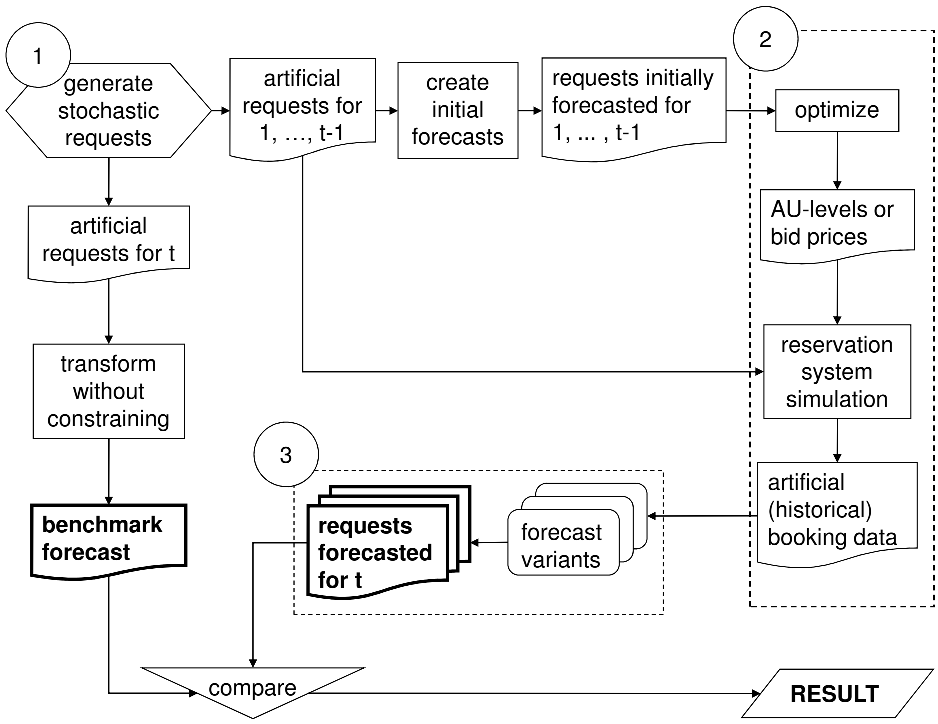

Figure 6.3.: Evaluating the Forecast Component

An approach based on a simulation system as presented in Figure 6.3 may avoid these catches. Once more, artificial demand is generated in step (1). Requests are generated for a time-line from 1 to $t-1$, to ensure a test of the forecast's ability to pick up trends and seasonality. Requests for periods 1 to $t-1$ are used to build a history of bookings for the forecast methods to be evaluated in step (2). In step (3), the different forecast methods are tested by providing the same data basis in the form of historical bookings and for the task of forecasting unconstrained bookings for period t.

Finally, the resulting forecasts can be evaluated by comparing them to the actual requests as generated for period t. The forecast that did best in recognizing the underlying patterns and unconstraining sales data should yield the result with the smallest deviation to the actual requests. As different approaches to forecasting may offer different forms of information, the requests need to be transformed in order to make a comparison feasible. This means that a kind of *perfect or psychic forecast* needs to be created, employing the same information and data format as the evaluated forecast method.

Of course, aspects of the process generating a psychic forecast may influence the test results. However, a consistent approach in transforming requests, such as the general principle of always exploiting maximum willingness to pay, can still enable a consistent comparison of diverse methods.

The comparison can follow standard statistical procedures. This includes the calculation of key indicators as described in Chapter 4.

6.3. Demand Volume Aspect

Forecasting for revenue management consists of two linked tasks. Historical booking information needs to be transformed in order to learn about past demand (*unconstraining*). Historical patterns need to be recognized and extrapolated in order to provide predictions toward future demand (*time series aspect*). When evaluating forecasts, it can make sense to separately consider a method's ability to fulfill both tasks. This may provide insight to build better approaches by combining existing concepts. While this section describes a simulation-based evaluation of the time series aspect, the next section will consider how the unconstraining aspect may be isolated and analyzed.

While unconstraining seems to be connected to the special nature of revenue management, the consideration of time series is a problem common to all kinds of forecasting. Once more, three aspects arise: Demand seasonality, long-term trends, and arrival timing.

Over long or short terms, developments depending on time can include recurring patterns. The fluctuations based on these patterns are called *seasonality*. Apart from actual seasons such as spring or summer, demand may also fluctuate periodically with regard to the day of week or the time of day. By recognizing seasonality and accurately predicting it, a forecast can use the knowledge of the time a flight departs to estimate expected demand.

While economic cycles may be considered as the outcome of a seasonality spreading over the course of several years, it is more common to view them as a separate aspect, called the trend. Trends provide information on the general development of demand volume over several seasons. Rather than assuming that the second week day of the tenth week of every year will see the exact same demand, different overall demand level for the current situation based on recent developments is predicted. Seasonality is helpful to predict patterns within this development.

Last but not least, demand arrival within the sales period also follows certain patterns that may depend on the customers requesting demand. With regard to the timing of a departure, demand arrival for the different booking classes or fares offered needs to be predicted in order to allow for a successful optimization. This, too, is time series forecasting.

Existing approaches to measuring a method's ability to predict time series tend to do so by providing a set of data that is not constrained. In revenue management, that means relying on historical bookings that occurred in classes that were always available. In practice, such data can only be the result of consistently low demand. Basing a forecast evaluation on it means basing it on a special demand situation – one cannot say how the forecast would perform given high levels of demand or high fluctuations in demand.

A simulation approach to evaluating the time series aspect can provide booking data that has not been constrained by inventory controls. In such a model, capacities can be set to infinity. In order to not turn even one customer away, capacity may be neglected entirely: By closing booking classes, sales constrained and no longer represent the complete demand. Without fixed capacities, unconstraining is necessary no more. Demand, no matter how high or low and no matter how volatile, directly translates into bookings that can be used as a basis for forecasts.

Based on this idea, requests are generated for departures $t = 1..T$ in step (1) of Figure 6.4. The requests for $1..T-1$ are turned into historical bookings by assigning them to

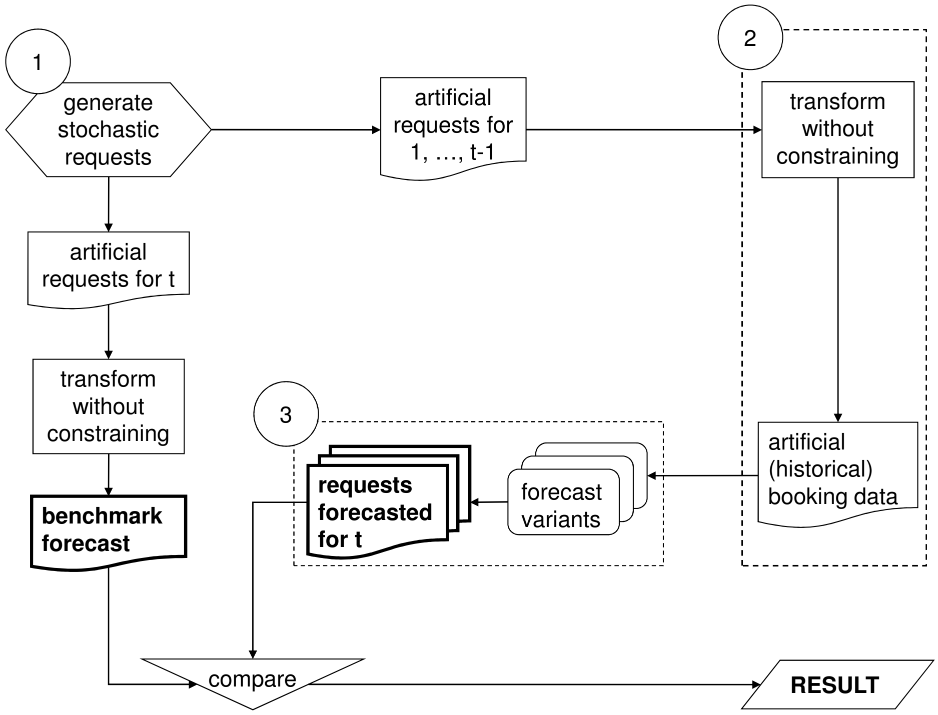

Figure 6.4.: Evaluating the Trend Component

booking classes without constraining them in step (2). Forecast methods are provided historical data and predict demand for the departure $t = T$ in step (3). The forecast is by design unconstrained and can therefore be easily compared to the requests generated. Any flaws in the prediction are due to demand fluctuations caused by seasonality or trends.

As with traditional methods, the three aspects of revenue management time series may be heeded by different indicators used in the error measurements. A forecasts ability to pick up trends can be observed by comparing the overall predicted demand to the overall amount of requests. Seasonality patterns included in forecasts may be evaluated by forecasting and evaluating requests for a number of departures $t = T - n..T$ rather than for just one departure. The correct identification of arrival patterns can be checked by comparing the predicted demand arrival rates for a departure to the demand arrival rates inherent to the original requests.

This section explained ways of evaluating the time series aspect alone based on the simulation concept. Such is also the focus of much research on forecasts and forecast evaluation as presented in Armstrong (2001). The further text will focus mainly on evaluating the performance of unconstraining as well as conclusions toward dependent demand behavior. This allows for an examination of problems particular to revenue management forecasting.

6.4. Unconstraining Aspect

Traditionally, a forecast's ability to deduct demand from sales data has been evaluated using historical data with as little overall fluctuations as possible. For example, seasonality may be filtered including only departures within one season, day of the week and time slot in the analysis. Trend may be excluded by considering a very uniform market with little fluctuations. By separately considering bookings for time slices before departure, the necessity of anticipating arrival rates can be neglected. All this makes the evaluation of a method possible only under very specific circumstances. The challenges a forecast faces for example when faced with a market with volatile demand are not considered.

A new approach to evaluating the unconstraining component is illustrated in Figure 6.5. In a simulation, requests can be generated repeatedly (including a normally distributed distortion) for only one point of time, t, in step (1). They can be turned into bookings by an optimization based on an initial forecast in step (2). With a history of bookings for one departure, the forecast methods tested only have to unconstrain the sales data in order to predict demand in step (3).

Unconstraining is only ever feasible within the range of available data. This means that if a class was never open, no method will be able to guess whether there was demand in a market for this class and if so, how much. A perfect forecast is likely to only yield one kind of (restrictive) inventory controls. Therefore, depending on how the demand model is calibrated, none of the applied methods are likely to correctly estimate all demand. They can, however, be compared with one another given the same degree of information. In order to further analyze methods, they might be compared based on different approaches to initial optimization. For example, a first-come-first-served optimization without a

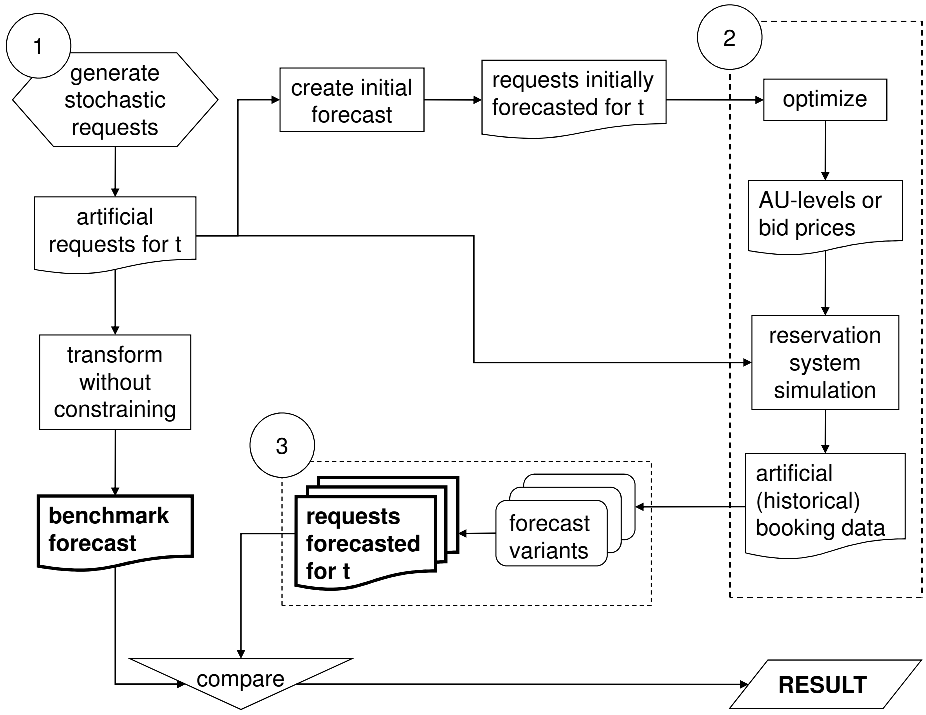

Figure 6.5.: Evaluating the Unconstraining Component

forecast may be offered in step (2). The forecast outcomes given the data basis may yield information about its reliability.

Diverse long-term effects are also conceivable. If, for example, overall demand is underestimated on the long run, this can lead to a spiral-down effect: Capacity could be reserved for expected high-value demand is allocated and sold at low fares. Customers with a high willingness to pay may still book cheap tickets if those are available. As a result, even fewer bookings are recorded in the high-value classes and even lower forecasts are calculated. Such effects may be analyzed by letting a simulation run repeatedly on the same, flexible customer model. If the customer behavior is stable, no changes in revenue outcomes and bookings should occur. Any trend then points toward systematic flaw in the uncontraining component.

Systematic flaws may be inherent to a forecast method. For example, as already mentioned in Section 3, Internet usage has lead to a shift in customer behavior. Customers book flexibly and bookings in one class or even on one flight or itinerary may be dependent on current availabilities. To derive informations on customer behavior from historical bookings and possibly additional information has become a new task of forecasting. The next section offers some ideas of how to evaluate the success in this regard.

6.5. Demand Behavior Aspects

The idea of dependent customer behavior as opposed to a model in which customers statically demand tickets in one booking class is still relatively new. Customer choice may depend on the lowest price available, flight or itinerary alternatives, or even product preferences as opposed to price sensitivity. In any case, availabilities have an influence on what customers book, and no availability in one specific booking class and for one specific flight does not automatically mean that a customer will not book anything else. Literature describing approaches to incorporating this idea into revenue management has been listed in Section 3.3.

Traditionally, models incorporating dependent demand have been evaluated by comparing them to concepts relying on static demand models. In practice, expected revenue improvements are calculated based on expected up-sell. The accuracy of a method's description of a customer model as derived from historical bookings, though, has rarely been emphasized. For example, when considering classical customer choice models such as introduced in Ben-Akiva & Lerman (1985), the assumption is that customer surveys establish reasonable estimates of customer choice factors.

A simulation with a demand model including customer choice offers the possibility of further analysis. By separately setting up the request generation including a cost function as well as price and product restrictions, mixed situations can be generated and the results of forecast methods applied can be compared. The actual factors of the cost function are known in the simulation.

The comparisons of forecasts including dependent aspects within a simulation may be conducted on any of the levels described in this section. Important is the consideration that the inclusion of a dependency assumption all by itself may not automatically lead to more accurate forecasts or even higher revenues. Once more, a simulation offers the opportunity of testing the concept within a freely manageable environment and on a stable data basis.

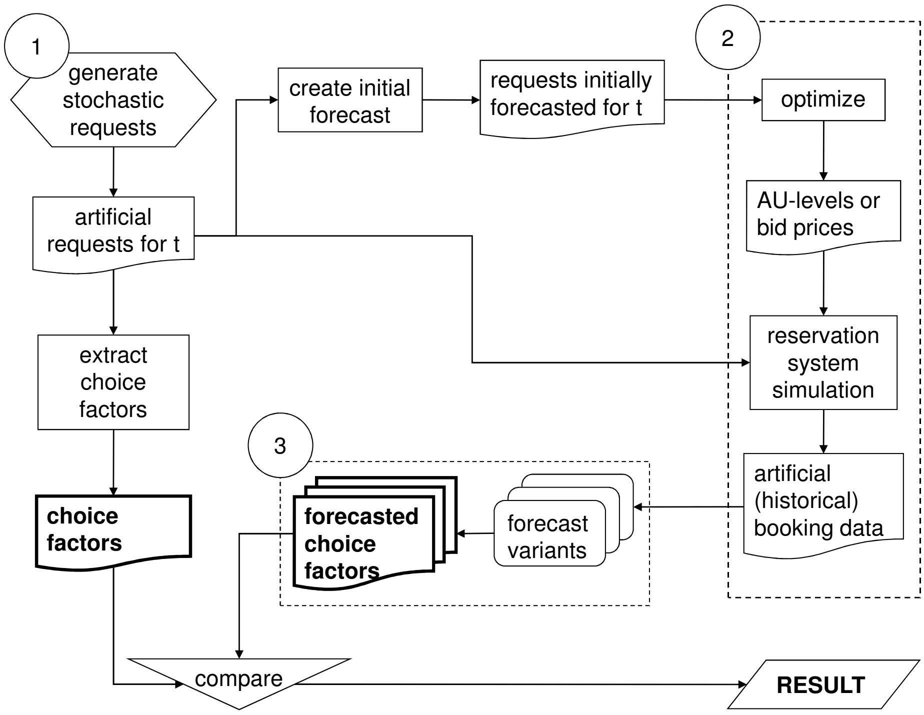

Figure 6.6.: Evaluating the Choice Component

As Figure 6.6 shows, the evaluation of the choice component differs from the evaluation of the unconstraining component only in so far as that not merely overall demand volumes are compared. Instead, the choice factors predicted by the forecast are extracted from the demand model and used for the evaluation.

7. Simulation Environment for Revenue Management

This section describes the revenue management simulation implemented and applied for the consideration of demand forecast evaluations. The simulation system in place can be divided into three major components. One is the simulation control, functioning as a framework that controls the event-based timing of a simulation as well as the management of information and the triggering of reporting functionalities. The choice-based demand model is a complex component necessary to make the simulation realistic. Finally, the part of the simulation that is modeled after an actual revenue management system: modules for forecasting, optimization, and inventory control. The components interact as simulation control triggers customer requests, forecasts and optimization, customer requests interact with the inventory control to create bookings, the results of this interaction are saved and processed, new forecasts and optimized seat allocations are calculated based on historical bookings.

Figure 7.1 shows the simulation system in the context of set-up and analysis. Supply and demand are prepared from schedule, fares and input parameters. The system transforms this data into result indicators such as bookings and revenue according to parameters such as required confidence intervals. The analysis component puts data into context, generating information that conclusions may be drawn from.

In Section 7.1, the details of the simulation control as well as the data and parameters involved are described. Subsequently, a description of the supply and demand model used in the simulation is provided in Section 7.2. Finally, in Section 7.3, the revenue management methods implemented are outlined.

Setting up a simulation system in order to evaluate forecasts, the evaluation approach of replicating outputs as described in Armstrong (2001) is included to some extent. The importance of a sufficiently large population of predictions and observations to measure errors on is pointed out in Armstrong & Collopy (1992). In the case of a simulation

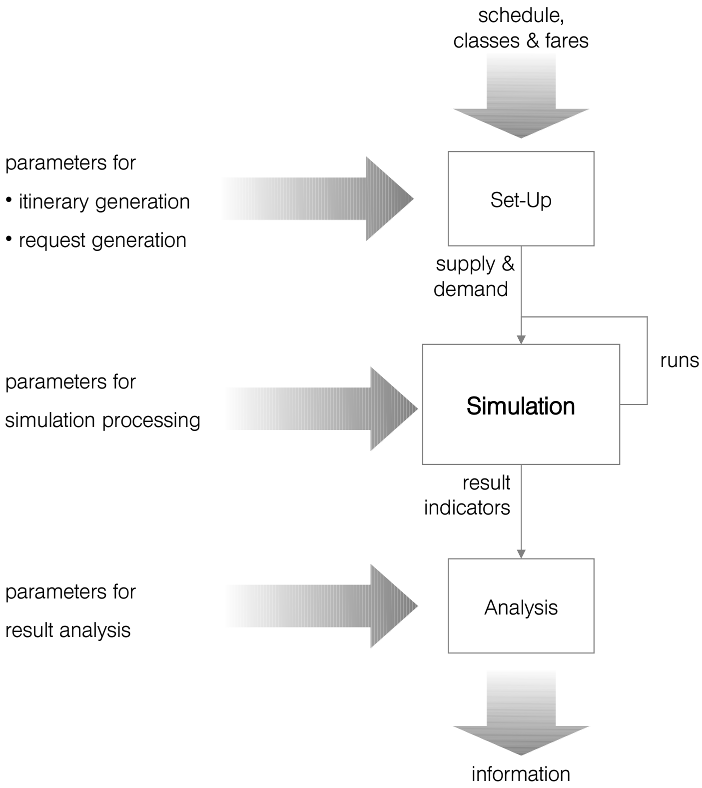

Figure 7.1.: The Simulation Cycle

system, this pertains the number of runs evaluated as well as the number of scenarios set up. Both are within the influence of the researcher but constrained by limits in time and effort. Evaluating a method by comparing its output to actual knowledge about what is to be predicted may yield misleading results as stochastic elements can lead to lucky flukes. However, repeating the stochastic process over the course of several simulation runs, such outliers can be neutralized: Results that are averaged over a number of replications differing only in stochastic error terms tend to be more significant.

The system presented here follows the theory of stochastic simulation as laid down in Law & Kelton (1997). In addition, the guidelines offered in Frank et al. (2008) are considered.

7.1. Simulation Control

Contemplating the design of a revenue management simulation system, several aspects can be identified. Each simulation experiment needs to be prepared: Data has to be provided and the nature of some processes as well as the reporting required need to be specified. An experiment consists of several simulation runs; these have to be initialized and from one run to the other, data needs to be updated or reset. Within each simulation run, points of time within a booking horizon pass and events need to be handled as customers request and book tickets. Finally, the simulation results need to be stored and processed.

7.1.1. Data Management

All data is stored as lists of records, with each record having a set of properties distinguishing it. The following overview outlines the information stored for each item; the algorithms used to process or create the data will be described in more detail in the next sections.

At the start of the simulation, supply and demand data is read into the system cache for fast access. As the simulation proceeds, simulation control provides access to the required data. After a simulation experiment has finished, reporting information stored in the result records is processed and written out (see Section 7.1.3).

As illustrated by Figure 7.1, data from outside the system enters the simulation at several moments. This data can be split into different sets as shown in the following list.

- Schedule, booking classes and fares: This information can be taken from the real world and may be based on the actual schedule as well as a selection of the actual fares
- Parameters for itinerary generation: This input describes how combinations of origin and destination and itineraries connecting the two are generated from the schedule.
- Parameters for request generation: This input describes how demand is to be generated for the simulation.
- Parameters for simulation processing: These parameters describe how the simulation is to be executed.

List	Properties	Description
Airports	*name, latitude, longitude, traffic area*	Airports connected by flights included in the schedule used in the simulation.
Flights	*carrier, departure airport, arrival airport, flight number, departure time, duration, days of operation, capacity*	Flights from one airport to another.
Legs	*origin, destination, traffic area, distance*	Connections; at least one flight between the two airports is needed to justify a leg, but one leg can include several flights at different times of day.
Booking Classes	*name, carrier, traffic area, one boolean value describing whether each of the included product characteristics applies*	Classes offered in the simulation.
Fares	*carrier, class name, pairing, price*	Prices of tickets in one class on one flight.
Pairings	*origin, destination, traffic area, request share, customer mix*	Origin and destination combination offered to customers.
Itineraries	*origin, destination, included flights, travel time*	Travel itineraries linking two airports in a pairing.

Table 7.1.: Simulation Environment: Supply Lists

List	Properties	Description
Customer Types	*name, request distribution, departure time distributions, cost function, willingness to pay, maximum accepted travel time, accepted deviation from the preferred departure time, a boolean value describing whether each of the included product characteristics is accepted by this customer, error term used for the individual distortion of requests*	Input for demand generation: Types of customers that may request tickets throughout the simulation.
Requests	*run, time of request, preferred departure time, pairing requested, actual cost function, actual price and product preferences*	Output of demand generation: Customers that request tickets during the simulation.

Table 7.2.: Simulation Environment: Demand Lists

- Parameters for result analysis: These parameters describe how the basic result indicators generated in the simulation can be refined and put into context.

The following parameters are required to prepare the generation of pairings of origin and destinations and itineraries connecting them. They define which connections are acceptable for travel and will become part of supply:

- minimum and maximum connection time,
- minimum and maximum travel distance,
- maximum alternative itineraries per pairing,
- maximum transfers.

The data input needed to set up a supply scenario and the results of the connection builder process are listed in Table 7.1.

Table 7.2 describes the customer types needed for setting up demand generation and the requests that are its result. To assign demand to markets, a customer mix needs to be provided as input parameter, defining a distribution over pairings and customer types.

Parameters influencing the analysis may be based on the motivation of the simulation experiment. Output of the simulation experiments will be described in Section 7.1.3.

7.1.2. Simulation Runs and Lists of Events

Apart from providing data access and reporting facilities, two tasks of simulation control remain. On the one hand, between simulation runs, preparatory measures and the evaluation of stochastic confidence need to be controlled. On the other hand, within each run, a list of events needs to be processed.

Between Simulation Runs

The *simulation* includes a number of repetitions (*runs*) as defined during set-up. As described in Law & Kelton (1997), a certain number of runs is necessary to make the outcome statistically significant. Within the simulation, any number of modules communicate with each other and access the data, and any number of processes can be triggered. Depending on the requirements as declared during set-up or hard-coded into the system, information on the simulation process is written out. Before the first run, an initialization needs to be performed. Before every successive run, data from previous runs needs to be processed and made available.

The initial run of any simulation experiment lacks historical information from preceding runs. If a forecast requiring such information is included in the model, a substitute needs to be provided in the first run. The task of simulation control is to recognize to first run of an experiment and to generate an initial forecast that may be methodically different from the forecasts in subsequent runs. Alternatively, before the first run, historical data might be imported from an external data set and provided for the forecast.

After the initial run, data created in the preceding run may be used in the upcoming run. This can be the case if the updating of the forecast method applied in the simulation experiment depends on historical bookings. The task of simulation control is then to store the booking data, reset the inventory, and trigger a forecast update at the beginning of the new run. After every run, some amount of information on the interaction of the demand model and revenue management processes needs to be permanently stored – this is the task of result reporting as documented in Section 7.1.3. Variables or data structures intended to hold only the data created within one run need to be reset.

The overall number of runs N^s included in one simulation experiment may be set before the simulation starts. Alternatively, it can be set to depend on a defined confidence interval.

The convergence of the sum of bookings and the sum of revenue toward the true mean is tested after every run based on a confidence interval. This interval is determined via the *student (t) distribution* and parameters, α and δ. Inequality (7.1) is used to test whether the probability that the expected value does not differ from the true mean by more than a percentage δ is equal to or lesser than α.

To describe the test for confidence, the following notation is needed.

- Let s be number of runs already processed.
- Let $\overline{x}$ be the average of the indicator (sum of bookings or revenue) over n.
- Let σ be the deviation of $\overline{x}$.
- Let α be the acceptable probability of error.
- Let δ be the slice of $\overline{x}$ that is acceptable as a confidence interval.
- Let $\hat{N}^s$ be the number of runs recommended to reach the confidence interval.

$$\left| t_{n-1,1-\frac{\alpha}{2}} \cdot \frac{\sigma}{\sqrt{n}} \right| \leq |\delta \cdot \overline{x}| \tag{7.1}$$

If Inequality (7.1) is not fulfilled within the set maximum number of overall runs N^s, the number of runs recommended to reach the confidence interval, $\hat{N}^s$ is calculated for informational purposes as shown in Definition (7.2). If deemed reasonable, the experiment can be repeated with the required number of runs.

$$\hat{n} := \frac{\left(t_{n-1,1-\frac{\alpha}{2}} \cdot \sigma\right)^2}{(\delta \cdot \overline{x})^2} \tag{7.2}$$

When the indicators can be expected to develop systematically during a simulation experiment, testing for a confidence interval is less useful. However, in this case, a number

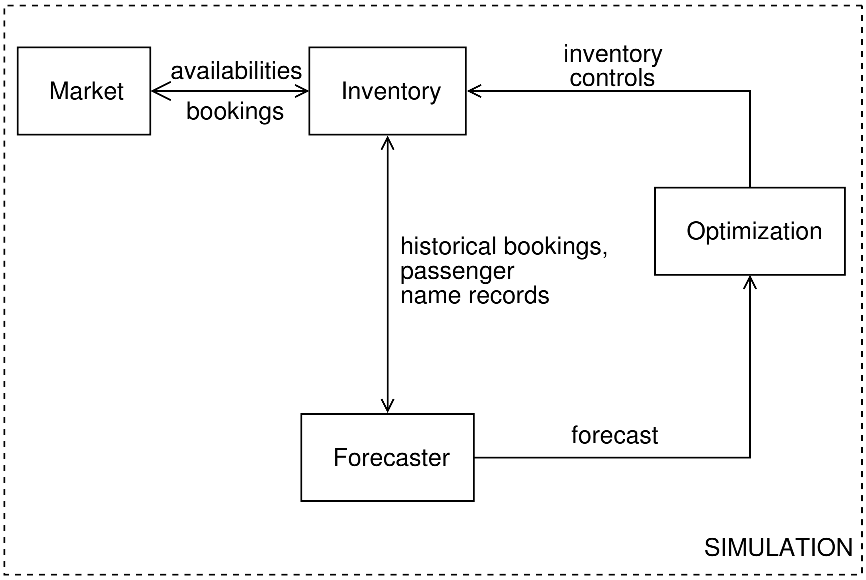

Figure 7.2.: Revenue Management Simulation

of initial runs may be exempted from the test (and the calculation of $\overline{x}$). The system can then be tested for stability under an expected dynamic behavior.

Within Simulation Runs

A revenue management simulation can be set up to answer a multitude of questions. Regardless of its purpose, certain modules seem to be compulsory and are depicted in Figure 7.2. These modules interact and thereby create or update data to be analyzed later.

Within each run, simulation control has to keep track of time and trigger various processes. This happens based on *events*. An event may be the arrival of a customer request, the end of the booking horizon and the departure of a flight, or the necessity for an update of forecast and optimization within the booking horizon. Event-controlled simulation allows for an efficient way of implementation: The alternative would be to account for each minuscule time slice within the booking horizon to test whether any action is necessary.

As a run begins, simulation control first triggers the update or the initialization of the forecast module. Based on this, inventory allocation is optimized for each flight included in the experiment. The list of customer requests is processed in the order of their arrival. Each request is scheduled for a specific time in the booking horizon. If an update was scheduled for a point of time that lies between two subsequent request, this update is triggered by the simulation control before the next request is processed.

Data is generated as the revenue management system starts working. This includes forecasts, optimization results, availabilities, and bookings as well as possibly details on customer choice behavior shown in the course of the booking process. Simulation control keeps track of this data as it is updated within the run.

The simulation includes the basic parts needed for revenue management: A forecast determines demand to come and thereby provides an objective function for revenue to be maximized by the optimization. That, in turn, yields a set of inventory controls intended to result in maximum revenue. The algorithms and the data used in these processes depend on the methods included as documented in Section 7.3. The resulting availabilities are managed by an inventory system. The inventory acts as the interface to the market.

A simulation could also be set up to examine the effects of competition with minimal effort. For this, several parallel revenue management systems would be implemented. This means that several inventories based on separate forecast and optimization modules are kept up. The required data can still be stored in the same component but is distributed according to airline. All airline inventories interface with the same demand model, the market. Keeping track of different methods and therefore modules assigned to different airlines is also within the responsibility of simulation control.

7.1.3. Reporting

Once a simulation experiment has been concluded, a broad data basis can be analyzed. This information may be used to evaluate the general performance of a system set-up, test the functioning of an implemented simulation, or even calculate error measurements considering the level of accuracy of a given forecast. The information needed depends on the specification of the simulation experiment considered. The data that has to be collected and processed in order to analyze the result of simulation experiments is presented in Table 7.3.

Bookings:	Per run, flight, class, point of time before departure.
Revenue:	Per run, flight, class, point of time before departure.
Yield:	Per run, flight, class, point of time before departure.
Available seats:	Per run, flight, class, point of time before departure.
Forecast Error:	Per run and point of time before departure, aggregated over flight, class, time before departure.
Buy-Down and Sell-Up:	Per customer request, aggregated over run and time before departure.
Denied and Accepted Requests:	Per customer request, aggregated over run and time before departure.

Table 7.3.: Output of Simulation Experiments

7.2. Supply and Demand Data

The goal of revenue management is to sell a perishable product at the right price. In the airline industry, the product is a seat on a flight differentiated by characteristics including the price, restrictions and features such as comfort and flexibility. Customers choose what to buy based on their preference. When the flight departs, the seats on it lose their value.

In order to implement a realistic revenue management simulation, both product (*supply*) and customer (*demand*) need to be defined in detail. Offering routes through a network as well as a range of booking classes tied to diverse restrictions and features confronts customers with alternative products. They choose from these according to a rational choice model including a cost function and a set of maximum values and boolean acceptance rules.

7.2.1. Supply Information

The *supply* of an airline consists of seats on flights, represented by tickets. The flights are described in the airline's schedule. The capacity of each flight is defined by the aircraft assigned during the fleet assignment process. By transferring from one flight to another, customers travel from origin to destination according to itineraries. These itineraries are calculated by booking engines provided by the airline itself or third parties. They are an integral part of a simulation that models flexible customers. Finally, tickets are sold for different booking classes. Each class defines a set of restrictions, features and the price of the ticket.

Flights and Itineraries

Flights are described by a departure airport, an arrival airport, and a departure date and time. They are further distinguished by a flight number and a carrier code. A combination of two airports connected by one or more flights is referred to as a *leg*. Based on airports as vertices and legs as edges, networks can be defined.

As described in Section 2, state of the art research tends to consider network rather than flight views. Especially with the advent of Internet booking portals, customers no longer book single flights but rather decide for or against itineraries leading them on a path through a network. While the decision of whether to include a network model in forecasting and optimization is a methodological one, a realistic simulation should model both customer choice behavior and the product range based on a network view.

In the revenue management simulation system presented, a network is created from the legs included during simulation set-up. In order to do so, two data structures are required. A list of *pairings* describes combinations of origin and destination linked by one or more legs. Further information that can be provided for each pairing includes the knowledge whether a direct flight from origin to destination is available, the traffic area and the geographical distance involved. For each pairing, a set of *itineraries* can be defined. These itineraries describe the legs that a customer might book tickets for in order to travel from the origin of a pairing to its destination. The complete process is presented in Figure 7.3; those parts of the model that enable the network view are marked bold.

The *connection builder* derives pairings and itineraries from a set of flights. For this purpose, a shortest path algorithm needs to be implemented. As indicated in Figure 7.3, a modified *Dijkstra* offering the n-shortest paths is used in this case. How many paths a customer may chose from and what itineraries are regarded as valid depends on a range of settings. For example, while it may make sense to travel from Munich to New York via Frankfurt, the customer should not be offered a trip from Munich to Frankfurt via

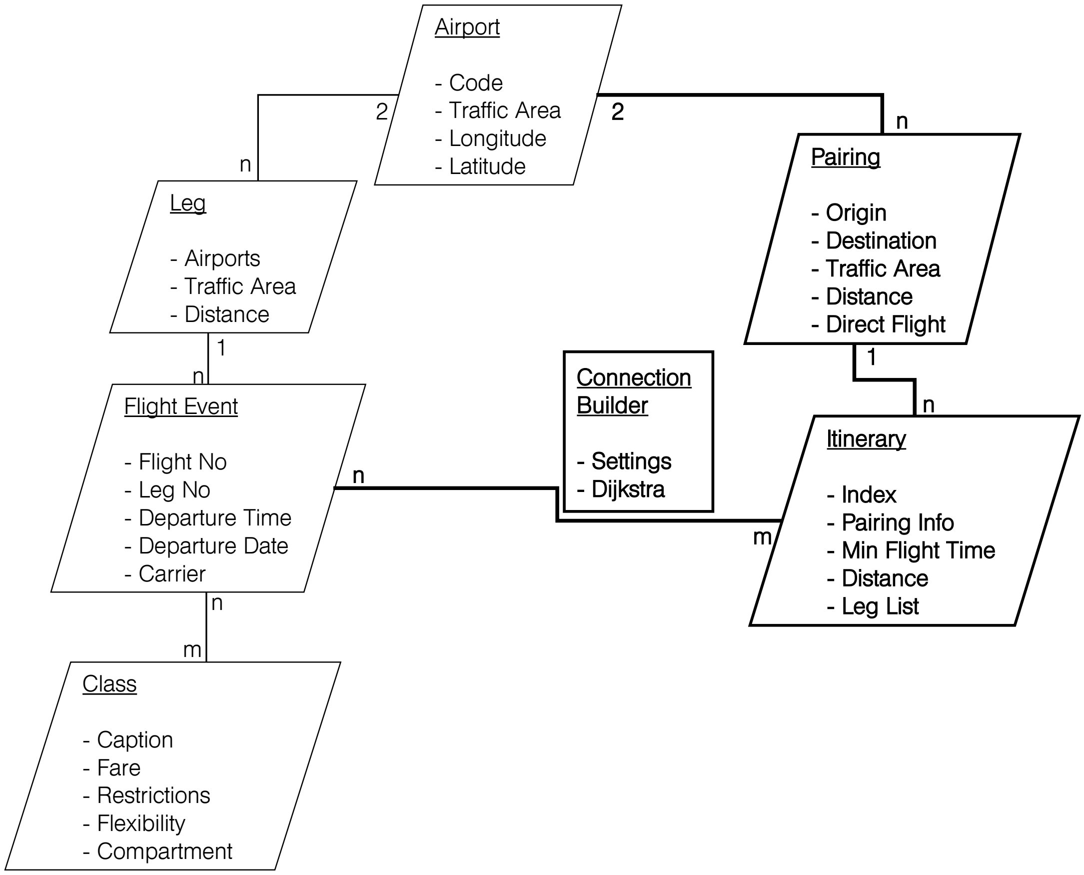

Figure 7.3.: Defining the Product

New York. In order to avoid such invalid itineraries, constraints on the overall duration of travel,the amount of transfers and connecting times are set by parameters.

The relationships of flights, pairings and itineraries are as presented in the following formulas.

- Let $f \in F$ be the flights included in the schedule.
- Let Θ represent the parameters defining maximum duration of travel, amount of transfers and connecting times.
- Let $q \in Q$ be the pairings that are possible based on a given set of airports.
- Let $\Delta(\hat{q}, F, \Theta)$ be the Dijkstra function defining the itineraries connecting origin and destination of pairing $\hat{q}$ based on a set of flights F and a set of parameters Θ.

- Let $i \in I$ be the itineraries derived from the schedule and the parameters.
- Let I_q be the set of itineraries derived from the schedule and the parameters connecting origin and destination of pairing q.
- Let F_i be the set of flights included in itinerary i.

Function (7.3) shows the generation of itineraries from pairings. If no itineraries can be computed according to the parameters, the pairing is removed from the set of pairings. New itineraries that are found for a pairing are added to the set of itineraries.

$$\Delta\left(\hat{q}, F, \Theta\right) = \begin{cases} I_{\hat{q}} & \rightarrow I := I \cup I_{\hat{q}} \\ \{\emptyset\} & \rightarrow Q := Q \backslash \hat{q} \end{cases} \qquad \forall\ \hat{q} \in Q \tag{7.3}$$

According to this process, Inequality (7.4) states that every pairing included in the set of offered pairings needs to be associated to one or more itineraries.

$$|I_q| \geq 1\ \forall\ q \in Q \tag{7.4}$$

Equation (7.5) shows that the set of offered itineraries is made up by subsets of itineraries offered for each pairing.

$$I = \bigcup_{q \in Q} I_q \tag{7.5}$$

Classes and Fares

Tickets are categorized by booking classes characterized by a set of restrictions or features. The price of the ticket (*fare*) is defined by a function over the booking class and the chosen flight. In practice, fare classes present an additional differentiation, defining diverse tariffs for the same booking class. The simplified model implemented in the simulation assumes that every booking classes represents exactly one fare class.

Every booking class has distinguishing characteristics: A caption naming it, a set of restrictions, and a set of product features. Examples of possible restrictions are a weekend or a minimum stay. Example of features are special flexibility or the seating in the business compartment. A lack of features may also be modelled as a restriction.

Each flight is assigned a set of booking classes. It seems intuitive that flights spanning greater distances are connected to a higher fare than short flights, even if the same booking class is chosen. Furthermore, it common that the fare increases with increasing features and decreasing restrictions as presented in Inequality (7.6). Therefore, the function defining the price of a ticket takes into account the booking class and the flight's traffic area.

- Let $c \in C$ be the classes offered, ordered by decreasing features and increasing restrictions.
- Let $f \in F$ be the flights offered.
- Let $p(f, c)$ define the price of booking class c for flight f.

$$p(f, c) > p(f, c+1) \ \ \forall\ f \in F;\ \ c \in C \tag{7.6}$$

7.2.2. Demand Model

If there were no customers demanding tickets, revenue management would be senseless. Therefore, any revenue management simulation needs to include a demand model. However, the degree of sophistication of this artificial demand depends on the simulation requirements. Each part of the simulation component needs input and processes and stores updated or additional data.

In order to provide challenges for current forecasting methods, the customer model should include as few simplifying assumptions as possible. For example, customers should aspire to travel via a network from origin to destination rather than statically booking single flights. Furthermore, the customer model should allow for flexible decisions based on individual availability situations as they are possible in the age of comparison shopping.

Finally, while the demand model implemented here strives to include new ideas on flexible customer behavior, one caveat needs to be mentioned. As flexible customer choice was excluded by the assumption of independence in early demand forecasts for revenue management, other aspects of customer behavior may not be included in the model implemented here. Such implicit assumptions may lead to results that deviate from real-life observations, but cannot be avoided whenever a model of reality is designed based on current knowledge.

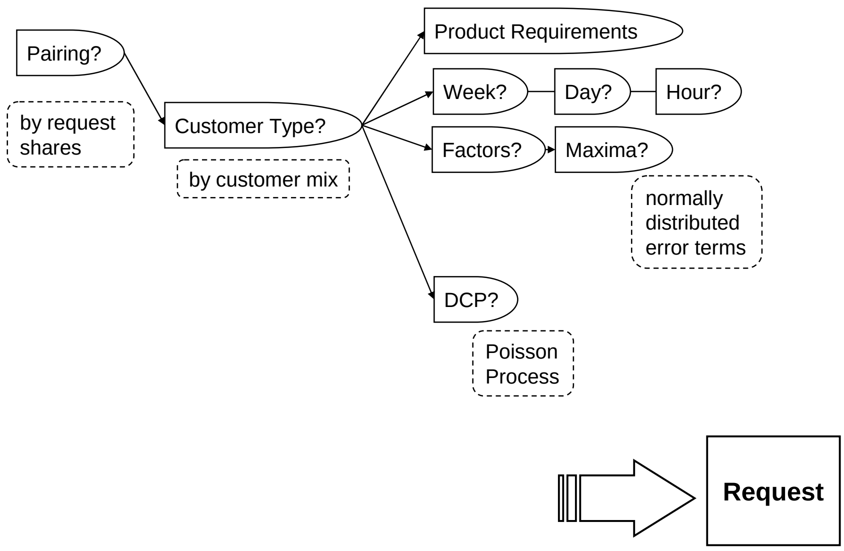

Figure 7.4.: Request Generation

Demand Volume and Arrival

The overall amount of requests is approximated within the set-up parameters. It is assigned as a percentage of the amount of seats offered based on the supply model. A demand volume setting of 150% will therefore result in about 1.5 times as many requests as there are seats. This is an approximation in two regards: Some of these requests can end up as bookings on more than one flight, and the actual demand volume is based on a *inhomogenous Poisson process* as described below. The demand volume is distorted by an error term for each run and used as an input for the intensity of a Poisson process. This avoids a too deterministic model of demand.

Definition (7.7) formalizes the distortion of overall demand volumes.

- Let $\overline{R}$ be the parameter defining the average amount of requests to be scheduled per simulation run of the experiment.

- Let ϵ_s be the error term of demand for simulation run s drawn from the normal distribution with expected value 0 and a deviation defined by an input parameter.
- Let R_s be the set of requests that is generated for simulation run s.

$$|R_s| := \overline{R} + \epsilon_s \; \forall \; s = 1, ..., N^s \tag{7.7}$$

Depending on the number of pairings included in the simulation, the share of requests allocated to each can be fixed manually or automatically. In an aggregated, automated approach, request shares depend on characteristics such as traffic area or the existence of direct flights.

Equation (7.8) presents the underlying constraints.

- Let $q \in Q$ be the pairings offered.
- Let $\gamma(q) \in [0, 1]$ be the share of overall requests allocated to pairing q.

$$\sum_{q \in Q} \gamma(q) = 1 \tag{7.8}$$

In order to allow for sufficiently complex patterns in customer arrival and behavior, the simulation is based on the concept of customer types. The *customer type* describes the factors applied for a choice of itinerary given a combination of origin and destination, the requirements of a booking class, a customer's price-sensitivity and the arrival distribution. The share of the requests assigned to one pairing that is connected to one customer type is determined by the pairing's *customer mix.* Again, this customer mix (a distribution over the existing customer types) may be assigned manually per pairing or automatically based on pairing characteristics.

Depending on the motivation for implementing a simulation system, the featured demand model can include different degrees of complexity. In order to test the workings of mathematical methods especially with regard to optimization, demand is frequently modeled based on stochastic distributions. Demand arrival in the sense of the implementation presented here is based on a Poisson process with separate parameters for each customer type.

Depending on the number of pairings included in the simulation, the request share can be fixed manually or automatically. In an aggregated, automated approach, pairings can be assigned shares depending on characteristics such as traffic area or the existence of direct flights.

Every pairing has a customer mix associated to it, which includes one or more customer types. Equation (7.9) presents the underlying constraints.

- Let $m \in M$ be the index of customer types.
- Let $\eta\left(q,m\right)) \in [0,1]$ be the share of requests allocated to pairing q that is to be based on customer type m.

$$\sum_{m \in M} \eta\left(q,m\right) = 1 \;\forall\; q \in Q \tag{7.9}$$

The absolute amount of requests based on one customer type for one pairing can be calculated from the request share and the customer mix as shown in Definition (7.10).

$$|R_s^{q,m}| := |R_s| \cdot \gamma\left(q\right) \cdot \eta\left(q,m\right) \tag{7.10}$$

Based on this, an inhomogeneous Poisson Process can be generated using the arrival distribution defined for the customer type.

- Let $\lambda_{q,m}$ be the overall intensity of the Poisson process for pairing q and customer type m throughout the booking horizon.
- Let $P\left[N\left(q,m,t+\tau\right) - N\left(q,m,t\right) = k\right]$ be the probability of k requests based on customer type m to arrive for pairing q in the time slice t to $t+\tau$ of the booking horizon of simulation run s.

This means that the Poisson Process $\left(X_{i,q,m}\right)_{i \in \mathbb{N}}$ is defined by $X_{i,q,m}$ is distributed according to $\exp(\lambda_{q,m})$. It defines $P\left[N\left(q,m,t+\tau\right) - N\left(q,m,t\right) = k\right]$.

When requests are generated from customer types, a normally distributed error term is added to the cost function. A function that computes requests from customer types therefore needs input variables as presented in Function (7.11).

- Let $|R_s^{q,m,N^t}|$ be the number of requests based on customer type m planned to arrive for pairing q within the booking horizon of simulation run s.
- Let σ_ϵ be the deviation of the normal distribution that error terms ϵ_r are drawn from.

- Let $\Psi\left(m, q, |R_s^{q,m,N^t}|, \epsilon_r\right)$ be the function that defines $|R_s^{q,m,N^t}|$ individual requests for pairing q based on customer types m and a given error deviation ϵ_r.

$$\begin{aligned} \Psi\left(m, q, |R_s^{q,m,N^t}|, \sigma_\epsilon\right) = R_s^{m,q,N^t} \\ \forall\ m \in M, q \in Q, s = 1, ..., N^s \end{aligned} \tag{7.11}$$

As shown in Equation (7.12), the overall set of requests generated for one simulation run can be divided into subsets of requests derived from specific customer types for specific pairings.

$$R_s = \bigcup_{m \in M} \bigcup_{q \in Q} R_s^{q,m,N^t}\ \forall\ s = 1, ..., N^s \tag{7.12}$$

Itinerary Choice

Given a departure day and a pairing, the choice of itinerary is based on the *cost function* and *product acceptance* stored in each request. A discussion of the decision factors follows. More factors are imaginable and could be implemented without much additional effort.

As this section formally describes the preference for one itinerary over the other given by each request, some notation is required. First, there are some additional features of itineraries and pairings:

- Let $q \in Q$ be the index of all pairings offered in the simulation.
- Let $\nu^{\text{dist}}(q)$ be the minimum distance between origin and destination airports of the pairing q.
- Let $\nu^{\text{dur}}(q)$ be the minimum travel time required by pairing q
- Let $i \in I$ be the index of all itineraries offered in the simulation.
- Let I_q be the set of all itineraries connecting origin and destination of pairing q.
- Let $I_{i,q}$ be a boolean matrix indicating whether itinerary i connects pairing q as shown in Definition (7.13).
- Let $x^{\text{dep}}(i)$ be the departure time of itinerary i.
- Let $x^{\text{dur}}(i)$ be the travel time attached to itinerary i.

- Let $x^{\text{trans}}(i)$ be the number of transfers in itinerary i.

$$I_{i,q} := \begin{cases} 1 & i \in I_q \\ 0 & i \notin I_q \end{cases} \forall\, i \in I;\ q \in Q \tag{7.13}$$

Next, the factors of the cost function and the product acceptance defined in the customer type:

- Let $\beta^{\text{dep}}(m)$ be the weight of the deviation from the preferred departure time in the cost function of the customer type m.
- Let $\beta^{\text{dur}}(m)$ be the weight of the difference between actual and minimum travel time in the cost function of the customer type m.
- Let $\beta^{\text{trans}}(m)$ be the weight of the number of transfers included in the chosen itinerary in the cost function of the customer type m.
- Let $\beta^{\text{car}}(m)$ be the cost factor attached to any itinerary that is not provided by the preferred carrier of customer type m.
- Let $\delta^{\text{dep}}(m)$ be the factor for maximum acceptable deviation from $w^{d}(r)$, defined by the customer type m.
- Let $\delta^{\text{dur}}(m)$ be the factor for maximum acceptable travel time, defined by the customer type m.
- Let σ_{ϵ} be the deviation of the normal distribution that the error terms are drawn from.

Finally, some more characteristics of requests:

- Let $r \in R$ be the index of requests included in the simulation.
- Let m_r be the customer type that request r was generated from.
- Let q_r be the pairing that request r was generated for.
- Let $w^{\text{dep}}(r)$ be the preferred departure time of request r.
- Let ϵ_r be the actual error term drawn from the normal distribution and attached to the cost function of request r.

- Let $\hat{C}(i,r)$ be the cost of itinerary i considered by request r, without regard for the actual ticket price (given the assumption that all itineraries cost the same).

With regard to the set of itineraries acceptable according to one request, note the relationship displayed in Definition (7.14). All sets of alternative itineraries are subsets of the quantity of itineraries offered. In order to belong to the set of acceptable itineraries of one request, the considered alternative needs to belong to the set of itineraries associated with the pairing the request is targeted at.

$$I_r \subset I_{q_r} \subset I \ \ \forall\, r \in R \tag{7.14}$$

Departure Deviation: The time of day that is the preferred departure time $w^{\text{dep}}(r)$ depends on a daily distribution from which the preferred hour of departure is drawn when the request is generated. The factor $\beta^{\text{dep}}(m_r)$ weights the difference between $w^{\text{dep}}(r)$ and the departure time of the itinerary considered, $x^{\text{dep}}(i)$, in the cost function. Only itineraries that fulfill the constraint shown in Inequality (7.15) are considered. This includes a maximum acceptable deviation from the preferred departure time that depends on a parameter $\delta^{\text{dep}}(m_r)$ as well as the distance covered by the pairing q_r, $\nu(q_r)$.

$$\begin{aligned} \left|w^{\text{dep}}(r) - x^{\text{dep}}(i)\right| &\leq \delta^{\text{dep}}(m_r) \cdot \sqrt{\nu(q_r)} \\ &\forall\, r \in R;\ i \in I_{q_r} \end{aligned} \tag{7.15}$$

Travel Time: The factor $\beta^{\text{dur}}(m_r)$ weights the difference between the minimum travel time $\nu^{\text{dur}}(q_r)$ connected to the pairing q_r requested and the travel time of the itinerary considered, $x^{\text{dur}}(i)$, in the cost function. Only itineraries that fulfill the constraint shown in Inequality (7.16) are considered. This constraint includes a maximum acceptable deviation from the minimum travel time that depends on a parameter $\delta^{\text{dur}}(m_r)$ as well as the distance covered by the pairing q_r, $\nu(q_r)$.

$$\begin{aligned} \left|\nu^{\text{dur}}(q_r) - x^{\text{dur}}(i)\right| &\leq \delta^{\text{dur}}(m_r) \cdot \sqrt{\nu(q_r)} \\ &\forall\, r\ \in R;\ i \in I_{q_r} \end{aligned} \tag{7.16}$$

Transfers: The factor $\beta^{\text{trans}}(m_r)$ weights the number of transfers included in the considered itinerary, $x^{\text{trans}}(i)$, in the cost function. Only itineraries that fulfill the constraint shown in Inequality (7.17) are considered; this constraint includes a maximum acceptable number of transfers that depends on a parameter $\delta^{\text{trans}}(m_r)$.

$$x^{\text{trans}}(i) \le \delta^{\text{trans}}(m_r) \qquad \forall\, r \in R;\ i \in I_{q_r} \tag{7.17}$$

Brand Preference: If a carrier is preferred by a customer type and several carriers are included in the simulation experiment, a factor $\beta^{\text{car}}(m_r)$ is included in the cost function. It adds a constant additional cost to any brand that is not the preferred carrier.

Cost Function Without Price: The cost function without regard for the price is shown in Equation (7.18). Note that a normally distributed error depending on the customer type, ϵ_r is added to each request's cost function to individualize it. By minimizing the cost, the first choice itinerary for every request can be determined from the quantity of acceptable itineraries according to the constraints shown above. However, when a customer makes a booking decision, he can only really consider itineraries for which tickets are available.

$$\begin{aligned}\hat{C}(i,r) = \beta^{\text{dep}}(m_r)\left(\left|\nu^{\text{dep}}(q_r) - x^{\text{dep}}(i)\right|\right) + \beta^{\text{dur}}(m_r)\left(\left|\nu^{\text{dur}}(q_r) - x^{\text{dur}}(i)\right|\right)\\ + \beta^{\text{trans}}(m_r)\cdot x^{\text{trans}}(i) + \beta^{\text{car}}(m_r) + \epsilon_r\\ \forall\, r \in R;\ i \in I_{q_r}\end{aligned} \tag{7.18}$$

Simplifying Model Assumptions: In the revenue management simulation presented in this text, some transformations are based on pragmatic assumptions. One is the association between maximum acceptable departure time deviations and travel times: A connection to the overall distance traveled is known, but the precise functional form is not established. For the data used in this simulation, the product of a factor and the square-root of the distance worked well. However, different functional forms are conceivable. The same transformation is used to make the customers' willingness to pay dependent on the distance traveled. Finally, the form of the cost function is assumed to be linear. Much more complex functional forms are conceivable. Implementing them is easily possible, but a comparison of cost functions was not the focus of this work. In order to model flexible customer choice behavior, the condition for bookings depending on availabilities, the linear cost function is sufficient.

Class Choice

In order to formally express the logic of class choice, additional variables have to be defined:

- Let $f \in F$ be the index of flights.
- Let F_i be the set of flights included in itinerary i.
- Let $c \in C$ be the index of booking classes offered.
- Let $p\left(f, c\right)$ be the fare associated to a ticket for flight f in booking class c.
- Let $z \in Z$ be the set of restrictions of booking classes - the absence of a feature, such as comfort seating, is modeled as a restriction.
- Let Z_c be the restrictions included in booking class c.
- Let Z_r be the restrictions accepted by request r.
- Let δ_r^{price} be the factor defining maximum willingness to pay for customer type m.
- Let δ_r^{price} be the factor defining maximum willingness to pay for request r.
- Let β_m^{price} be the weight of fare in the cost function of the customer type m.
- Let $C(i, r)$ be the cost of itinerary i as defined by the cost function of request r when a lowest available fare has been found.
- Let $\delta_{r,c}^{\text{product}}$ be a boolean matrix defining whether the product represented by booking class c is acceptable according to the product requirements of request r.
- Let $\delta_{r,c,i}^{\text{wtp}}$ be a boolean matrix defining whether the price of class c on itinerary i is acceptable according to the willingness to pay of request r.
- Let c_r^{min} be the cheapest acceptable class according to the product requirements of request r.

Product Characteristics: Every customer type includes a list of acceptable restrictions and required features for classes. These correspond to the restrictions and features that are used to describe the booking classes in the supply model. Whether one of the classes available for an itinerary is acceptable becomes a matter of aligning class characteristics and customer requirements. Class c is acceptable for request r if it fulfills the condition presented in Inequality (7.19):

$$\delta_{r,c}^{\text{product}} := \begin{cases} 1 & Z_c \subset Z_r \\ 0 & \text{else.} \end{cases} \qquad \forall\, r \in R;\ c \in C \tag{7.19}$$

Based on product acceptance, the acceptable class with the lowest price can already been defined. This is based on the model limitation that classes have the same descending order of price and restrictions on all flights. The cheapest acceptable class of a request r is determined according to Definition (7.20).

$$c_r^{\min} := c' \mid |p(\circ, c') = \min\left(p(\circ, c)\ \forall\, c \in \left\{C \times \delta_{r,c}^{\text{product}}\right\}\right) \qquad \forall\, r \in R \tag{7.20}$$

Price Characteristics: Included in the definition of the customer type is the maximum willingness to pay. It depends on the parameter δ_m^{price} and the distance covered by the pairing q_r, $\nu(q_r)$. δ_r^{price} is calculated by distorting the underlying $\delta_{m_r}^{\text{price}}$ with the normally distributed error term ϵ_r as shown in Definition (7.21). The choice of a combination of class c and itinerary i is only acceptable if the overall fare fulfills the restriction shown in Inequality (7.22).

$$\delta_r^{\text{price}} := \delta_{m_r}^{\text{price}} + \epsilon_r\ \forall\, r \in R \tag{7.21}$$

$$\delta_{r,i}^{\text{wtp}} := \begin{cases} 1 & \sum_{f \in F_i} p\left(f, c_r^{\min}\right) \leq \delta_r^{\text{price}} \cdot \sqrt{\nu(q_r)} \\ 0 & \text{else.} \end{cases} \qquad \forall\, r \in R;\ i \in I_{q_r} \tag{7.22}$$

The demand model is based on the assumption that every customer will accept a class that includes more features or less restrictions than required, if that class is offered at a fare lower than the maximum willingness to pay. A customer will always chose the cheapest of all acceptable classes of one itinerary. The cost of tickets in the cheapest acceptable class of the flights included is added to each itinerary's cost function. The resulting costs for the considered alternatives are compared and the itinerary with the lowest cost is chosen. The cheapest available and acceptable class on this itinerary is booked.

Cost Function With Price: The cost function including the price is shown in Equation (7.23). The minimum available and acceptable price is added to the function with a weight factor $\beta_{m_r}^{\text{price}}$.

$$C(i,r) = \hat{C}(i,r) + \beta_{m_r}^{\text{price}} \cdot \sum_{f \in F_i} p\left(f, c_r^{\min}\right) \qquad \forall\ r \in R;\ i \in I_{q_r} \tag{7.23}$$

If no acceptable class is available, the itinerary is not considered. If no acceptable class could be determined on any of the acceptable itineraries, no booking takes place.

7.2.3. Exemplary Scenario

An example may be useful in illustrating how a simulation scenario can be generated. This section will describe the design of supply data and the generation of customer requests.

Picture a small network consisting of four airports (*vertices*): The domestic airports FRA (Frankfurt) and HAM (Hamburg) as well as the intercontinental airports JFK (New York) and BKK (Bangkok). The four airports are connected by three vice-versa flight legs (*edges*): FRA-HAM/HAM-FRA, FRA-BKK/BKK-FRA, HAM-JFK/JFK-HAM.

As shown in Figure 7.5, part (a), such a network offers a maximum of six possible pairings. If no other constraints are considered, a customer can travel from any of the airports to any of the other airports. An example is traveling from HAM to BKK via FRA or directly, routes described by the dotted line in part (b).

A connection builder including constraints such as a maximum number of acceptable transfers and a minimum as well as a maximum connecting time, however, may enforce a limitation of offered pairings. In the given example, traveling from BKK to JFK is only possible by transfers at both FRA and HAM. With a restriction to a maximum of one transfer per itinerary, the pairing BKK-JFK is no longer considered, with the result shown in part (c).

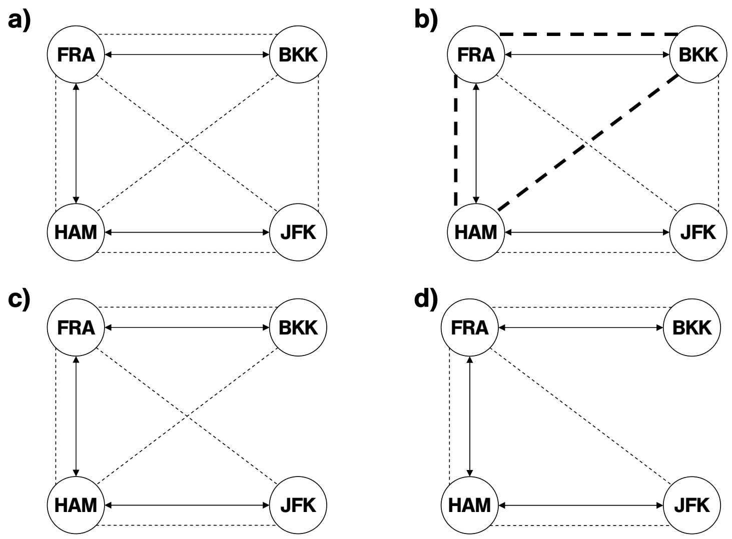

Figure 7.5.: Example – Product

Other pairings, in the given example HAM-BKK, may be excluded due to connecting times exceeding the maximum given by input parameters. The network presented in part (d) of Figure 7.5 would remain.

In order to provide complete itineraries for customers to chose from, a connection builder combines the existing flights to form paths through the network. In keeping with the current example, four pairings and eight directed itineraries emerge: FRA-HAM (direct), HAM-FRA (direct), FRA-BKK (direct), BKK-FRA (direct), HAM-JFK (direct), JFK-HAM (direct), and FRA-JFK (via HAM) as well as JFK-FRA (via HAM). Other theoretically possible paths such as FRA-JFK-FRA-HAM are excluded by restrictions of the connection builder.

Before the three existing pairings can be assigned shares of customer types, these need to be defined. This small example contains only two types of customers: Business travelers and tourists.

The timing of requests for tickets needs to be fixed per customer type. Keeping in line with classical assumptions, tourists plan their trips a long time before departure, whereas business travelers spontaneously decide to travel. Accordingly, the bulk of tourist

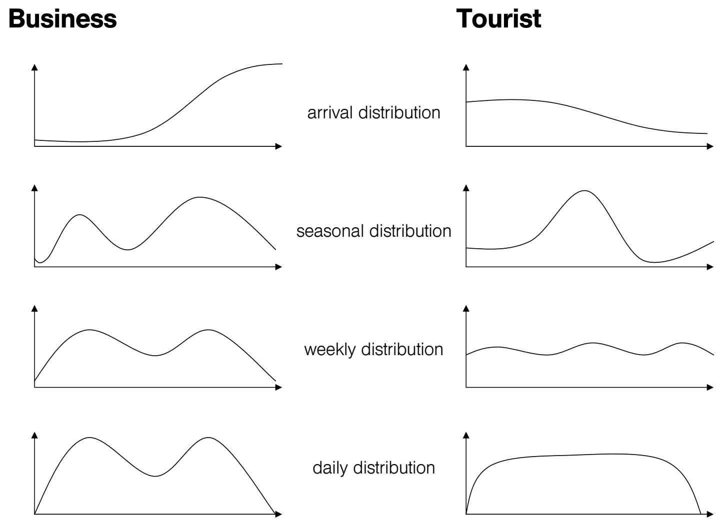

Figure 7.6.: Example – Customer Types

customers arrives during the first two thirds of the booking horizon, whereas most business customers arrive during the last third. As presented in Section 7.2.2, this is modeled as changes in the intensity of the Poisson process over time. The different distributions with regard to arrival and desired departure time can be seen in Figure 7.6.

Tourists may include different factors in their cost functions and tend to display a lower maximum willingness to pay than business travelers. In this example, the customer type representing tourists is willing to accept a weekend stay as well as a minimum stay of five days. The customer type representing business travelers requires extra flexibility and seating in the business compartment.

Pairings are assigned a general share of requests as well as a customer mix. To keep a simulation realistic, the amount of customer requests is calibrated to match productivity indicators such as seat load factors. An orientation for the assignment of shares to customer types may be taken from actual booking data for classes aimed at different product segments.

For this example, HAM-JFK (vice-versa) is assigned a 40% share of pairings whereas FRA-BKK (vice versa) gets a 20% share. Assuming an overall demand of a thousand requests, this means 400 requests for HAM-JFK and 200 requests for FRA-BKK.

One of the pairings is a major business route while the other is set to be a predominantly tourist market. The consequence is to assign HAM-JFK a customer mix that determines an 80% share of classical business customers and a 20% share of tourist customers, whereas FRA-BKK gets exactly the opposite, 20% business customers and 80% tourists.

When requests are generated, the request shares of the pairings are regarded as fixed. For 200 requests, the fact that the customer will want to travel from FRA to BKK is certain. For 40 of these, the customer type will be "business" whereas for the rest (160), the customer type will be "tourist". A random element enters the model when these numbers are used as the intensity of the respective Poisson processes for tourist and business customer types. The request arrivals generated from these Poisson processes can still differ from the expected intensity.

The distributions underlying the customer types are presented in Figure 7.6. The desired departure time is drawn from three distributions: First the week, then the week day, finally the preferred hour of departure are drawn. The result might look something like this: A business customer requests a flight from FRA to BKK, leaving in the tenth week of the year, on Tuesday, at 9 a.m. – the request arrives ten days before departure. The request will include the product restrictions defined in the type description, for example an exclusive acceptance of seats in the business compartment combined with an intolerance for weekend stays.

The underlying cost function is taken from the customer type. It is distorted by an error term drawn from a normal distribution for each request.

7.3. Revenue Management Components

Apart from supply and demand data, a revenue management simulation needs to include a model of the systems actually in use in airline revenue management. Required are a forecast, an optimization, and an inventory control. The implementation details for these modules are described in the following text.

7.3.1. Forecast

Following basic approaches to demand forecasting as described in Section 3, two forecast methods have been implemented. The exponential smoothing forecast follows the traditional assumption of static demand by observing historical bookings for one class to predict demand for the same class. The price sensitive forecast alternative considers possible buy-down and sell-up to occur between classes that are regarded as interchangeable by the customers.

Initial Forecasts

Every forecast method needs past observations to predict the future. In the first run of a simulation experiment, such information is not yet available. Several ways of handling this are conceivable:

Zero Forecast: When this forecast is chosen as the initialization method, all demand to come is set to zero. As actual bookings are observed, the forecast is expected to increase.

Random Forecast: When this forecast is chosen as the initialization method, a parameter is provided to indicate the predicted seat load factor. From this, the absolute number of demand is derived via the capacity. This demand is then predicted equally distributed over the booking classes.

Real Forecast: Using the real-world forecast for the supply included in the simulation could be a way of ensuring a realistic initial status. However, the condition for this is a demand model that is accurately calibrated to match real-world demand. As full information on real-world customer choice behavior is not available, this is condition cannot be fulfilled.

Psychic Forecast: This is a forecast that is based on knowledge of demand as inherent in the simulation. Different methods of computing it are conceivable and may lead to different results.

Exponential Smoothing

The exponential smoothing forecast predicts demand volume based on historical observations. By predicting demand for time slices before departure, the forecast considers demand arrival patterns. By predicting demand for each flight, departure time patterns are included.

The inclusion of seasonal patterns in the forecast could be implemented by adding methods taken from time series forecasting. However, as the focus of the simulation experiments conducted is on the evaluation of the quality of unconstraining and recognition of demand behavior, this has not been implemented.

In the exponential smoothing forecast the unconstraining aspect is included via *additive pick-up*. Historical bookings are transformed unless the class in which they occurred was available throughout the considered time slice. If it was available, the observation is added to the history of bookings as shown in Definition (7.24). If it was not available, the number of bookings observed is compared to that observed in previous runs while the class was open. The higher value is used. This process is formally expressed in Definition (7.25).

- Let $s = 2, ..., N^s$ be the runs included in a simulation for which a forecast update is performed. $s + 1$ occurs chronologically after s and can be based on historical data derived from s. For $s = 1$, an initial forecast is supplied.
- Let $t = 0, ..., N^t$ be points of time in booking horizon, with $t = 0$ designating the start of the booking horizon and $t = N^t$ being the time of departure.
- Let $c \in C$ be the booking classes ordered by descending price.
- Let $f \in F$ be the flights included in the schedule.
- Let $b(f, c, t, s)$ be the bookings observed for flight f, class c, between points of time before departure $t - 1$ and t, in simulation run s.
- Let $\hat{b}(f, c, t, s)$ be the average of historical bookings during the runs 1 to s that occurred on flight f between points of time t and $t - 1$ while booking class c was available.

- Let $b^{\text{unc}}(f,c,t,s)$ be the unconstrained bookings for flight f, class c, between points of time $t-1$ and t, in simulation run s.
- Let $a(f,c,t,s)$ be the seats available for sale for flight f, class c, at points of time before departure t in simulation run s.

First, the unconstrained bookings observed while the class was available need to be updated:

$$\hat{b}(f,c,t,s) := \begin{cases} \frac{\hat{b}(f,c,t,s-1)\cdot(s-1)+b(f,c,t,s)}{s} & a(f,c,t-1,s)>0,\ a(f,c,t,s)>0 \\ \hat{b}(f,c,t,s-1) & \text{else.} \end{cases} \quad (7.24)$$
$$\forall\ f \in F;\ c \in C;\ t=1,...,N^t;\ s=2,...,N^s$$

Next, the observed bookings need to be unconstrained if the class was not available throughout the time slice:

$$b^{\text{unc}}(f,c,t,s) := \begin{cases} b(f,c,t,s) & a(f,c,t-1,s)>0;\ a(f,c,t,s)>0 \\ \max\left(\hat{b}(f,c,t,s), b(f,c,t,s)\right) & \text{else.} \end{cases}$$
$$\forall\ f \in F;\ c \in C;\ t=1,...,N^t;\ s=2,...,N^s \quad (7.25)$$

The forecast is updated with the help of the unconstrained bookings. The emphasis given to new bookings over the existent forecast is influenced by a parameter α. Definition (7.26) shows how the forecast is updated after each run.

- Let $f^{\text{unc}}(f,c,t,s)$ be the predicted demand to arrive for flight f between points of time $t-1$ and t per class c and simulation run s.
- Let $f^{\text{unc}}(t,c,t,1)$ be the initial forecast.
- Let α^{exp} be the weight of new bookings in the calculation of the updated unconstrained forecast.

$$f^{\text{unc}}(f,c,t,s) := (1-\alpha^{\text{exp}})\cdot f^{\text{unc}}(f,c,t,s-1) + \lambda^{\text{exp}}\cdot b^{\text{unc}}(f,c,t,s)$$
$$\forall\ f \in F;\ c \in C;\ t=1,...,N^t;\ s=2,...,N^s;\ \alpha^{\text{exp}} \in [0,1] \quad (7.26)$$

A variation of the exponential smoothing method inlcudes updates of the forecast within the booking horizon of one simulation run. A comparison of actual and predicted bookings is used to determine the current difference in terms of a factor. This factor is applied to future bookings, as shown in Definition (7.27):

$$f^{\mathrm{unc}}(f,c,t',s) := f^{\mathrm{unc}}(f,c,t',s) \cdot \left(1 + \frac{b^{\mathrm{unc}}(f,c,t,s) - f^{\mathrm{unc}}(f,c,t,s)}{f^{\mathrm{unc}}(f,c,t,s)}\right)$$
$$\forall\, t' = t+1,...,N^t;\ f \in F;\ c \in C;\ s = 2,...,N^s \tag{7.27}$$
$$\text{if } f^{\mathrm{unc}}(f,c,t,s) > 0,\ a(f,c,t-1,s) > 0,\ a(f,c,t,s) > 0$$

Price-Sensitive Estimators

Alternative forecasting method uses price-sensitive estimators. The underlying assumption is that every customer will buy the cheapest alternative if classes are only differentiated by price. In the simulation implemented, this is true for booking classes that offer the same set of characteristics. In that case, demand is influenced by two factors, the price of the cheapest available booking class and the time before departure.

- Let $c \in C$ be a set of booking classes sharing the same set of restrictions, ordered by descending price – this means N^c is the booking class with the lowest price.
- Let $b(f,c,t,s)$ be the bookings that were observed for flight f in class c between points of time $t-1$ and t of simulation run s.
- Let $\overline{o}(f,c,t,s)$ be a boolean matrix indicating the lowest available class for flight f between points of time $t-1$ and t of simulation run s.
- Let $\omega^T_{f,t,s}$ be the vector of time elasticity for flight f depending on the point of time t before departure of run s.
- Let $\omega^P_{f,c,s}$ be the vector of price elasticity for flight f depending on the class c for run s.
- Let $u^T(f,c,t,s)$ be the time-based estimator before departure for flight f, class c, and point of time t of run s.
- Let $u^P(f,c,t,s)$ be the price-based estimator for flight f, class c, and point of time t of run s.
- Let α^T be the weight of the time-based estimator in the joint estimator.

- Let α^P be the weight of the estimator based on price in the joint estimator.
- Let $u^J(f,c,t,s)$ be the joint estimator for flight f, class c, point of time t and run s.

First of all, the cheapest class available needs to be determined as shown in Definition (7.28). This information is stored as a boolean flag per time slice before departure, class, and run.

$$\overline{o}(f,c,t,s) := \begin{cases} 1 & a(f,c,t,s) > 0;\ a(f,c',t,s) = 0\ \forall\ c' = c+1,...,N^c \\ 1 & c = 1;\ a(f,c',t,s) = 0\ \forall\ c' \in C \\ 0 & \text{else.} \end{cases} \tag{7.28}$$
$$\forall\ f \in F;\ c \in C;\ t = 1,...,N^t;\ s = 1,...,N^s - 1$$

According to the assumption that customers will book the cheapest available class, bookings can only be expected for this class. Therefore, the observed bookings in the class indicated by $\overline{a}(f,c,t,s) = 1$ enter both the price and the time estimator matrices. For those classes that were not the cheapest available during time slice t, values are derived via price-elasticity and time-elasticity vectors. The rules according to which the price and time estimator matrices are filled are formalized in Definitions (7.29) and (7.30).

$$u^P(f,c,t,s) := \begin{cases} b(f,c,t,s-1) & \overline{o}(f,c,t,s-1) = 1 \\ u^P(f,c-1,t,s)/\omega^P(f,c-1,s-1) & \sum_{c'=1}^{c} \overline{o}(f,c',t,s-1) = 1 \\ u^P(f,c+1,t,s) \cdot \omega^P(f,c,s-1) & \sum_{c'=c+1}^{N^c} \overline{o}(f,c',t,s-1) = 1 \\ e^P(f,c,t,s-1) & \text{else.} \end{cases}$$
$$\forall\ f \in F;\ c \in C;\ t = 1,...,N^t;\ s = 2,...,N^s \tag{7.29}$$

$$u^T(f,c,t,s) := \begin{cases} b(f,c,t,s-1) & \overline{o}(f,c,t,s-1) = 1 \\ u^T(f,c,t+1,s-1) \cdot \frac{\omega^T(f,t,s-1)}{\omega^T(f,t+1,s-1)} & \sum_{t'=t+1}^{N^t} \overline{o}(f,c,t',s-1) = 1 \\ u^T(f,c,t-1,s-1) \cdot \frac{\omega^T(f,t,s-1)}{\omega^T(f,t-1,s-1)} & \sum_{t'=1}^{t-1} \overline{o}(f,c,t',s-1) = 1 \\ u^T(f,c,t,s-1) & \text{else.} \end{cases}$$
$$\forall\ f\ \in F;\ c \in C;\ t = 1,...,N^t;\ s = 2,...,N^s \tag{7.30}$$

The joint estimator matrix is computed as shown in (7.31) using the parameters α^{exp} (as introduced for exponential smoothing), α^T and α^P. Note that demand for one class c is set to be higher or equal to that for the next, cheaper class $c+1$.

$$\begin{aligned} u^J(f,c,t,s) &:= \\ &(1-\alpha^{\text{exp}}) \cdot u^J(f,c,t,s-1) \\ &+ \min\left(\alpha^{\text{exp}} \cdot \left(\alpha^T \cdot u^T(f,c,t,s) + \alpha^P \cdot u^P(f,c,t,s)\right), u^J(f,c+1,t,s)\right) \\ &\forall\ \alpha^{\text{exp}}, \alpha^T, \alpha^P \in [0,1];\ \alpha^T + \alpha^P = 1;\ f \in F;\ c \in C;\ t = 1,...,N^t;\ s = 2,...,N^s \end{aligned} \tag{7.31}$$

The price-elasticity and time-elasticity vectors are updated based on the new joint estimator matrix. Considered are those points of time up to which a class was the lowest available and from which on it was no longer available. If no such change took place, the elasticity vectors are not updated. This is formalized by Definitions (7.32) and (7.33):

$$\begin{aligned} \omega^P(f,c,s) := \sum_{t=1}^{N^t} \left(\overline{o}(f,c,t-1,s-1) - \overline{o}(f,c,t,s-1)\right) \cdot \frac{u^J(t,c,s)}{u^J(t-1,c+1,s)} \\ \forall\ f \in F;\ c \in C;\ s = 2,...,N^s \end{aligned} \tag{7.32}$$

$$\begin{aligned} \omega^T(f,t,s) := \begin{cases} \sum_{c \in C} \left(\overline{o}(f,c,t,s-1) \cdot u^J(f,c,t,s)\right) & \\ \cdot \prod_{c \in C} \left(\sum_{c'=1}^{c} \overline{o}(f,c',t,s-1)\right) \cdot \frac{1}{\omega^P_{c,s}} & \sum_{c \in C} \overline{o}(f,c,t,s-1) > 0 \\ \omega^T(f,t,s-1) & \sum_{c \in C} \overline{o}(f,c,t,s-1) = 0 \end{cases} \\ \forall\ f \in F;\ t = 1,...,N^t;\ s = 2,...,N^s \end{aligned} \tag{7.33}$$

In order to make this forecast a valid input for the implemented EMSR-b optimization algorithm, it needs to be transformed into a pseudo-static version. For this purpose, each customer is assumed to request the booking class that corresponds to his highest willingness to pay. The transformation is formally described by Definition (7.34):

$$\begin{aligned} f^{\text{unc}}(f,c,t,s) := \begin{cases} u^J(f,c,t,s) - u^J(f,c+1,t,s) & c < N^c \\ u^J(f,c,t,s) & c = N^c \end{cases} \\ \forall\ f \in F;\ c \in C;\ t = 1,...,N^t;\ s = 2,...,N^s \end{aligned} \tag{7.34}$$

7.3.2. Optimization

The optimization method implemented is EMSR-b as first presented in P. Belobaba (1987b) (see also Section 2.2). As the focus of this work is foremost on forecast evaluation, EMSR-b was chosen for its adaptability to different forecast methods. This way, the optimization method could be kept stable as forecasts vary and are transformed to match its demands.

EMSR-b uses the accumulated forecast to allocate seats to booking classes. The decision of whether or not to reserve capacity for more expensive classes is based on the *expected marginal seat revenue*. This is calculated as a ratio of expected demand to arrive in one class and the next cheaper class multiplied by their respective value.

When forecasts are updated within the booking horizon, the optimization can be updated as well. In order to include this option, the point of time t before departure is included in all calculations described here. If $t = 0$, the optimization takes place before any bookings were observed. If availabilities are calculated again after a forecast update, $t > 0$ and bookings may have already taken place. Whether or not to update the forecast and the availabilities is a methodological decision.

- Let $c \in C$ be a set of booking classes sharing the same set of restrictions, ordered by descending price – this means N^c is the booking class with the lowest price.
- Let $\sum_{t=1}^{N^t} f^{\text{unc}}(f, c, t, s)$ be the sum of the unconstrained demand to arrive for flight f in class c until point of time t in the booking horizon of run s.
- Let $\sigma\left(f^{\text{unc}}(f, c, t, s)\right)$ be the standard deviation of the forecast of demand for flight f in class c at point of time t in run s.
- Let $b(f, c, t, s)$ be the bookings that arrived for flight f in class c in run s between points of time $t - 1$ and t.
- Let $K(f, t, s)$ be the available capacity of the flight f the point of time t before departure of simulation run s.
- Let $p(f, c)$ be the price of class c for the flight f.
- Let $\overline{p}(f, c, t, s)$ be the expected marginal seat revenue for a seat in class c, flight f at point of time t of simulation run s.

- Let U^{-1} be the inverse of the normal distribution.
- Let $S(f, c, t, s)$ be a "safety margin" of demand expected for class c of flight f during the time between $t-1$ and t, based on the standard deviation of predicted demand and the inverse of the normal distribution.
- Let $\hat{a}(f, c, t, s)$ be the protected seats assigned to flight f in class c in simulation run s at the point of time t.

At the beginning of the booking horizon, revenue-maximizing availabilities are computed based on the demand forecast. Protected seats are calculated as the part of the capacity that is reserved for one particular booking class. In a nested structure, these seats can be used for bookings in the class they are reserved for or any more expensive class. Before the booking horizon has started, at $t = 0$, capacity $K(f, 0, s)$ is equal to the overall capacity of the aircraft assigned to the considered flight.

To err on the side of caution, the average price of the considered class c and all the more expensive classes offered is calculated. This average is weighted by the predicted demand as shown in Definition (7.35). The result is the *expected marginal seat revenue*:

$$\overline{p}(f, c, t, s) := \frac{\sum_{c'=1}^{c} \left(\left(\sum_{t'=t+1}^{N^t} f^{unc}(f, c', t', s) \right) \cdot p(f, c') \right)}{\sum_{c'=1}^{c} \left(\sum_{t'=t+1}^{N^t} f^{unc}(t', c', s) \right)} \qquad \forall\ f \in F;\ c \in C;\ t = 0, ..., N^t - 1;\ s = 1, ..., N^s \tag{7.35}$$

The standard deviation of the forecast up to class c is computed as shown in Definition (7.36):

$$\hat{\sigma}\left(f^{\mathrm{unc}}(f, c, t, s)\right) := \sqrt{\sum_{c'=c}^{N^c} \sigma\left(f^{\mathrm{unc}}(f, c, t, s)\right)^2} \qquad \forall\ f\ \in F;\ c \in C;\ t = 1, ..., N^t;\ s = 1, ..., N^s \tag{7.36}$$

A safety margin $S(f, c, t, s)$ is calculated using the inverse of the normal distribution, U^{-1} and the standard deviation of the forecast.

$$S(f, c, t, s) := U^{-1}\left(1 - \frac{p(f, c)}{\overline{p}(f, c, t, s)}\right) \cdot \hat{\sigma}\left(f^{\mathrm{unc}}(f, c, t, s)\right) \qquad \forall\ f \in F;\ c \in C;\ t = 0, ..., N^t - 1;\ s = 1, ..., N^s \tag{7.37}$$

Finally, the protection levels are assigned as shown in Definition (7.38). Note that in the case of availabilities being updated throughout the booking horizon, bookings that already took place automatically become part of the protection levels. Protected seats are allocated from top down: First, the most expensive class gets its share, left-over capacity is assigned to the cheapest class.

$$K^* (f, c, t, s) := K (f, t, s) - \sum_{c'=1}^{c-1} (\hat{a} (f, c', t, s))$$
$$\forall\ f \in F;\ c \in C;\ t = 0, ..., N^t;\ s = 1, ..., N^s$$

$$F^* (f, t, c, s) :=$$
$$\left(\sum_{c'=c}^{N^c} \left(\sum_{t'=t+1}^{N^t} f^{\text{unc}} (f, c', t', s) \right) + S (f, c, t, s) \right)$$
$$- \left(\sum_{c'=c+1}^{N^c} \left(\sum_{t'=t+1}^{N^t} f^{\text{unc}} (f, c', t', s) \right) + S (f, c+1, t, s) \right)$$
$$\forall\ f \in F;\ c \in C;\ t = 0, ..., N^t - 1;\ s = 1, ..., N^s$$

$$\hat{a} (f, c, t, s) := \begin{cases} \min \left(K^* (f, c, t, s), F^* (f, t, c, s) \right) + \sum_{t'=0}^{t} (b (f, c, t', s)) & c < N^c \\ K^* (f, c, t, s) + \sum_{t'=0}^{t} (b (f, c, t', s)) & c = N^c \end{cases} \tag{7.38}$$
$$\forall\ f \in F;\ c \in C;\ t = 0, ..., N^t - 1;\ s = 1, ..., N^s$$

7.3.3. Inventory

A *reservation system* (*inventory*) is needed to calculate up-to-date availabilities and handle customer bookings. In the inventory, customer requests for tickets meet the class providing the desired product characteristics at the lowest available fare. Availabilities are either calculated based on authorization levels and a nesting structure or by comparing the value of a booking class to the current bid price.

In the case of authorization levels, each booking class offered is assigned a certain amount of protected seats. An example for an optimization method yielding results in this form is EMSR-b as presented in Section 7.3.2. In order to make the system more robust and allow for more than one booking class to be sold at any given time, the concept of *nesting* as described for example in Talluri & Van Ryzin (2004b) is applied. As protected seats are calculated using the EMSR-b method in the implemented simulation system, this system uses a nested fare structure with availabilities that are equal to or exceed the protected seats.

Any class is available for as long as the overall available seats exceed those seats protected for more expensive classes. The relationship is presented in Definition (7.39) according to the notation introduced in Section 7.3.2.

$$a\left(f,c,t,s\right) := \min\left(K\left(f,t,s\right) - \left(\sum_{c'=1}^{c-1}\left(\hat{a}\left(f,t,c',s\right)\right) - \sum_{c'=1}^{c-1}\left(\sum_{t'=1}^{t} b\left(f,t',c',s\right)\right)\right),0\right) \tag{7.39}$$
$$\forall\ f\in F;\ c\in C;\ t=0,...,N^t-1;\ s=1,...,N^s$$

According to this concept, once the protected seats reserved for one class are used up, customers wanting to book a ticket in this class can access left-over protected seats in any of the lower classes. Depending on the precise implementation, sold seats are either subtracted from the lowest class's protected seats or only from the protected seats of the class directly nested under the desired class. Availabilities for every class are calculated by adding up its protected seats as well as all the protected seats for the lower-nested classes. Depending on whether the separation between the compartments is treated as flexible or as fixed, a nesting structure can include all the classes offered or only the classes offered within one compartment.

A graphic illustration of this concept is provided by Figure 7.7. Three classes are included in the example shown. Class 1 is the most expensive one with the highest nesting position. While 50 seats are protected for this class alone, all 100 seats making up overall capacity may be sold in this class if none of the cheaper ones is booked. Class 2 represents an intermediate value. Based on unconstrained forecasts, 10 protected seats have been calculated for this class. As 50 seats are protected for class 1, the remaining 50 seats may be sold in class 2 if the cheaper class is not booked. Class 3, the cheapest option, is allocated the "left over" seats after protected seats for all more expensive classes

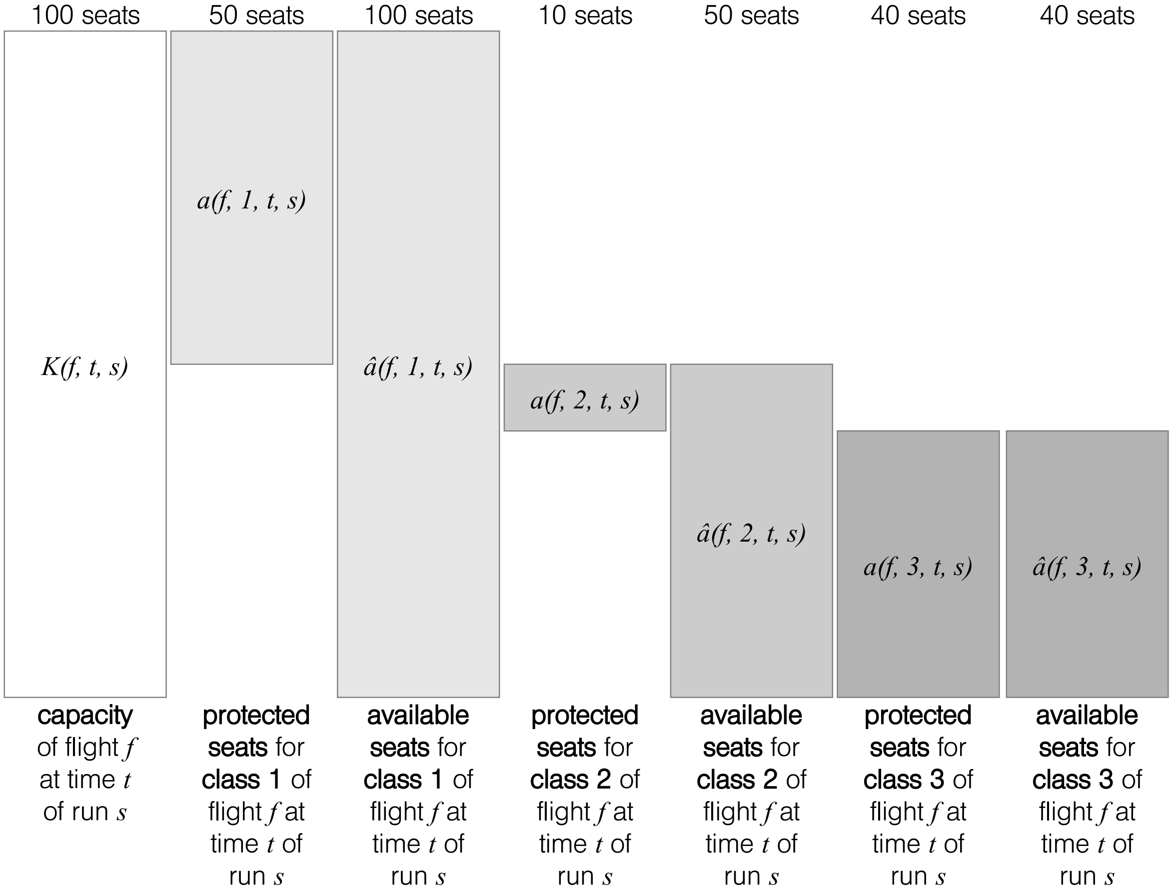

Figure 7.7.: Inventory: Protected and Available Seats

have been reserved. These 40 seats may be sold in class 3. As the cheapest class, it cannot access availabilities in any of the other classes.

If bid price controls are implemented, calculating availabilities becomes more straight forward. Each booking class is assigned a value based on its fare and, possibly, the buy-down it triggers according to a forecast including such information. This value, also referred to as *capacity allocation value*, is compared to the bid price assigned by the optimization component. If the class's value is lower than the bid price, it is not available, if it is equal or higher than the bid price, it is available. The bid price changes flexibly with regard to the seats already sold and the customers expected to request seats in the future.

When a customer request arrives, the lowest available classes complying to the customer's product preferences are returned. Now, the customer can chose whether to book a ticket; he will automatically book the alternative minimizing his cost function if his maximum willingness to pay has not been excelled by all offers. If a reservation is affirmed, the inventory updates the availabilities. In the case of a nesting structure, this includes updating the protected seats of the lower-priced classes if necessary. In the case of bid price controls, this includes updating the bid price according to the specifications given by the bid price curve provided by the optimization.

7.4. Market Implementations

As the basis of the simulation experiments presented in this thesis, one particular instance of the demand model is continually used. This means the creation of a number of specific customer types, their assembly to a customer mix, and the confrontation of this demand with some variations of supply.

To ensure consistency, demand is varied along two clearly defined parameters and confronted three alternative supply structures. These combinations can be used to evaluate the application of a number of forecast and forecast evaluation methods.

7.4.1. Demand Variations

A number of different customer types were implemented to allow for a variety of product-oriented and price-oriented behaviors. Underlying is the assumption of rational choice behavior: Customers will always try to minimize cost in terms of price and preference factors. For the mode, this means that all customer types will buy the cheapest ticket available given all other (product) conditions are equal. Furthermore, they will choose the itinerary that offers the best conditions according to the factors of their cost function.

Only one basic demand model was realized for the simulation experiments conducted. This is based on the theory that revenue management does not change basic market characteristics but instead utilizes different market segments by developing a targeted set of offers. According to this, the same demand model should lead to quite different result indicators depending on both the fare structures and the inventory controls in place.

To allow for scenarios that are purely price-driven, purely product-driven, and hybrid, requests based on customer choice behaviors including either priority need to be generated. The implemented customer types are:

- No Frills 1: Main focus on price, will only accept "no-refund" restriction, maximum willingness-to-pay is about 10% of the price span offered.
- No Frills 2: Main focus on price, will only accept "no-refund" restriction, maximum willingness-to-pay is about 30% of the price span offered.
- Tourist Weekend: Main focus on price, will accept "no-refund" and "weekend-stay" restriction, maximum willingness-to-pay is about 50 % of the price span offered.
- Tourist Medium: Main focus on price, will accept all restrictions, maximum willingness-to-pay is about 60% of the price span offered.
- Tourist Lux: Main focus is on travel time and departure time, will accept all restrictions, maximum willingness-to-pay is over 100 % of the price span offered.
- Business Long: Main focus is on travel time and departure time, will accept "no-refund" and "minimum stay" restrictions, maximum willingness-to-pay is about 90% of the price span offered.
- Business Flex: Main focus is on flexibility, travel time and departure time, will not accept restrictions, maximum willingness-to-pay is about 95% of the price span offered.
- Business Pure: Main focus is on flexibility, business compartment seats, travel time and departure time, will not accept restrictions, maximum willingness-to-pay is over 100% of the price span offered.

Figure 7.8 shows the distribution of these customer types in the mix applied to all markets in the simulation. Depending on the error term added to the individual origin-destination combinations, the customer mix presented may vary. This error term is based

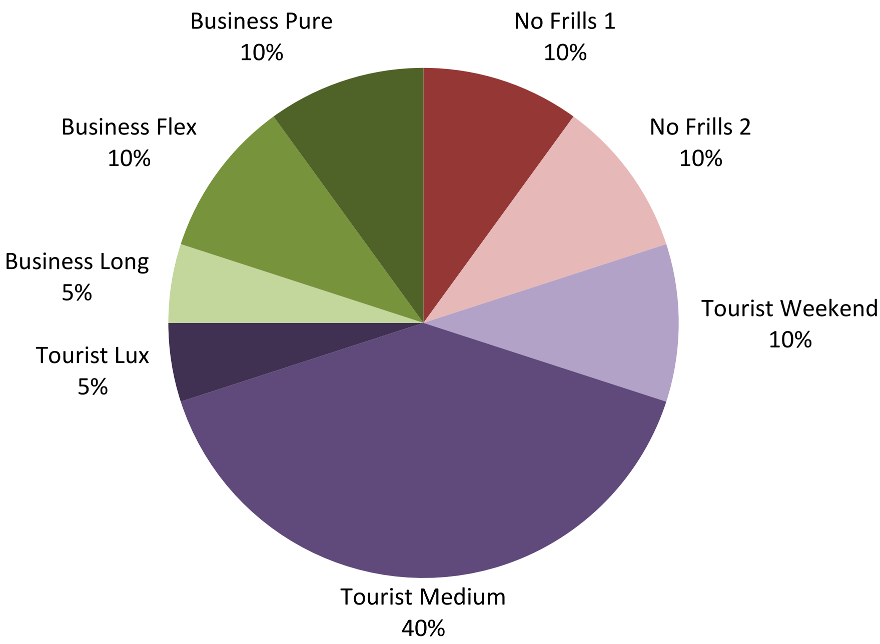

Figure 7.8.: Mix of Customer Types

on an input parameter defining the *standard deviation* of the normal distribution the error term is drawn from.

Overall demand was varied according to two parameters, demand volume and the deviation of the error term. While volume was controlled over the overall amount of requests generated, the error term drawn from the normal distribution with a zero average and the set standard deviation is included in multiple parts of the model. This way, the input parameter deviation influences the uncertainty of demand.

The realized variations are listed below. The abbreviation “Vol.” describes the number of requests in a percentage relation to the number of seats included in the simulation. For example, “Vol.050” indicates that the number of generated requests equals 50% of the number of seats. The abbreviation “Dev.” describes the standard deviation set for the normal distribution the error term is drawn from. In the case of "Dev.00", demand only varies from one run to the other due to the effects of drawing from a Poisson distribution.

- Vol.050 Dev.00: 3.000 requests are scheduled to arrive. The standard deviation of error terms is 0.
- Vol.050 Dev.01: 3.000 requests are scheduled to arrive. The standard deviation of error terms is 1.
- Vol.050 Dev.05: 3.000 requests are scheduled to arrive. The standard deviation of error terms is 5.
- Vol.050 Dev.10: 3.000 requests are scheduled to arrive. The standard deviation of error terms is 10.
- Vol.050 Dev.20: 3.000 requests are scheduled to arrive. The standard deviation of error terms is 20.
- Vol.100 Dev.00: 6.000 requests are scheduled to arrive. The standard deviation of error terms is 0.
- Vol.100 Dev.01: 6.000 requests are scheduled to arrive. The standard deviation of error terms is 1.
- Vol.100 Dev.05: 6.000 requests are scheduled to arrive. The standard deviation of error terms is 5.
- Vol.100 Dev.10: 6.000 requests are scheduled to arrive. The standard deviation of error terms is 10.
- Vol.100 Dev.20: 6.000 requests are scheduled to arrive. The standard deviation of error terms is of 20.

The consequences of varied volume and error deviation are presented in Figure 7.9. Shown are the results of first-come-first-serve inventory controls averaged over 50 simulation runs. As expected, the average seat load factor is higher for "Vol.100". In addition, it decreases slightly with increasing error deviation. The resulting deviation of the indicator seat load factor increases with increasing error deviation. Average revenue is lower for "Vol.100": With more overall requests, more requests for cheap booking classes are

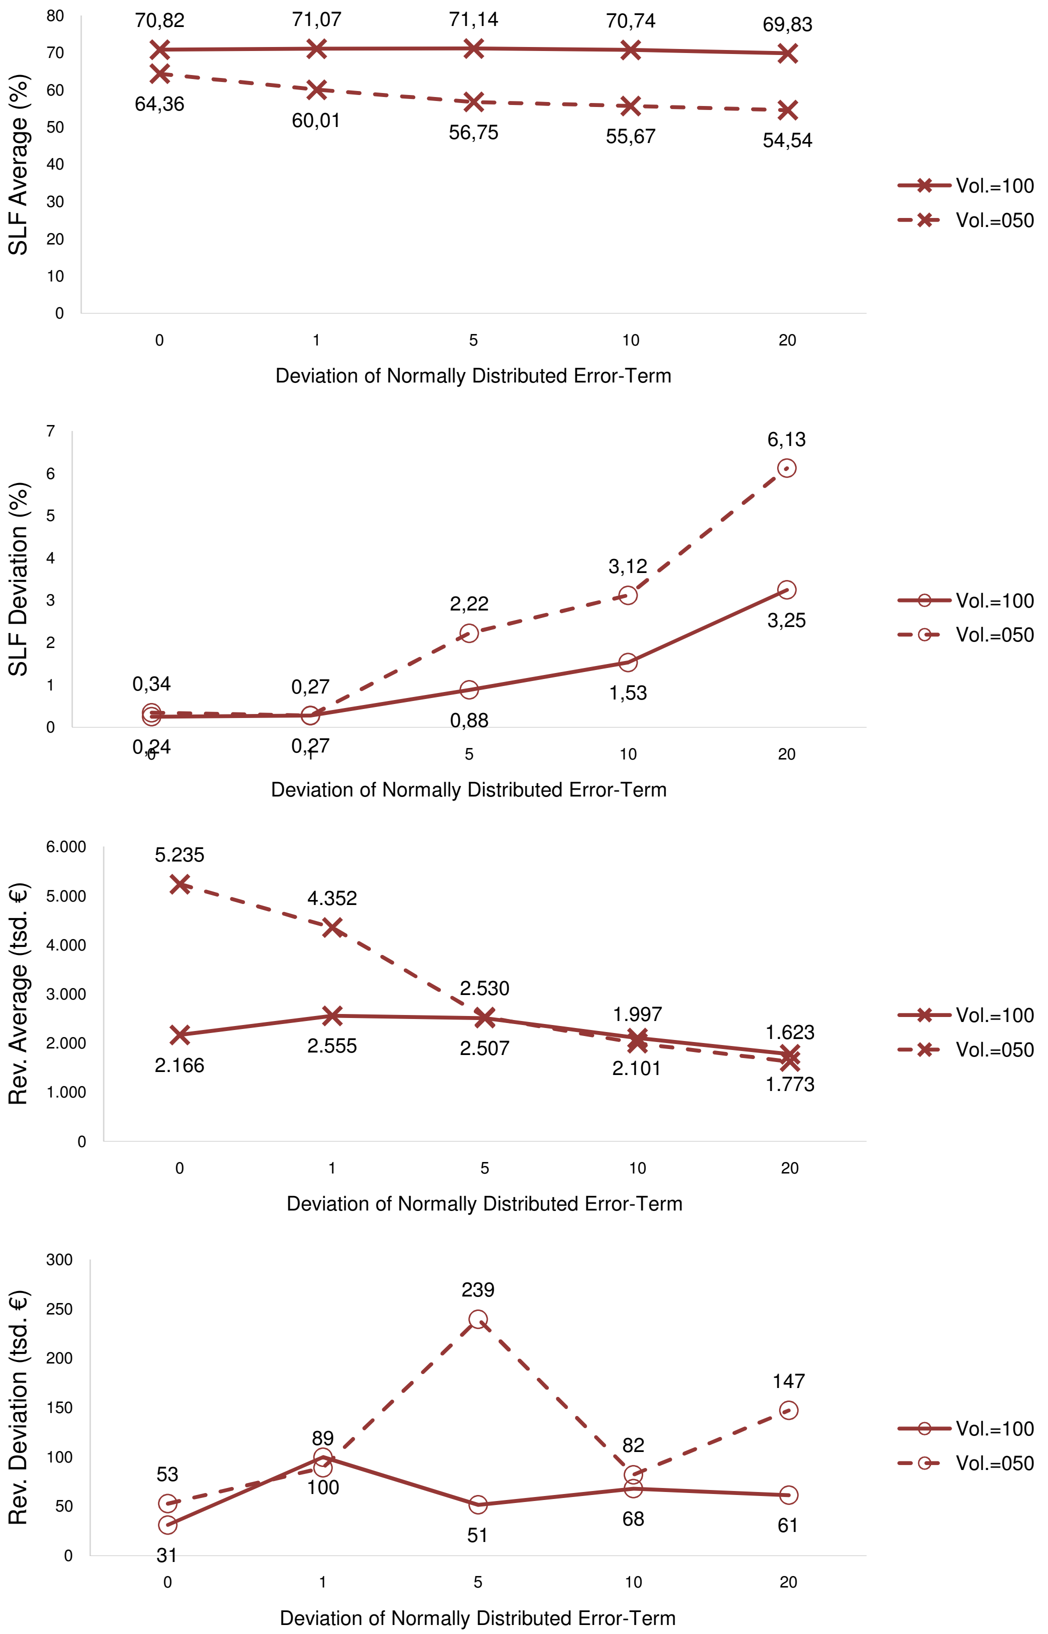

Figure 7.9.: SLF Average and Deviation depending on Error Term Deviation

Caption	Description	Fare Near	Fare Far
C	seat in business compartment	350	700
Y	refundable ticket	250	500
G	base fare	175	350
M	minimum stay restriction	150	300
W	weekend stay restriction	125	250

Table 7.4.: Booking Classes Differentiated by Product-Feature

accepted before valuable customers arrive. Average revenue also decreases with increasing error deviation. The resulting deviation of revenue shows no clear trend over increasing error deviation, however, low request volume react stronger to changes in error deviation. This seems to hold true over all four observed indicators.

7.4.2. Supply Variations

As described in the previous section, the customer behavior is set to be stable throughout the experiments. However, three representative supply variations are realized to model three supply strategies found in applied airline revenue management: product-based differentiation, price-based differentiation, and hybrid differentiation.

Table 7.4 lists the booking classes that present product-based differentiation. Table 7.5 lists the booking classes that present price-based differentiation. Table 7.6 lists the booking classes that present hybrid differentiation. Classes are presented from top to bottom in their nesting order. Two fares are assigned for flights spanning one ("Fare Near") or two traffic areas ("Fare Far").

Parallel to the assumption of rational customer choice behavior described in Section 7.2.2, an assumption of rational supply planning is held with regard to the calibration of the market scenarios realized. As it would not make sense in terms of economic rationality for a customer to buy an expensive ticket if a cheap ticket that corresponds to his or her product preferences was available, it seems unreasonable for an airline to offer a product that does not trigger demand from a customer segment specific to this product.

Caption	**Description**	**Fare Near**	**Fare Far**
G	base fare	175	350
K	reduced fare	100	200
L	reduced fare	75	150
T	reduced fare	50	60
E	reduced fare	20	20

Table 7.5.: Booking Classes Differentiated by Price

Caption	**Description**	**Fare Near**	**Fare Far**
C	seat in business compartment	350	700
Y	refundable ticket	250	500
G	base fare	175	350
M	minimum stay restriction	150	300
W	weekend stay restriction	125	250
K	reduced fare	100	200
L	reduced fare	75	150
T	reduced fare	50	60
E	reduced fare	20	20

Table 7.6.: Booking Classes Differentiated by Product Characteristics and Price (Hybrid Differentiation)

If the risk of buy-down is considered a cost, any new product differing from a base fare and standard features is costly or neutral. Product characteristics such as a business compartment differing in terms of physical comfort obviously gain added value through higher production cost. Features such as the possibility of refund add uncertainty to the airline's plans. Tickets that are sold at rates lower than that of the highest rate offered may cause buy-down and therefore include the risk of lost revenue.

Therefore, demand was calibrated according to the assumption that the airline would not offer a new booking class if there was no reason to expect additional demand. Using an expanding fare structure, a number of simulation experiments based on the variation *Vol.050 Dev.00* of the demand model described above were conducted. The fare structures develop as follows:

- G: This is the version including a single, non-refundable base fare representing fare structures in a time when neither product nor price differentiation was applied.
- G+Y: With the class Y, a more flexible class that is refundable is added at a price that exceeds that of G.
- G Y+C: With the class C, an option of comfortable seating in a business class is added at a price that exceeds that of Y.
- G Y C + M: The class M is the first reduced-fare class that is added to the supply – however, to avoid buy-down, a "minimum-stay" restriction is included.
- G Y C M + W *(product-based scenario)*: The class W is another reduced-fare class that is added to the supply, with a fare that is lower than that of M – however, to avoid buy-down, a "weekend-stay" restriction is included.
- G Y C M W + K, L, T, E *(hybrid scenario)*: In the tradition of no-frills airline, more classes are introduced at low fares that are not differentiated by restrictions but only differ in price. The underlying hope is that with good inventory controls, the gain through increased bookings will exceed the loss due to buy-down behavior.
- G K L T E *(price-based scenario)*: Finally, to include an option that represents that of low-fare or no-frills airlines, all classes based on product-differentiation are excluded from the scenario once more and only classes differentiated by price are offered.

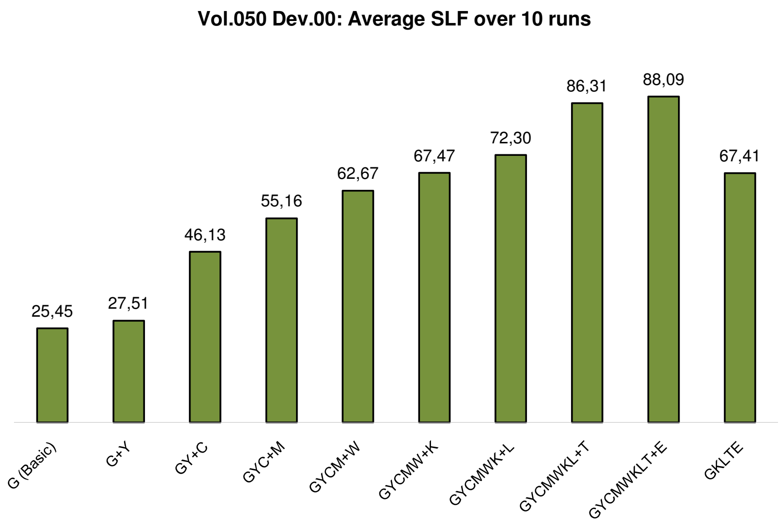

Figure 7.10.: Increase in Bookings by Additional Classes

The results of differentiation are illustrated in Figure 7.10. In this diagram, every column presents seat load factors under first-come-first-serve seat allocation. The demand model used was “Vol.050 Dev.00”. With every added class, seat load factors increase as a new customer segment is addressed with a new product. The last column shows seat load factors under a solely price-differentiated environment: Low prices activate customer segments with a low willingness to pay, but a lack of product features such as business compartment seats drives away a different customer segment.

In this section, market implementations for a revenue management simulation have been presented. Using this data, the simulation including flexible demand, and the concept for decomposition introduced earlier, simulation experiments can be conducted. These may be used to analyze statements regarding forecast performance.

Part III.

Experiments and Conclusions

With the help of the simulation system described in Chapter 7 and the concept of decomposition outlined in Chapter 6, experiments can be designed and executed to analyze forecast performance and its evaluation. Common evaluation methods can be combined with knowledge of the implemented demand model. The results of this are presented in Chapter 8 as combinations of formal statements and simulation results visualized by graphs.

In Chapter 9, a summary of the outcomes of this thesis is presented. This chapter recaptures the goals as first introduced in Section 1.2 and detailed in Section 5. It explains the actions that were taken in order to fulfill the goals and their results. A list of recommendations compiled from the thoughts documented Chapter 8 is included. Finally, more ideas on how to apply and extend the concept and the simulation environment documented in Part II are offered.

8. Simulation Based Analysis of Forecast Performance

This chapter lists ideas on forecast performance and forecast evaluation methods. These ideas are used to design simulation experiments that illustrate their ramification and consequences. Considered are the long-term effect of revenue management methods, aspects of error measurements, the use of psychic forecasts, uncertainty of demand, and the inclusion of price-sensitivity in forecasts.

8.1. Observations on Long-Term Effects of Forecast Methods

When forecasts use historical data generated under their own influence, the repeated application of revenue management methods leads to an evolving dynamic. A common example for this is the so-called *spiral-down effect*. Its theoretical background has been described in Cooper et al. (2006). In this section, it is used as an example for long-term effects of revenue management methods.

The spiral-down effect can be expected when forecasting methods based on the assumption of independent demand meet flexible customers. Such customers tend to buy the cheapest acceptable class available. As a consequence, the forecast method will systematically predict more demand for low-fare classes and less demand for valuable classes. An optimization using this forecast reserves less seats for valuable classes and allows more availability in low-fare classes. The forecast becomes a self-fulfilling prophecy as the increased availability is used by flexible customers and more bookings are observed in cheap classes.

This effect is the result of a combination of methods and demand model. Therefore, the spiral-down effect can be observed best when the complete system is analyzed as proposed in Granger & Pesaran (2000).

The simulation system offers ways of evaluating the development over time in fast-forward mode. As the demand model can be kept stable, the effects of the cycle of bookings-forecast-optimization-bookings-etc. can be observed over dozens of simulation runs. The results of the complete system and the accuracy of the forecast component are evaluated.

Revenue Management Configurations: To validate statements on long-term effects, four forecast methods are applied to the price-based scenario with customers choosing between restriction-free classes. Implemented are three variations of an exponential smoothing forecast in combination with EMSR-b. They differ in the weight that is attached to new observations. While *"Exp025"* includes these new observation with a smoothing factor of $\alpha = 0.25$, *"Exp050"* uses a smoothing factor of $\alpha = 0.5$ and *"Exp075"* applies a smoothing factor of $\alpha = 0.75$. To provide a lower boundary, a first-come first-serve strategy is provided and referred to as *"FCFS"*.

Given the combination of demand, supply, and methods, over the course of several simulation runs, shifts in bookings and availability as well as plunging yield and revenue should manifest. This does hold true as shown by the simulation results depicted in the further text and figures. The spiral-down effect can be observed best during the first simulation runs – for this reason, further examinations focus on $s = 1..20$. In consecutive runs, the development slows down as it approaches a steady state.

Indicators: Several indicators can be used to describe the spiral-down effect. It has consequences for:

- predicted demand (forecast),
- protected seats (availabilities),
- bookings,
- revenue,
- yield,
- forecast evaluations (error measurements).

With regard to their expected development, these indicators and their formulaic expression are listed in the following paragraphs. The outcome of the respective experiment is used to illustrate the ideas presented in this section.

- Let $c \in C$ be the index of restriction-free booking classes ordered descending by their price.
- Let $f \in F$ be the index of flights.
- Let $t = 0, ..., N^t$ be points of time before departure; demand arrives after $t = 0$, $t = N^t$ is the time of departure.
- Let $s = 1, ..., N^s$ be the runs included in a simulation. $s + 1$ occurs chronologically after s and can be based on historical data derived from s.
- Let $f^{\text{unc}}(f, c, t, s)$ be the unconstrained demand forecast per class c on flight f for the time between $t - 1$ and t of simulation run s.
- Let $\hat{a}(f, c, t, s)$ be the protected seats per class c on flight f at point of time t in simulation run s.
- Let $a(f, c, t, s)$ be the available seats per class c on flight f at point of time t in simulation run s.
- Let $b(f, c, t, s)$ be the bookings per class c on flight f that arrived between points of time $t - 1$ and t of simulation run s.
- Let R_s be the set of requests assigned to run s in the simulation demand model. For every simulation run, the demand model is equivalent, even though individual requests differ due to the error terms included: $R_s \equiv R_{s+1}$
- Let $r(s)$ be the overall revenue generated in simulation run s.

Forecast: When the forecast is based on historical bookings and customers book the cheapest tickets available, the amount of predicted requests for valuable classes decreases while that for cheap classes increases. In the real world, rather than being updated for simulation runs s, the forecast is computed per departure. As no historical data is available in the first run ($s = 1$), a psychic forecast is used to generate the first prediction. Over consecutive runs, the forecast is updated using exponential smoothing.

- Let the psychic forecast be described by a function F^{psy}.
- Let the exponential smoothing forecast be described by a function F^{hist}.

Definition (8.1) describes how forecasts are generated:

$$f^{\text{unc}}(f,c,t,s) = \begin{cases} F^{\text{psy}}(R_s, F, C, t, 1) & s = 1 \\ F^{\text{hist}}(b(f,c,t,s-1), a(f,c,t,s-1), f^{\text{unc}}(f,c,t,s-1)) & s > 1 \end{cases}$$
$$\forall\ f \in F;\ c \in C;\ t = 1,...,N^t;\ s = 1,...,N^s \qquad (8.1)$$

If classes C are ordered in the descending order of their value, class 1 is the most expensive class while class N^c is the cheapest. Given the conditions of the spiral-down effect, the following development is expected: Forecasts for valuable classes decrease, while forecasts for cheap classes increase. This can be expressed formally as in Hypothesis (8.2):

$$\begin{aligned} &\lim_{s\to\infty} f^{\text{unc}}(f,c,t,s) = 0 \quad c < N^c \\ &\exists\ n_s \in \mathbb{N} \mid f^{\text{unc}}(f,N^c,t,s) \leq f^{\text{unc}}(f,N^c,t,s+n_s) \\ &\forall\ f \in F;\ c \in C;\ t = 1,...,N^t;\ s = 1,...,N^s - 1 \end{aligned} \qquad (8.2)$$

An indicator independent of the overall amount of predicted demand can be derived by computing the percentage of forecasted requests per class. This forecast-mix $f^{\text{unc}\ \%}$ is expected to show the behavior described in Hypothesis (8.2), normalized to 100%.

$$\begin{aligned} &f^{\text{unc}\ \%}(f,c,t,s) = \frac{f^{\text{unc}}\ (f,c,t,s)}{\sum_{c\in C} f^{\text{unc}}\ (f,c,t,s)} \cdot 100 \\ &\forall\ f \in F;\ c \in C;\ t = 1,...,N^t;\ s = 1,...,N^s \end{aligned} \qquad (8.3)$$

$$\begin{aligned} &\lim_{s\to\infty} f^{\text{unc}\%}(f,c,t,s) = \begin{cases} 0 & c < N^c \\ 100 & c = N^c \end{cases} \\ &\forall\ f \in F;\ c \in C;\ t = 1,...,N^t \end{aligned} \qquad (8.4)$$

Figure 8.1 shows the amount of demand predicted in the five classes as method *"Exp050"* is applied. In order to make the development of the forecast-mix comparable over the scenario variations, forecasts are expressed as percentages of overall predicted demand,

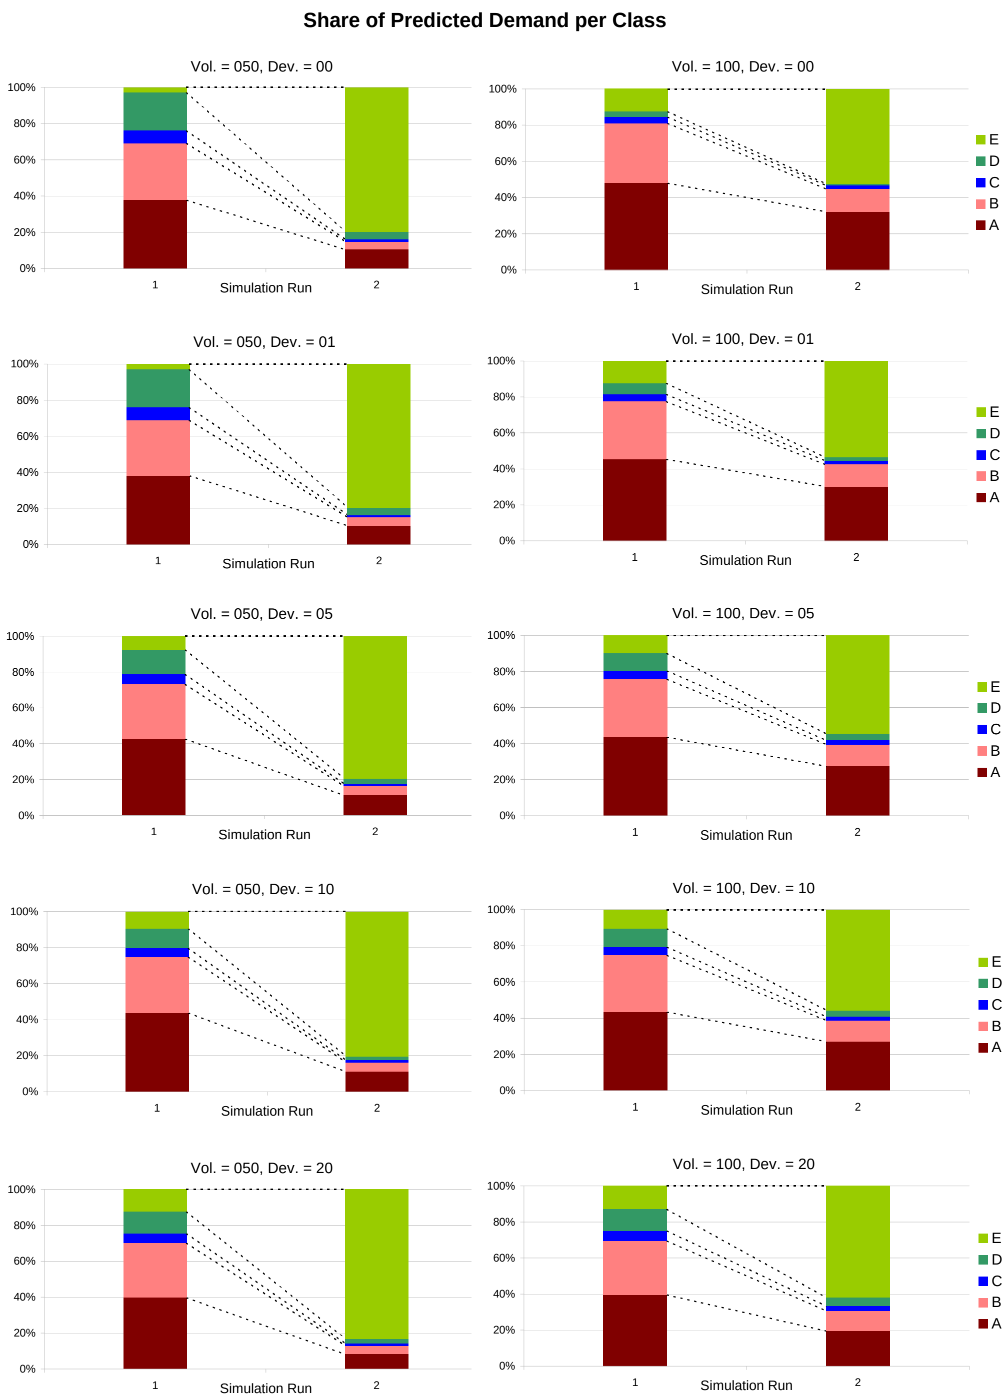

Figure 8.1.: Predicted Demand per Class with Exp050

$f^{\text{unc}\,\%}$ as described in Definition (8.3). As can be seen, forecasts for the most valuable class, *"A"*, decrease over the course of 50 runs as those for the cheaper classes offered, *"B", "C", "D", and "E"* increase.

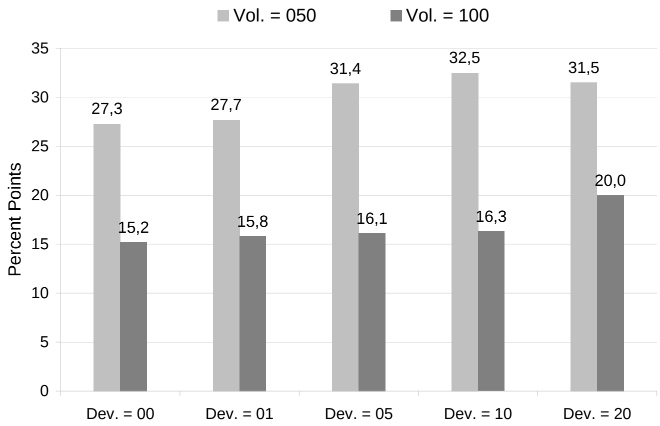

Figure 8.2.: Decrease in Demand Predicted for Class "A"

The degree of this development differs over the different markets. When demand volume is high as in *"Vol. = 100"*, it is not as steep as when observing markets with low demand volume as in *"Vol. = 050"*. Additionally, the decrease grows stronger with increasing deviation of the error term distribution. The decrease of the share of predicted demand for the class *"A"* in percent points is illustrated by Figure 8.2. However, the trend described is always the same and can be observed in all the variations of the price-based scenario.

Availabilities: Availabilities are the result of the optimization, which is based on the demand forecast. When EMSR-b is applied, the result is expressed as protections, seats reserved for valuable classes. In a system of nested classes as described in Section 7.3.3, protected seats are used to compute authorization levels, ensuring a minimum of availability for valuable classes.

As shown in Definition (8.5), the function A^{emsrb} uses the EMSR-b algorithm to generate protections for class c on flight f and run s for the point of time before the booking horizon starts, $t = 0$, $\hat{a}(f, c, 0, s)$, from the forecast $f^{\mathrm{unc}}(f, c, t, s)$ and the prices per class c on flight f, $p(f, c)$:

$$\hat{a}(f, c, 0, s) := A^{\mathrm{emsrb}}\left(\sum_{t=1}^{N^t} f^{\mathrm{unc}}(f, c, t, s), p(f, c)\right) \qquad \forall\ f \in F;\ c \in C;\ s = 1, ..., N^s \tag{8.5}$$

Given the conditions of the spiral-down effect, the following development is expected: As forecasts for valuable classes decrease, so do the protected seats computed for these classes. This can be expressed formally as presented in Hypothesis (8.6).

$$\begin{gathered} \lim_{s \to \infty} \hat{a}(f, c, 0, s) = 0 \quad c < N^c \\ \exists\ n_s \in \mathbb{N} \mid \hat{a}(f, N^c, 0, s) \leq \hat{a}(f, N^c, 0, s + n_s) \\ \forall\ f \in F;\ c \in C;\ s = 1, ..., N^s + 1 \end{gathered} \tag{8.6}$$

An indicator that is independent of the overall capacity can be derived by computing the percentage of seats allocated to each class. This availabilities-mix $\hat{a}^{\%}$ is expected to show the same behavior described in Hypothesis (8.6), normalized to 100%.

$$\begin{gathered} \hat{a}^{\%}(f, c, t, s) = \frac{\hat{a}(f, c, t, s)}{\sum_{c \in C} \hat{a}(f, c, t, s)} \cdot 100 \\ \forall\ f \in F;\ t = 1, ..., N^t;\ s = 1, ..., N^s \end{gathered} \tag{8.7}$$

$$\begin{gathered} \lim_{s \to \infty} \hat{a}^{\%}(f, c, 0, s) = \begin{cases} 0 & c < N^c \\ 100 & c = N^c \end{cases} \\ \forall\ f \in F;\ c \in C \end{gathered} \tag{8.8}$$

Figure 8.3 shows the amount of seats protected in the five classes as method *"Exp050"* is applied. In order to make the development of the protection-mix comparable over the scenario variations, protected seats are expressed as percentages of overall capacity, $a^{\%}$ as described in Definition (8.7). Protected seats for the most valuable class, *"A"*, decrease over the course of 50 runs as those for the cheaper classes offered, *"B", "C", "D", and "E"* increase.

The degree of this development differs over the different markets. When demand volume is high as in *"Vol. = 100"*, it is not as steep as when observing markets with low demand

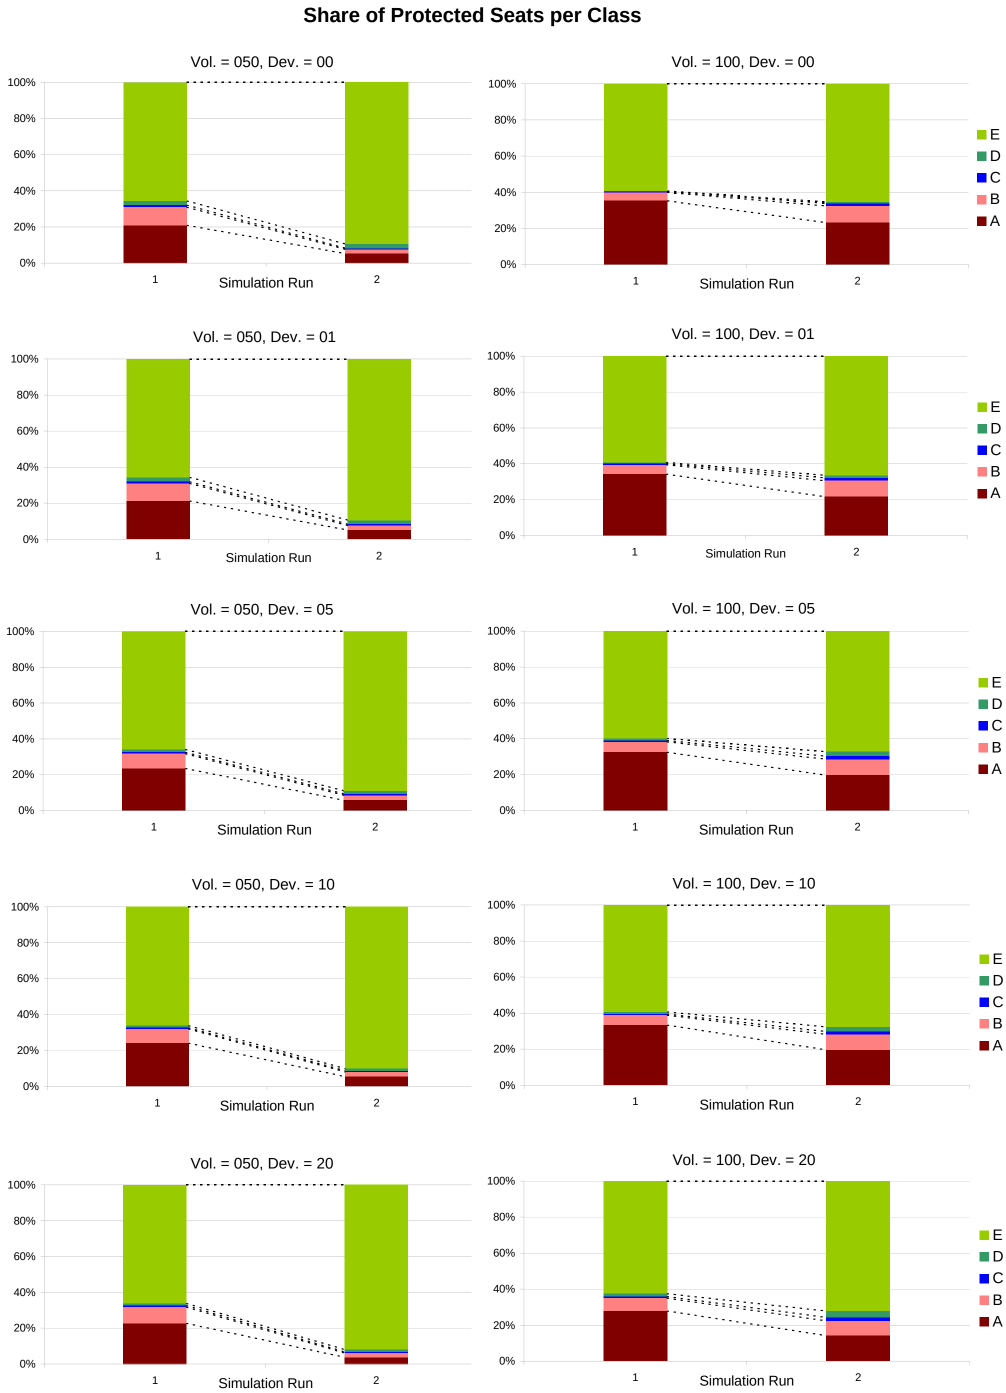

Figure 8.3.: Protected Seats per Class with Exp050

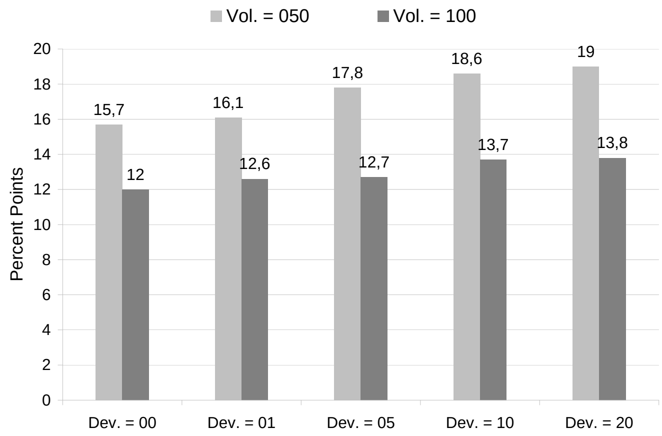

Figure 8.4.: Decrease in Seats Protected for Class "A"

volume as in *"Vol. = 050"*. Additionally, the decrease grows stronger with increasing deviation of the error term distribution. The decrease of the share of protected seats for the class *"A"* in percent points is illustrated by Figure 8.4. However, the trend described is always the same and can be observed in all the variations of the price-based scenario.

Bookings: In the demand model implemented, given the same product, customers always book the cheapest class available. For a set of restriction-free classes $c = 1, ..., N^c$ with similar product characteristics and increasing value, if overall demand volume stays constant, this means that the amount of bookings $b(f, c, t, s)$ depends solely on the availabilities $a(f, c, t, s)$. As availabilities for cheap classes increase, so do the bookings in these classes. Formally, this expectation can be expressed as follows:

$$\begin{aligned} \lim_{s \to \infty} b\left(f, c, t, s\right) = 0 \quad c < N^c \\ \exists\ n_s \in \mathbb{N} \mid b\left(f, N^c, t, s\right) \leq b\left(f, N^c, t, s + n_s\right) \\ \forall\ f \in F;\ \ c \in C;\ \ t = 1, ..., N^t;\ \ s = 1, ..., N^s - 1 \end{aligned} \tag{8.9}$$

An indicator that is independent of the overall bookings can be derived by computing the percentage of bookings for each class. This bookings-mix $b^{\%}$ is expected to show the same behavior described in Hypothesis (8.9), normalized to 100%.

$$b^{\%}(f,c,t,s) = \frac{b(f,c,t,s)}{\sum_{c\in C} b(f,c,t,s)} \cdot 100 \qquad \forall\ f \in F;\ c \in C;\ t = 1,...,N^t;\ s = 1,...,N^s \tag{8.10}$$

$$\lim_{s\to\infty} b^{\%}(f,c,t,s) = \begin{cases} 0 & c < N^c \\ 100 & c = N^c \end{cases} \qquad \forall\ f \in F;\ c \in C;\ t = 1,...,N^t \tag{8.11}$$

Figure 8.5 shows the amount of seats booked in the five classes as method *"Exp050"* is applied. In order to make the development of the booking-mix comparable over the scenario variations, bookings per class are expressed as percentages of overall bookings, $b^{\%}$ as described in Definition (8.10). As can be seen over all scenario variations, bookings for the most valuable class, *"A"*, decrease over the course of 50 runs as those for the cheaper classes offered, *"B", "C", "D", and "E"* increase.

The degree of this development differs over the different markets. When the deviation of the error term distribution is low, it is stronger for markets with high demand. When the deviation of the error term is high, the opposite seems to be true. The decrease grows stronger with increasing deviation of the error term distribution for *"Vol. 050"*, but not for *"Vol. 100"*. The decrease of the share of bookings for the class *"A"* in percent points is compared is illustrated by Figure 8.6. The trend described is always the same and can be observed in all the variations of the price-based scenario.

Revenue: The major indicator in revenue management, overall revenue per run, $r(s)$, is computed as the sum over the product of bookings $b(f,c,t,s)$ and the price of classes $p(f,c)$. Outside a simulation, this may indicate the revenue generated by all flights on one departure day.

$$r(s) = \sum_{f\in F}\sum_{c\in C}\left(p(f,c)\cdot\sum_{t=1}^{N^t} b(f,c,t,s)\right) \qquad \forall\ s = 1,...,N^s \tag{8.12}$$

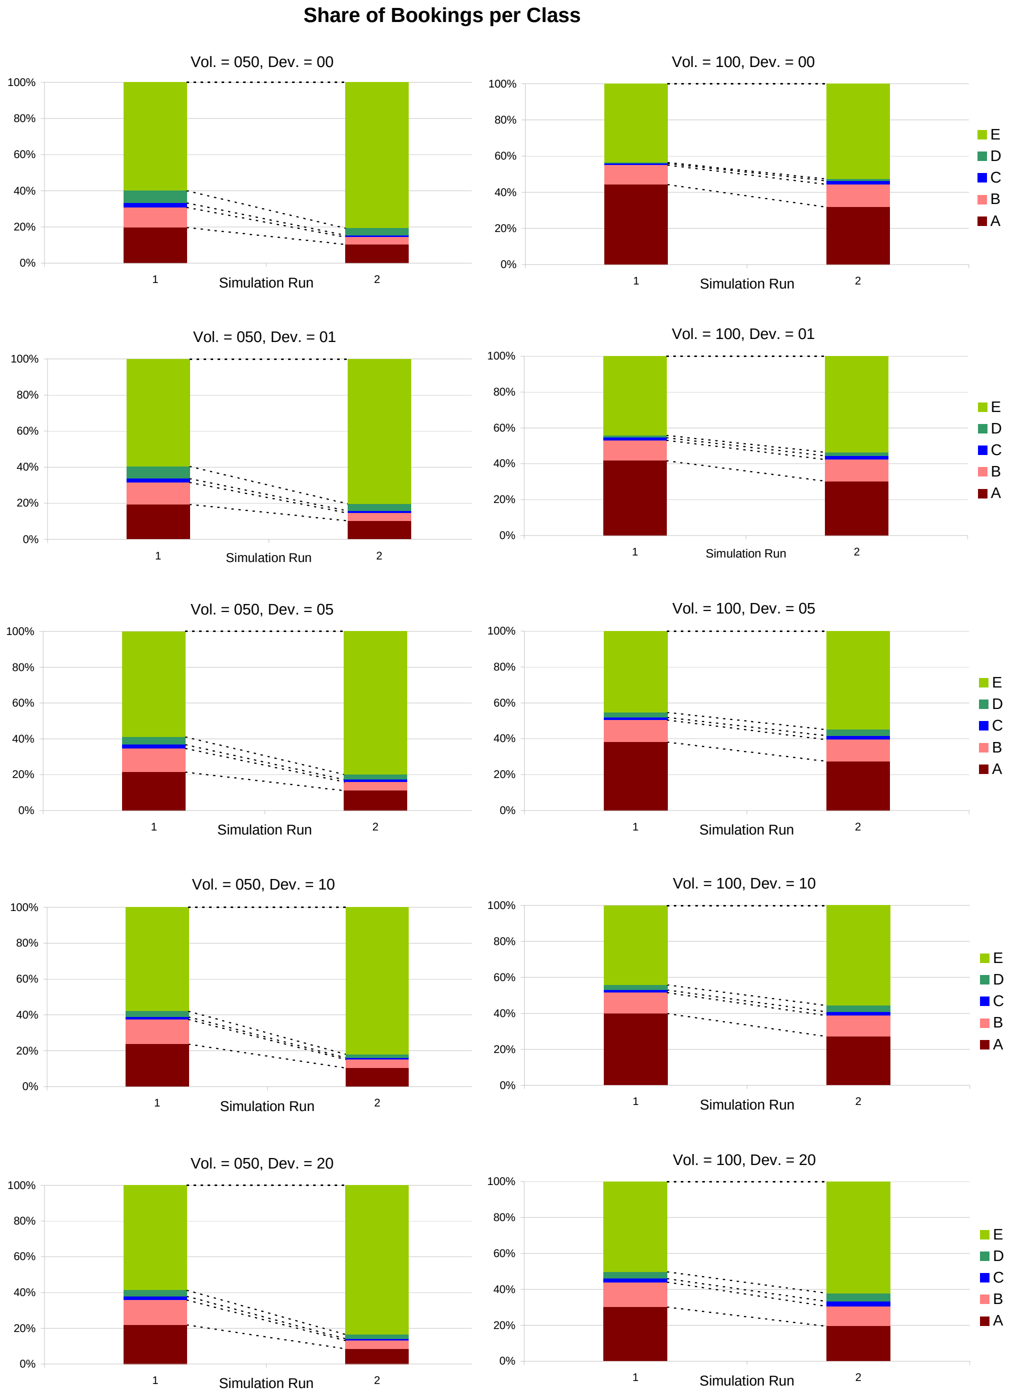

Figure 8.5.: Observed Bookings per Class with Exp050

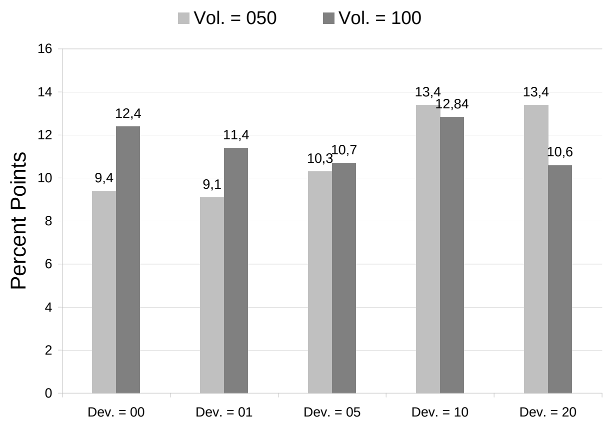

Figure 8.6.: Decrease in the Share of Bookings Observed for Class "A"

If overall demand volume stays constant and the booking mix changes according to Hypothesis (8.9), less bookings in valuable classes and more bookings in expensive classes lead to decreasing overall revenue. Formally, this can be expressed as follows:

$$\begin{aligned}\sum_{f\in F}\sum_{c\in C}\sum_{t=1}^{N^t} b\left(f,c,t,s\right) \equiv \sum_{f\in F}\sum c \in C\sum_{t=1}^{N^t} b\left(f,c,t,s+1\right)\\ \rightarrow r\left(s\right) \leq r\left(s+1\right)\\ \forall\ s=1,...,N^s-1\end{aligned} \tag{8.13}$$

However, the spiral-down effect may be beneficial for revenue if the gain in bookings due to increased availabilities for cheap classes compensates for the loss in bookings in valuable classes. There is actually a break-even point from which on the revenue lost to buy-down is compensated by that gained through low-fare acquisition:

$$\begin{aligned}&\sum_{t=1}^{N^t}\left(b\left(f,N^c,t,s\right)-b\left(f,N^c,t,s+1\right)\right)\cdot p\left(f,N^c\right)\\ &\leq \sum_{t=1}^{N^t}\sum_{c=1}^{N^c-1}\left(b\left(f,c,t,s+1\right)-b\left(f,c,t,s\right)\cdot p\left(f,c\right)\right)\\ &\qquad\forall\ f\in F;\ s=1,...,N^s-1\end{aligned} \tag{8.14}$$

To observe the development of revenue over time, a percentage indicator may be calculated. Given an initial simulation run $s = 1$, revenue for all future $s = 2, ..., N^s$ may be converted to a percentage $r^{\%}$ of the revenue earned during the initial run.

$$r^{\%}(s) = \frac{r(s)}{r(1)} \cdot 100 \quad \forall\ s = 1, ..., N^s \tag{8.15}$$

$$\sum_{f \in F} \sum_{c \in C} \sum_{t=1}^{N^t} b^{\%}(f,c,t,s) \equiv \sum_{f \in F} \sum_{c \in C} \sum_{t=1}^{N^t} b^{\%}(f,c,t,s+1) \rightarrow r^{\%}(s) \leq r^{\%}(s+1) \quad \forall\ s = 1, ..., N^s - 1 \tag{8.16}$$

Figure 8.7 shows the revenue earned as exponential smoothing methods "Exp025", "Exp050", and "Exp075" are applied. In order to make the development of revenue comparable over the scenario variations, it is expressed as percentages of the revenue earned in run 1, $r^{\%}$ as described in Definition (8.15). As can be seen over all scenario variations, revenue decreases over the course of 50 runs.

The form of this development differs over the different markets. With high demand volume *"Vol. = 100"*, as the conditions described by Hypothesis (8.14) do manifest during the first runs, a small revenue increase can be observed in the beginning of the simulation experiment. However, soon the loss of bookings in valuable classes stops being compensated by the gain of bookings in low-fare classes and overall revenue decreases. With low demand volume *"Vol. = 050"*, the decrease of revenue starts immediately after the first run – its consequences are also more severe. When deviation is high as in *"Dev. = 20"*, the development is not as straightforward as revenue shifts with volatile demand, yet it is even steeper. The trend described is always the same and can be observed in all the variations of the price-based scenario.

After a number of runs, revenue reaches a plateau that still exceeds what is earned with first-come-first-serve controls. Inventory controls that were originally based on the psychic forecast initialization keep some seats protected for valuable customers – and as long as these customers request tickets before cheaper classes are available, bookings will be observed in these classes. This leads to a halt in the spiral-down effect. Volatile

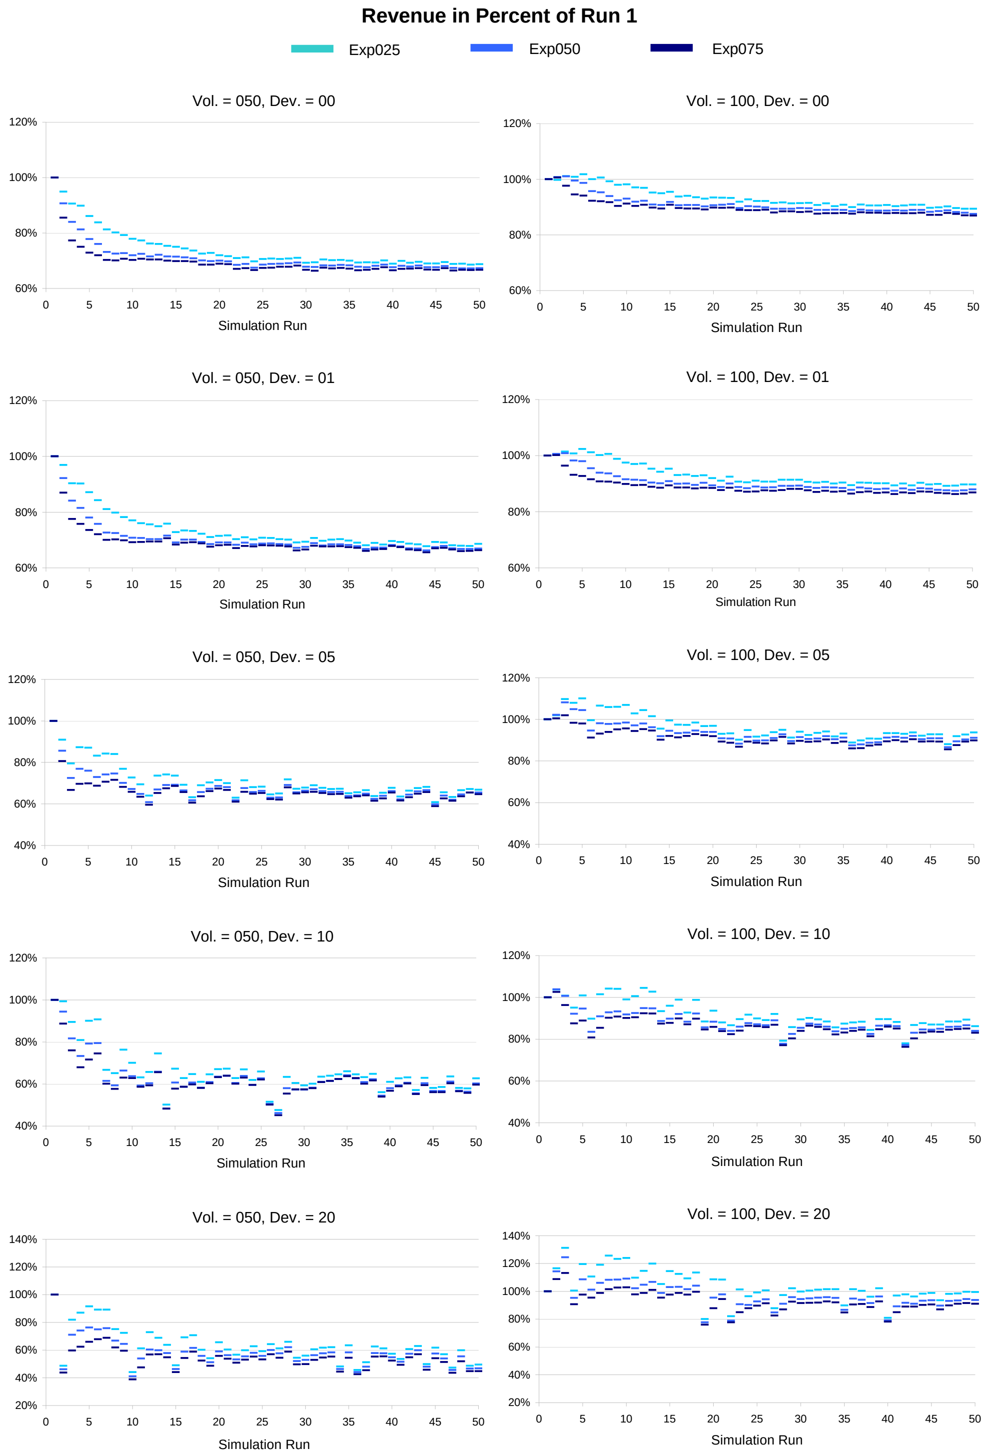

Figure 8.7.: Revenue in Percent of Revenue Earned in Run 1

customer arrival patterns, however, can still lead to a further shift toward cheap classes, while no shift toward expensive classes can occur in a purely price-based market.

When several methods are compared, the percentage difference to a lowest benchmark may be computed. A possible benchmark is the result of a first-come-first-serve seat allocation, $r^{\text{fcfs}}(s)$. The indicator $r^{\text{fcfs}\,\%}(s)$ is computed as the percentage by which the revenue of the considered method exceeds that gained when first-come-first-serve was applied.

$$r^{\text{fcfs}\,\%}(s) = \frac{r(s) - r^{\text{fcfs}}(s)}{r^{\text{fcfs}}(s)} \cdot 100 \qquad \forall\ s = 1, ..., N^s \tag{8.17}$$

Based on forecasts becoming self-fulfilling prophecies, the spiral-down effect is even more severe when the forecast method picks up new data quickly. In Figure 8.7 the development of revenue given the application of the three exponential smoothing methods is shown. The indicator used to depict the implications of spiral-down for revenue depending on the method used is $r^{\text{fcfs}\,\%}$ as described by Definition (8.17). Revenue decreases more quickly when the $\alpha^{\exp}$ employed to weight new values in the forecast is higher.

Yield: Yield is the average revenue gained per booking. It can be computed as shown by dividing the overall sum of revenue by the overall sum of bookings.

$$y(s) = \frac{\sum_{f \in F} \sum_{c \in C} \left(p(f,c) \cdot \sum_{t=1}^{N^t} b(f,c,t,s) \right)}{\sum_{f \in F} \sum_{c \in C} \sum_{t=1}^{N^t} b(f,c,t,s)} \qquad \forall\ s = 1, ..., N^s \tag{8.18}$$

The development predicted for revenue in case of spiral-down can be applied to yield as well. However, while extreme gains in low-fare bookings may lead to an overall revenue compensation, the yield decrease can be expected with certainty. As only the cheapest class, $c = 1$, is booked any more, the yield ends up as the price of this class.

$$\lim_{s \to \infty} y(s) = \frac{\sum_{f \in F} \left(p(f, N^c) \cdot \sum_{t=1}^{N^t} b(f, N^c, t, s) \right)}{\sum_{f \in F} \sum_{t=1}^{N^t} b(f, N^c, t, s)} \qquad \forall\ s = 1, ..., N^s \tag{8.19}$$

To better observe the development of yield over time, a percentage indicator may be calculated. Given an initial simulation run $s = 1$, yield for all future $s = 2..N^s$ may be converted to a percentage $y^{\%}$ of the yield observed during the initial run.

$$y^{\%}(s) = \frac{y(s)}{y(1)} \cdot 100 \tag{8.20}$$

$$\begin{aligned} &\sum_{f \in F} \sum_{t=1}^{N^t} b(f, N^c, t, s) \leq \sum_{f \in F} \sum_{t=1}^{N^t} b(f, N^c, t, s+1) \\ \text{and } &\sum_{f \in F} \sum_{t=1}^{N^t} b(f, 1, t, s) \geq \sum_{f \in F} \sum_{t=1}^{N^t} b(f, 1, t, s+1) \\ &\rightarrow y^{\%}(s) \leq y^{\%}(s+1) \\ &\forall\ s = 1, ..., N^s - 1 \end{aligned} \tag{8.21}$$

Figure 8.8 shows the amount yield earned as method *"Exp050"* is applied. In order to make the development of yield comparable over the scenario variations, it is expressed as percentages of the yield earned in run 1, $y^{\%}$ as described in Definition (8.20). As can be seen over all scenario variations, yield decreases over the course of 50 runs.

The form of this development differs over the different markets. As the yield is independent of the amount of overall bookings, the observed decrease in yield is much smoother than the decrease observed with regard to revenue. With low demand volume *"Vol. = 050"*, yield is lower even in the first run compared to high demand volume. In addition, the decrease of yield is steeper. A high deviation of the error term distribution contributes to this effect.

When several methods are compared, the percentage difference to a lowest benchmark may be computed. A possible benchmark in the simulation is the yield resulting from a first-come-first-serve seat allocation, $y^{\text{fcfs}}(s)$. The indicator $y^{\text{fcfs}\ \%}(s)$ is computed as the percentage by which the average yield of the considered method exceeds that observed when first-come-first-serve was applied.

$$\begin{aligned} y^{\text{fcfs}\ \%}(s) &= \frac{y^{\text{fcfs}}(s) - y(s)}{y^{\text{fcfs}}(s)} \cdot 100 \\ &\forall\ s = 1, ..., N^s \end{aligned} \tag{8.22}$$

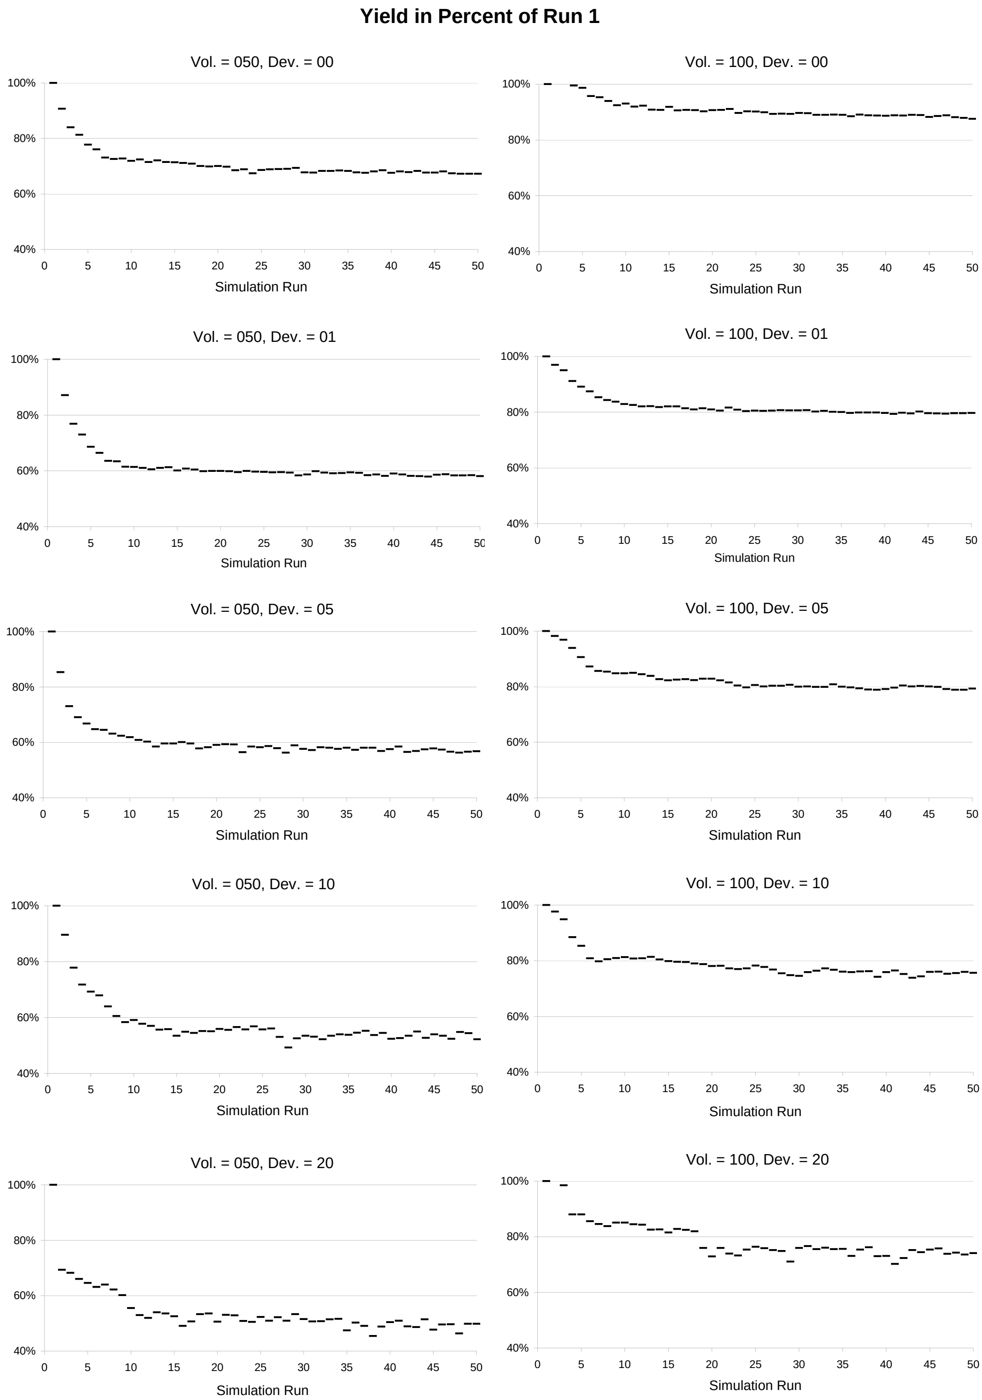

Figure 8.8.: Yield in Percent of Yield Earned in Run 1

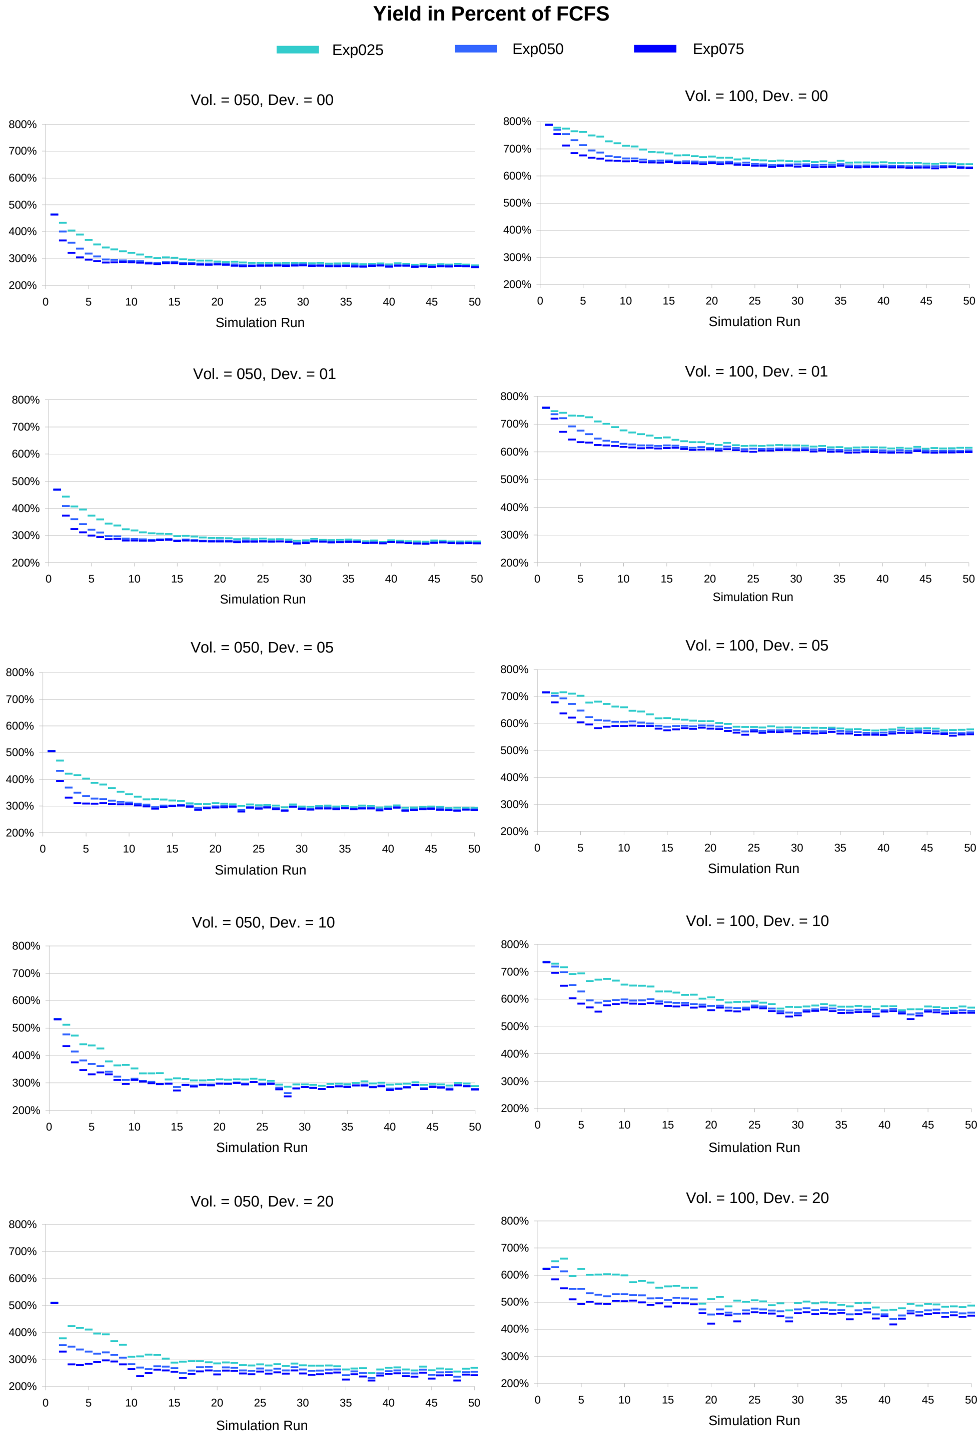

Figure 8.9.: Yield in Percent of Yield Earned with First-Come-First-Serve

In Figure 8.9, presents the development of yield given the application of the three exponential smoothing methods *"Exp025"*, *"Exp050"*, and *"Exp075"*. The indicator used to depict the implications of spiral-down for revenue depending on the method used is $y^{\text{fcfs}\,\%}$ as described by Definition (8.22). As the diagrams show, yield decreases more quickly when the α^{exp} employed to weight new values in the forecast is higher.

Forecast Evaluations: The expected development of forecast quality under the conditions of the spiral-down effect is the same for all error measurements based on the comparison of forecasted demand and bookings as outlined in Chapter 4.

- Let $e_{\circ}^{c-c}(s)$ be the average forecast error computed for simulation run s based on some to be defined method $\circ$ comparing observed bookings $b(f,c,t,s)$ and the constrained demand forecast $f^{const}(f,c,t,s)$.

As the forecast based on bookings predicts more demand to come for cheap classes and availabilities based on the forecast allow for this demand to realize in more bookings, the forecast becomes a self-fulfilling prophecy. This systematic flaw is interpreted as an improvement of forecast quality:

$$\begin{gathered} e_{\circ}^{c-c}(s) \geq e_{\circ}^{c-c}(s+1) \\ \forall\ s = 1, ..., N^s - 1 \end{gathered} \tag{8.23}$$

When the conditions of the spiral-down effect are fulfilled and forecasts are evaluated based on comparisons to actual bookings, their quality seems to improve. This can be validated by observing the development of MAD, RMSE, MAPE and U2 applied to the comparison of the constrained forecast f^{const} and observed bookings b over the course of 50 runs. All indicators show a decrease over time as forecast and bookings converge due to the spiral-down effect.

Figure 8.10 presents the development of MAD as three exponential smoothing methods are applied. Figure 8.11 presents the development of RMSE as three exponential smoothing methods are applied. Figure 8.12 presents the development of MAPE as three exponential smoothing methods are applied. Figure 8.13 presents the development of U2 as three exponential smoothing methods are applied.

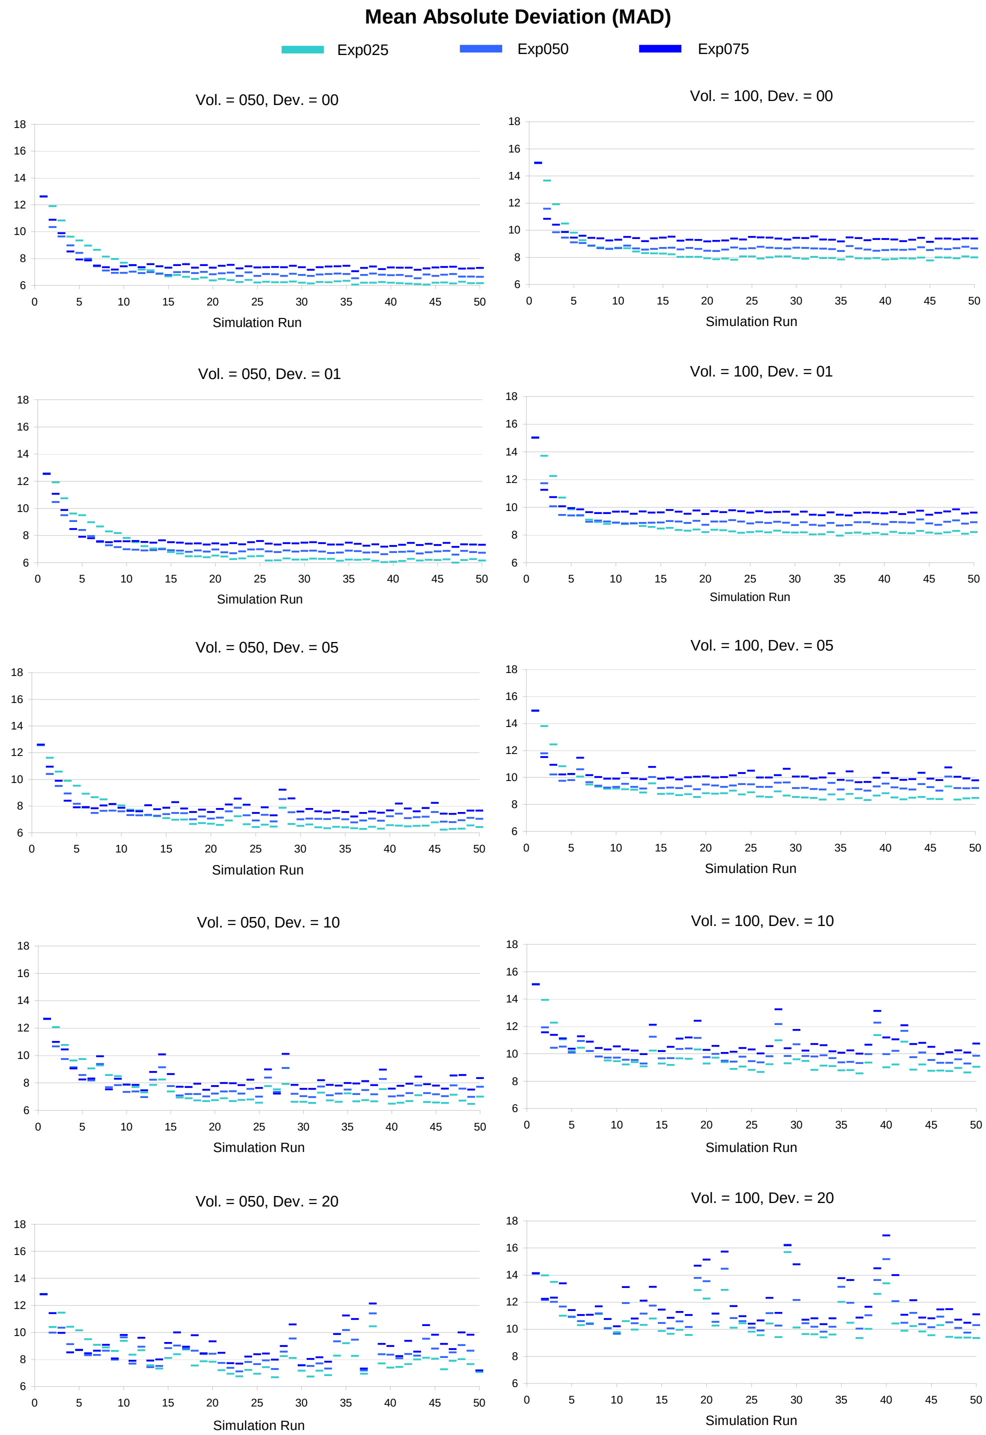

Figure 8.10.: Mean Absolute Deviation (MAD): Constrained FC from Observed BKD

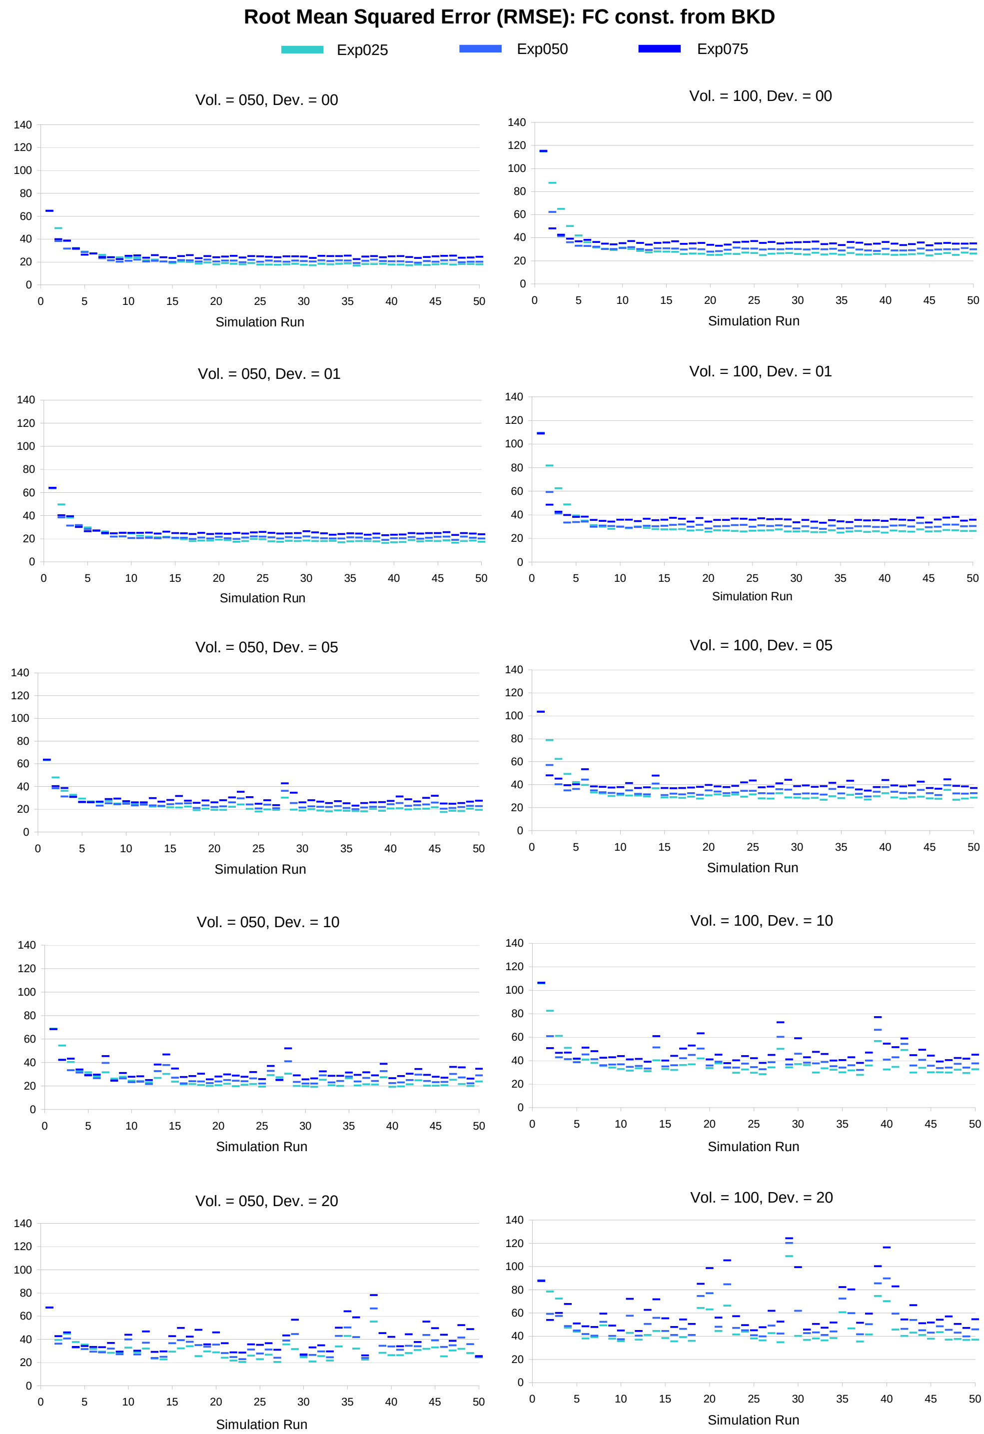

Figure 8.11.: Root Mean Squared Error (RMSE): : Constrained FC from Observed BKD

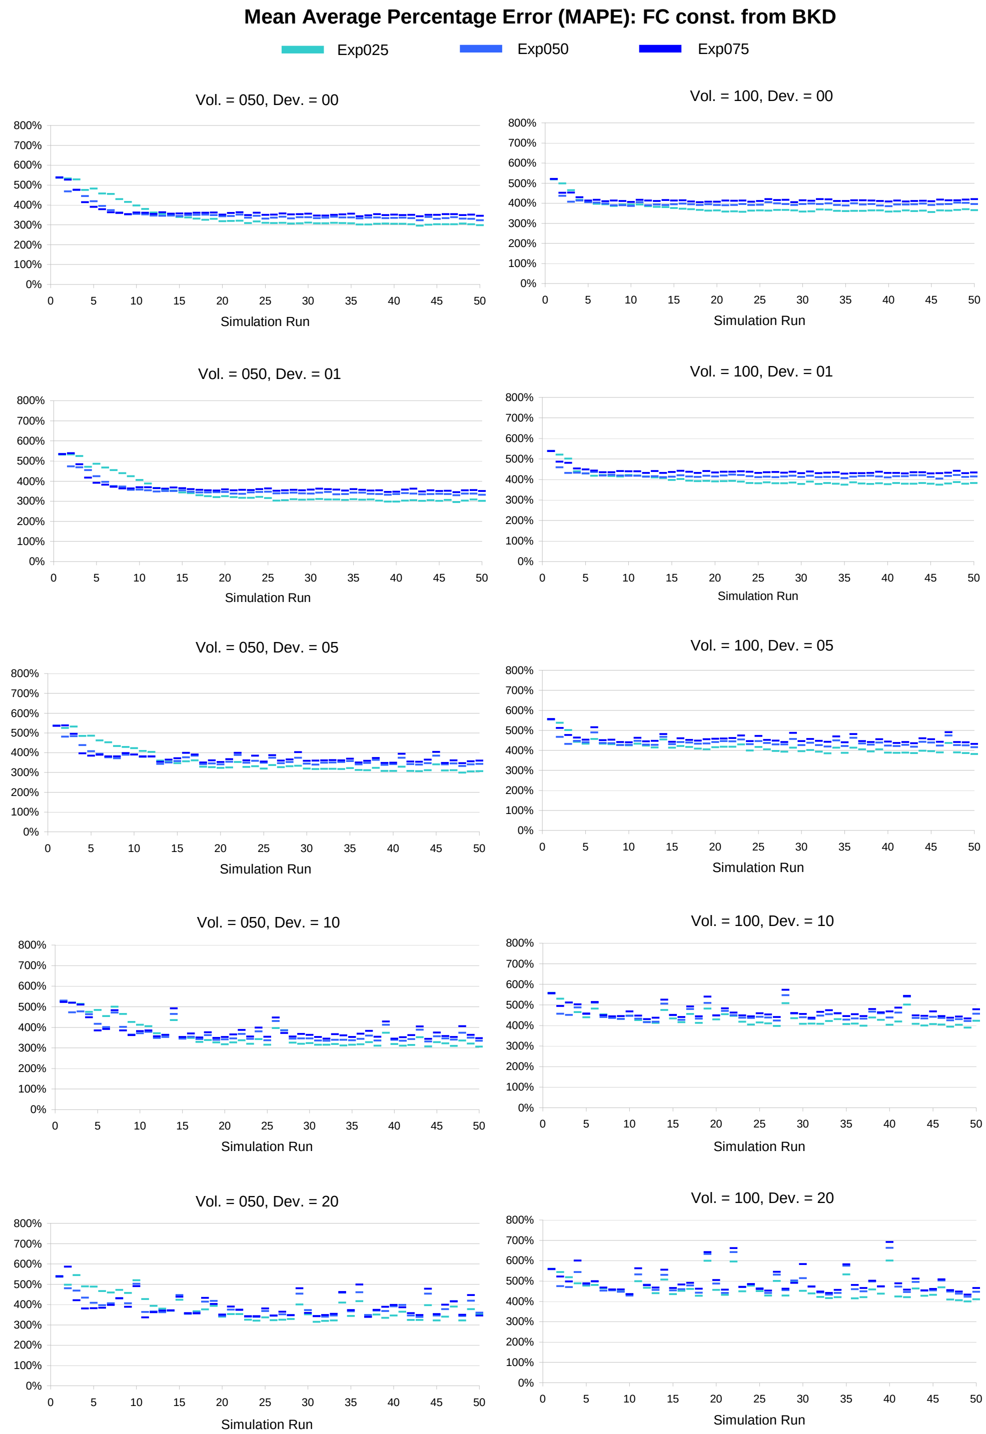

Figure 8.12.: Mean Avg. Percentage Error (MAPE): Constrained FC from Observed BKD

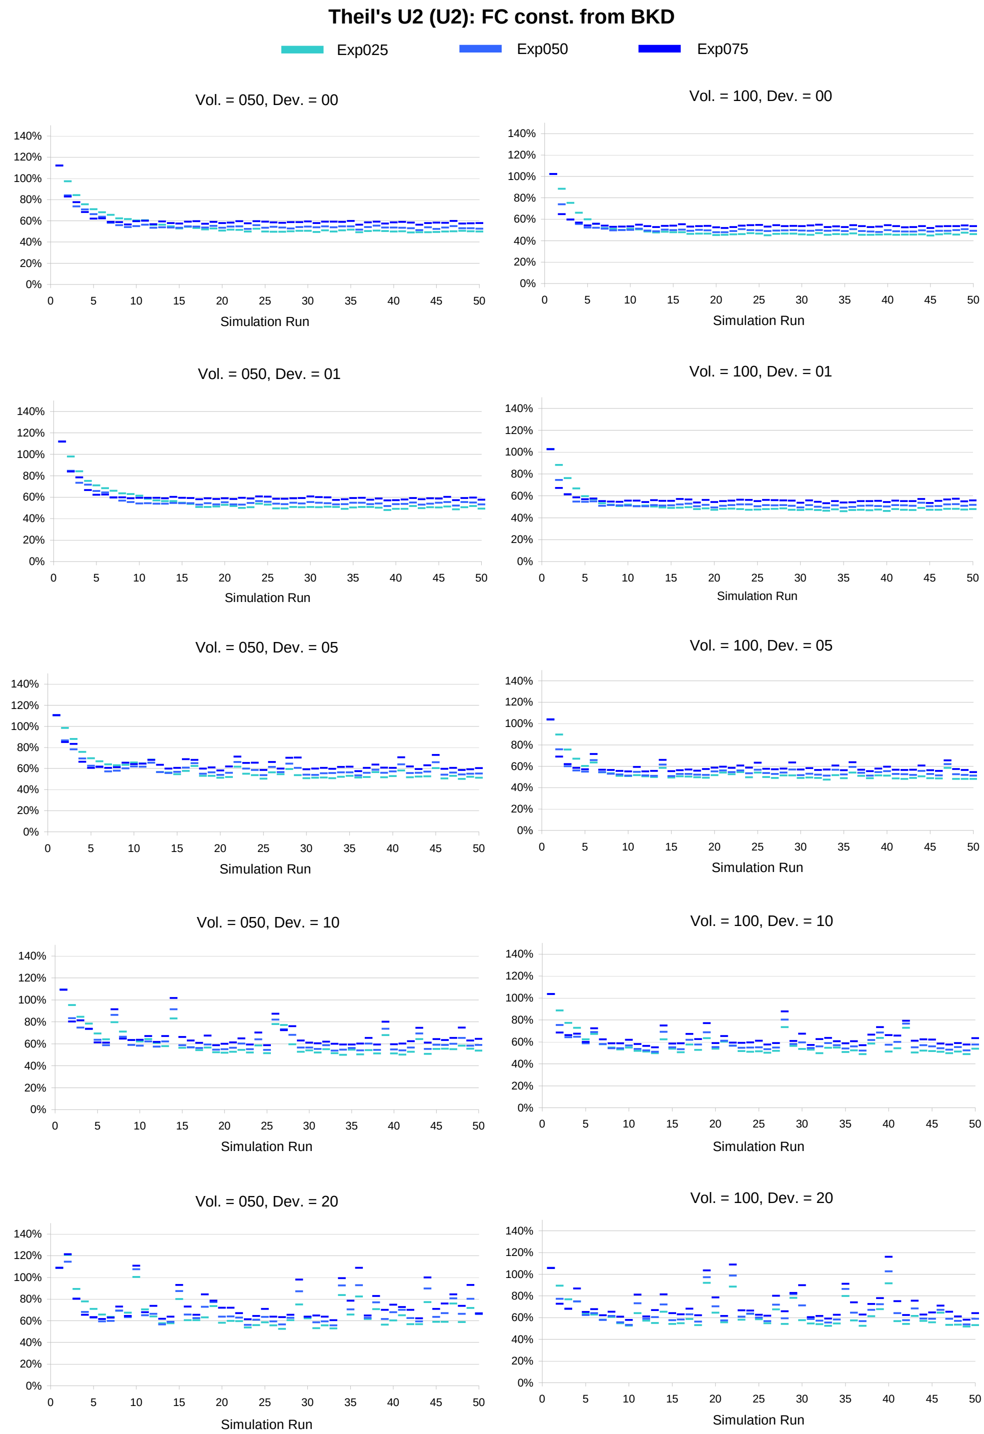

Figure 8.13.: Theil's U2 (U2): Constrained FC from Observed BKD

For all indicators, the trend of forecast quality over the course of a number of simulation runs is more volatile if the deviation of the error term distribution is higher. As demand volume shifts more strongly from one run to the next due to this, it becomes harder to predict and a self-fulfilling prophecy takes longer to manifest. As for other indicators, the development is steeper for markets with a low overall demand volume.

All indicators also show that just as the decrease of yield is steeper when new values are weighted stronger, forecast quality seems to improve quicker. When the new observations based on availabilities that are already influenced by the spiral-down effect are weighted heavier, the forecast turns into a self-fulfilling prophecy even more quickly.

Effects of Updates within the Booking Horizon: Finally, the spiral-down effect can even be observed within the booking horizon of a single run if forecasts are updated. Figure 8.14 shows a comparison of revenue over 50 runs for all price-based scenarios when exponential smoothing with a smoothing factor of 0.5 is applied. With *"Exp050"*, the forecast is not updated throughout the booking horizon as bookings are observed. With *"Exp050upd"*, the forecast is updated as described in Section 7.3.1. Revenue decreases with every run when the forecast is updated. This is due to availabilities updated based on decreasing shares of predicted demand for the most expensive classes as described earlier in this section.

However, as presented in Figure 8.15, this updating of the forecast does not even lead to better results concerning traditional forecast evaluation. Using MAD as an exemplary error measurement, it becomes clear that the computed error of the forecast decreases as the booking horizon progresses. This is due to the decreasing remaining time span for which the forecast is valid. However, while this development can be observed both for *"Exp050"* and *"Exp050upd"*, the overall level of mean absolute deviation is lower when the forecast is not updated based on observed bookings.

The graph is shown only for those variations of demand where the deviation of the error term is 0 ("Dev. = 00"). Not much variation can be observed as volume and deviation change – this is due to both the fact that constrained values are compared and that the comparison is limited to the first run, initialized using the psychic forecast.

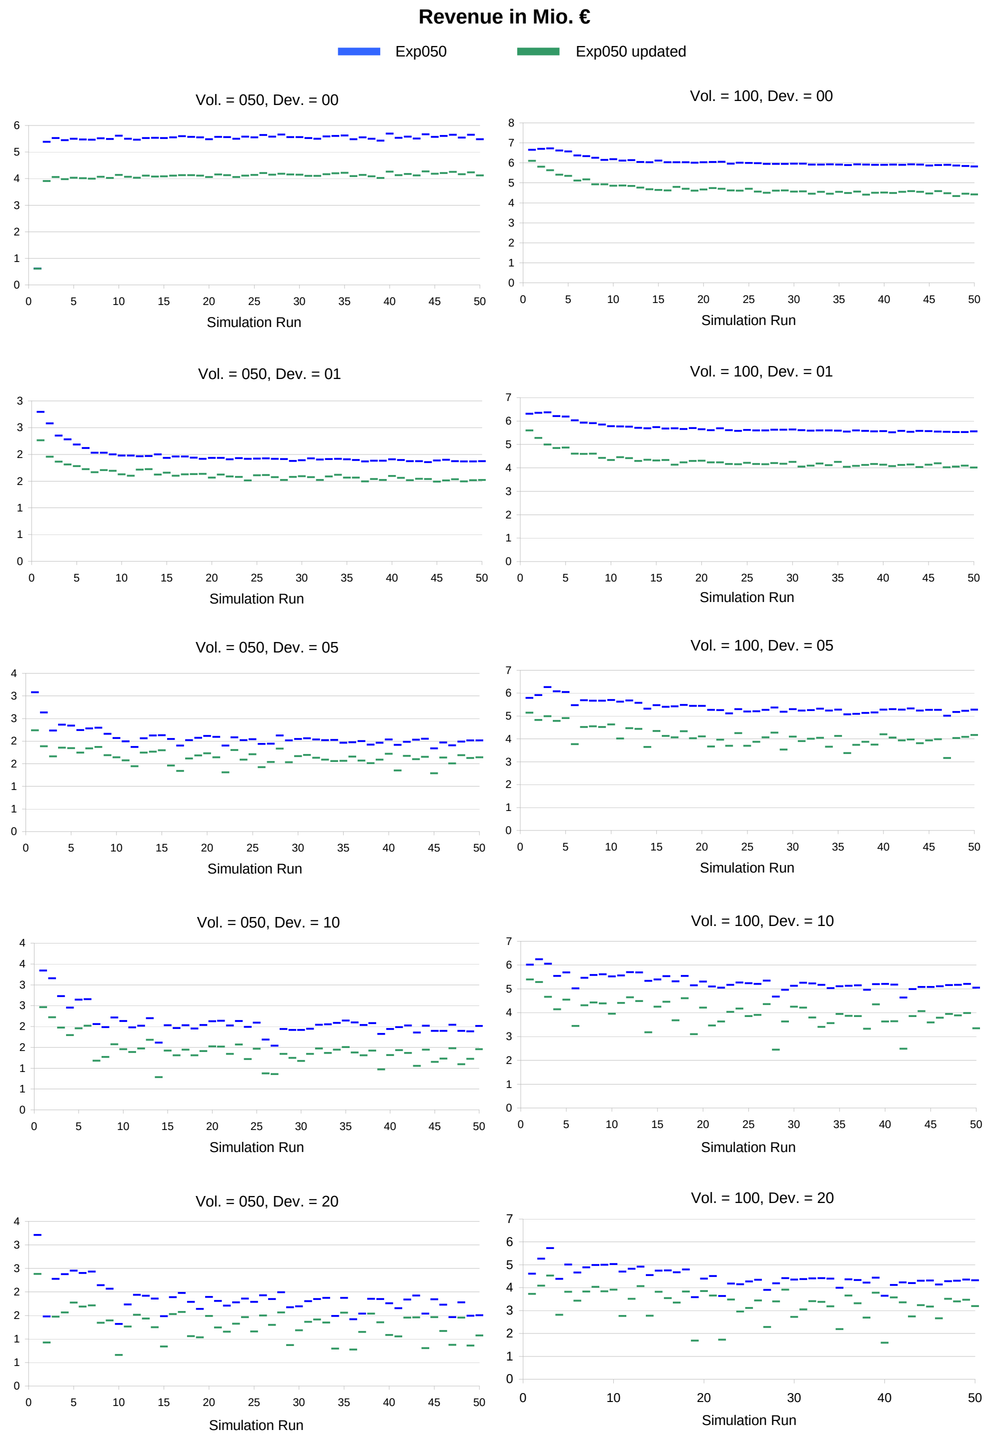

Figure 8.14.: Revenue Resulting from Exp050 and Exp050upd

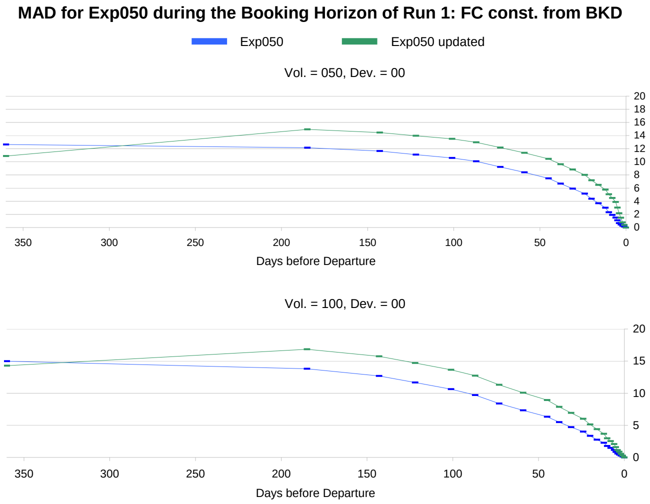

Figure 8.15.: MAD during the Booking Horizon of Run 1

Conclusion: When applied over longer terms, flawed assumptions included in adaptive forecast methods can cause systematic trends in inventory controls and revenue development. In addition, when causing a spiral-down effect, forecast methods can become self-fulfilling prophecies.

8.2. Consequences of Possible Definitions of Psychic Forecasts

In the previous section, the spiral-down effect was demonstrated by providing an initial forecast and then applying an adaptive method (exponential smoothing). The initial forecast was implicitly assumed to present a more accurate prediction of demand, with the adaptive method based on a model of static demand leading to a spiral-down effect when confronted with flexible customer behavior. The initial forecast was based on knowledge of the demand model implemented in the simulation, which is not available in the real world. This class of forecasts will be referred to as *psychic forecasts* in the further text.

The concept of psychic forecasts will be the focus of this section. When evaluating approaches to demand forecasting based on a simulation system, there are two applications for psychic forecasts: They may be used as initialization to observe the development of indicators when adaptive methods are applied or as a benchmark to compare a method's ability of picking up the characteristics of the demand model.

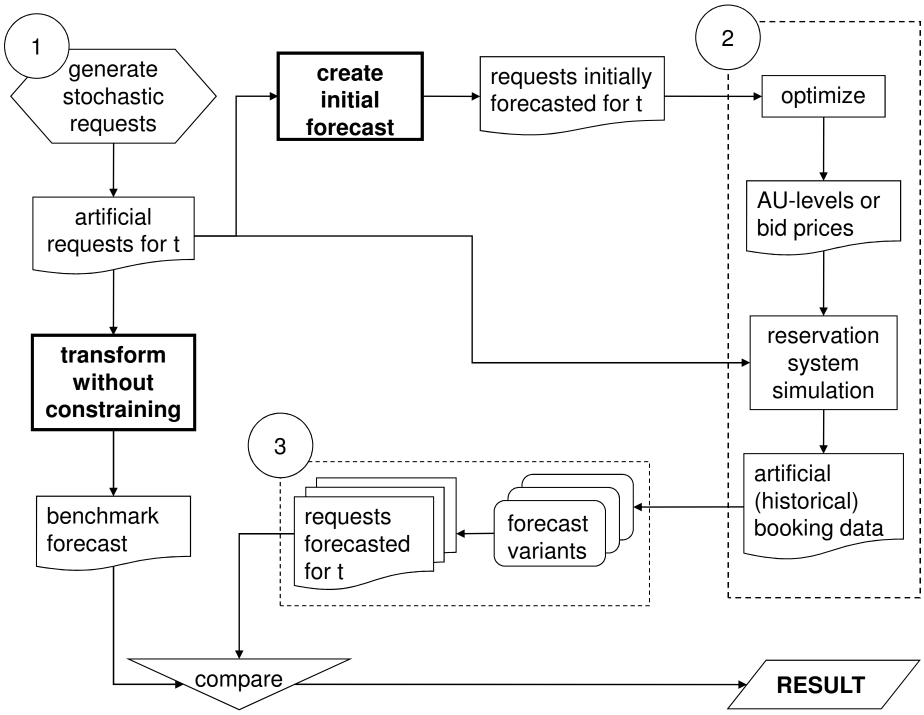

Figure 8.16.: Uses of the Psychic Forecast in the Simulation

Figure 8.16 emphasizes those parts of the system that have not been further examined so far. For example when evaluating the unconstraing aspect as described in Section 6.4, the previous text has regarded the psychic forecast as a constant method. It is employed in two places, both marked in bold.

Several approaches are conceivable to calculate a psychic forecast. When used as an initialization or as a benchmark for methods based on a static model of demand, the psychic forecast can be varied along two parameters: *Class choice* and *itinerary choice*. These choices relate to the ways in which customers decide which classes to book (class choice) on which flights (itinerary choice) in order to reach their destination.

	preferred itinerary	**alternative itineraries**
maximum price	c-max 1-i	c-max 2-i
minimum price	c-min 1-i	c-min 2-i
acceptable prices	c-frac 1-i	c-frac 2-i

Table 8.1.: Possible Variations of Choice in Psychic Forecasts

Two variations of *itinerary choice* and three variations of *class choice* as well as their combinations will be examined in this section. Together with their abbreviations, they can be found in Table 8.1.

In order to express the options of psychic forecasting and their consequences formally, additional notation is required. The following list recaptures some of the variables introduced in Section 7.2.2.

- Let $r \in R$ be the index of requests.
- Let R_s be the set of requests included in the run s.
- Let t_r indicate the point of time t at which request r arrives.
- Let $\delta^{\text{time}}_{r,t}$ be a boolean matrix indicating whether request r arrives at point of time t: $t = t_r$.
- Let s_r indicate the simulation run s for which request r was generated.
- Let $i \in I$ be the index of itineraries offered.
- Let I_r be the set of itineraries acceptable according to the product requirements specified for request r (i.e. origin, destination, travel time, departure day and time, transfers).
- Let F_i be the set of flights included in itinerary i.
- Let $I^{\text{it}}_{f,i}$ be a boolean matrix indicating whether flight $f \in F_i$ for any f and i.
- Let $\delta^{\text{product}}_{r,c}$ be a boolean matrix defining whether the product represented by booking class c is acceptable according to the product requirements of request r (see Definition (7.19).

- Let $\delta^{\text{wtp}}_{r,c,i}$ be a boolean matrix defining whether the price of class c on itinerary i is acceptable according to the willingness to pay of request r (see Definition (7.21)).
- Let $p\,(f,c)$ be the price of booking class c on flight f.
- Let $f^{\text{unc}}(f,c,t,s)$ be the unconstrained demand forecast for flight f, class c, and point of time t of simulation run s.
- Let $\hat{C}(i,r)$ be the cost of itinerary i considered by request r, without regard for the actual ticket price (given the assumption that all itineraries cost the same).

Preferred Itinerary (1-i): Psychic forecasts predicting demand to arrive only for the preferred itinerary use knowledge of the customers' cost function given by the simulation. Demand is expected to arrive in a given booking class of the itinerary that would be chosen if all itineraries were available at the same price. The hypothesis is that the demand is most likely to manifest in the itinerary that presents the first choice according to the cost function. The formal expression for the choice of itinerary in this psychic forecast is given in Definitions (8.24).

- Let $\delta^{\text{i-choice}}_{i,r}$ be a boolean matrix indicating the chosen itinerary based on the cost function of request r.
- Let n^i be the number of itineraries that are chosen according to the method of psychic forecast. For "1-i", $n^i = 1$.

$$\delta^{\text{i-choice}}_{i,r} := \begin{cases} 1 & i = i' |\max_{i' \in I_r} \hat{C}\,(i',r) \\ 0 & \text{else.} \end{cases} \qquad \forall\, r \in R;\; i \in I_r \tag{8.24}$$

Alternative Itineraries (n-i): Psychic forecasts that predict demand for two or more itineraries use knowledge of the customers' cost function. Demand is expected to arrive in a given booking class of those itineraries that make the "top list" if all itineraries were available at the same price. The number of itineraries n^i that demand is distributed over and the weight α^{it} given to each itinerary are parameters of the method. The hypothesis is that demand is most likely to manifest in itineraries presenting good choices according to the cost function, and that customers will divert from their first choice if availability is lacking. Itineraries are defined as chosen in an iterative process, starting with the first choice. The formal expression for the choice of itinerary in this psychic forecast is given in Definitions (8.25).

$$\delta_{i,r}^{\text{i-choice}} := \begin{cases} 1 & i = i' | \max_{i' \in I_r} \left(\delta_{i',r}^{\text{i-choice}} - 1 \right)^2 \cdot \hat{C}\left(i', r\right) \\ 0 & \text{else.} \end{cases} \tag{8.25}$$

$$\forall\; r \in R;\; i \in I_r$$

Maximum Price (c-max): Psychic forecasts predicting maximum prices use knowledge of the customers' maximum willingness to pay given in the simulation. They expect demand to arrive in the *most expensive* acceptable booking class of a given itinerary. The hypothesis is that when the forecast predicts customers to arrive according to their highest willingness to pay, inventory controls will be restrictive enough to prevent much of buy-down. The formal expression for this psychic forecast is given in Definitions (8.26) and (8.27).

- Let $\delta_{r,c,i}^{\max}$ be a boolean matrix indicating the most expensive booking class available on all flights of the given itinerary i and acceptable for r .

$$\delta_{r,c,i}^{\max} := \begin{cases} 1 & c = c' | \max_{c' \in C} \left(\delta_{r,c'}^{\text{product}} \cdot \delta_{r,c',i}^{\text{price}} \cdot \sum_{f \in F_i} p\left(f, c\right) \right) \\ 0 & \text{else.} \end{cases} \tag{8.26}$$

$$\forall\; r \in R;\; I \in I_r$$

$$f^{\text{unc}}\left(f, c, t, s\right) := \sum_{r \in R_s} \delta_{r,t}^{\text{time}} \cdot \frac{\sum_{i \in I_r} I_{f,i}^{\text{it}} \cdot \delta_{i,r}^{\text{i-choice}} \cdot \delta_{r,c,i}^{\max}}{n^i} \tag{8.27}$$

$$\forall\; f \in F;\; c \in C;\; t = 1, ..., N^t;\; s = 1 ... N^s$$

Minimum Price (c-min): Psychic forecasts predicting minimum prices use knowledge of the customers' acceptance of booking classes given by the simulation. Demand is expected to arrive in the *cheapest* acceptable booking class of a given itinerary. The hypothesis is that when the forecast predicts customers to arrive in the cheapest class according to their acceptance of product characteristics given by booking classes, the worst case of buy-down is already included in the forecast. The formal expression for this psychic forecast is given in Definitions (8.28) and (8.29).

- Let $\delta^{\text{min}}_{r,c,i}$ be a boolean matrix indicating the cheapest booking class available on all flights of the given itinerary i and acceptable for r .

$$\delta^{\text{min}}_{r,c,i} := \begin{cases} 1 & c = c' | \min_{c' \in C} \left(\delta^{\text{product}}_{r,c'} \cdot \delta^{\text{price}}_{r,c',i} \cdot \sum_{f \in F_i} p\left(f,c\right) \right) > 0 \\ 0 & \text{else.} \end{cases} \tag{8.28}$$

$$\forall\ r \in R;\ I \in I_r$$

$$f^{\text{unc}}\left(f,c,t,s\right) := \sum_{r \in R_s} \delta^{\text{time}}_{r,t} \cdot \frac{\sum_{i \in I_r} I^{\text{it}}_{f,i} \cdot \delta^{\text{i-choice}}_{i,r} \cdot \delta^{\text{min}}_{r,c,i}}{n^i} \tag{8.29}$$

$$\forall\ f \in F;\ c \in C;\ t = 1, ..., N^t;\ s = 1...N^s$$

Acceptable Prices (c-frac): Psychic forecasts predicting all acceptable prices use the knowledge of customers' maximum willingness to pay and acceptance of booking classes given by the simulation. Fractional demand is expected to arrive in *all acceptable* booking classes of a given itinerary. The hypothesis is that demand will manifest in one of these classes respectively and the forecast strives to get a reasonable estimate including the possibility of buy-down. The formal expression for this psychic forecast is given in Definition (8.30).

$$f^{\text{unc}}\left(f,c,t,s\right) := \sum_{r \in R_s} \delta^{\text{time}}_{r,t} \cdot \sum_{i \in I_r} \left(I^{\text{it}}_{f,i} \cdot \delta^{\text{i-choice}}_{i,r} \cdot \frac{\delta^{\text{product}}_{r,c} \cdot \delta^{\text{price}}_{r,c,i}}{\sum_{c \in C} \delta^{\text{product}}_{r,c} \cdot \delta^{\text{price}}_{r,c,i}} \right) / n^i \tag{8.30}$$

$$\forall\ f \in F;\ c \in C;\ t = 1, ..., N^t;\ s = 1...N^s$$

Psychic Forecasts as Benchmarks: To present the consequences of the use of the different psychic forecasts, a normalization to the effects of first-come-first-serve inventory controls is performed. These effects can be calculated by applying the controls to the same demand that is later used to evaluate psychic forecasts. First-come-first-serve controls provide a benchmark as they are the simplest alternative to applying any forecast at all.

- Let $\sum_{f \in F} K\left(f,0,s\right)$ be the overall available capacity at the beginning of run s.

- Let $\sum_{f\in F}\sum_{c\in C}\sum_{t=1}^{N^t} b\left(f,c,t,s\right)$ be the overall bookings generated in simulation run s.
- Let $r\left(s\right)$ be the overall revenue generated in simulation run s.
- Let $r^{\text{fcfs}}\left(s\right)$ be the overall revenue generated in simulation run s given first-come-first-serve inventory controls.
- Let $r_{\text{2-i c-max}}^{\text{\% fcfs}}\left(s\right)$ be the overall revenue generated in simulation run s using (e.g.) psychic forecast method "2-i c-max", as a percentage of the revenue resulting from the application of first-come-first-serve controls.
- Let $y^{\text{fcfs}}\left(s\right)$ be the average yield generated in simulation run s given first-come-first-serve inventory controls.
- Let $y_{\text{2-i c-max}}^{\text{\% fcfs}}\left(s\right)$ be the average yield generated in simulation run s using (e.g.) psychic forecast method "2-i c-max", as a percentage of the revenue resulting from the application of first-come-first-serve controls.
- Let $l\left(s\right)$ be the average seat load factor generated in simulation run s.
- Let $l^{\text{fcfs}}\left(s\right)$ be the average seat load factor generated in simulation run s given first-come-first-serve inventory controls.
- Let $l_{\text{2-i c-max}}^{\text{\% fcfs}}\left(s\right)$ be the average seat load factor generated in simulation run s using (e.g.) psychic forecast method "2-i c-max", as a percentage of $l^{\text{fcfs}}\left(s\right)$.

Revenue is computed as described by Definition (8.12). The normalization to the benchmark generated by first-come-first-serve controls is shown in Definition (8.31). To present the result of a complete simulation experiment in a single indicator, the normalized revenue may be averaged over all runs as presented in Definition (8.32).

$$r^{\text{\% fcfs}}\left(s\right) := \frac{r\left(s\right)}{r^{\text{fcfs}}\left(s\right)}\cdot 100 \qquad \forall\ s=1,...,N^s \tag{8.31}$$

$$\hat{r}^{\text{\% fcfs}} := \frac{\sum_{s=1}^{N^s} r^{\text{\% fcfs}}\left(s\right)}{N^s} \tag{8.32}$$

Yield is computed as described by Definition (8.18). The normalization to the benchmark generated by first-come-first-serve controls is shown in Definition (8.33). To present the result of a complete simulation experiment in a single indicator, the normalized revenue may be averaged over all runs as presented in Definition (8.34).

$$y^{\%\ \text{fcfs}}(s) := \frac{y(s)}{y^{\text{fcfs}}(s)} \cdot 100 \qquad \forall\ s = 1, ..., N^s \tag{8.33}$$

$$\hat{y}^{\%\ \text{fcfs}} := \frac{\sum_{s=1}^{N^s} y^{\%\ \text{fcfs}}(s)}{N^s} \tag{8.34}$$

The seat load factor is computed as a function of capacity and bookings, as presented in Definition (8.35). It can be normalized, as Definition (8.36) shows. To present the result of a complete simulation experiment in a single indicator, the normalized seat load factor may be averaged over all runs as presented in Definition (8.37).

$$l(s) := \frac{\sum_{f \in F} \sum_{c \in C} \sum_{t=1}^{N^t} b(f, c, t, s)}{\sum_{f \in F} K(f, 0, s)} \cdot 100 \qquad \forall\ s = 1, ..., N^s \tag{8.35}$$

$$l^{\%\ \text{fcfs}}(s) := \frac{l(s)}{l^{\text{fcfs}}(s)} \cdot 100 \qquad \forall\ s = 1, ..., N^s \tag{8.36}$$

$$\hat{l}^{\%\ \text{fcfs}} := \frac{\sum_{s=1}^{N^s} l^{\%\ \text{fcfs}}(s)}{N^s} \tag{8.37}$$

The simulation experiments conducted to analyze the effect of psychic forecasts are based on the “hybrid” market scenario described in Section 7.4. As in this scenario both product- and price-oriented supply and demand are included, the broadest variety of effects is expected. Variations of demand volume and deviation are included in the data and will be pointed out whenever helpful.

Given the different approaches to translating willingness to pay and the acceptance of booking classes into psychic forecasts, the assumption is that consequences of the choice of class-forecast will become clear when considering revenue. The precise expectation is formalized in Hypothesis (8.38)

$$\hat{r}^{\%\ \text{fcfs}}_{\text{c-max}} \geq \hat{r}^{\%\ \text{fcfs}}_{\text{c-frac}} \geq \hat{r}^{\%\ \text{fcfs}}_{\text{c-min}} \geq 100 \tag{8.38}$$

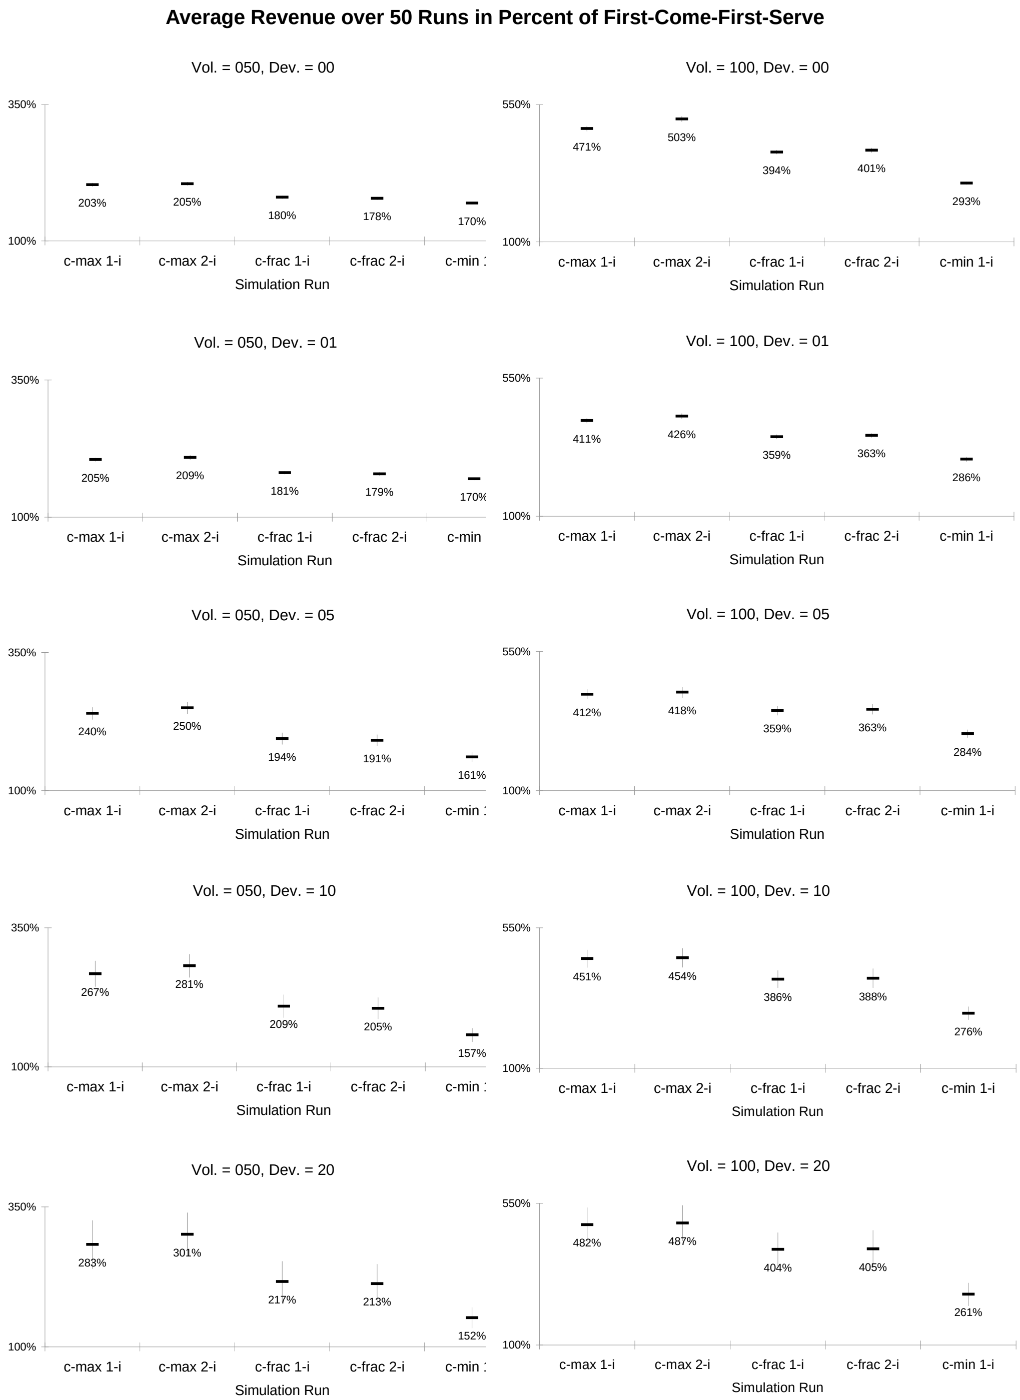

Figure 8.17.: Average Revenue over 50 Runs in Percent of First-Come-First-Serve

Figure 8.17 shows the average revenue over 50 runs resulting from the application of the listed variations of psychic forecasts over all runs. Revenue is expressed as $\hat{r}^{\%\ \mathrm{fcfs}}$ as formalized in Definition (8.32). In the diagram, the deviation of the average revenue in both directions is presented by a grey vertical line. As expected, the deviation of the average revenue grows as the deviation of the error term distribution grows and demand volume is more volatile between runs.

In all cases, as predicted by Hypothesis (8.38), the psychic forecasts based on "c-max" result in higher revenue than those based on "c-frac". At the same time, "c-frac" still leads to higher revenues than "c-min". As the simulation experiment was based on a hybrid scenario including some customers that do not accept the cheapest booking class, even "c-min" always resulted in revenues exceeding those observed when first-come-first-serve controls were applied.

In addition, it is remarkable that revenue was always highest when the forecast was not only based on "c-max" but also took into account two alternative preferred itineraries as in the variant "2-i". This effect is strongest in the case of "Vol. = 100, Dev. = 00", when demand volume is high and the deviation of the error term distribution is low, and in the case of "Vol. = 050, Dev. = 20", when demand volume is low and the deviation of the error term is high. In both cases, the probability for a few itineraries being requested significantly more often than others seems to be higher. In such cases, it is advantageous to include demand predicted also for the second-choice itineraries in the optimization.

The effect can even be observed in markets with a high demand volume when "c-frac" is applied. However, the revenue observed in markets with a high error term deviation is not much higher for "2-i" than it is for "1-i". The reason for this becomes clear when considering yield and seat load factor.

As demand is distributed more equally among itineraries with alternative itinerary forecasts, availability control may be expected to turn more restrictive. The consequence of this would be rising yields but also declining bookings. The precise expectation is formalized in Hypotheses (8.39) and (8.40).

$$\hat{l}_{\text{2-i}}^{\%\ \mathrm{fcfs}} \leq \hat{l}_{\text{1-i}}^{\%\ \mathrm{fcfs}} \leq 100 \tag{8.39}$$

$$\hat{y}_{\text{2-i}}^{\%\ \mathrm{fcfs}} \geq \hat{y}_{\text{1-i}}^{\%\ \mathrm{fcfs}} \geq 100 \tag{8.40}$$

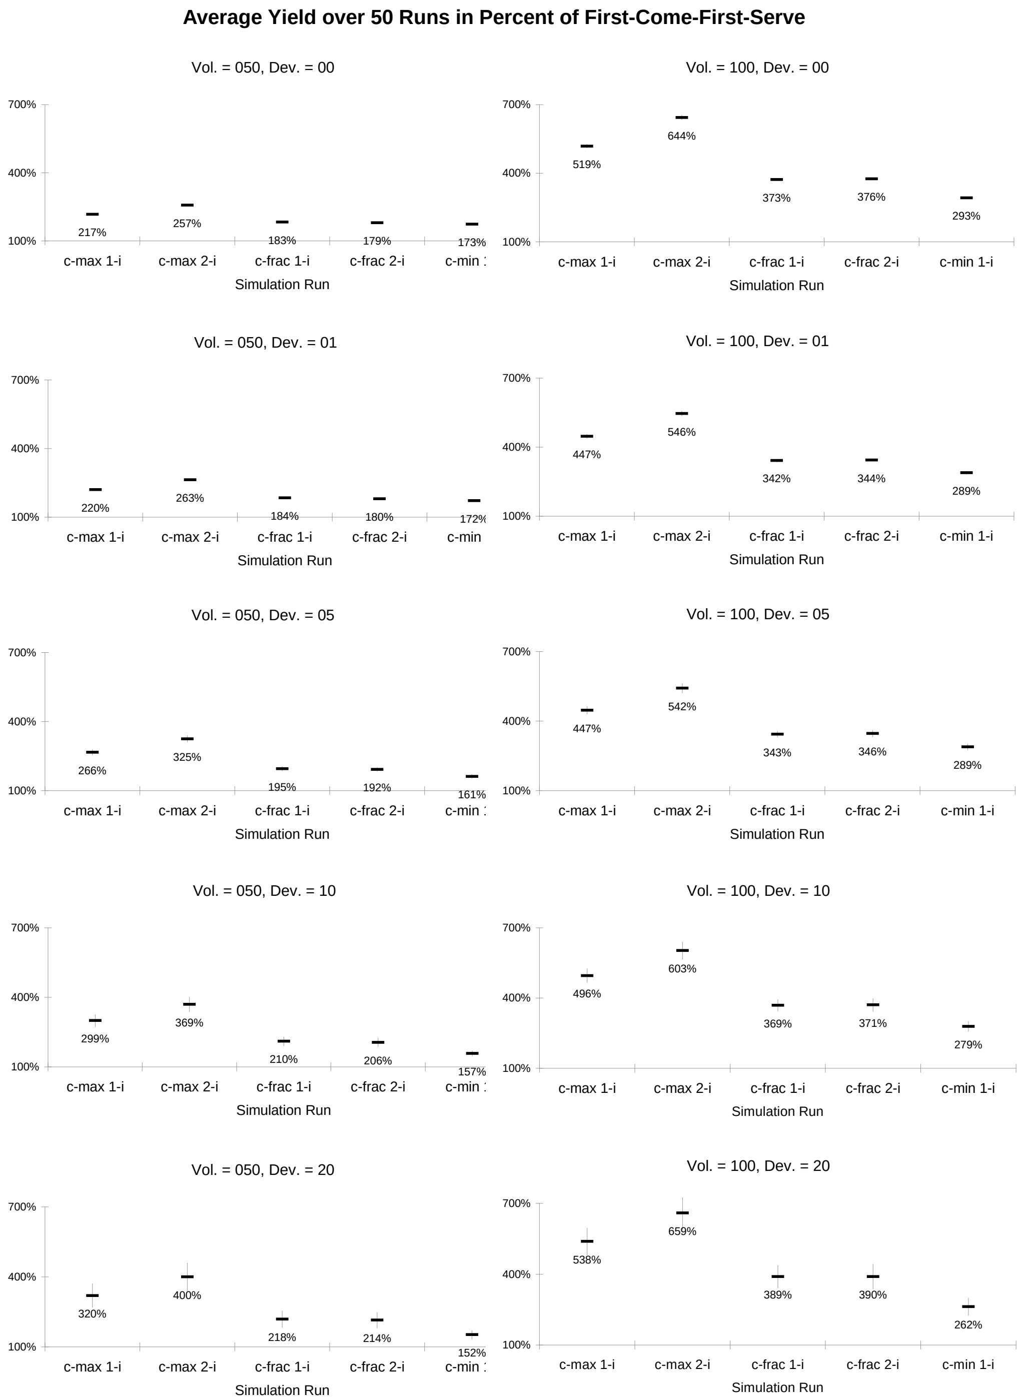

Figure 8.18.: Average Yield over 50 Runs in Percent of First-Come-First-Serve

Figure 8.18 shows the average yield over 50 runs resulting from the application of the listed variations of psychic forecasts over all runs. Yield is expressed as $\hat{y}^{\%\ \text{fcfs}}$ as shown in Definition (8.34). In the diagram, the deviation of the average yield in both directions is presented by a grey vertical line. As can be expected, the deviation of the average yield grows as the deviation of the error term distribution grows and demand volume is more volatile between runs.

In all cases, the psychic forecasts based on "c-max" result in higher Yield than those based on "c-frac". At the same time, "c-frac" still leads to higher yield than "c-min". As the simulation experiment was based on a hybrid scenario including some customers that do not accept the cheapest booking class, even "c-min" always resulted in yields exceeding those observed when first-come-first-serve controls were applied.

As already remarked with regard to revenue and predicted by Hypothesis (8.40), yield is always highest when the forecast is not only based on "c-max" but also takes into account two alternative preferred itineraries as in the variant "2-i". This effect is strongest in the case of "Vol. = 100, Dev. = 20", when demand volume is high and the deviation of the error term distribution is also high. The consequence of the high deviation of the error distribution is a less homogeneous distribution of demand volume over the existing itineraries. This leads to more itineraries being especially desirable and more itineraries being "second-choice". Without a "2-i" forecast, the second-best alternatives are assigned less restrictive inventory controls. However, in the case of "Vol. = 100, Dev. = 20", the consideration of second-best alternatives seems to have lead to losses in bookings: While yield strongly exceeds the yield gained by "1-i", the revenue as shown in Figure 8.17 was not much higher.

Figure 8.19 shows the average seat load factor over 50 runs resulting from the application of the listed variations of psychic forecasts over all runs. Seat load factor (SLF) is expressed as $\hat{l}^{\%\ \text{fcfs}}$ as shown in Definition (8.37). In the diagram, the deviation of the average SLF in both directions is presented by a grey vertical line. However, in contrast to what was observed for yield and revenue previously, the deviation of this indicator is small across all the markets.

In all cases, the psychic forecasts based on "c-max" result in lower SLF than those based on "c-frac". At the same time, "c-frac" still leads to lower SLF than "c-min". The reason

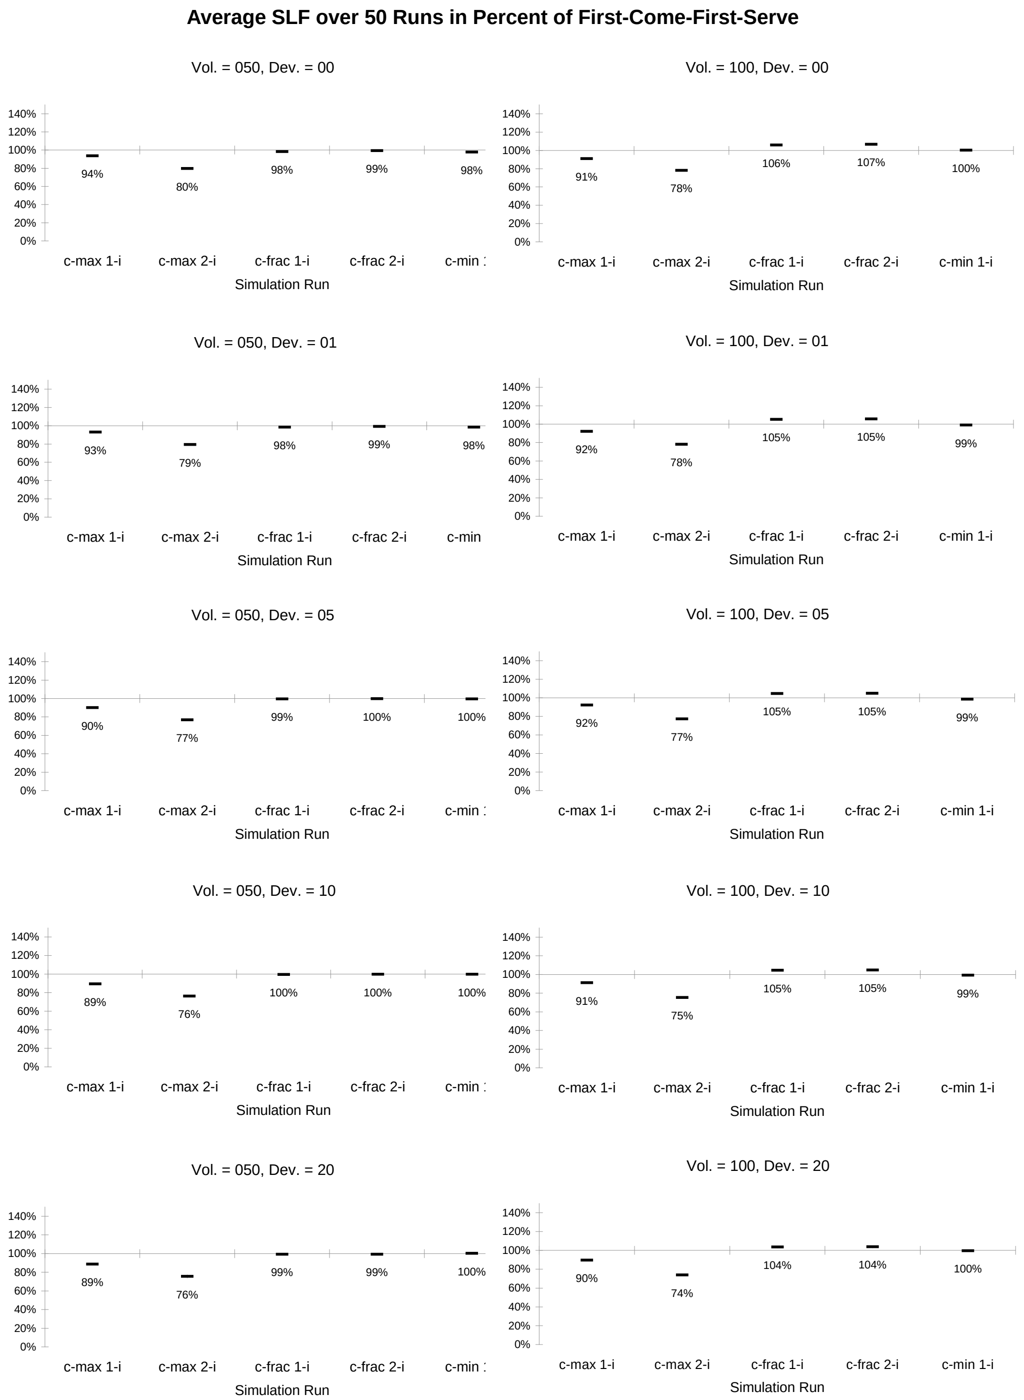

Figure 8.19.: Average SLF over 50 Runs in Percent of First-Come-First-Serve

for this are more restrictive inventory controls resulting from demand being predicted to arrive in classes that are not the cheapest. In the case of high demand volume, *"Vol. = 100"*, the psychic forecasts of the variant "c-min" lead to SLF that are even higher than those achieved with first-come-first-serve controls. This is due to more bookings being accepted as availabilities are optimized.

As already remarked with regard to yield, the itinerary choice has a direct effect on SLF. They are even lower when the forecast is not only based on "c-max" but also takes into account two alternative preferred itineraries as in the variant "2-i". This effect can be observed regardless of demand volume. It is stronger when the deviation of error term distribution is high. The consequence of the high deviation of the error distribution is a less homogeneous distribution of demand volume over the existing itineraries. This leads to more itineraries being especially desirable and more itineraries being "second-choice". With a "2-i" forecast, the second-best alternatives are assigned more restrictive inventory controls.

The consideration of second-choice alternatives leads to losses in bookings as more seats are reserved for valuable customers that book on their preferred itinerary. This is the reason for the effect observed with regard to revenue: While yield strongly exceeds the yield gained by "1-i", the revenue as shown in Figure 8.17 was not much higher.

Psychic Forecasts as Initialization: Used as an initialization method, psychic forecasts influence the first set of historical bookings that adaptive methods such as exponential smoothing can be based on. The degree of their influence may be derived from a measurement that describes in how far the results of the simulation experiments based on different initial forecasts but adapted according to the same method diverge. Therefore, new indicators need to be introduced.

- Let $\{\text{c-max 1-i}, \text{c-frac 1-i}, \text{c-min 1-i}, \text{c-max 2-i}, \text{c-frac 2-i}\}$ be the set of available psychic forecasts.
- Let $r^{\%fcfs}(s)$ be the revenue in run s as a percentage of the revenue gained with first-come-first-serve controls.
- Let $\sigma(r(s))$ be the deviation of revenues at run s over all simulation experiments considered.

As the same adaptive method is applied in a range of simulation experiments starting out with different initial forecasts, a development in the deviation of results can be expected. The assumed relationship is formalized in Hypothesis (8.41).

$$\lim_{s \to \infty} \sigma\left(r\left(s\right)\right) = 0 \tag{8.41}$$

Figure 8.20 shows $r^{\%fcfs}\left(s\right)$ over the course of 50 runs for different initializations of Exp050. When "zero FC" is used as initialization method, the forecast is set to be zero for all flights and classes in the first run. It is then updated based on observed unconstrained bookings using the exponential smoothing method "Exp050". The other options shown correspond to the psychic forecast variants presented previously.

The "zero FC" initialization leads to first-come-first-serve controls being applied in the first run. As no demand is predicted to arrive in any class, no protected seats are computed by the EMSR-b optimization. A slight spiral-up effect can be observed in later runs: As the exponential smoothing method picks up on product-based demand, protected seats are introduced and the inventory controls are no longer merely based on first-come-first-serve.

When used as an initialization method, "c-max 2-i" seems superior to "c-max 1-i". In contrast to the constant use of the psychic forecast that predicts demand to arrive in part for second-choice itineraries, the exponential smoothing avoids the effect of overly restrictive inventory controls. Both methods that are based on predicting demand to arrive according to customers' highest willingness to pay result in constantly higher revenue than the alternatives.

While leading to lower revenue than the "c-max" and even "c-frac" alternatives, "c-min" is still more successful than the "zero FC" initialization. As the market that is being analyzed in these simulation experiments includes both a class structure and demand that is price- as well as product-oriented, a share of customers accord to the static demand assumption. An accurate prediction of this demand even in the cheapest class customers are willing to buy therefore still results in protected seats for more valuable classes. The initialization method "c-min" may also be regarded to present the state a complete spiral-down effect would finally lead to in a hybrid market: All customers are predicted to arrive in the cheapest class they are willing to buy, which is not necessarily the cheapest class available.

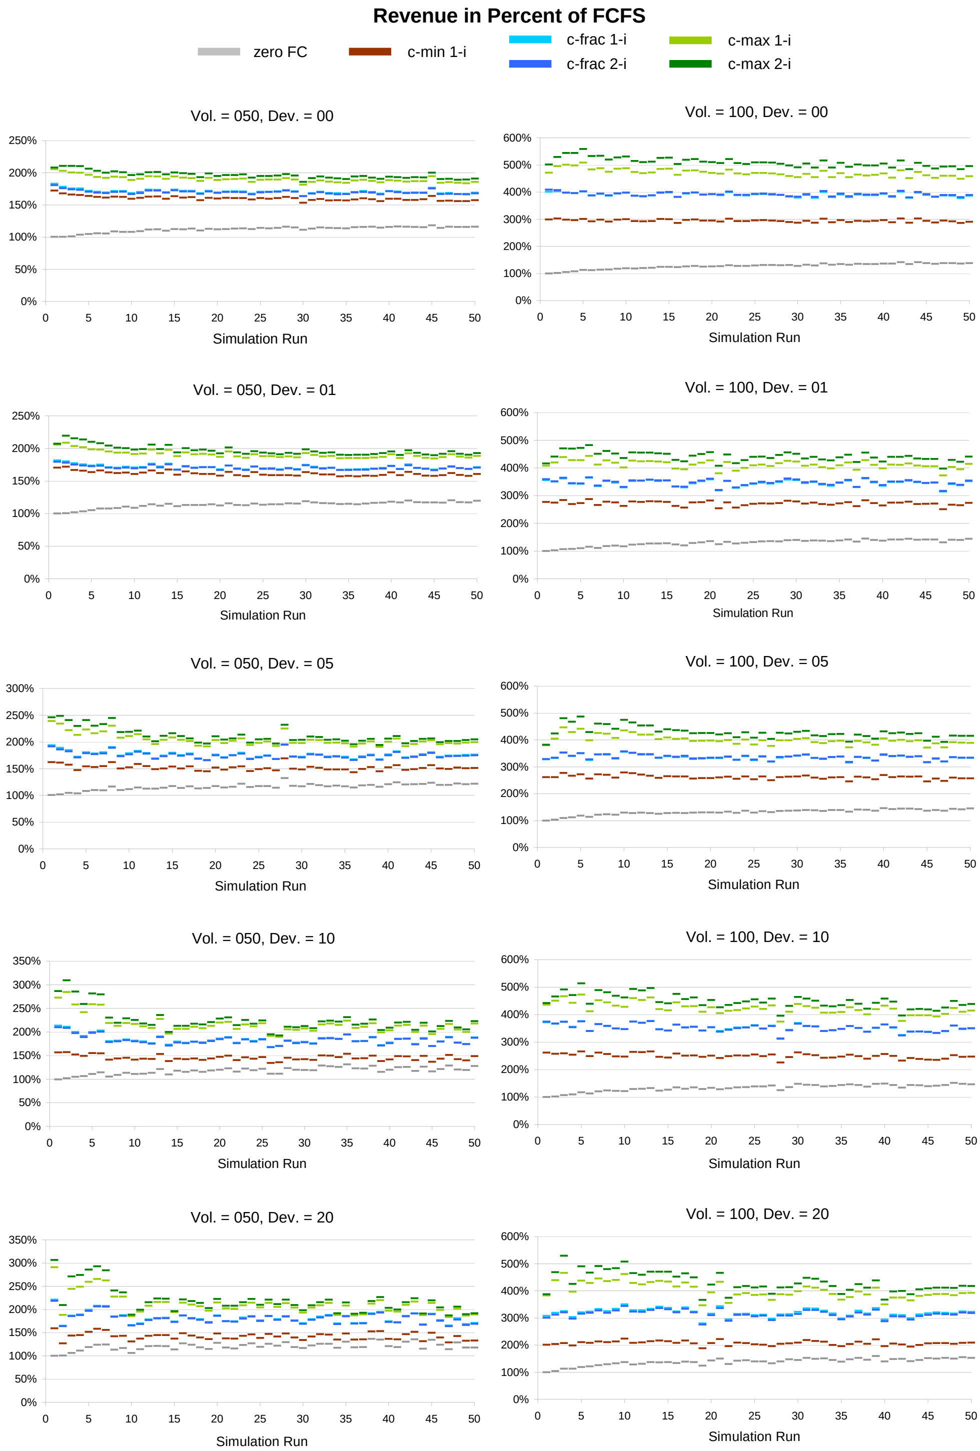

Figure 8.20.: Revenue in Percent of First-Come-First-Serve

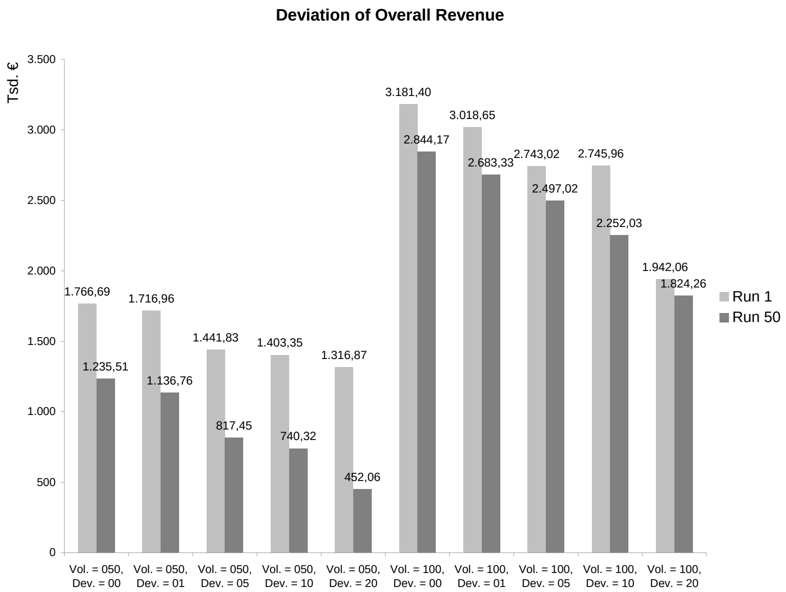

Figure 8.21.: Deviation of Revenue between Simulation Experimenets

Figure 8.21 shows the development of $\sigma\left(r\left(s\right)\right)$ from run 1 to run 50, given different forecast initializations and the adaptive method Exp050 as already discussed with regard to Figure 8.20. As can be seen, the prediction formalized in Hypothesis (8.41) holds true: The deviation of revenue between simulation experiments decreases as the same exponential smoothing method is applied to different initial forecasts.

The difference in deviation between run 1 and run 50 is greater for markets including low demand volume ("Vol. = 050") and high deviation ("Dev. = 20"). This is due to the fact that the spiral-down effect that leads to similar (low) revenue results takes place faster under these conditions. This has been demonstrated in Section 8.1.

In all market variations including high demand volume ("Vol. = 100"), the initial deviation is higher than when demand volume is low. Due to the high number of requests, inventory control can have a wider range of success as the maximum of revenue that may be gained is greater. As the spiral-down effect does not immediately lead to lower revenues in such cases (this also been explained in Section 8.1), revenues do not decrease from their initial state as much as they do in markets with low demand volume. For this reason, the deviation does not decrease in a similar way, either.

Conclusion: The concept of psychic forecasts can be used for two purposes in a simulation environment: To initialize adaptive forecast methods lacking historical bookings and to serve as a benchmark for the evaluation of other forecast methods. However, the choice of psychic forecast has a great influence on its success. The way in which the psychic forecast is defined and computed from knowledge of the demand model has direct consequences in the form of the resulting revenue. It influences both the success of adaptive methods that use it as an initialization method and the outcome of evaluations based on comparisons with it.

8.3. Evaluation of Standard Accuracy Indicators

The concept of simulation offers the opportunity of considering long-term developments and comparisons with exclusive knowledge of the demand model. In the previous two sections, these views have been applied to forecast methods. This section focuses on the ways of evaluating forecasts as presented in Section 4.1 with the help of a simulation system. It strives to analyze the advantages and fallacies of common approaches to forecast evaluation.

In order to compare the effects of different methods, forecast evaluation needs to be defined more clearly. Three dimensions can be identified that describe traditional forecast error measurements: the *objects* of comparison, the *level* of comparison, and the *method* of comparison.

Objects of Comparison: In general, the objects of comparison are always bookings and forecasts. However, in the case of demand forecasting for revenue management, a differentiation between the states "constrained" and "unconstrained" applies to both of these indicators as already explained in Section 6.2. This view results in four indicators that may be compared, two considering bookings and two considering forecasts. In addition, the unconstrained or constrained psychic forecast available in the simulation may also be used as a benchmark to compare the results of evaluation.

- Let $b(c, f, t, s)$ be the bookings observed in class c of flight f between points of time $t-1$ and t in the booking horizon of simulation run s.

- Let $b^{\text{unc}}(c, f, t, s)$ be the *unconstrained* bookings. The transformation from actual to unconstrained bookings is implicit in the forecast method applied.
- Let $f^{\text{unc}}(c, f, t, s)$ be the demand predicted to arrive for class c of itinerary f between points of time $t-1$ and t in the booking horizon of simulation run s. This is the output of the forecast method and the input for the optimization method.
- Let $f^{\text{const}}(c, f, t, s)$ be the *constrained* predicted demand. The transformation from unconstrained to constrained is handled over the inventory controls that were applied during the time period considered.
- Let $f^{\text{unc}}_{\text{1-i c-max}}(c, f, t, s) = f^{\text{unc}}_{\text{psy}}(c, f, t, s)$ be the psychic forecast according to the method using the preferred itinerary and maximum willingness to pay described in Section 8.2. As it was shown to be a good upper benchmark in most scenarios, this psychic forecast will be used representatively for all psychic forecasts in this section.
- Let $f^{\text{const}}_{\text{1-i c-max}}(c, f, t, s) = f^{\text{const}}_{\text{psy}}(c, f, t, s)$ be the constrained version of the psychic forecast. The transformation from unconstrained to constrained is handled over the inventory controls calculated on the basis of the psychic forecast.

Intuitively, it appears sensible to compare only indicators that are of the same transformation. This would mean comparing only unconstrained bookings with the original unconstrained forecast and only comparing actual bookings with a constrained forecast. However, as bookings and forecasts appear originally in different states, an evaluation that considers the difference between actual bookings and the original, unconstrained forecast may also be of interest. In contrast to this, there seems to be no point in transforming both indicators from their original state and comparing unconstrained bookings with constrained forecasts.

The resulting options for comparison are:

- actual bookings vs. constrained forecast, e^{c-c};
- unconstrained bookings vs. unconstrained forecast, e^{u-u};
- actual bookings vs. unconstrained forecast, e^{c-u};
- unconstrained forecast (e.g. a psychic variant) vs. unconstrained forecast, $e^{\widetilde{u}-u}$.

Levels of Comparison: Different aggregation levels may be considered as well. Before comparing them, indicators may be averaged over classes, itineraries or points of time before departure. Summing up over classes or itineraries does not seem useful with regard to forecast evaluation: If this was done, errors that occur with regard to the distribution of predicted demand over classes or itineraries may be compensated. However, the distribution of demand over classes and itineraries is a vital information for revenue management.

The question of whether or not to sum up forecasts over points of time before departure to evaluate them is not that trivial. The EMSR-b heuristic implemented in the revenue management system modeled does not consider the order of arrival of demand. Does this mean the information is irrelevant? In the course of this section, two aggregation levels will therefore be analyzed:

- Let $\hat{e}(s)$ be the series error calculated for run s comparing bookings and forecast per flight, class and point of time.
- Let $\bar{\hat{e}}(s)$ be the series error calculated for run s comparing bookings and forecast summed up over points of time per flight and class.

Methods of Comparison: Finally, four ways of calculating error measurements have been described in Section 4.1: The absolute measurements MAD (mean absolute deviation) and RMSE (root mean squared error) as well as the percentage error measurements MAPE (mean average percentage error) and U2 (Theil's U2). To specify the error calculated, further notation is needed.

- Let the usage of MAD be indicated by e_{MAD}.
- Let the usage of RMSE be indicated by e_{RMSE}.
- Let the usage of MAPE be indicated by e_{MAPE}.
- Let the usage of U2 be indicated by e_{U2}.

From these options, a tree of possible forecast error measurements emerges. It is outlined in Figure 8.22. According to it, using the range presented so far, 32 different error measurements may be calculated from the combination of different forecasts or bookings.

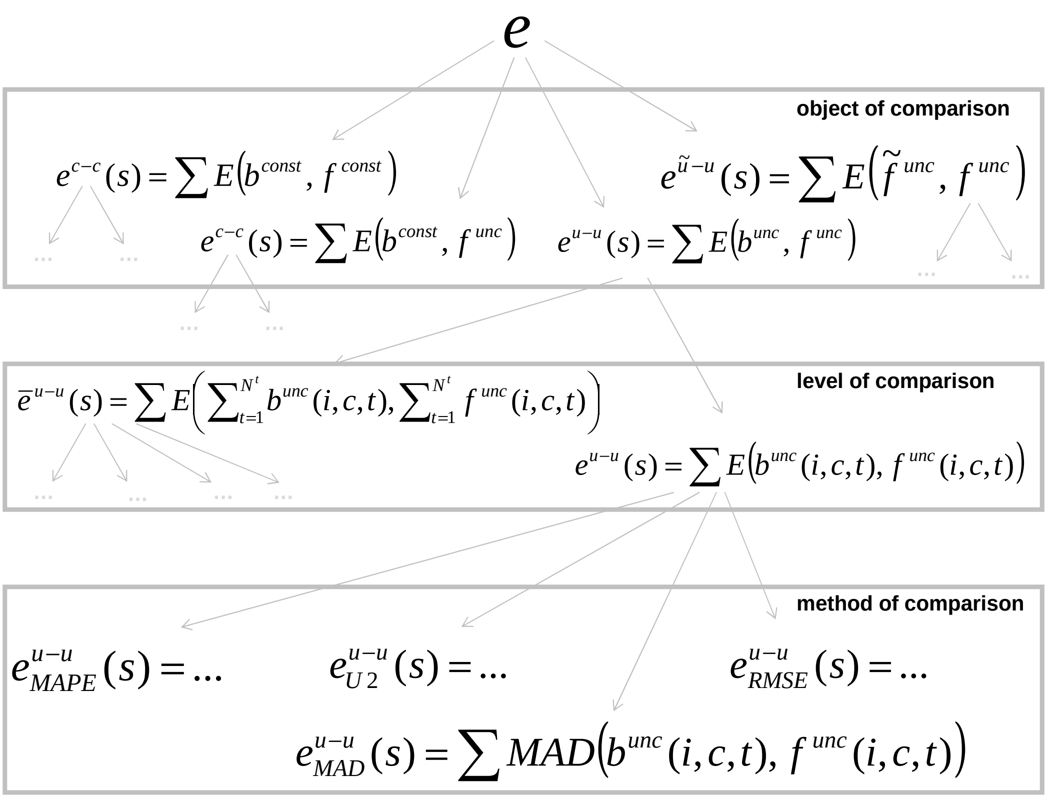

Figure 8.22.: Possible Error Measurements

By applying the alternative error measurements described, the consequences of choosing one over the other can be tested using the simulation system. Applying different measurements to the same scenario and combining it with knowledge about the actual demand and quality of forecast methods applied allows for a systematic evaluation of three fields: the choice of indicators to compare, the choice of aggregation level, and the choice of error to compute.

Consequences of the Choice of Object: As documented in Section 4.2, it is common to use the comparisons e^{u-u} and e^{c-c} for forecast evaluation. The condition for this is the equivalence of both indicators. While the absolute numbers depend on the type of comparison, their conclusion should be the same. This may be tested by comparing the ranking of methods based on each error measurement as shown in Definition (8.42).

- Let $R(\circ, e, s)$ be the rank of forecasting method $\circ$ resulting from its evaluation via the error measurement e in simulation run s.

$$e^{u-u} \equiv e^{c-c} := R\left(\circ, e^{u-u}, s\right) = R\left(\circ, e^{c-c}, s\right) \tag{8.42}$$

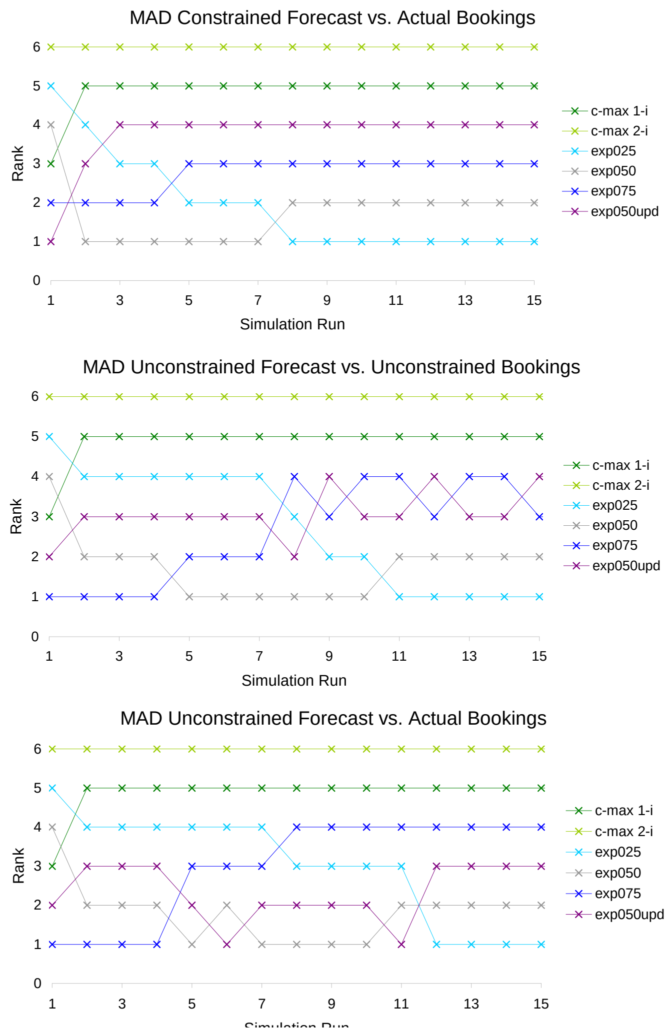

Figure 8.23.: Rank of Methods according to MAD in Product-Based Scenario with "Vol. = 050, Dev. = 00"

Figure 8.23 shows the resulting ranks of the computation of e^{u-u}_{MAD}, e^{c-c}_{MAD}, and e^{c-u}_{MAD} for the forecast methods "1-i c-max", "2-i c-max", "exp025", "exp050", "exp075" and the version of "exp050" that is updated throughout the booking horizon, "exp050-upd". The methods are ranked based on observations made when applying them to a product-sensitive market with low demand volume, "Vol. = 050", and low deviation of the error distribution "Dev. = 00". This market shows the clearest results, however, similar trends emerge when analyzing different error term deviations. As the ranks are constant after the tenth run, only the first ten runs of the simulation experiment are shown in the graph.

As can be seen, after a number of runs, all error measurements included in the analysis lead to the same result in terms of rank: While the Exp025 forecast is judged to be the most accurate, the two psychic forecast options occupy the last places. The relationship stated by Hypothesis (8.42) is confirmed.

Given a view of ranks rather than absolute errors, a comparison of actual bookings and unconstrained forecasts becomes possible. While higher absolute quantities can be expected due to the inherent quantitative difference of bookings and demand, intuitively rankings may still be equivalent. A successful optimization should always exclude bookings in cheap classes if enough demand for more expensive classes is available, this may not be true for scenarios including a high volume of demand. Here, a systematic (and desirable) difference between demand and forecast may lead to a disadvantage for methods that correctly predict the demand that never gets to manifest as bookings. The expectation according to this logic is that possible equivalences disappear as absolute demand increases. This is formally expressed in Hypothesis (8.43).

$$\begin{aligned}
\left|\lim_{|R|\to\infty} R\left(\circ, e^{u-u}, s\right) - R\left(\circ, e^{c-u}, s\right)\right| &\geq 1 \\
\left|\lim_{|R|\to\infty} R\left(\circ, e^{c-c}, s\right) - R\left(\circ, e^{c-u}, s\right)\right| &\geq 1 \\
\forall\ s = 1, ..., N^s &
\end{aligned} \tag{8.43}$$

Figure 8.24 shows the resulting ranks of the computation of e^{u-u}_{MAD}, e^{c-c}_{MAD}, and e^{c-u}_{MAD} for the forecast methods "1-i c-max", "2-i c-max", "exp025", "exp050", "exp075" and the version of "exp050" that is updated throughout the booking horizon, "exp050-upd". The methods are ranked based on observations made when applying them to a product-sensitive market with high demand volume, "Vol. = 050", and low deviation of the error

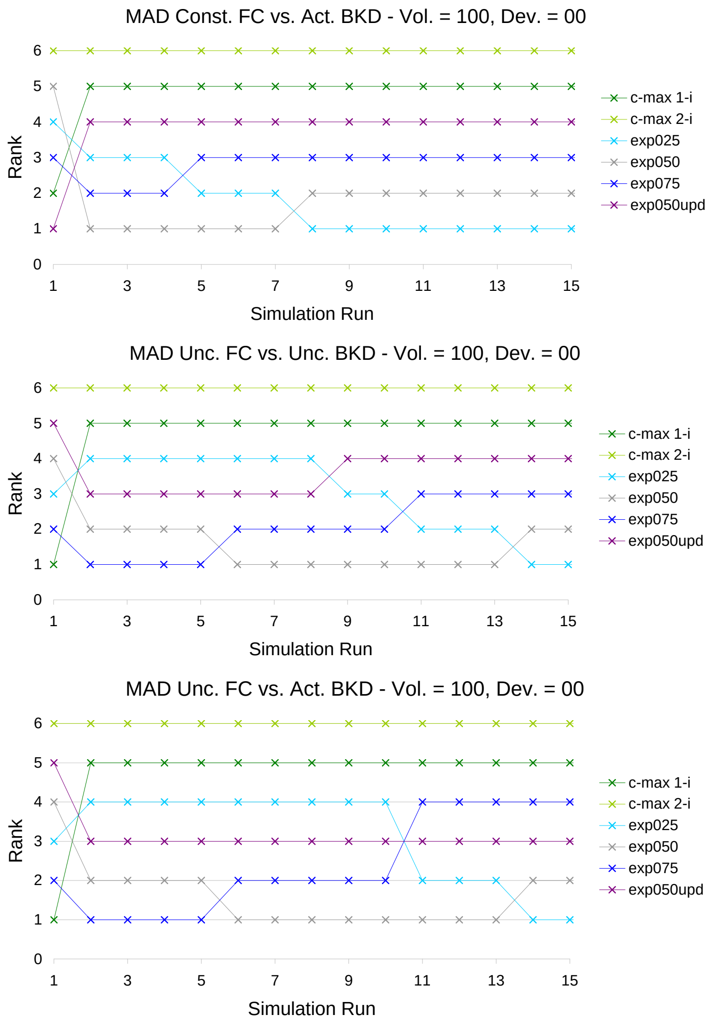

Figure 8.24.: Rank of Methods according to MAD in Product-Based Scenario with "Vol. = 100, Dev. = 00"

distribution "Dev. = 00". This market shows the clearest results, however, similar trends emerge when analyzing different error term deviations. As the ranks are constant after the tenth run, only the first ten runs of the simulation experiment are shown in the graph.

Contrary to what could be expected based on Hypothesis (8.43), there is no difference to be found between the ranking of forecasts after 15 runs – and from then on, it stays constant for "Dev. = 00". However, in getting to this state, especially during the first two or three runs, the ranks seem to change quicker when demand volume is high. This may be due to the higher pressure caused by more valuable demand being in the market.

In the simulation system, a psychic forecast includes the most accurate information. Accordingly, correct error measurements should rate it as the best of all options. As can be seen in Figure 8.23 and Figure 8.24, this is not the case. With regard to e^{u-u}, the reason for this lies in deficiencies in the unconstraining methods applied by the adaptive forecasts and used to transform bookings for the comparison. These do not only lead to a spiral-down-effect but also make forecasts that predict a smaller overall demand volume seemingly more attractive. With regard to e^{c-c}, the buy-down that is still possible in spite of restrictive inventory controls leads to a deviation from the predicted maximum willingness to pay.

As it incorporates buy-down and does not leave much space for spiral-down, $f^{\text{unc}}_{1-ic-min}$ can be expected to perform better according to e^{c-c} than $f^{\text{unc}}_{1-ic-max}$ does. This logic would predict $f^{\text{unc}}_{1-ic-frac}$ to end up ranked between the two other alternatives. Hypothesis (8.44) formalizes this expectation.

$$R\left(\text{1-i c-min}, e^{c-c}, s\right) \leq R\left(\text{1-i c-frac}, e^{c-c}, s\right) \leq R\left(\text{1-i c-max}, e^{c-c}, s\right) \quad \forall\ s = 1, ..., N^s \tag{8.44}$$

Figure 8.25 shows that the prediction formalized in Hypothesis (8.44) does come true when MAD is applied to variants of psychic forecasts in the product-sensitive scenario. The graph illustrates the value of MAD averaged over 50 runs. The variant "1-i c-min", as predicted, always achieves the lowest deviation from actual bookings while "1-i c-max" leads to the highest deviation. One exception to the rule is the case of high demand volume "Vol. = 100" and low deviation "Dev. = 00" – while the order of "1-i c-max" and "1-i c-min" stays the same, "1-i c-frac" is rated worst in this case.

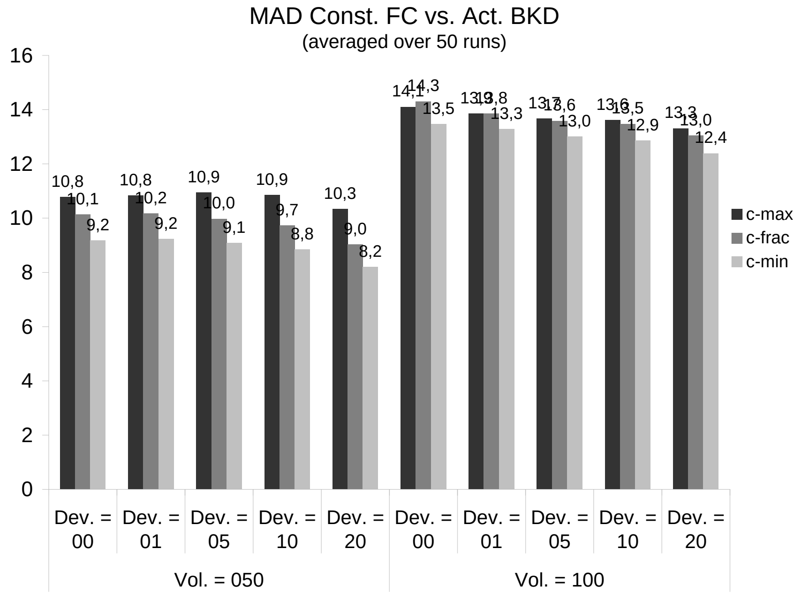

Figure 8.25.: MAD: Constrained Psychic Forecasts vs. Actual Bookings in the Product-Based Scenarios

Finally, if nothing changes in customer behavior over time, the repeated application of adaptive forecasts should lead to an improvement of forecast quality. In addition to e^{c-c} and e^{u-u}, the error measurement alternative including psychic forecasts, $e^{\widetilde{u}-u}$, can be used to test for this. The expected forecast behavior can be found Hypotheses (8.45).

$$\begin{aligned}
&\lim_{s\to\infty} e^{\widetilde{u}-u}(s) = 0 \\
&\lim_{s\to\infty} e^{u-u}(s) = 0 \\
&\lim_{s\to\infty} e^{c-c}(s) = 0
\end{aligned} \tag{8.45}$$

However, as shown in Figure 8.26, Hypothesis (8.45) does not hold for the adaptive forecast methods based on exponential smoothing implemented here. In contrast to that, the prediction of Hypothesis (8.45) can be observed to come true over the course of the fifty simulation runs included in the experiment. The conclusion drawn from this observation is that the seemingly adaptive effect that leads to an improvement of e^{u-u} and e^{c-c} is not in fact due to an adaption to the real demand. Instead, it is the result of a spiral-down effect as described in Section 8.1.

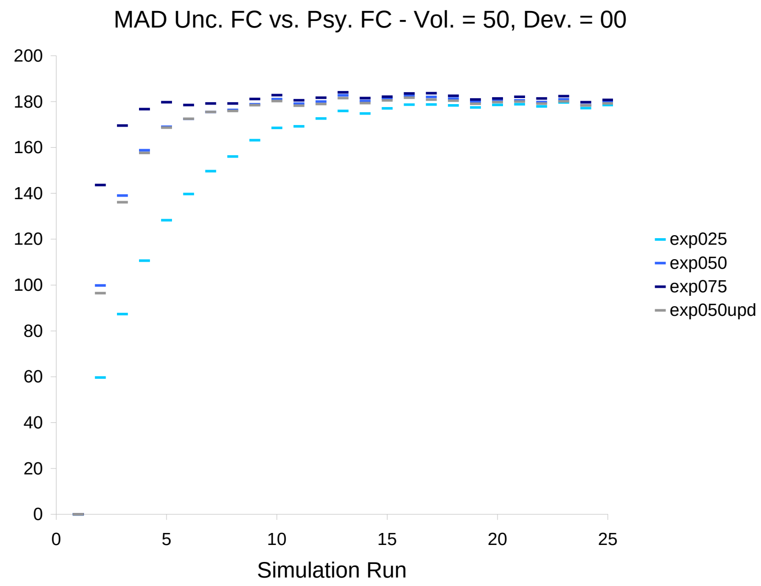

Figure 8.26.: MAD of Unconstrained Forecasts from Psychic Forecast in Product-Based Scenario with "Vol. = 050, Dev. = 00"

Consequences of the Level of Comparison: Whether or not to aggregate bookings and forecasts over time before departure before calculating error measurements can also be considered on the basis of a simulation experiment. For this, the computation of a reverse psychic forecast is introduced.

- Let $f^{\text{unc}}_{\text{1-i c-max}}(f, c, t, s)$ be the psychic forecast predicting demand to arrive for class c and flight f between points of time $t-1$ and t in the booking horizon of simulation run s.
- Let $\overline{f}^{\text{unc}}_{\text{1-i c-max}}(f, c, t, s)$ be the mirror psychic forecast as presented in Definition (8.46).

$$\overline{f}^{\text{unc}}_{\text{1-i c-max}}(f, c, t, s) := f^{\text{unc}}_{\text{1-i c-max}}\left(f, c, N^t - t + 1, s\right) \quad \forall\ f \in F;\ \ c \in C;\ \ t = 1, ..., N^t;\ \ s = 1, ..., N^s \tag{8.46}$$

Considering that EMSR-b does not take into account the order of demand arrival, $\overline{f}^{\text{unc}}$ and f^{unc} may be expected to be equivalent with regard to their consequences. Using revenue $r(s)$ as an indicator of overall system performance, this expectation has been formalized in Hypothesis (8.47).

- Let r be the overall revenue resulting from the application of a forecast f^{unc} to one scenario in a simulation experiment.
- Let $\overline{r}$ be the overall revenue resulting from the application of a mirrored forecast $\overline{f}^{\text{unc}}$ to the same scenario.

$$f^{\text{unc}} \equiv \overline{f}^{\text{unc}} \rightarrow r = \overline{r} \tag{8.47}$$

However, many revenue management systems include re-optimization routines within the booking horizon as described in Section 7.3. These are modifications of forecasts and inventory controls based on the deviation of already observed bookings from the predicted demand.

- Let f'^{unc} be a forecast that is updated throughout the booking horizon.
- Let $\overline{f'}^{\text{unc}}$ be a mirror forecast that is updated throughout the booking horizon.
- Let $\overline{r'}$ be the overall revenue resulting from the application of a mirror forecast that is updated.

As updating is applied, a difference in revenue can be expected. As seen before, adaptation to observed values does not always have a positive effect. For this reason, the expectation phrased by Hypothesis (8.48) is neutral.

$$f'^{\text{unc}} \neq \overline{f'}^{\text{unc}} \rightarrow r' \neq \overline{r'} \tag{8.48}$$

Figure 8.27 shows the revenue earned when inventory controls based on the psychic forecast are re-optimized, a mirrored psychic forecast is applied, or inventory controls based on the mirrored psychic forecast are re-optimized. In order to render the difference clearly visible without regard to variations of overall demand volume and error term deviation, revenue is expressed as a percentage of the revenue earned under the same conditions with the psychic forecast.

As can be seen, Hypothesis (8.47) holds true: The mirrored psychic forecast without re-optimization leads to a result that is equivalent to that of the psychic forecast that was not mirrored. The revenue percentage displayed is 100 %.

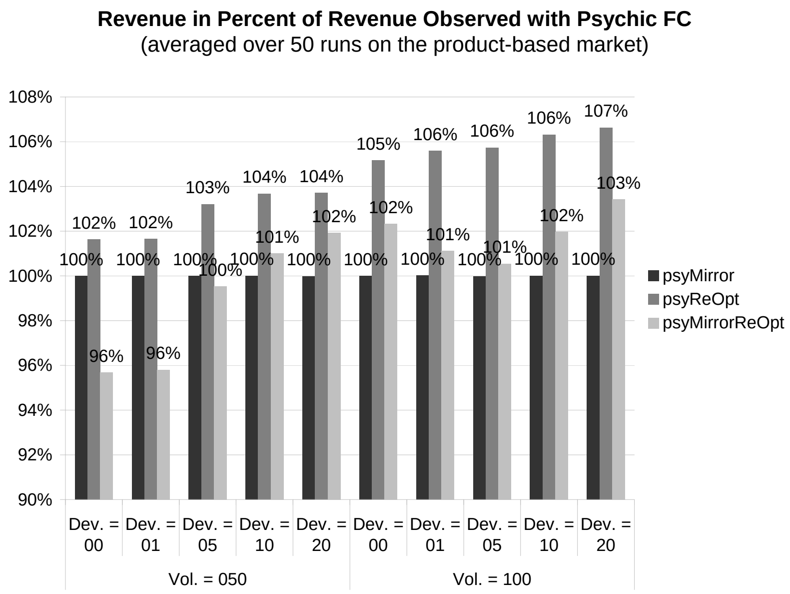

Figure 8.27.: Revenue in Percent of Revenue Earned by Psychic Forecast – Product-Based Scenario

The consequences of updating as predicted by Hypothesis (8.48) can also be observed in Figure 8.27. As shown, re-optimization combined with the psychic forecast leads to higher revenue in all market variations included. This is due to the overly restrictive inventory controls of the psychic forecast being corrected by re-optimization: While yield decreases, the increase in the number of bookings is strong enough to compensate.

In situations with low demand volume and low error term deviation, combining the mirrored psychic forecast with a re-optimization leads to lower revenue. As valuable demand is expected to arrive at incorrect times, the re-optimization does not compute inventory controls that succeed at reserving seats. Instead, the number of additional bookings gained when cheap tickets are available does not compensate for the loss in yield. However, in situations with high demand volume or a high error term deviation leading to a non-homogeneous distribution of demand over itineraries, even the re-optimized mirrored psychic forecast leads to higher revenue than the original psychic forecast. Again, this is due to an increase in bookings that compensates for losses in yield.

Consequences of Choice of Method: So far, all simulation experiments have focused on the use of mean absolute deviation, MAD, to compute errors. However, three other methods of computing error have been mentioned earlier. Of these, two (MAPE and U2) are based on percentage rather than absolute errors. RMSE is another absolute error measurement.

Two assumptions may be held about these indicators. One is the equivalence of absolute error measurements as well as that of percentage error measurements. The assumption has been taken from the literature presented in Section 4.1. It states that errors based on a percentage calculation are preferable as they do take into account the overall amount of demand.

Equivalence of MAD to RMSE and MAPE to U2 can be tested for based on a set-up similar to that applied when considering the equivalence of object choices. If MAD and RMSE are equivalent, this is not necessarily a question of their absolute value but one of the resulting rankings for different forecast methods. The same is true for MAPE and U2. This statement is formalized in Hypothesis (8.49).

$$\begin{aligned} e^{c-c}_{\text{MAD}} \equiv e^{c-c}_{\text{RMSE}} &\rightarrow R\left(\text{MAD}, e^{c-c}, s\right) = R\left(\text{RMSE}, e^{c-c}, s\right) \\ e^{c-c}_{\text{MAPE}} \equiv e^{c-c}_{\text{U2}} &\rightarrow R\left(\text{MAPE}, e^{c-c}, s\right) = R\left(\text{U2}, e^{c-c}, s\right) \end{aligned} \tag{8.49}$$

If error measurements based on percentage values were more accurate than those based on absolute values, one consequence would be a general difference in ratings. This expectation can be expressed in similar terms as the expected equivalence of methods described by Hypothesis (8.49). It is presented in Hypothesis (8.50).

$$\begin{aligned} e^{c-c}_{\text{MAD}} \neq e^{c-c}_{\text{MAPE}} &\rightarrow R\left(\text{MAD}, e^{c-c}, s\right) \neq R\left(\text{MAPE}, e^{c-c}, s\right) \\ e^{c-c}_{\text{RMSE}} \neq e^{c-c}_{\text{U2}} &\rightarrow R\left(\text{RMSE}, e^{c-c}, s\right) \neq R\left(\text{U2}, e^{c-c}, s\right) \end{aligned} \tag{8.50}$$

Figure 8.28 shows that the prediction does hold true: Different rankings result from the application of absolute or percentage values. No statement about the quality of these measurements can be made yet. One way of evaluating the quality of the error measurements analyzed here has already been introduced – the comparison to a psychic forecast. It may offer insights as the quality of the psychic forecast is known to be superior to that of the adaptive alternatives.

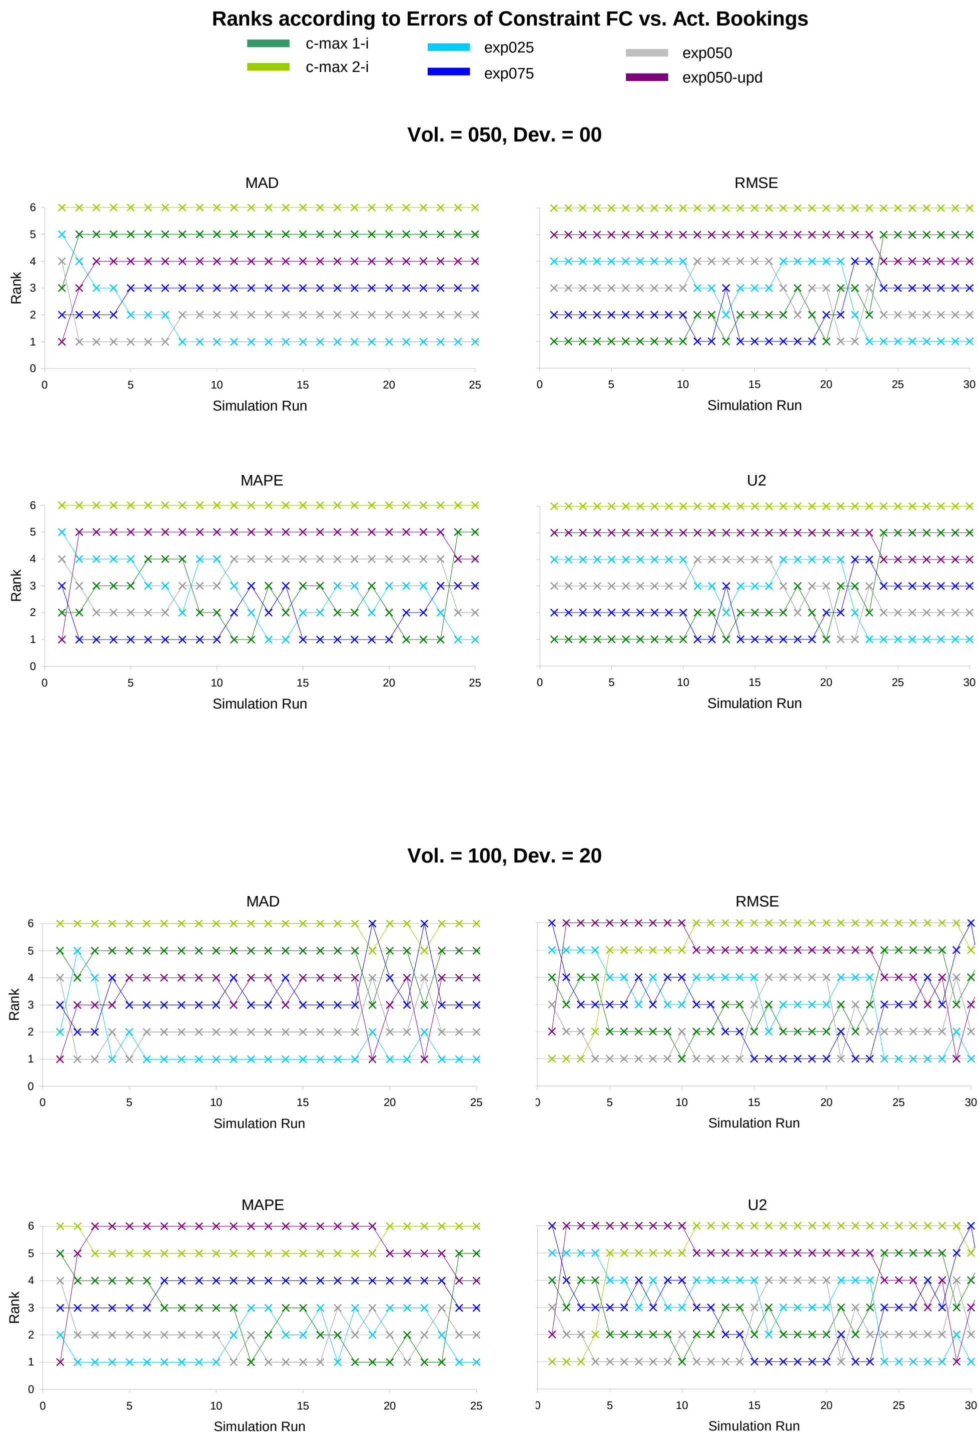

Figure 8.28.: Ranks According to Different Error Measurements

Conclusion: Different error measurements may result in different evaluations of the accuracy of forecasts. A consequence of the choice of objects of comparison, level of comparison, and method of comparison may be quite different rankings of the method compared. Additionally, some simulation experiments involving mirrored psychic forecasts have indicated that an accurate forecast is not necessarily the most successful forecast in terms of resulting revenue.

8.4. Definitions and Effects of Uncertainty of Demand

The previous sections have considered the effects of the application of certain forecast methods over a long term, the consequences of psychic initializations and benchmarks and the characteristics of error measurements. All statements about the analysis of forecast performance under these aspects have been illustrated by the results of simulation experiments over an array of market variations. Depending on the volume of demand included and the deviation of the error term involved in distorting demand over several simulation runs, different results could be observed. The goal of this section is to examine in how far the uncertainty included in a market influences the forecast quality and how this aspect can be included in forecast performance evaluation.

The concept that some markets may include more uncertainty than others and therefore be more difficult to predict has already been mentioned in Section 4.1 with special regard to Diebold & Lopez (1996). When the customer behavior is random and volatile, no forecast can achieve good results. Uncertainty includes two dimensions: information on customer behavior is not available (*apparent* randomness) and the uncertainty of demand is high. Both dimensions are to be analyzed.

Uncertainty due to Lacking Information: A way of comparing forecasts given different amounts of information about customer behavior is to change the decision parameters included in the market. An example of this can be found in the comparison of the product-oriented to the price-oriented scenario. Customers confronted with a class structure clearly differentiated by restrictions and basing their decision mainly on their product acceptance tend to book the same classes whenever they are available. Customers confronted with a product structure that is only differentiated by price and basing their decision mainly on the price tend to book the cheapest class available. A forecast that assumes static demand includes the necessary information to predict product-oriented, but not price-oriented demand. The consequence is a higher degree of uncertainty due to information (on buy-down behavior, in this case) not being available. This lack of information becomes even more influential to the success of the forecast in terms of revenue when overall demand volume is low, as under such circumstances buy-down takes place more frequently.

- Let S^{product} indicate a combination of supply and demand that is based on a differentiation by product characteristics (product-sensitive market).
- Let S^{price} indicate a combination of supply and demand that is based on a differentiation by price (price-sensitive market).
- Let S^{hybrid} indicate a combination of supply and demand that is based on a differentiation by price and product characteristics (hybrid market).
- Let $S = \left\{S^{\text{product}}, S^{\text{price}}, S^{\text{hybrid}}\right\}$ be the set of market scenarios available.
- Let $\Psi\left(S, \sigma_\epsilon, V, F^\circ\right)$ describe a simulation experiment based on a market structure S, a deviation of the error term of demand $\sigma\left(\epsilon\right)$, overall demand volume V and a forecast F°.
- Let $y^{\text{fcfs}\ \%}\left(\Psi\left(S, \sigma_\epsilon, V, F^\circ\right)\right)$ be the average yield generated in a simulation experiment as the percentage of the yield generated when inventory controls based first-come-first-serve were applied.
- Let $r^{\text{psy}\ \%}\left(\Psi\left(S, \sigma_\epsilon, V, F^\circ\right)\right)$ be the average revenue generated in a simulation experiment as the percentage of the revenue generated when inventory controls based first-come-first-serve were applied.
- Let $y^{\text{psy}\ \%}\left(\Psi\left(S, \sigma_\epsilon, V, F^\circ\right)\right)$ be the average yield generated in a simulation experiment as the percentage of the yield generated when inventory controls based on a "c-max 1-i" psychic forecast were applied.
- Let $r^{\text{psy}\ \%}\left(\Psi\left(S, \sigma_\epsilon, V, F^\circ\right)\right)$ be the average revenue generated in a simulation experiment as the percentage of the revenue generated when "c-max 1-i" psychic under first-come-first-serve inventory controls were applied.

- Let $e^{\text{c-c}}_{\text{MAD}}\left(\Psi\left(S,\sigma_\epsilon,V,F^\circ\right)\right)$ be the mean average deviation of the constrained forecast compared to the bookings observed in simulation experiment averaged over all runs.

In Hypothesis (8.51), the following expectation for simulation experiments using a psychic forecast “c-max, 1-i” is formalized: If the market is product-based, the decrease of yield resulting from a decrease in demand volume will not be as steep as when it is price based.

$$\frac{y^{\text{fcfs }\%}\left(\Psi\left(S^{\text{product}},\sigma_\epsilon,V,F^{\text{c-max, 1-i}}\right)\right)}{y^{\text{fcfs }\%}\left(\Psi\left(S^{\text{product}},\sigma_\epsilon,V',F^{\text{c-max, 1-i}}\right)\right)} \leq \frac{y^{\text{fcfs }\%}\left(\Psi\left(S^{\text{price}},\sigma_\epsilon,V,F^{\text{c-max, 1-i}}\right)\right)}{y^{\text{fcfs }\%}\left(\Psi\left(S^{\text{price}},\sigma_\epsilon,V',F^{\text{c-max, 1-i}}\right)\right)} \quad \forall\ V' \leq V; \sigma_\epsilon \in \mathbb{R} \tag{8.51}$$

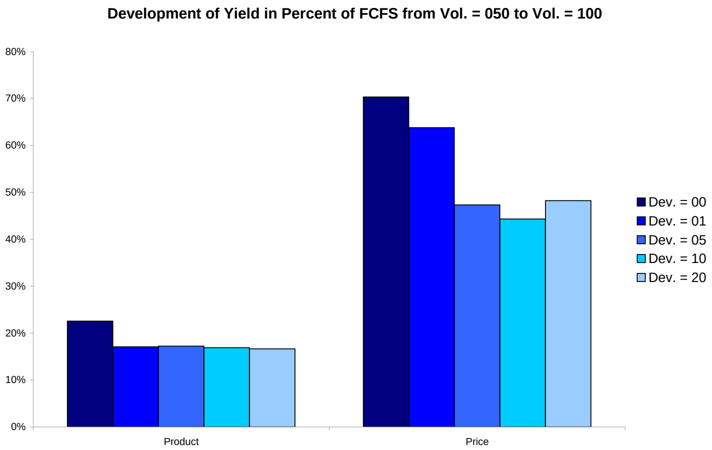

Figure 8.29.: Increase of Yield in Percent of FCFS from Vol. = 050 to Vol. 100

Figure 8.29 shows the increase of yield indicated by $y^{\text{fcfs}\%}$ as demand volume grows from “Vol. = 050” to “Vol. = 100” The increase is depicted in percent of $y^{\text{fcfs}\%}$ given “Vol. = 050”. While growing yield can also be observed in a product-based market, the increase is higher on price-based markets. On a product-based market, yield increases in yield as with absolute demand volume being higher, the absolute share of customers that exclusively request more valuable classes is also higher. On a price-based market, yield increases as the absolute number of customers with a high willingness to pay increases and inventory controls are become more restrictive.

The effect of yield also seems to depend on the deviation given in the market ("Dev. = 00" vs. "Dev. = 20"). In general, it can be stated that deviation leads to a weaker growth in yield as demand volume increases.

Hypothesis (8.52) formalizes the expectation that in a price-based scenario, a static forecast will be less successful at maximizing revenue than in a product-based scenario even when it is based on "psychic" knowledge of maximum willingness to pay as described by Section 8.2.

$$r^{\text{fcfs }\%}\left(\Psi\left(S^{\text{price}},\sigma_\epsilon,V,,F^{\text{c-max, 1-i}}\right)\right) \leq r^{\text{fcfs }\%}\left(\Psi\left(S^{\text{product}},\sigma_\epsilon,V,F^{\text{c-max, 1-i}}\right)\right) \quad \forall\ \sigma_\epsilon \in \mathbb{R}; V \in \mathbb{N} \tag{8.52}$$

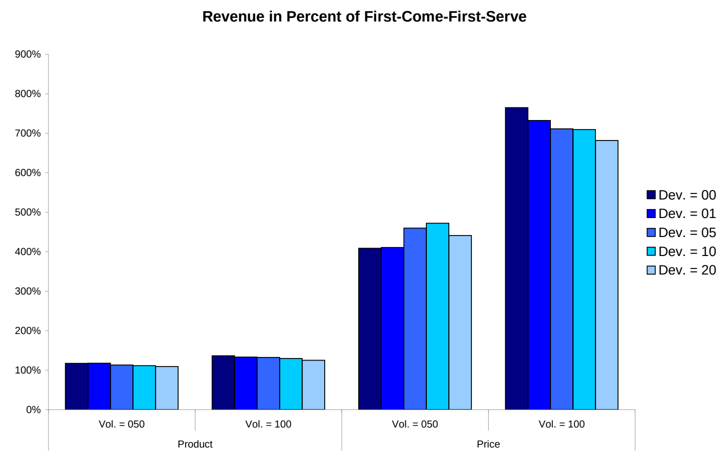

Figure 8.30.: Revenue in Percent of FCFS on Price- and Product-Based Markets

Figure 8.30 shows a comparison of the revenue as a percentage of the average yield observed under first-come-first-serve inventory controls over price- and product-based markets. Obviously, the expectation stated in Hypothesis (8.52) is confirmed.

In the product-based markets, the difference in revenue between first-come-first-serve controls and the "c-max 1-i" psychic forecast is not large to start with: At most, the psychic forecast exceeds first-come-first-serve by 37 % in cases with low deviation ("Dev. = 00") and high demand volume. As customers request tickets for specific classes, and nested inventory controls ensure availability in the most expensive classes, even well-computed protection levels can not improve revenue much if enough capacity is available to accept most requests. This is the reason for a difference of only 9 % for low demand volume and high deviation ("Vol. = 050, Dev. = 20"). For product-based markets, inventory controls are most useful for high demand, when bookings in cheap classes threaten to use capacity that could be sold at higher prices.

The case is quite different when the price-based markets are considered. On these, the difference between first-come-first-serve and the psychic forecast in terms of revenue is at least 300%: Four times as much revenue can be gained when applying the inventory controls based on "c-max 1-i" to low deviation and low demand ("Vol. = 050, Dev. = 00"). As deviation rises, the distribution of demand becomes less homogenuous and on some flights, demand exceeds capacity. In these cases, as already described for the product-based scenarios, inventory controls based on the psychic forecast avoid using the sparse capacity for customers requesting cheap tickets even on markets with low overall demand. Revenue based on the psychic forecast exceeds that based on the first-come-first-serve controls by up to 370 %. When demand volume is high, the effect of high deviation does not add to the revenue. Instead, revenue is at its maximum when deviation is low, exceeding what was earned with first-come-first-serve by 665 %.

Hypothesis (8.53) formalizes the expectation that in any scenario, a static forecast based on exponential smoothing will be less successful at maximizing revenue when the deviation of the distribution the error term is drawn from is high.

$$\begin{gathered}\lim_{\sigma_\epsilon \to \infty} r^{\text{fcfs}\ \%}\left(\Psi\left(S^{\text{price}}, \sigma_\epsilon, V, F^{\circ}\right)\right) = 100 \\ \forall\ S \in \left\{S^{\text{product}}, S^{\text{price}}, S^{\text{hybrid}}\right\}; V \in \mathbb{N}; \left\{F_{\text{exp25}}, F_{\text{exp50}}, F^{\text{exp75}}\right\}\end{gathered} \tag{8.53}$$

As could be seen in Figure 8.30, this statement only holds true for high demand volume. When overall demand volume is low, a high deviation of the error term distribution can have positive effects that may lead to an increase in revenue: As the distribution of demand over itineraries becomes less homogeneous, in spite of low demand volumes, inventory controls become more effective at reserving seats for valuable customers.

Hypothesis (8.54) formalizes the expectation that in a price-based scenario, a static forecast will be less successful at predicting demand accurately than in a product-based scenario. The quality of demand forecast is measured as the mean absolute deviation between the constrained forecast and the observed bookings.

$$e^{\text{c-c}}_{\text{MAD}}\left(\Psi\left(S^{\text{price}}, \sigma_\epsilon, V, , F^{\text{c-max, 1-i}}\right)\right) \geq e^{\text{c-c}}_{\text{MAD}}\left(\Psi\left(S^{\text{product}}, \sigma_\epsilon, V, F^{\text{c-max, 1-i}}\right)\right) \quad \forall\ \sigma_\epsilon \in \mathbb{R}; V \in \mathbb{N} \tag{8.54}$$

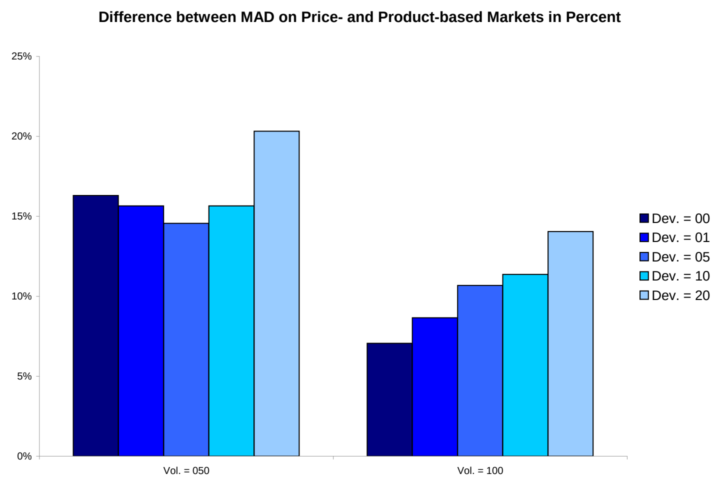

Figure 8.31.: Difference in MAD (Price-Based - Product-Based Market) in Percent

Figure 8.31 shows the positive difference between the MAD observed on price- and product-based markets. It indicates that the MAD computed from the constrained forecast and actual bookings is always higher for price-based markets. As capacity and the overall demand volume were kept constant for all simulation experiments, this fulfills the prediction stated in Hypothesis (8.54).

Uncertainty due to Market Volatility: In the simulation, the uncertainty created by the volatility of a market is modeled as the deviation of the distribution the error term is drawn from. A potentially high error term can result in overall demand volume changing strongly from one run to the other. Furthermore, the distribution of demand over different routes becomes more skewed.

$$\begin{aligned} \sigma_\epsilon &\leq \sigma'_\epsilon \\ \rightarrow e^{\text{c-c}}_{\text{MAPE}}\left(\Psi\left(S, \sigma_\epsilon, V, F^\circ\right)\right) &\leq e^{\text{c-c}}_{\text{MAPE}}\left(\Psi\left(S, \sigma'_\epsilon, V, F^\circ\right)\right) \\ \forall\ S \in \left\{S^{\text{product}}, S^{\text{price}}, S^{\text{hybrid}}\right\}; V \in \mathbb{N}; &\left\{F_{\text{exp25}}, F_{\text{exp50}}, F_{\text{exp75}}\right\} \end{aligned} \tag{8.55}$$

In any scenario a static forecast will be less successful accurately forecasting demand when the deviation of the distribution the error term is drawn from is high. Standard error indicators will interpret this as a decrease in forecast performance, even though not the forecast's abilities but the potential for accurate forecasts has been diminished. For the forecast error MAPE, this expectation is formalized by Hypothesis (8.55).

Figure 8.32 presents the value of $e^{\text{u-u}}_{\text{MAPE}}\left(\Psi\left(S, \sigma_\epsilon, V, F^\circ\right)\right)$ averaged over 50 runs. The error indicator is displayed for simulation experiments applying the naive forecast F^{naive} as well as exponential smoothing methods initialized by the psychic forecast "c-max 1-i"; F^{exp025}, F^{exp050}, and F_{exp075}; to markets S^{product} and S^{price}. Overall demand volume changes from "Vol. = 050" to "Vol. = 100" – the four possible combinations of market-type and demand volume have been used to generate one graph each. The deviation of the error term distribution changes from "Dev. = 00" via "Dev. = 01", "Dev. = 05", and "Dev. = 10" to "Dev. = 20". The mean average percentage error (MAPE) is illustrated by a group of bars for each method, with the color changing from dark blue (low deviation) to light blue (high deviation).

As can be seen in Figure 8.32, the forecast error tends to increase as the deviation of the error term distribution increases. Its level is higher for the exponential smoothing methods than for the naive forecast. This is due to the initialization of the exponential smoothing methods based on the psychic forecast: As explained in Section 8.3, when evaluated by standard error indicators, the psychic forecast leads to high errors. Furthermore, the spiral-down effect that takes place on price-based markets as explained in Section 8.1 leads to smaller errors on price-based markets.

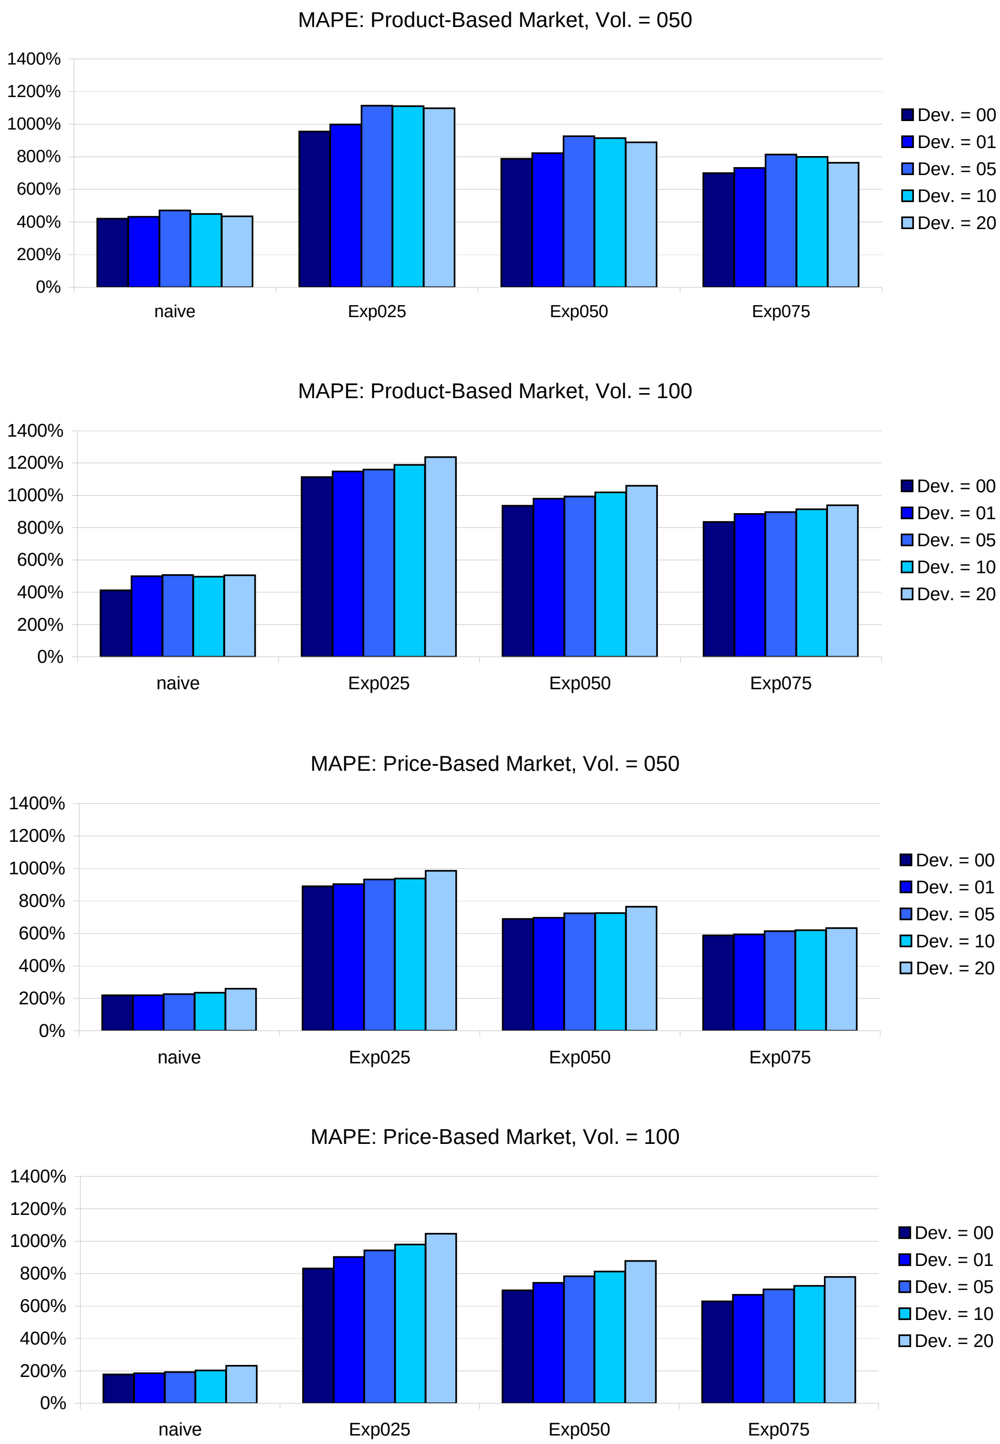

Figure 8.32.: MAPE Averaged over 50 Runs

Indicators of Uncertainty: To normalize the development of the forecast error to the degree of market uncertainty, a new indicator is introduced.

- Let $e_{\text{PB}}^{\text{u-u}}\left(\Psi\left(S, \sigma_\epsilon, V, , F^\circ\right)\right)$ be the percentage of cases in which deviation of unconstrained forecast F used in simulation experiment S from the observed bookings was smaller than that of the naive forecast.

The performance of the naive forecast depends on the volatility of the market: If the same amount of requests for the same classes arrives in every run, it is very successful. While in such a situation, forecasting is also easier for other methods, the benchmark of the naive forecast becomes harder to reach. If high market uncertainty leads to volatile request volumes, the difference between the naive forecast and the actual observations grows. While this makes forecasting harder for other methods, as well, the benchmark of the naive forecast becomes easier to reach. This way, forecast performance is normalized to the situation.

Figure 8.33 presents the value of $e_{\text{PB}}^{\text{u-u}}\left(\Psi\left(S, \sigma_\epsilon, V, F^\circ\right)\right)$ averaged over 50 runs. The performance indicator "Percent Better" is displayed for simulation experiments applying the naive forecast F^{naive} as well as exponential smoothing methods initialized by the psychic forecast "c-max 1-i"; F^{exp025}, F^{exp050}, and F^{exp075}; to markets S^{product} and S^{price}. Overall demand volume changes from "Vol. = 050" to "Vol. = 100" – the four possible combinations of market-type and demand volume have been used to generate one graph each. The deviation of the error term distribution changes from "Dev. = 00" via "Dev. = 01", "Dev. = 05", and "Dev. = 10" to "Dev. = 20". The percentage of cases for which the respective method performed better than the naive method (PB) is illustrated by a group of bars for each method applied, with the color changing from dark blue (low deviation) to light blue (high deviation).

The introduction of a percent-better indicator does normalize the evaluation of forecast accuracy according to the level of uncertainty included in a market. The number of cases in which the evaluated method performed better than the naive forecast does not seem to clearly depend on the deviation of the error distribution. While it grows with rising deviation in the product-based market with high demand volume, it does not show a trend for the other combinations of market and demand volume shown.

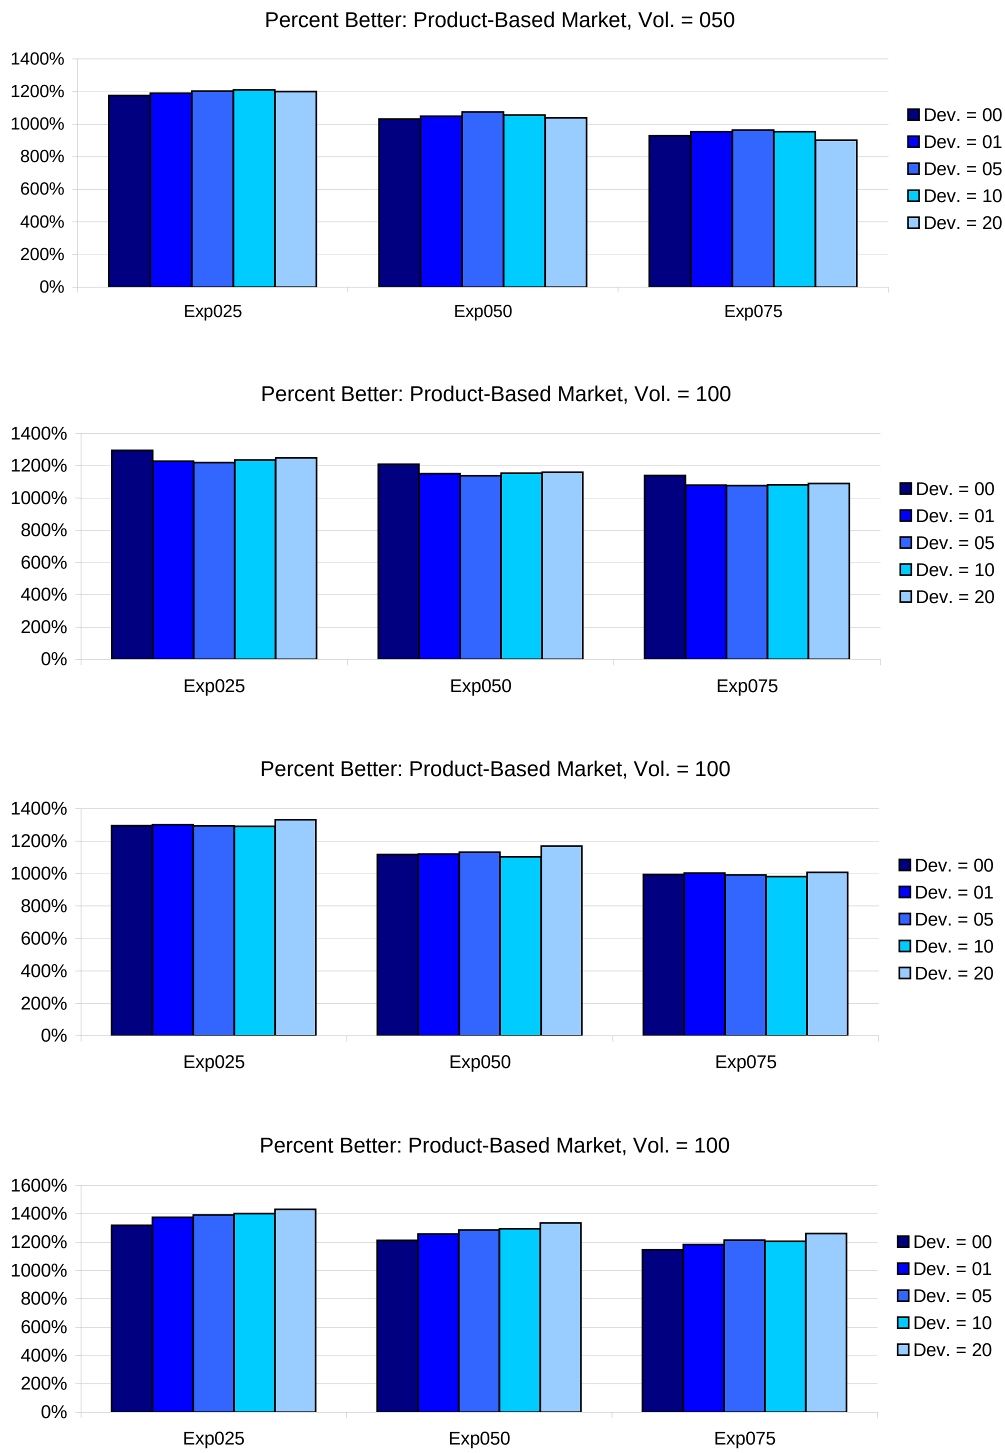

Figure 8.33.: “Percent Better” Averaged over 50 Runs

According to what was shown so far, a general indicator of the uncertainty included in a market may be the quality of the naive forecast. In a very stable environment, a prediction based on the values observed during the last run can be expected to be accurate if no trend influences demand. With demand fluctuating from one run to the other as the deviation of error terms increases, this changes. Hypothesis (8.56) formalizes this expectation.

$$\sigma_\epsilon \leq \sigma'_\epsilon \rightarrow e^{\text{c-c}}_{\text{MAD}} \left(\Psi\left(S, \sigma_\epsilon, V, F^{\text{naive}}\right)\right) \leq e^{\text{c-c}}_{\text{MAD}} \left(\Psi\left(S, \sigma'_\epsilon, V, F^{\text{naive}}\right)\right) \\ \forall\ S \in \left\{S^{\text{product}}, S^{\text{price}}, S^{\text{hybrid}}\right\}; V \in \mathbb{N} \tag{8.56}$$

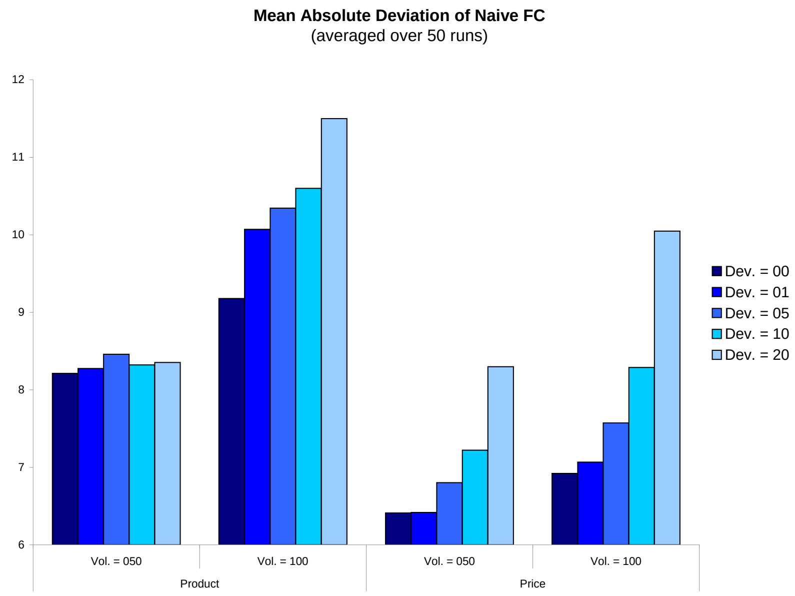

Figure 8.34.: MAD for Naive Forecast Averaged over 50 Runs

Figure 8.34 shows the development of $e^{\text{c-c}}_{\text{MAD}}\left(\Psi\left(S, \sigma_\epsilon, V, F^{\text{naive}}\right)\right)$ over all markets as σ_ϵ increases. As was already observed to some extend with regard to the error indicator MAPE, the mean absolute deviation of the constrained naive forecast from observed bookings increases with an increase in the deviation of the error term distribution. The exception is the product-based market with low overall demand volume: Even with the resulting high error term, product-based customers apparently stay equally predictable as long as demand does not exceed capacity.

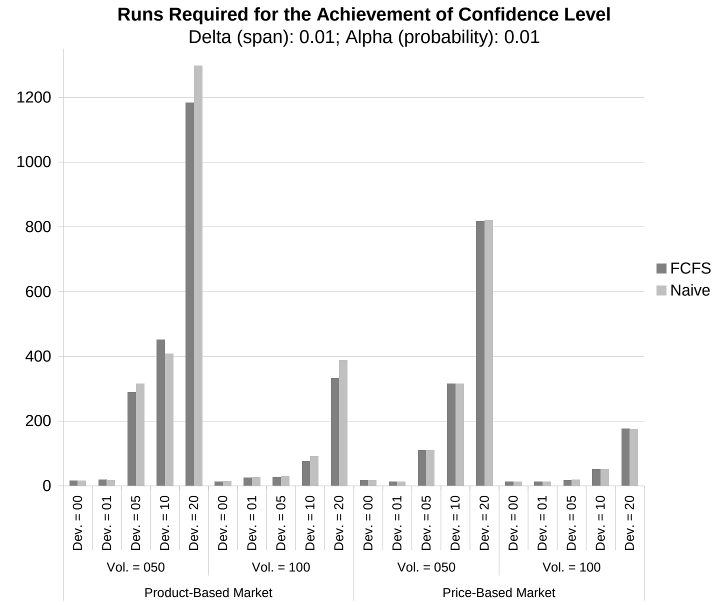

Figure 8.35.: Runs Required for Confidence Level

In a simulation environment, another indication of market uncertainty is the number of runs needed to reach the desired confidence level. Based on the computation of confidence described in Section 7.1, the number of runs a simulation experiment increases with the deviation of the distribution error terms. This observation is illustrated by Figure 8.35. The graph shows the development of the number of runs required for simulation experiments including first-come-first-serve controls or a naive forecast and EMSR-b, based on price- or product-based markets with low or high demand volume and a range of deviations of the error term distribution.

As can be seen, both methods applied need roughly an equal amount of runs to reach the required confidence level. Note that on price-based markets, a naive forecast that is initialized with a zero forecast and first-come-first-serve lead to similar inventory controls and to bookings only in the cheapest class included in the scenario. Adaptive methods such as psychic forecasts and exponential smoothing need many more runs as they interact with the volatility of the market. However, in both cases the number of runs required increases with the deviation of the distribution that the demand error term is drawn from.

Without having to rely on information exclusive to simulation environments, the error variance of the naive forecast can also indicate uncertainty. The more that the fluctuation of demand is influenced by a high error term and thereby randomized, the more random also the quality of a naive forecast. Hypothesis (8.57) formalizes this expectation.

- Let $\sigma^2\left(e^{\text{c-c}}_{\text{MAD}}\left(\Psi\left(S, \sigma_\epsilon, V, F^{\text{naive}}\right)\right)\right)$ be the variation of the mean absolute deviation of the constrained naive forecast from observed bookings over the runs of simulation experiment Ψ.

$$\sigma_\epsilon \leq \sigma'_\epsilon \rightarrow \sigma^2\left(e^{\text{c-c}}_{\text{MAD}}\left(\Psi\left(S, \sigma_\epsilon, V, F^{\text{naive}}\right)\right)\right) \leq \sigma^2\left(e^{\text{c-c}}_{\text{MAD}}\left(\Psi\left(S, \sigma'_\epsilon, V, F^{\text{naive}}\right)\right)\right) \\ \forall\ S \in \left\{S^{\text{product}}, S^{\text{price}}, S^{\text{hybrid}}\right\}; V \in \mathbb{N} \tag{8.57}$$

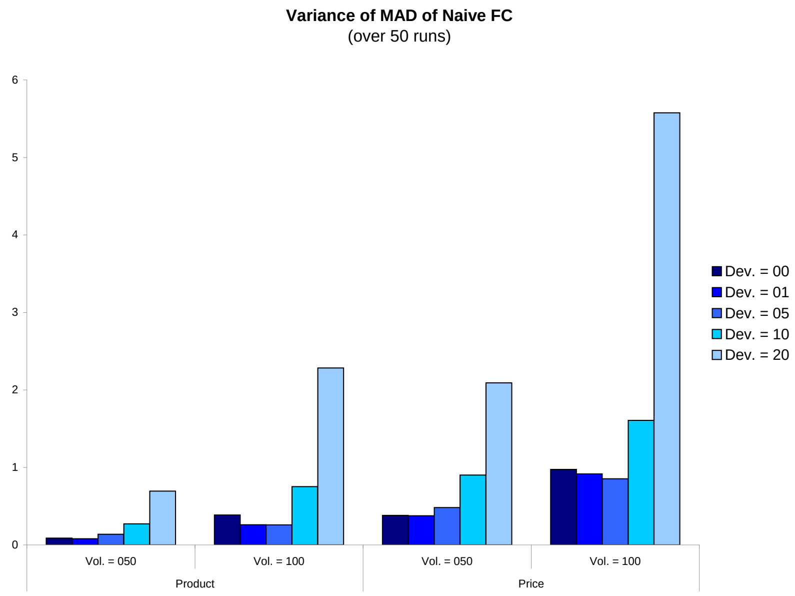

Figure 8.36.: Variance of MAD for Naive Forecast over 50 Runs

Figure 8.36 shows the development of $\sigma^2\left(e^{\text{c-c}}_{\text{MAD}}\left(\Psi\left(S, \sigma_\epsilon, V, F^{\text{naive}}\right)\right)\right)$ as σ_ϵ increases over all markets. As predicted in Hypothesis (8.57), the variance of error increases rapidly as the variance of the error term distribution increases.

Indicators of Robustness: Observing the development of revenue and forecast quality as the deviation of the distribution of error terms increases can also give an indication of the *robustness* of a forecast method. In contrast to an observation of forecast quality under static conditions, robustness provides information on how well a forecast reacts to changes in the market place. Hypothesis (8.58) formalizes the description of a measurement of robustness based on the comparison of revenue as uncertainty represented by the deviation of the distribution that error terms are drawn from increases.

- Let $O^{\text{rev}}\left(F^{\circ}, \sigma_{\epsilon}, \sigma'_{\epsilon}\right)$ be a revenue-based indicator of the robustness of a forecast method F° given an increase in uncertainty from σ_{ϵ} to σ'_{ϵ}

$$O^{\text{rev}}\left(F^{\circ}, \sigma_{\epsilon}, \sigma'_{\epsilon}\right) := \frac{r\left(\Psi\left(S, \sigma'_{\epsilon}, V, F^{\circ}\right)\right)}{r\left(\Psi\left(S, \sigma_{\epsilon}, V, F^{\circ}\right)\right)} \quad \forall\ S \in \left\{S^{\text{product}}, S^{\text{price}}, S^{\text{hybrid}}\right\}; V \in \mathbb{N} \tag{8.58}$$

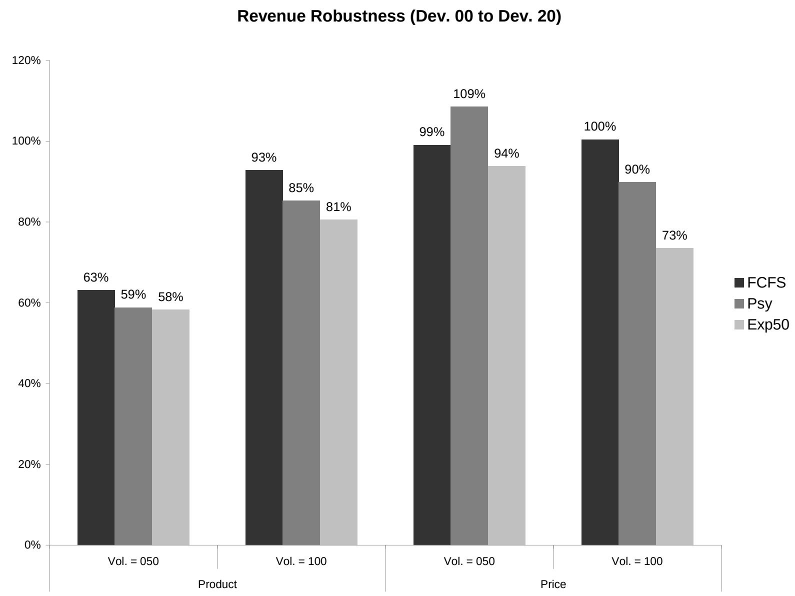

Figure 8.37.: Revenue Robustness based on Rev. Averaged over 50 Runs

Figure 8.37 shows the revenue-based robustness of first-come-first-serve controls, the psychic forecast based on "c-max 1-i", and the exponential smoothing method Exp050 for $\sigma'_{\epsilon} = 0$ and $\sigma_{\epsilon} = 20$. As may be expected, the first-come-first-serve controls are the most robust choice on most market variations: First-come-first-serve inventory controls are not changed based on forecasts that may be mislead by demand with a high deviation of the error term distribution. However, based on these controls, the lowest absolute revenue is generated in most cases. The psychic forecast, which anticipates at least part of the demand volatility caused by the error term, comes in on second place with regard to robustness.

Conclusion: Some combination market characteristics, labeled uncertainty, has an effect on the potential success of any forecast method. In the simulation environment used for experiments, uncertainty is modeled by an error term influencing the volatility of customer behavior. As the deviation of the normal distribution this error term is drawn from increases and thereby uncertainty rises, forecasts perform worse both in terms of accuracy and in terms of resulting revenue. Finally, some indicators based on first-come-first-serve controls or the naive forecast, by which to estimate the level of uncertainty included in a market, have been introduced.

8.5. Evaluation Approaches for Price-Sensitive Forecasts

As some markets include more uncertainty than others, some forecast methods include more information than others. The previous section presented simulation experiments designed to show how markets where customers make decisions based on price-sensitivity pose a problem for static forecasts not taking into account information such as which class is the lowest available.

This section intends to demonstrate the advantages of additional demand information considered in both the forecast and its evaluation. As an example for additional information, the price sensitivity of customers is considered and a forecast method including predictions on customers' buy-down behavior is applied to the price-sensitive scenario.

The new forecast is based on price-sensitive estimators depending on the cheapest class available. Details of its implementation are described in Section 7.3.1. It is influenced by a weight parameter α^P balancing the influence of price to that of time before departure. In the further text, the results of the use of this forecast are marked as "price09" - the parameter that determines the weight of the price-estimators was set to 0.9 for the experiments presented here. The exponential smoothing parameter α^{exp} can also be varied. However, as the general consequences of this parameter have already been presented in

Section 8.1, it is kept at $\alpha^{\text{exp}} = 0.5$ for all experiments presented here. For this reason, the exponential smoothing forecast used for comparison is also "Exp050".

To maximize the observed effects, the price-oriented market scenario is used in all simulation experiments with the price-sensitivity forecast. When booking classes are only differentiated by price and customers base their decisions on the information which class is the cheapest available, the new forecast should fully realize its potential.

Revenue: First, it needs to be established that the new forecast method actually is advantageous compared to exponential smoothing approaches. Hypothesis (8.59) presents the expectation that at least one parametrization of the price-sensitive forecast performs better in term of revenue than the exponential smoothing alternative.

- Let $\Psi\left(S^{\text{price}}, \sigma_\epsilon, V, F^\circ\right)$ be a simulation experiment based on the price-sensitive market scenario S^{price} with error deviation σ_ϵ, demand volume V and the forecast method F°.
- Let $r_{\Psi(S^{\text{price}}, \sigma_\epsilon, V, F^\circ)}(s)$ be the revenue generated in run $s = 1, ..., N^s$ of the above described simulation experiment.

$$\begin{aligned} \exists\ F^{\text{price}} \rightarrow r_{\Psi(S^{\text{price}}, \sigma_\epsilon, V, F^{exp050})}(s) \leq r_{\Psi(S^{\text{price}}, \sigma_\epsilon, V, F^{\text{price}})}(s) \\ \forall\ \sigma_\epsilon \in \mathbb{R};\ V \in \mathbb{N} \end{aligned} \tag{8.59}$$

Figure 8.38 shows the results of the application of the psychic forecast "c-max, i-1", the exponential smoothing method "exp050" and the price-sensitive forecast "market09" to the price sensitive market variations over volume and the deviation of the error term. To provide a measure of comparison, revenue is presented as a percentage what was earned when first-come-first-serve inventory controls are applied. All methods are initialized with the $F_{\text{c-max, 1-i}}$ psychic method.

Accuracy: With new parameters and variables included, new forecast evaluation methods become available. In addition to comparing observed bookings to the predicted demand, observed customer behavior may now also be transformed to make it comparable to the estimators used in the new forecast.

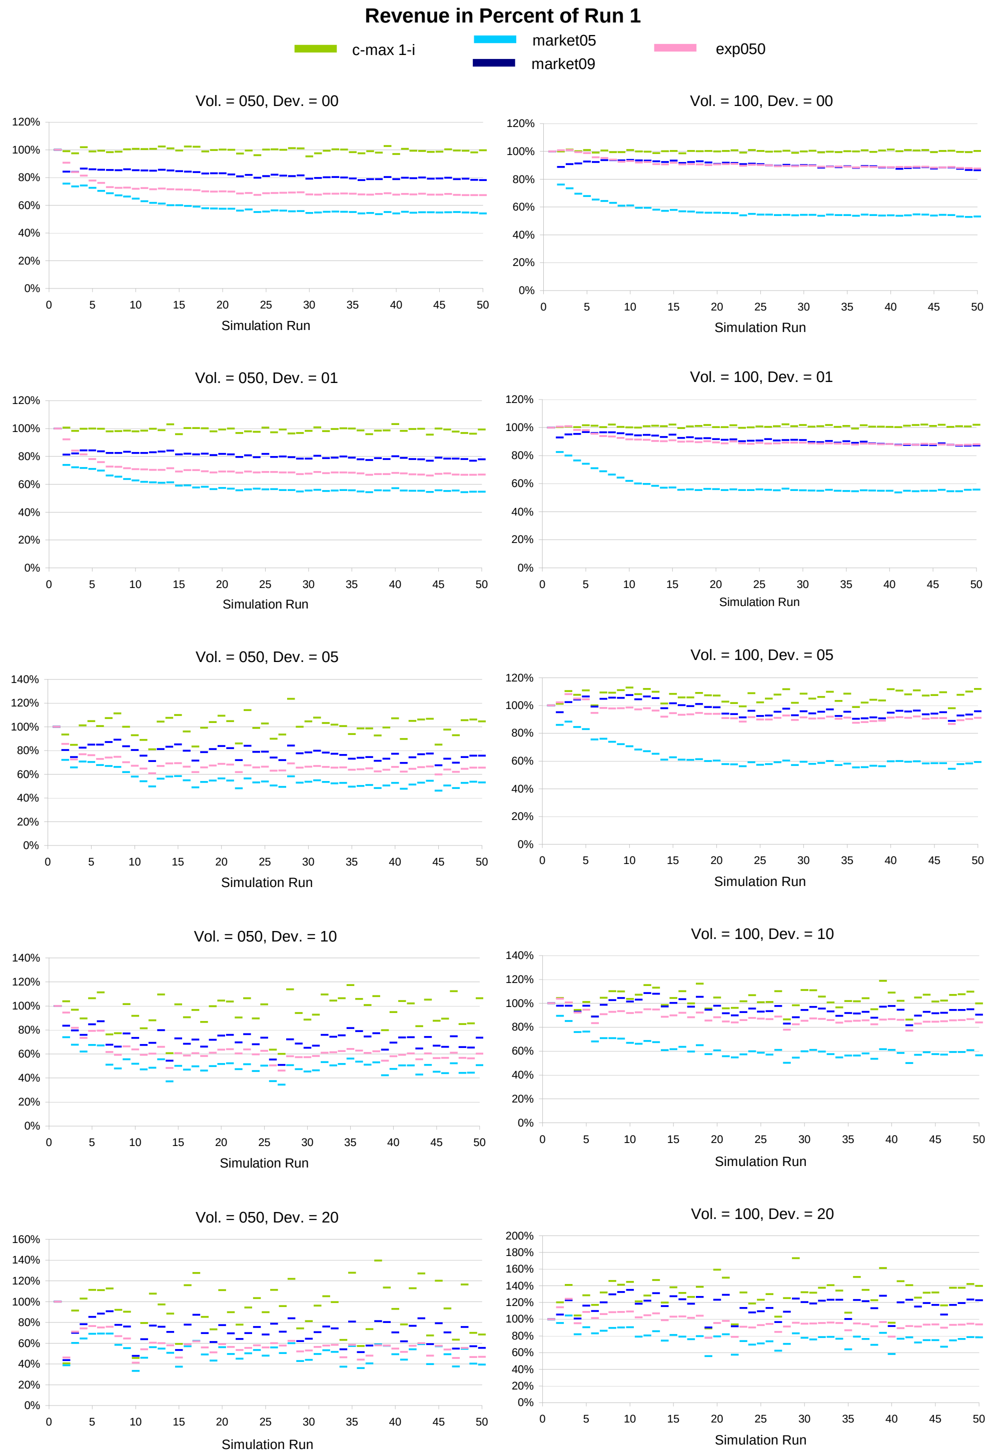

Figure 8.38.: Revenue in Percent of Revenue Earned in Run 1

- Let $b^{\text{unc}}(f, c, t, s)$ be the unconstrained bookings for flight f observed in class c during time slice $t-1$ to t before departure in simulation run s.
- Let $\omega_F^{P_{f,c,s}}$ be the vector of price elasticity for flight f depending on the class c for run s according to the forecast F.
- Let $\omega_b^{P_{f,c,s}}$ be the vector of price elasticity observed for flight f depending on the class c for run s according to observed bookings.
- Let $\hat{e}_{\omega\text{MAD}}(s)$ be the mean absolute deviation of the forecasted price elasticity from the observed price elasticity during run s.

Definition (8.60) shows the computation of the price elasticity from observed bookings. As no consideration is given to time before departure and the assumption of the price-sensitive forecast is that all customers always buy the lowest class available, elasticity can be computed as the share of bookings in one class compared to the overall bookings.

$$\omega_{\text{b}}^{P_{f,c,s}} = \frac{\sum_{t=1}^{N^t} b^{\text{unc}}(f, c, t, s)}{\sum_{c=1}^{N^c} \sum_{t=1}^{N^t} b^{\text{unc}}(f, c, t, s)} \quad \forall\ f \in F;\ c \in \{2, ..., N^c\};\ s = 1, ..., N^s \tag{8.60}$$

Definition (8.61) shows the computation of the mean absolute deviation of forecasted elasticity from observed elasticity. The only difference to the computation of the previously used MAD indicator is in the indicators used for comparison.

$$\hat{e}_{\omega\text{MAD}}(s) := \frac{\sum_{f \in F} \sum_{c=2}^{N^c} \left| \omega_F^{P_{f,c,s}} - \omega_b^{P_{f,c,s}} \right|}{|F| \cdot (|C| - 1)} \quad \forall\ s = 1, ..., N^s \tag{8.61}$$

After spiral-down is completed, as presented in Section 8.1, exponential smoothing methods assuming static demand will predict all demand to request the cheapest booking class in a purely price-based scenario. According to this logic, the elasticity vector of "Exp050" is filled with zeros as no sell-up is predicted.

In Figure 8.39, the development of $e_{\omega\text{MAD}}$ is shown for two experiments with price-sensitive forecasts, "market05" and "market09", over the course of 50 simulation runs.

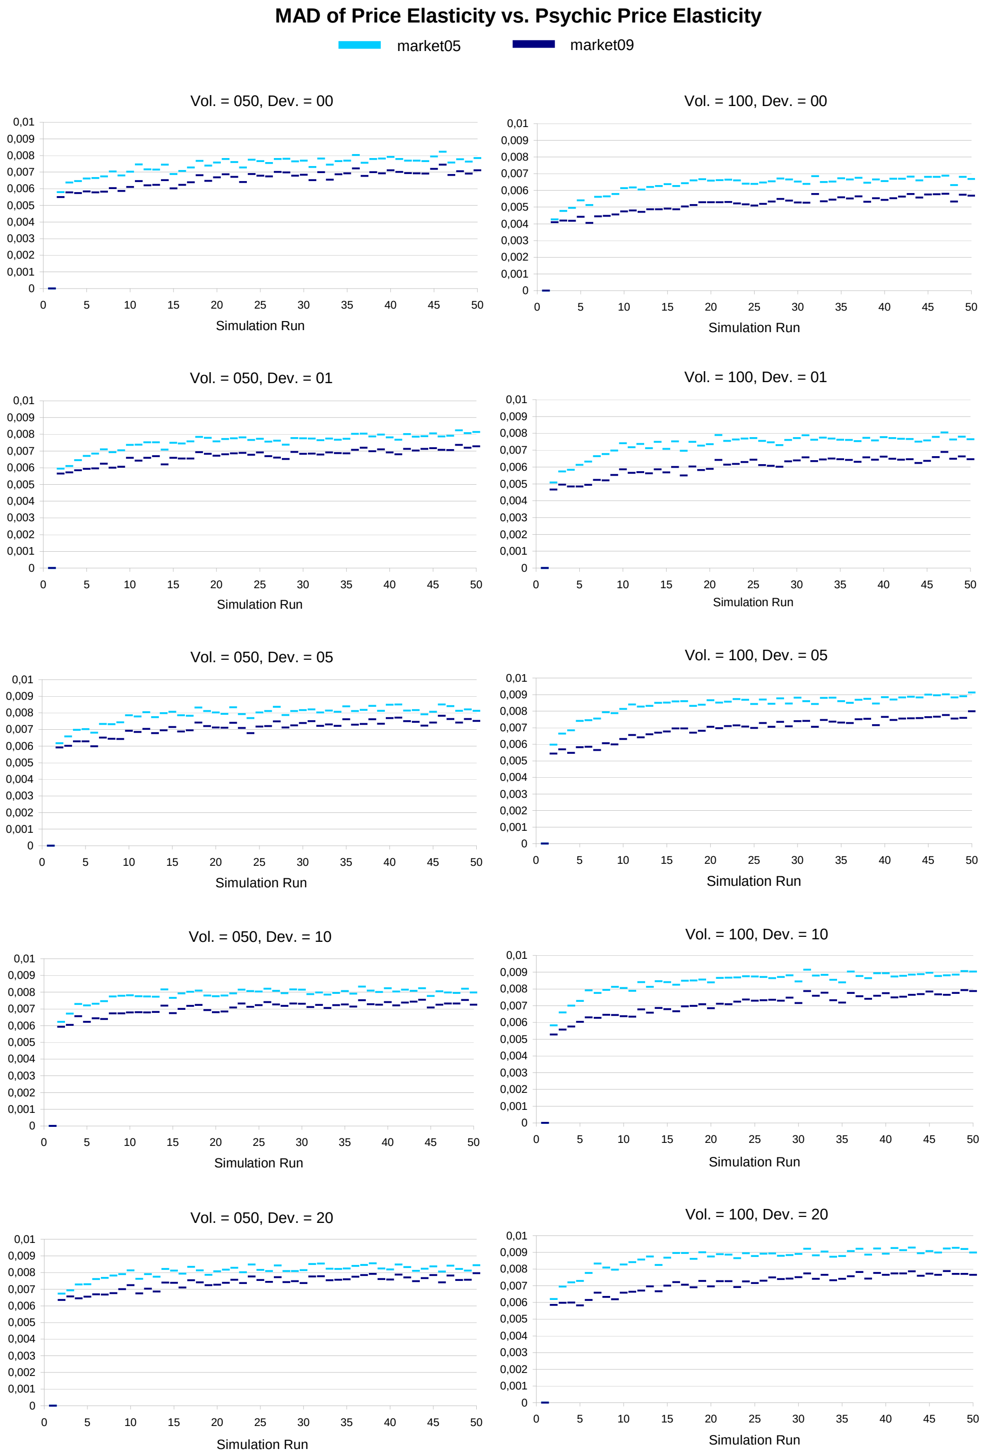

Figure 8.39.: MAD of Elasticity vs. Psychic Elasticity

As in the case of the price-sensitive forecast “market05”, the update includes more time-sensitive information, the spiral-down effect occurs. The result of this could already been seen in Figure 8.38, as revenue declined. In the case of the mean absolute deviation, this results in higher errors compared to psychic knowledge of price elasticity.

Conclusion: The advantages of forecasts that include a model of demand that is not static when confronting flexible demand structures have been pointed out. As a conclusion, the need for new error measurements including objects of comparison that differ from the traditional set of bookings and forecasted demand volume may be stressed.

8.6. Simulation-Based Findings Recaptured

A range of statements on the relationship between forecast performance, forecast accuracy, and forecast evaluation methods have been formalized. Expectations concerning the long-term effects of methods were drawn from existing research with regard to recent challenges as described in Section 2.3. Possible consequences of definitions of psychic forecasts were highlighted as the by-product of thoughts on customer-choice arising when designing the implementation of the demand model as described in Section 7.2.2. The standard accuracy indicators evaluated have been introduced previously in Section 4.1. The considerations of the effect of uncertainty of demand were listed as the consequence of observations on results depending on the deviation of the error term. Finally, some evaluation approaches for price-sensitive forecasts were offered using the forecast method described in Section 7.3.1.

Theoretical statements were used to design experiments using the simulation system described in Chapter 7. Based on the decomposition concept introduced in Chapter 6, for the first time such theories were analyzed ceteris-paribus in a system that include a volatile, flexible customer model and a sophisticated range of supply. Processing the output of the simulation experiments lead to information concerning the conditions under which the theoretical statements apply.

A number of findings has been verified using the concept for the decomposition and evaluation of forecasts as documented in Chapter 6 based on the simulation environment presented in Chapter 7. The following observations have been made with regard to the long-term application of methods in revenue management as described in Section 8.1:

- When forecasts based on static demand assumptions are confronted with a flexible, price-sensitive customer model and a restriction-free class structure, the so-called spiral-down effect can be observed (see Cooper et al. (2006).
- The spiral-down effect leads to decreasing forecasts for more valuable classes, decreasing amounts of protected seats for more valuable classes, and decreasing numbers of bookings in the valuable classes. At the same time, forecasts, availability and bookings for the cheapest class increase.
- The ultimate consequence of the spiral-down effect is a loss in revenue. However, in markets with high demand volume, the increase in bookings can temporarily compensate for a loss in yield, resulting in increasing revenue for a limited time.
- The spiral-down effect is more drastic in markets with low demand volume and volatile customer behavior.
- For adaptive forecast methods based on static demand that include a higher adaptive weight, the spiral-down effect is more drastic than for those including a low adaptive weight.

The following observations have been made with regard to the consequences of possible definitions of psychic forecasts as described in detail in Section 8.2:

- In a simulation system, psychic forecasts may be used to initialize or to evaluate other forecast methods.
- When predicting demand to arrive per booking class and itinerary, psychic forecasts may differ with regard to the interpretation of customers' willingness to pay and preferred itineraries.
- The method chosen for generating the psychic forecast has consequences with regard to its revenue performance, ...
- ... with regard to the revenue performance that adaptive methods based on it show, ...

- ... and with regard to the evaluation of accuracy of forecast methods compared to it.

The following observations have been made with regard to the evaluation of standard accuracy indicators for demand forecasts in revenue management as described in Section 8.3:

- Forecast error measurements may lead to different results depending on the objects, the level, and the method of comparison.
- Error measurements including different units or operate on different levels of aggregation may be compared by comparing the resulting ranks of evaluated methods.
- With regard to the object of comparison, the decision whether to compare constrained forecasts to actual bookings, unconstrained forecasts to unconstrained bookings, or unconstrained forecasts to actual bookings needs to be taken.
- With regard to the level of comparison, even when an optimization method does not explicitly use information on the timing of demand arrival, it can be useful to include timing in the evaluation of forecast accuracy. The information available on this level may be used later in the process, for example when updating the forecast within the booking horizon.
- The psychic forecast can be used as a benchmark for the self-improving effects of forecasts that trigger a spiral-down effect.

The following observations have been made with regard to the definition and the effect of uncertainty of demand as described in Section 8.4:

- The potential for forecast success in terms of accuracy and revenue depends on the level of uncertainty included in a market.
- Uncertainty may be measured by the success of the naive forecast.
- In a simulation system, the robustness of forecast methods with regard to uncertainty included in a market may be measured by comparing the results of the forecast when normalized to the results of first-come-first-serve inventory controls at different levels of uncertainty.

The following observations have been made with regard to the evaluation of price-sensitive forecast methods as described in detail in Section 8.5:

- In price-sensitive markets including a restriction-free product structure and customers that base their decisions on price, forecasts that include a non-static model of demand allow only for a weak spiral-down-effect.
- Forecasts that predict price-sensitive behavior require new evaluation methods based on other objects of comparison than overall bookings and predicted demand volume.

According to the method applied in this chapter, more theories may be formalized and tested. After a summary of the ideas introduced in this thesis, the final chapter will provide a number of suggestions concerning further fields of inquiry.

This chapter described analyses applying the simulation environment for revenue management to the decomposition and evaluation of demand forecasts. The documentation of findings that have been made possible by the concept and the implementation of this system represents the closing argument of this text.

9. Conclusion

The following sections first summarizes the ideas discussed in this text. Finally, an outlook of future possible research and open questions is provided.

9.1. Summary

In this thesis, a new concept to evaluate demand forecast methods for revenue management based on a decomposed view of the system was introduced. With the help of such a concept, the components of revenue management and demand forecasts can be isolated and evaluated. Interferences stemming from the market environment and the interaction of the parts of the system can be controlled and analyzed: When evaluated in the presented framework rather than on an actual market, the success of a method can be isolated from possible economic trends. When the output of individual components is analyzed, the performance of specific methods can be separated from their fit with the system. These merits were demonstrated using a simulation environment. This chapter summarizes the steps taken.

Background: Having introduced revenue management in general and the problems of forecast performance evaluation in particular, a list of tasks was formulated in the initial chapter, 1. Consecutively, these tasks have been made seized and set as goals for research. The results have been documented in the previous text.

Categorization of Forecast Methods: In order to provide a sound basis for a new concept of forecast evaluation, existing approaches to demand forecasting for revenue management were listed in Chapter 3. Research on demand forecasting was categorized by three aspects: the prediction of *volume*, the *unconstraining* of bookings to compute historical demand figures, and *customer behavior*. With regard to all categories, the influence of forecasting on the outcome of revenue management was emphasized. The categorization serves as the basis for a decomposition approach.

Characterization of Forecast Evaluation Methods: In Chapter 4, demand forecast performance measurements are listed and characterized by the object, the level, and the method of evaluation. A theoretical background of existing indicators is provided and instances of forecast performance measurement applied in existing publications are listed.

In addition, in Chapter 4, research opportunities related to the special difficulties of forecast evaluation for revenue management were highlighted. Based on the established knowledge of forecast methods and approaches to forecast evaluation, the need for further investigation into the decomposed evaluation was further justified. The consequences of different error measurements are analyzed in Chapter 8.

Conceptualization of the Decomposition of Revenue Management Systems: To introduce such a decomposed evaluation, Chapter 6 describes revenue management systems in the terms of separate forecast, optimization and inventory modules. The forecast module is further decomposed along the lines of the categorization introduced in Chapter 3. This results in separate components for the prediction of overall demand volume, the unconstraining and the prediction of demand behavior. Such a framework enables the isolated evaluation and comparison of forecast methods ceteris-paribus.

Development of Processes for the Evaluation of Demand Forecast Components: In Chapter 6, a new concept for the decomposition and evaluation of revenue management systems is presented. In the framework described, the special characteristics of demand forecasting for revenue management can be confronted by isolating aspects and using knowledge of the customer demand model to compute new benchmarks. Detailed descriptions of the use of this system to analyze a complete revenue management system, the whole forecasting component as well as the aspects of predicting volume, unconstraining bookings to compute historical demand and predicting customer behavior are offered.

Implementation of a Simulation Environment: In Chapter 7, a simulation environment for revenue management that enables the decomposed view is documented. This includes a description of simulation control as well as of the supply information and the demand model included. Furthermore, the revenue management components implemented are formally introduced. Market implementations available based on the supply and demand model are also outlined. The simulation-based approach provides access to information on customer behavior not available in the real world and allows for the accelerated observation of long-term developments. Given this view, the problem caused by forecasts' influence on observed bookings that later serve as a quality benchmark can be quantified and avoided.

Formalization of Statements on Forecast Performance Evaluation: A number of statements on forecast performance are formally expressed in terms of the decomposed concept in Chapter 8. This includes a long-term view of the effects of methods applied, a discussion of alternative methods of computing forecasts from available knowledge of the demand model, a comparison of existing key performance indicators for forecast evaluation in a simulation, an analysis of the concept of uncertainty of markets and finally an introduction to the advantages and additional features of price-sensitive forecasting.

Analysis of Simulation Experiments: Based on statements on forecast performance evaluation, simulation experiments have been designed, conducted and documented in Chapter 8. For example, the difficulties stemming from traditional forecast evaluation with regard to the spiral-down effect are demonstrated. Open decisions concerning the transformation of knowledge on the demand model are analyzed. The effects of differences in the market place with special regard to the concept of uncertainty are demonstrated and quantified.

9.2. Outlook

Some opportunities for future research in revenue management become obvious without special regard to forecast evaluation. New forecasting methods taking into account flexible customer behavior based on an ever more transparent market are still required and so are optimization algorithms to efficiently compute inventory controls that maximize revenue based on this information.

With regard to topics of this thesis, three topics that further research may expand on emerge. These are the decomposition concept presented in Chapter 6, the simulation environment introduced in Chapter 7, and applications of the combination of concept and environment to the analysis of forecast evaluation documented in Chapter 8.

First of all, as a theoretical approach, the idea of decomposing the revenue management system to evaluate its components separately and ceteris-paribus may not only be applied to forecasting. Optimization methods and inventory solutions can also be analyzed in this way. The following list offers a number of problems that may be considered in a similar way as forecasting has been analyzed so far:

- Given a fixed market and forecast, what revenue management optimization method works best?
- Given a fixed market, how can pricing and product design be improved?
- Given certain conditions of demand, how can scheduling and flight planning be evaluated?
- Given methods of revenue management, how can fleet assignment be designed to use opportunities of synergy?

Secondly, on a larger scale, different parts of the airline planning problem could be modeled as a process built from separate components that may be interlinked to different degrees.

- Given a choice behavior of crew, how can crew scheduling be processed to maximize gain both for the employees and the airline?
- Given certain options of action and probabilities of disturbance, how can operations handling be designed in a robust fashion?

Thirdly, apart from including more aspects of the airline planning problem, the simulation environment presented as a method of realizing the concept of decomposition could be expanded to present a more realistic model of revenue management. From the customer cost function, which is strictly linear so far, to the inclusion of cancellations and no-shows, a whole array of improvements is conceivable. The following list offers a number of features that may be included in future versions of such a simulation environment:

- More methods for demand forecasting (network-based forecasts, market-sensitivity, estimation maximization, neural networks);
- more methods for revenue optimization (network-based optimization, linear and dynamic programming, new heuristics);
- real-time dynamic planning to realize strategic goals by applying rule-based systems in the inventory;
- expansions to the customer model: more decision factors, different functional forms for the cost function, cancellations, group bookings;
- expansions to the time horizon: seasonal departures, updating of forecasting within simulation runs for consecutive departures.

When the simulation environment is considered as an object of further research, the existence of implicit assumptions within and their consequences for evaluations based on it have to be kept in mind. For example, any aspect of customer choice that is influential in the real world but has not been identified in the model so far may distort the findings. Therefore, whenever additional layers are added to this system, the risk of implicit assumptions needs to be noted.

Fourthly, there are still many open questions that may be approached empirically with the help of a concept for decomposition based on a simulation environment. Some of these problems could already be tackled by the existing system. The following list includes topics that are closely related to those presented in the last chapter.

- In how far can a simulation system be calibrated to present realistic market behavior?
- Can a calibrated simulation system be used to generate new boundaries for revenue opportunity?
- Can more aspects of market uncertainty be isolated and quantified?
- Can forecasts be evaluated based on the number of dimensions of demand they predict?

Finally, a more abstract idea may be drawn from the concept and experiments presented here. Using a simulation system to evaluate the interaction of a clearly defined set of rules with a number of entities basing their actions on a cost function is a concept that can be applied to many other fields of research. As decribed with regard to customer choice behavior in the context of airline revenue management, the same conditions can be replicated and confronted with different strategies in any simulation system. The challenge is always the decomposition of the system, the efficient definition of the boundaries of the model included in the simulation, and the formulation of effective analyses and experiments.

Bibliography

Abdelghany, A., & Abdelghany, K. (2007). Evaluating airlines ticket distribution strategies: a simulation-based approach. *International Journal of Revenue Management*, *1*(3), 231–246.

Abdelghany, A., & Abdelghany, K. (2008). A micro-simulation approach for Airline Competition Analysis and Demand Modelling. *International Journal of Revenue Management*, *2*(3), 287–306.

Abed, S., Ba-Fail, A., & S.M., J. (2001). An econometric analysis of international air travel demand in Saudi Arabia. *Journal of Air Transport Management*, *7*(3), 143–148.

Aburto, L., & Weber, R. (2007). A sequential hybrid forecasting system for demand prediction. *Lecture Notes in Computer Science*, *4571*, 518–532.

Algers, S., & Beser, M. (2001). Modelling choice of flight and booking class-a study using Stated Preference and Revealed Preference data. *International Journal of Services Technology and Management*, *2*(1), 28–45.

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

American Airlines. (1987). The art of managing yield. *American Airlines Annual Report*, 22–25.

Anderson, C., & Wilson, J. (2003). Wait or buy? The strategic consumer: Pricing and profit implications. *Journal of the Operations Research Society*, *54*(3), 299–306.

Andersson, S. (2001). Forecasting and optimisation for airline operational planning. *International Journal of Services Technology and Management*, *2*(1), 161–172.

Armstrong, J. (2001). *Principles of forecasting: A handbook for researchers and practitioners*. Kluwer Academic Publishers.

Armstrong, J., & Collopy, F. (1992). Error measures for generalizing about forecasting methods: Empirical comparisons. *International Journal of Forecasting*, *8*(1), 69–80.

Armstrong, J., Collopy, F., & Yokum, J. (2004). Decomposition by causal forces: A procedure for forecasting complex time series. *International Journal of Forecasting*, *17*, 143–157.

Armstrong, J., & Fildes, R. (1995). On the selection of error measures for comparisons among forecasting methods. *Journal of Forecasting*, *14*(1), 67–71.

Ball, M. O., & Queyranne, M. (2006). *Toward robust revenue management: Competitive analysis of online booking.*

Barnhart, C., Belobaba, P., & Odoni, A. (2003). Applications of Operations Research in the Air Transport Industry. *Transportation Science*, *37*(4), 368.

Battersby, B. (2005). *Consumer demand theory and regional air travel: An integrated economic and econometric approach*. Charles Stuart University.

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

Belobaba, P. (1987a). Airline yield management: An overview of seat inventory control. *Transportation Science*, *21*(2), 63–73.

Belobaba, P. (1987b). *Air travel demand and airline seat inventory management*. Thèse de doctorat non publiée, Massachusetts Institute of Technology.

Belobaba, P. (1989). Application of a Probabilistic Decision Model to Airline Seat Inventory Control. *Operations Research*, *37*(2), 183–197.

Belobaba, P., & Gorin, T. (2004). Revenue management performance in a low-fare airline environment: insights from the Passenger Origin-Destination Simulator. *Journal of Revenue and Pricing Management*, *3*(3), 215.

Belobaba, P., & Hopperstad, C. (2004). Algorithms for revenue management in unrestricted fare markets. In *AGIFORS Reservations and Yield Management Study Group Annual Meeting Proceedings.* (Auckland, New Zealand)

Belobaba, P., & Weatherford, L. (1996). Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations. *Decision Sciences*, *27*(2), 343–363.

Belobaba, P. P. (1998). Pods results update: Impacts of forecasting on o-d control methods. In *AGIFORS Reservations and Yield Management Study Group Annual Meeting Proceedings.* (Melbourne, Australia)

Ben-Akiva, M., & Lerman, S. (1985). *Discrete Choice Analysis: Theory and Application to Travel Demand.* MIT Press.

Bertsimas, D., & de Boer, S. (2005). Simulation-based booking limits for airline revenue management. *Operations Research*, *53*(1), 90–106.

Bhadra, D. (2003). Demand for air travel in the United States: Bottom-up econometric estimation and implications for forecasts by origin-destination pairs. *Journal of Air Transportation*, *8*(2), 19–56.

Bhatia, A., & Parekh, S. (1973). Optimal Allocation of Seats by Fare. *AGIFORS Reservations and Yield Management Study Group Meeting*.

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

Bliemel, F. (1973). Theil's forecast accuracy coefficient: A clarification. *Journal of Marketing Research*, *10*(4), 444–446.

Bodily, S., & Weatherford, L. (1995). Perishable-asset Revenue Management: Generic and Multiple-price Yield Management with Diversion. *Omega*, *23*(2), 173–185.

Bohutinsky, C. (1990). *The Sell Up Potential of Airline Demand.* Massachusetts Institute of Technology, Dept. of Aeronautics & Astronautics, Flight Transportation Laboratory.

Botimer, T., & Belobaba, P. (1999). Airline pricing and fare product differentiation: A new theoretical framework. *Journal of the Operational Research Society*, *50*(11).

Boyd, E. (2004). Dramatic changes in distribution will require renewed focus on pricing and revenue management models. *Journal of Revenue and Pricing Management*, *3*(1), 100–103.

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

Boyd, E., & Kallesen, R. (2004). The science of revenue management when passengers purchase the lowest available fare. *Journal of Revenue and Pricing Management*, *3*(2), 171–177.

Brons, M., Pels, E., Nijkamp, P., & Rietveld, P. (2002). Price Elasticities of Demand for Passenger Air Travel: A Meta-Analysis. *Journal of Air Transport Management*, *8*, 165–175.

Brumelle, S., & McGill, J. (1993). Airline seat allocation with multiple nested fare classes. *Operations Research*, *41*(1), 127–137.

Brumelle, S., McGill, J., Oum, T., Sawaki, K., & Tretheway, M. (1990). Allocation of airline seats between stochastically dependent demands. *Transportation Science*, *24*(3), 183–192.

Brumelle, S., & Walczak, D. (1997). Dynamic Allocation of Airline Seat Inventory with Batch Arrivals. *NASA*(19980018508).

Buhr, J. (1982). Optimal Sales Limits for Two-Sector Flights. *AGIFORS Symposium Proceedings*, *22*.

Calder, S. (2006). *No frills: The truth behind the low-cost revolution in the skies.* Virgin Books.

Carrier, E. (2003). Market Based Revenue Management with Competitor Spill-out. In *PODS Consortium Meeting.* (Zürich, Switzerland)

Castelli, L., Pesenti, R., & Ukovich, W. (2003). An airline-based multilevel analysis of airfare elasticity for passenger demand. In *Proceeding of the 7 thATRS conference.*

Chatfield, C. (1988). Apples, oranges and mean square error. *International Journal of Forecasting*, *4*(4), 515–518.

Chen, V., Gunther, D., & Johnson, E. (2003). Solving for an optimal airline yield management policy via statistical learning. *Journal of the Royal Statistical Society Series C (Applied Statistics)*, *52*(1), 19–30.

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

Cho, M., Fan, M., & Zhou, Y. (2007). *A study of an online fare-track and order system.* (Working Paper, University of Washington Business School, Seattle, WA)

Cléaz-Savoyen, R. (2005). *Airline revenue management methods for less restricted fare structures.* Massachusetts Institute of Technology.

Coldren, G., & Koppelman, F. (2005). Modeling the Competition among Air Travel Itinerary Shares: GEV Model Development. *Transportation Research Part A*, *39*, 345–365.

Cooper, W., & Mello, T. Homem-de. (2002). *Revenue Management Using Sampling-Based Optimization and Markov Decision Processes* (Rapport technique). Working paper, University of Minnesota, Minneapolis, MN & Ohio State University, Columbus, OH.

Cooper, W., & Mello, T. Homem-de. (2003). *A class of hybrid methods for revenue management.* (Working paper 03-015, Department of IE/MS, Northwestern University, Illinois)

Cooper, W., Mello, T. Homem-de, & Kleywegt, A. (2006). Models of the Spiral-Down Effect in Revenue Management. *Operations Research*, *54*(5), 968–987.

Cross, R. G. (1997). *Revenue management: hard-core tactics for market domination.* Broadway Books, NY.

Cunningham, S., & De Haan, A. (2006). Long-term forecasting for sustainable development: Air travel demand for 2050. *International Journal of Environment and Sustainable Development*, *5*(3), 297–314.

Curry, R. E. (1990). Optimum Airline Seat Allocation With Fare Classes Nested by Origin and Destination. *Transportation Science*, *24*(3), 193–204.

Cusano, A. (2003). *Airline revenue management under alternative fare structures.* Massachusetts Institute of Technology.

Cutshall, C., & Weisbrodt, J. (2006). Implementing an OD revenue management solution. *Journal of Revenue and Pricing Management*, *5*(2), 128–134.

de Boer, S. (2003). *Advances in Airline Revenue Management and Pricing.* Thèse de doctorat non publiée, Massachusetts Institute of Technology, Sloan School of Management, Operations Research Center.

Diebold, F., & Lopez, J. (1996). *Forecast Evaluation and Combination.* National Bureau of Economic Research Cambridge, Mass., USA.

Diebold, F., & Mariano, R. (1995). Comparing Predictive Accuracy. *Journal of Business and Economic Statistics*, *13*(3), 253–263.

Dror, M., Trudeau, P., & Ladany, S. (1988). Network models for seat allocation on flights. *TRANSP. RES.*, *22*(4), 239–250.

D'Sylva, E. (1982). O and D Seat Assignment to Maximize Expected Revenue. *Unpublished internal report, Boeing Commercial Airplane Company, Seattle, WA.*

Dunleavy, H., & Westermann, D. (2005). Future of revenue management: Future of airline revenue management. *Journal of Revenue and Pricing Management*, *3*(4), 380–383.

Ferguson, M., Crystal, C., Higbie, J., & Kapoor, R. (2007). *A comparison of unconstraining methods to improve revenue management systems (ed. 3).* (Working Paper, Georgia Institute of Technology)

Fiig, T., & Isler, K. (2004). SAS O&D Low Cost Project. In *PODS Consortium Meeting.* (Minneapolis, MN)

Fildes, R. (1992). The evaluation of extrapolative forecasting methods. *International Journal of Forecasting*, *8*(1), 81–98.

Fischer, T., & Kamerschen, D. (2003). Measuring competition in the us airline industry using the rosse-panzar test and cross-sectional regression analyses. *Journal of Applied Economics*, *6*(1), 73–93.

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

Frank, M., Friedemann, M., & Schröder, A. (2008). Principles for simulations in revenue management. *Journal of Revenue & Pricing Management*, *7*(1), 7–16.

Fudenberg, D., & Villas-Boas, J. (2006). Behavior-based price discrimination and customer recognition. In E. T.J. Hendershott (Ed.), *Handbook on economics and information systems* (pp. 377–436). Elsevier.

Gallego, G., & Hu, M. (2007). Dynamic pricing of perishable assets under competition. In *Agifors reservations and yield management study group annual meeting proceedings.* (Jeju Island, South Korea)

Gallego, G., Kou, S., & Phillips, R. (2008). Revenue Management of Callable Products. *Management Science*, *54*(3), 550.

Gallego, G., Krishnamoorthy, S., & Phillips, R. (2006). Dynamic revenue management games with forward and spot markets. *Journal of Revenue and Pricing Management*, *5*(1), 10–31.

Gallego, G., & Phillips, R. (2004). Revenue Management of Flexible Products. *Manufacturing & Service Operations Management*, *6*(4), 321.

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

Garcia-Diaz, A., & Kuyumcu, A. (1997). Cutting-plane procedure for maximizing revenues in yield management. *Computers & Industrial Engineering*, *33*(1), 51–54.

Garrow, L., Jones, S., & Parker, R. (2007). How much airline customers are willing to pay: An analysis of price sensitivity in online distribution channels. *Journal of Revenue and Pricing Management*, *5*(4), 271–290.

Gerchak, Y., & Parlar, M. (1987). A single period inventory problem with partially controllable demand. *Computers and Operations Research*, *14*(1), 1–9.

Gerchak, Y., Parlar, M., & Yee, T. (1985). Optimal rationing policies and production quantities for products with several demand classes. *Canadian Journal of Administrative Science*, *2*(1), 161–176.

Glover, F., Glover, R., Lorenzo, J., & McMillan, C. (1982). The passenger mix problem in the scheduled airlines. *Interfaces*, *12*(3), 73–80.

Gorin, T. (2000). *Airline revenue management: Sell-up and forecasting algorithms.* Massachusetts Institute of Technology.

Gorin, T. (2004). *Assessing low-fare entry in airline markets: impacts of revenue management and network flows.* Thèse de doctorat non publiée, Massachusetts Institute of Technology.

Granger, C., & Pesaran, M. (2000). Economic and statistical measures of forecast accuracy. *Journal of Forecasting*, *19*(7), 537–560.

Grosche, T., Rothlauf, F., & Heinzl, A. (2007). Gravity models for airline passenger volume estimation. *Journal of Air Transport Management*, *13*(4), 175–183.

Guo, C. (2007). Review: Methods for Estimating Sell-Up. In *PODS Consortium Meeting.* (Mainz, Germany)

Harvey, D. I. (1997). *The Evaluation of Economic Forecasts.* Thèse de doctorat non publiée, University of Nottingham.

Hersh, M., & Ladany, S. (1978). Optimal seat allocation for flights with one intermediate stop. *Computers and Operations Research*, *5*, 31-37.

Hopperstad, B. P., C.H. (2004). Q Investigations – Algorithms for Unrestricted Fare Classes. In *PODS Consortium Meeting.* (Amsterdam, Netherlands)

Hopperstad, C. (1994). The Application of Path Preference and Stochastic Demand Modelling to Market Based Forecasting. *AGIFORS Reservations and Yield Management Study Group Proc.*

Hopperstad, C. (2000). *Passenger Origin Destination Simulator Technical Specifications (Revision 1).* (November, Unpublished Working Paper, Seattle, WA)

Hopperstad, C. (2007). Market-based Forecasting Investigation I. In *PODS Consortium Meeting.* (Mainz, Germany)

Hyndman, R., & Koehler, A. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, *22*(4), 679–688.

Jung, J., & Fuji, E. (1976). The price elasticity of demand for air travel. *Journal of Transport Economics and Policy*, *10*, 257–262.

Kanafani, A., & Sadoulet, E. (1977). The partitioning of long haul air traffic- A study in multinomial choice. *Transportation Research*, *11*, 1–8.

Kayser, M. (2007). Use of Adjusted Fares to Come and Fare Adjustment Scaling. In *PODS Consortium Meeting.* (Mainz, Germany)

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

Kincaid, W., & Darling, D. (1963). An inventory pricing problem. *Journal of Mathematical Analysis and Applications*, *7*(2), 183–208.

Kraft, D., Oum, T., & Tretheway, M. (1986). Airline Seat Management. *Logistics and Transportation Review*, *22*(2), 115–130.

Ladany, S., & Bedi, D. (1977). Dynamic Rules for Flights with an Intermediate Stop. *Omega*, *5*, 721–730.

Lan, Y., & Gao, H. (2007). *Seat inventory control with limited demand information.* (Submission for AGIFORS Anna Valicek Medal)

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

Lautenbacher, C. J., & Stidham, S. (1999). The Underlying Markov Decision Process in the Single-Leg Airline Yield-Management Problem. *Transportation Science*, *33*(2), 136–146.

Law, A., & Kelton, W. (1997). *Simulation Modeling and Analysis.* McGraw-Hill Higher Education.

Lee, A. O. (1990). Airline reservations forecasting–probabilistic and statistical models of the booking process.

Lee, T. C., & Hersh, M. (1993). A Model for Dynamic Airline Seat Inventory Control with Multiple Seat Bookings. *Transportation Science*, *27*(3), 252–265.

Lennon, C. (1972). Market Reactions to Fare Constraints and Its Effects on Passenger Revenue. *AGIFORS Symposium Proc*, *12*.

Levin, Y., McGill, J., & Nediak, M. (2006). *Optimal dynamic pricing of perishable items by a monopolist facing strategic consumers.* (Submitted to Operations Research)

Li, M. (1997). On the Multi-fare Seat Allocation Problem. *The Conference Proceedings of the 1997 Air Transport Research Group,(ATRG), of the WCTR Society*.

Li, M. Z. F., & Oum, T. H. (1998). *Seat Allocation Game on Flights with Two Fares* (Rapport technique). Working Paper, Division of Applied Economics, Nanyang Business School, Nanyang Technological University, Singapore.

Li, M. Z. F., & Oum, T. H. (2002). A Note on the Single Leg, Multifare Seat Allocation Problem. *Transportation Science*, *36*(3), 349–353.

Littlewood, K. (1972). Forecasting and control of passenger bookings. *AGIFORS Symposium Proceedings*, *12*, 95–117.

Lua, F. (2007a). Lowest Competitor Class Open (loco) Closure Matching. In *PODS Consortium Meeting.* (Minneapolis, MN)

Lua, F. (2007b). Lowest Competitor Class Open (loco) Seat Availability Matching (Network S)). In *PODS Consortium Meeting.* (Mainz, Germany)

Lua, F. (2007c). Lowest Competitor Class Open (loco) Seat Availability Matching (Single Symmetric Market). In *PODS Consortium Meeting.* (Mainz, Germany)

Makridakis, S. (1986). The Art and Science of Forecasting: An Assessment and Future Directions. *International Journal of Forecasting*, *2*(1), 15–39.

Makridakis, S. (1993). Accuracy measures: theoretical and practical concerns. *International Journal of Forecasting*, *9*(4), 527–529.

Mayer, M. (1976). Seat Allocation, or a Simple Model of Seat Allocation via Sophisticated Ones. *AGIFORS Annual Symposium Proceedings*, *16*.

McGill, J. (1989). *Optimization and estimation problems in airline yield management.* Thèse de doctorat non publiée, University of British Columbia.

McGill, J. (1995). Censored regression analysis of multiclass passenger demand data subject to joint capacity constraints. *Annals of Operations Research*, *60*(1), 209–240.

McGill, J., & Ryzin, G. van. (1999). Revenue management: Research overview and prospects. *Transportation Science*, *33*(2), 233–256.

Möller, A., Romisch, W., & Weber, K. (2004). A new approach to O&D revenue management based on scenario trees. *Management*, *3*(3), 265–276.

Nahmias, S. (1994). Demand estimation in lost sales inventory systems. *Naval research logistics*, *41*(6), 739–757.

Nason, S. (2007). Forecasting the future of airline revenue management. *Journal of Revenue and Pricing Management*, *6*(1), 64–66.

Netessine, S., & Shumsky, R. (2005). Revenue management games: Horizontal and vertical competition. *Management Science*, *51*(5), 813–831.

Neuling, R., Riedel, S., & Kalka, K. (2004). New approaches to origin and destination and no-show forecasting: Excavating the passenger name records treasure. *Journal of Revenue and Pricing Management*, *3*(1), 62–73.

Njegovan, N. (2006). Elasticities of demand for leisure air travel: A system modelling approach. *Journal of Air Transport Management*, *12*, 13–30.

Ozdaryal, O., & Saranathan, K. (2004). Managing revenue in broken fare fence environment. In *AGIFORS Reservations and Yield Management Study Group Annual Meeting Proceedings.* (Auckland, New Zealand)

Pak, K., & Piersma, N. (2002). Overview of OR techniques for airline revenue management. *Statistica Neerlandica*, *56*(4), 480–496.

Pfeifer, P. (1989). The airline discount fare allocation problem. *Decision Sciences*, *20*(1), 149–157.

Phillips, R., Boyd, D., & Grossman, T. (1991). An algorithm for calculating consistent itinerary flows. *Transportation science*, *25*(3), 225–239.

Pölt, S. (1998). Forecasting is difficult – especially if it refers to the future. In *AGIFORS Reservations and Yield Management Study Group Annual Meeting Proceedings.* (Melbourne, Australia)

Pölt, S. (2000). From bookings to demand: The process of unconstraining. In *AGIFORS Reservations and Yield Management Study Group Annual Meeting Proceedings.* (New York City, NY)

Pratte, N. (1986). Automated Full Fare Protection System. *AGIFORS Symposium Proc*, *26*.

Rannou, B., & Melli, D. (2003). Measuring the impact of revenue management. *Journal of Revenue & Pricing Management*, *2*(3), 261–270.

Ratliff, R., & Vinod, B. (2005). Airline pricing and revenue management: A future outlook. *Journal of Revenue and Pricing Management*, *4*(3).

Ratliffe, R. (2008). Multi-flight demand untruncations: Simulation results. In *Agifors reservations and yield management study group annual meeting proceedings.* (Tahiti, French Polynesia)

Reyes, M. (2006). *Hybrid forecasting for airline revenue management in semi-restricted fare structures.* Massachusetts Institute of Technology.

Richter, H. (1982). The Differential Revenue Method to Determine Optimal Seat Allotments by Fare Type. *AGIFORS Symposium Proceedings*, *22*, 339–362.

Robinson, L. (1995). Optimal and Approximate Control Policies for Airline Booking with Sequential Nonmonotonic Fare Classes. *Operations Research*, *43*(2), 252–263.

Rockmann, J., & Alder, C. (2009). Revenue increase by o & d optimization at lufthansa. In *Agifors reservations and yield management study group annual meeting proceedings.* (Amsterdam, Netherlands)

Rothstein, M. (1968). *Stochastic models for airline booking policies.* Thèse de doctorat non publiée, New York: New York University.

Rothstein, M., & Stone, A. (1967). Passenger Booking Levels. *Proc. Seventh AGIFORS Sympos.*

Ryzin, G. van, & McGill, J. (2000). Revenue Management Without Forecasting or Optimization: An Adaptive Algorithm for Determining Airline Seat Protection Levels. *Management Science*, *46*(6), 760–775.

Ryzin, G. van, & Vulcano, G. (2006). Simulation-based optimization of virtual nesting controls for network revenue management. *Operations Research*.

Sa, J. (1987). Reservations forecasting in airline yield management.

Simpson, R. (1985). Setting Optimal Booking Levels for Flight Segments with Multi-Class, Multi-Market Traffic. *AGIFORS Symposium Proceedings*, *25*.

Smith, B., Leimkuhler, J., & Darrow, R. (1992). Yield Management at American Airlines. *Interfaces*, *22*(1), 8–31.

Smith, B., & Penn, C. (1988). Analysis of alternative origin-destination control strategies. *AGIFORS Annual Symposium Proceedings*.

Stefanescu, C., DeMiguel, V., Fridgeirsdottir, K., , & Zenios, S. (2004). Revenue management with correlated demand forecasting. In *Proceedings of the American Statistical Association, Business and Economics Statistics Section.* (Alexandria, VA)

Stone, R., & Diamond, M. (1992). *Optimal Inventory Control for a Single Flight Leg* (Rapport technique). Working paper, Northwest Airlines, Operations Research Division, Minneapolis, MN.

Su, X. (2007). Inter-temporal pricing with strategic customer behavior. *Management Science*, *53*(5), 726–741.

Subramanian, J., Stidham, S., & Lautenbacher, C. J. (1999). Airline Yield Management with Overbooking, Cancellations, and No-Shows. *Transportation Science*, *33*(2), 147–167.

Sullivan, R., Timmermann, A., & White, H. (2003). Forecast evaluation with shared data sets. *International Journal of Forecasting*, *19*(2), 217–227.

Swan, W. (1993a). *Forecasting for Revenue Management Systems* (Rapport technique). Working Paper, Boeing Commercial Aircraft, Seattle, WA.

Swan, W. (1993b). Modeling Variance for Yield Management. *prepared for AFIFORS Yield Management Group*, 2–5.

Swan, W. (1993c). *Revenue Management Nesting Structures* (Rapport technique). Working Paper, Boeing Commercial Aircraft Co., Seattle, WA.

Swan, W. (1999). Spill Modeling for Airlines. *World*.

Swan, W. (2002). Airline demand distributions: passenger revenue management and spill. *Transportation Research Part E: Logistics and Transportation Review*, *38*(3), 253–263.

Swan, W. M. (1990). *Revenue Management Forecasting Biases* (Rapport technique). Working Paper, Boeing Commercial Aircraft, Seattle, WA.

Swift, A. (2002). Renaissance: A revenue management change program at British Airways. *Journal of Revenue and Pricing Management*, *1*(2), 166–173.

Talluri, K. T. (1994). *Separable Convex Cost Flows and Revenue Management on a Line Network* (Rapport technique). working paper, USAir, Operations Research Department.

Talluri, K. T. (2001). Airline revenue management with passenger routing control: a new model with solution approaches. *International Journal of Services Technology and Management*, *2*(1), 102–115.

Talluri, K. T., & van Ryzin, G. (2000). *A discrete choice model of yield management.* Preprint.

Talluri, K. T., & Van Ryzin, G. J. (2004a). Revenue management under a general discrete choice model of consumer behavior. *Management Science*, *50*(1), 15–33.

Talluri, K. T., & Van Ryzin, G. J. (2004b). *Theory and practice of revenue management.* Kluwer Academic Publishers, Boston.

Taneja, N. (1978). Airline traffic forecasting: A regression analysis approach. *Lexington, Mass., D. C. Heath and Co., 1978. 245 p.*

Taylor, C. (1962). The Determination of Passenger Booking Levels. *AGIFORS Symposium Proc*, *2*.

Tayman, J., & Swanson, D. (1999). On the validity of MAPE as a measure of population forecast accuracy. *Population Research and Policy Review*, *18*(4), 299–322.

Theil, H. (1966). *Applied Economic Forecasting.* Amsterdam.

Theis, G., Adler, T., Clarke, J., & Ben-Akiva, M. (2006). Risk Aversion to Short Connections in Airline Itinerary Choice. *Transportation Research Record*, *1951*, 28–36.

Titze, B., & Griesshaber, R. (1983). Realistic Passenger Booking Behavior and the Simple Low-Fare/High-Fare Seat Allotment Model. *AGIFORS Symposium Proceedings*, *23*, 197–223.

van Ryzin, G. (2005). Future of revenue management: Models of demand. *Journal of Revenue and Pricing Management*, *4*(2), 204–209.

Vinod, B. (1989). *A Set Partitioning Algorithm for Virtual Nesting Indexing Using Dynamic Programming* (Rapport technique). Internal Technical Report, SABRE Decision Technologies.

Vinod, B. (1990). Reservation Inventory Control Techniques to Maximize Revenues. *IATA Third International Airline Yield Management Conference Proc.*

Vinod, B. (1991). *Stochastic network optimization – a nonlinear programming approach* (Rapport technique). Internal Technical Report, SABRE Decision Technologies, Fort Worth, TX.

Vinod, B. (1995). Origin-and-destination yield management. *The Handbook of Airline Economics*, 459–468.

Vinod, B. (2006). Advances in inventory control. *Journal of Revenue & Pricing Management*, *4*, 367–381.

Vinod, B., & Ratliff, R. (1990). A Discount Allocation Optimization Method Using Stochastic Linear Programming. *AGIFORS Reservations and Yield Management Study Group, Rome, Italy*.

Vulcano, G. (2006). Choice based revenue management: An empirical study of estimation and optimization. In *21st European Conference on Operational Research.* (Reykjavik, Iceland)

Walczak, D. (2005). Incorporating market information into pricing optimization. In *INFORMS.* (San Francisco, CA)

Walczak, D., & Brumelle, S. (2007). Semi-markov information model for revenue management and dynamic pricing. *OR Spectrum*, *29*(1), 61-84.

Wang, K. (1982). Modelling the Interaction between Payload Restriction, Passenger Demand and Reservation Booking Levels. *AGIFORS Symposium Proceedings*, *22*.

Wang, K. (1983). Optimum Seat Allocation for Multi-Leg Flights with Multiple Fare Types. *AGIFORS Symposium Proceedings*, *23*.

Weatherford, L. (1997a). Optimization of perishable-asset revenue management problems that allow prices as decision variables. *International Journal of Technology Management*, *12*.

Weatherford, L. (1997b). Using Prices More Realistically as Decision Variables in Perishable-Asset Revenue Management Problems. *Journal of Combinatorial Optimization*, *1*(3), 277–304.

Weatherford, L. (2000). Unconstraining methods. In *AGIFORS Reservations and Yield Management Study Group Annual Meeting Proceedings.* (New York City, NY)

Weatherford, L. (2001). Optimisation of perishable asset revenue management problems that allow prices as decision variables. *International Journal of Services Technology and Management*, *2*(1), 71–101.

Weatherford, L., & Belobaba, P. (2002). Revenue impacts of fare input and demand forecast accuracy in airline yield management. *Journal of the Operational Research Society*, *53*(8), 811–821.

Weatherford, L., & Bodily, S. (1992). A Taxonomy and Research Overview of Perishable-Asset Revenue Management: Yield Management, Overbooking, and Pricing. *Operations Research*, *40*(5), 831–844.

Weatherford, L., Bodily, S., & Pfeiffer, P. (1993). Modeling the customer arrival process and comparing decision rules in perishable asset revenue management situations: The airline industry. *Transportation science*, *27*(3), 239–251.

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.

Weber, K., & Thiel, R. (2004). Methodological issues in low cost revenue management. In *AGIFORS Reservations and Yield Management Study Group Annual Meeting Proceedings.* (Auckland, New Zealand)

Westermann, D. (2006). (Realtime) dynamic pricing in an integrated revenue management and pricing environment: An approach to handling undifferentiated fare structures in low-fare markets. *Journal of Revenue and Pricing Management*, *4*(4).

Williamson, E. (1988). *Comparison of optimization techniques for origin-destination seat inventory control.* Mémoire de Master non publié.

Williamson, E. (1992). *Airline Network Seat Inventory Control: Methodologies and Revenue Impacts.* Thèse de doctorat non publiée, Massachusetts Institute of Technology.

Wilson, J., Anderson, C., & Kim, S. (2006). Optimal booking limits in the presence of strategic consumer behavior. *International Transactions in Operational Research*, *13*(2), 99–110.

Wollmer, R. (1986a). An airline reservation model for opening and closing fare classes. *unpublished company report, Douglas Aircraft Company, McDonnell Douglas Corporation, Long Beach, CA.*

Wollmer, R. (1986b). A Hub-Spoke Seat Management Model. *Unpublished company report, Douglas Aircraft Company, McDonnell Douglas Corporation, Long Beach, CA.*

Wollmer, R. (1992). An Airline Seat Management Model for a Single Leg Route When Lower Fare Classes Book First. *Operations Research*, *40*(1), 26–37.

Wong, J. (1990). *Airline network seat allocation.* Thèse de doctorat non publiée, Northwestern University Evanston, Ill.

Wong, J., Koppelmann, F., & Daskin, M. (1993). Flexible assignment approach to itinerary seat allocation. *Transportation research. Part B: methodological*, *27*(1), 33–48.

Xu, X., & Hopp, W. (2005a). *Dynamic pricing and inventory control: The value of demand learning.* (Working Paper, Department of Industrial Engineering, Northwestern University, Illinois)

Xu, X., & Hopp, W. (2005b). *Implications of strategic customer behavior for revenue management.* (Working Paper, Department of Industrial Engineering, Northwestern University)

Young, Y., & van Slyke, R. (1994). *Stochastic Knapsack Models of Yield Management* (Rapport technique). Technical Report 94-76, Polytechnic University, New York, NY.

Young, Y., & van Slyke, R. (2000). Finite Horizon Stochastic Knapsacks with Applications to Yield Management. *Operations Research*, *48*(1), 155–172.

Zaki, H. (2000). Forecasting for airline revenue management. *Journal of Business Forecasting Methods and Systems*, *19*(1), 2–7.

Zeni, R. (2001a). *Improved forecast accuracy in revenue management by unconstraining demand estimates from censored data.* State University of New Jersey.

Zeni, R. (2001b). Improving forecast accuracy by unconstraining censored demand data. In *AGIFORS Reservations and Yield Management Study Group Annual Meeting Proceedings.* (Bangkok, Thailand)

Zhao, W., & Zheng, Y. (1998). *General Dynamic Models for Airline Seat Allocation* (Rapport technique). Working paper, The Wharton School, University of Pennsylvania, Philadelphia, PA.

Zhao, W., & Zheng, Y. (2000). Optimal Dynamic Pricing for Perishable Assets with Nonhomogeneous Demand. *Management Science*, *46*(3), 375–388.

Zhao, W., & Zheng, Y.-S. (2001). A dynamic model for airline seat allocation with passenger diversion and no-shows. *Transportation Science*, *35*(1), 80–98.

Zhu, J. (2006). Using turndowns to estimate the latent demand in a car rental unconstrained demand forecast. *Journal of Revenue and Pricing Management*, *4*(4), 344-353.

Notation used in Formulas

α^{exp}	Weight of new bookings in the calculation of forecasts based on exponential smoothing.
α^{P}	Weight of the estimator based on price in the joint estimator.
α^{T}	Weight of the time-based estimator in the joint estimator.
A^{emsrb}	Function that computes protected seats based on the EMSR-b algorithm.
$a(f, c, t, s)$	Seats available for class c on flight f, at point of time before departure t, in simulation run s.
$\hat{a}(f, c, t, s)$	Seats protected for class c on flight f, at point of time before departure t, in simulation run s.
$\beta^{\text{dep}}(m)$	Weight of the deviation from the preferred departure time in the cost function of the customer type m.
$\beta^{\text{dur}}(m)$	Weight of the difference between actual and minimum travel time in the cost function of the customer type m.
$\beta^{\text{car}}(m)$	Weight factor attached to any itinerary that is not provided by the preferred carrier of customer type m.
$\beta^{\text{price}}(m)$	Weight of the fare attached to the considered itinerary in the cost function of the customer type m.
$\beta^{\text{trans}}(m)$	Weight of the number of transfers included in the chosen itinerary in the cost function of the customer type m.
$b(f, c, t, s)$	Bookings observed for flight f, class c, between points of time before departure $t-1$ and t, in simulation run s.
$\hat{b}(f, c, t, s)$	Average of historical bookings during the runs 1 to s that occurred on flight f between points of time t and $t-1$ while booking class c was available.
$b^{\text{unc}}(f, c, t, s)$	Unconstrained bookings observed for flight f, class c, between points of time before departure $t-1$ and t, in simulation run s.

$\gamma(p) \in [0,1]$	Share of overall requests to be assigned to pairing p.
C	Set of booking classes offered, ordered by descending price.
$C(i,r)$	Cost function of request r considering itinerary i when a lowest available fare has been found.
$\hat{C}(i,r)$	Cost of itinerary i considered by request r, without regard for the actual ticket price (given the assumption that all itineraries cost the same).
C_r	Set of classes acceptable according to the criteria of request r, ordered by descending price.
$\delta^{\text{dep}}(m)$	Factor for maximum acceptable deviation from $w^d(r)$, defined by the customer type m.
$\delta^{\text{dur}}(m)$	Factor for maximum acceptable travel time, defined by the customer type m.
$\delta^{\text{price}}(m)$	Factor defining maximum willingness to pay for customer type m.
$\delta^{\text{price}}(r)$	Factor defining maximum willingness to pay for request r.
ϵ_r	Error term defining the distortion of request r.
ϵ_s	Error term defining the distortion of demand in run s.
$e_\circ^{c-c}(f,c,t,s)$; $e_\circ^{c-u}(f,c,t,s)$; $e_\circ^{u-u}(f,c,t,s)$	Error comparing actual or unconstrained bookings to constrained or unconstrained forecast based on method $\circ$ per flight f, class c and run s.
$\hat{e}_\circ^{c-c}(s)$; $e_\circ^{c-u}(s)$; $e_\circ^{u-u}(s)$	Series error comparing actual or unconstrained bookings to constrained or unconstrained forecast based on method $\circ$ per run s.
F	Set of flights included in the schedule.
F_i	Set of flights included in itinerary i.
$f^{\text{const}}(f,c,t,s)$	Constrained demand predicted to arrive for flight f, class c, between points of time before departure $t-1$ and t, in simulation run s.

$f^{\text{unc}}(f, c, t, s)$	Unconstrained demand predicted to arrive for flight f, class c, between points of time before departure $t-1$ and t, in simulation run s.
$f^{\text{unc}\,\%}(f, c, t, s)$	Forecasted demand per class as percentage of forecasted overall demand for the flight.
$\eta(p, m) \in [0, 1]$	Share of requests scheduled to arrive for pairing p based on customer type m.
I_q	Set of itineraries connecting the origin and destination included in pairing q.
I_r	Set of itineraries acceptable according to the criteria of request r
$i \in I$	Itineraries derived from the schedule.
$K(f, t, s)$	Available capacity of the flight f the point of time t before departure of simulation run s.
$\lambda_{p,m,s}$	Overall intensity of the Poisson process for pairing p and customer type m throughout the booking horizon of simulation run s
$\lambda_{p,m,s}^{t,t+\tau}$	Intensity of the Poisson process defining the arrival pattern of customer type m for pairing p in the slice of the booking horizon of simulation run s defined by the interval $[t, t+\tau]$
$l^{\text{fcfs}}(s)$	Average seat load factor generated in simulation run s given first-come-first-serve inventory controls.
$l_{\text{2-i c-max}}^{\%\ \text{fcfs}}(s)$	Average seat load factor generated in simulation run s as percentage of $l^{\text{fcfs}}(s)$.
$l(s)$	Average seat load factor generated in simulation run s.
M	Set of customer types included in the demand model.
$m(r) \rightarrow m$	Function resulting in the customer type from which a request r was created.
$\nu^{\text{dist}}(p)$	Minimum distance between origin and destination airports of the pairing p.
$\nu^{\text{dur}}(p)$	Minimum travel time required by pairing p

$\overline{o}\,(f, c, t, s)$	Matrix of boolean values indicating the lowest available class c for flight f between points of time $t-1$ and t of simulation run s.
$P_{p,m,s}[N\,(p, m, t+\tau)$ $-N\,(p, m, t) = k]$	Poisson probability of k requests based on customer type m to arrive for pairing p in time slice t to $t+\tau$ of simulation run s.
$p(f, c)$	Price of a ticket for flight f in booking class c.
$\overline{p}\,(f, c, t, s)$	Expected marginal seat revenue for a seat in class c, flight f at point of time t of simulation run s.
$q \in \{1, ..., N^p\}$	Pairings derived from the schedule.
$q\,(i) \rightarrow p$	Function resulting in the pairing for which an itinerary i was created.
R	Set of customer requests.
R_s	Customer requests created for simulation run s.
$R_s^{p,m,t}$	Requests based on customer type m that arrives for pairing p up to point of time t in the booking horizon of simulation run s.
$\overline{R}$	Input parameter defining the average number of requests to be scheduled per run.
$r\,(s)$	Overall revenue generated during simulation run s.
σ_ϵ	Deviation of the normal distribution that the error terms are drawn from.
$\sigma\,(f^{\text{unc}}\,(f, c, t, s))$	Standard deviation of the forecast of demand for flight f in class c at point of time t in run s.
S^{hybrid}	Hybrid market scenario.
S^{product}	Product-sensitive market scenario.
S^{price}	Price-sensitive market scenario.
$s = 1, ..., N^s$	Simulation runs included in the simulation experiment in chronological order with run $s-1$ to occur before run s.
$t = 0, ..., N^t$	Points of time before departure, demand arrives after $t = 0$, $t = N^t$ is the time of departure.

$u^P(f, c, t, s)$	Price-based estimator for flight f, class c, and point of time t of run s
$u^T(f, c, t, s)$	Time-based estimator for flight f, class c, and the point of time t of run s.
$u^J(f, c, t, s)$	Joint estimator for flight f, class c, point of time t and run s.
V	Indicator of the average demand volume per simulation run.
$w^{\text{dep}}(r)$	Preferred departure time of request r.
$x^{\text{dep}}(i)$	Departure time of itinerary i.
$x^{\text{dur}}(i)$	Travel time attached to itinerary i.
$x^{\text{trans}}(i)$	Number of transfers in itinerary i.
$y(s)$	Average yield generated during simulation run s.
$\Psi(M, \sigma_\epsilon, V, F)$	Simulation experiment based on a market structure M, a deviation of the error term of demand σ_ϵ, overall demand volume V, and a forecast method F.
$\omega^P_{f,c,s}$	Vector of price elasticity for flight f depending on the class c for run s.
$\omega^T_{f,t,s}$	Vector of time elasticity for flight f depending on the point of time t before departure of run s.
Z	Set of possible restrictions of booking classes - the absence of a feature, such as comfort seating, is modeled as a restriction.
$Z^{\text{class}}(z, c) \rightarrow \{0, 1\}$	Boolean function defining for every booking class c whether or not it includes restriction z.
$Z^{\text{request}}(z, r) \rightarrow \{0, 1\}$	Boolean function defining for every request r whether or not it accepts restriction z.

Glossary

arrival distribution	Timing of requests over the booking horizon of a flight.
authorization level	Inventory control: A class is available if the bookings already accepted do not exceed the authorization level.
bid price	Inventory control: The fare of each class is compared to a bid price, the class is available for sale if its fare exceeds the bid price.
booking class	Set of restrictions and features defining the conditions under which a ticket is sold by a carrier; defined by a caption.
booking horizon	Period of time before the departure of a flight during which tickets can be bought; also: reservation phase.
buy-down	Phenomenon of customers buying a class that is cheaper than the most expensive class acceptable according to their willingness to pay.
cancellation	Customers returning tickets before the departure day.
carrier	Airline offering flights and tickets.
connecting time	Time available between two connecting flights for customers to transfer.
connection builder	Function applied to generate pairings and itineraries given flights and set-up parameters.
cost function	Function weighting factors according to a customers preferences, used to choose between itineraries.

cost function factor	Weight of characteristics of itineraries and prices in the customers' cost function.
customer mix	Distribution over customer types for a given pairing.
customer segmentation	Concept used to enable revenue management; customers' are segmented according to their product requirements and price acceptance.
customer type	Template of customer characteristics including arrival time, product requirements, willingness to pay, and cost function.
denied boarding	Customers not being able to get a seat on a booked flight due to flaws in overbooking.
destination	Airport at which a customer desires to end the journey.
EMSR-b	Heuristic maximizing revenue in a flight-based revenue management system given static forecasts per flight and class.
error term	Stochastic distortion of request information, normally distributed.
Exp025	Exponential smoothing method applying the weight 0.25 to new data.
Exp050	Exponential smoothing method applying the weight 0.50 to new data.
Exp075	Exponential smoothing method applying the weight 0.75 to new data.
exponential smoothing	Forecast method extrapolating expected demand from historical bookings based on a static view of demand.
fare	Price of a ticket in a booking class on a specific flight.

feature	Positive condition attached to a ticket; examples can be flexible refund or seat in the business compartment.
flight	Direct connection between two airports; defined by a flight number, a carrier, a departure day and a departure time.
flight view	View of revenue management that strives to maximize revenue per flight.
itinerary	Way of traveling from one origin to one destination using one or more connecting flights.
leg	Combination of two airports for which one or more direct flights are offered.
MAD	Error measurement: Mean absolute deviation.
MAPE	Error measurement: Mean average percentage error.
naive forecast	Forecast method extrapolating expected demand from the bookings observed in the previous run based on a static view of demand.
network view	View of revenue management that strives to maximize revenue over a complete network.
no-frills airlines	Airlines offering restriction-free classes at usually low prices.
no-show	Customers not using their ticket on the departure day.
origin	Airport from which a customer desires to start the journey.

overbooking	Technique to compensate for no-shows and cancellations by selling more seats than available based on capacity.
pairing	Combination of origin and destination for which itineraries are offered.
PODS	Passenger Origin and Destination Simulator – MIT simulation used to evaluate revenue management strategies.
Poisson Process	Stochastic process used to describe the arrival distribution of customer types.
price acceptance	Customers' maximum willingness to pay.
product acceptance	Customers' requirements regarding restrictions and features of a booking class.
protected seats	Result of optimization defining the seats to be sold exclusively in a specific class; used to calculate authorization levels.
random walk	See: naive forecast.
request	Instance of a customer desiring to buy a ticket for a combination of origin and destination given certain requirements regarding product and price.
restriction	Negative condition attached to a ticket; examples can be minimum stay or a lack of features such as flexible refund.
RMSE	Error measurement: Root mean squared error.
sell-up	Phenomenon of customers buying a class that is more expensive than the cheapest class acceptable according to their product requirements.
simulation	...

simulation experiment	Simulation given specific data, methods, and set-up parameters.
simulation scenario	Data input for a simulation experiment.
traffic area	Geographical area to categorize flights, legs, itineraries, and pairings; example: continental vs. intercontinental.
U2	Error measurement: also called Theil's U2.
uncertainty	Characteristic describing the volatility and thereby the predictability of a market.