Airline Revenue Management with Dynamic Offers: Bundling Flights and Ancillary Services

by

Kevin K. Wang

M.Eng., Aerospace Engineering, University of Cambridge (2018)
B.A., Mechanical Engineering, University of Cambridge (2018)

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Master of Science in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2020

Author . . . .
Department of Aeronautics and Astronautics
May 19, 2020

Certified by . . . .
Peter P. Belobaba
Principal Research Scientist, Aeronautics and Astronautics
Thesis Supervisor

Accepted by . . . .
Sertac Karaman
Associate Professor of Aeronautics and Astronautics
Chair, Graduate Program Committee

Airline Revenue Management with Dynamic Offers: Bundling Flights and Ancillary Services

by

Kevin K. Wang

Submitted to the Department of Aeronautics and Astronautics
on May 19, 2020, in partial fulfillment of the
requirements for the degree of
Master of Science in Aeronautics and Astronautics

Abstract

Airline revenue management will have to adapt to a new world of airline retailing enabled by the New Distribution Capability. One area of interest is Dynamic Offer Generation (DOG), in which airlines respond to every booking request in real-time with a customized set of offers and prices. In these offers, ancillary services may be bundled with the flight. Selecting and pricing these offer sets represents a new joint pricing and assortment optimization problem in revenue management.

We propose a formulation for the dynamic offer generation problem and study the robustness of its solution. We derive conditions under which selling the flight in a bundle with an ancillary service increases total net revenues over selling the ancillary as an optional add-on. We show how this model integrates with traditional revenue management systems. We simulate DOG under competition in the Passenger Origin-Destination Simulator (PODS) to show the potential revenue benefits.

The simulation results show that bundling the flight with an ancillary service can generate higher revenues than selling both services separately. This is especially true when the ancillary service is highly valued by passengers, can be provided at low cost by the airline and passengers make purchase decisions rationally. We also show that price segmentation between passenger types can increase revenue and that there is a first-mover advantage for airlines to implement dynamic offer generation mechanisms.

When one of four airlines implements DOG, it can increase its total net revenue by up to 2.6% through ancillary bundling alone and up to 12% in combination with dynamic flight pricing. Most of these dynamic flight pricing gains are attributable to undercutting the existing fares offered by airlines with traditional RM systems. When all four airlines use DOG, their revenue increases by up to 0.9% through bundling alone and 7% with dynamic pricing. Under more realistic market conditions, the simulated net revenue gain of DOG reduces to 1.7% when all airlines implement it.

Thesis Supervisor: Peter P. Belobaba
Title: Principal Research Scientist, Aeronautics and Astronautics

Acknowledgments

I am extremely grateful to Peter Belobaba, my thesis advisor, for his endless support and dedication to his students, MIT and the airline industry as a whole. Through his instruction, I was able to gain deep insights into the airline industry, the field of revenue management and its business challenges. His mentorship has made me a better researcher and communicator.

As I am writing this thesis, I am witnessing the airline industry undergo an existential struggle of unprecedented proportions. I am extremely fortunate to have been academically and financially supported by the MIT PODS Consortium, whose members include some of the largest airline groups and airline IT providers worldwide. It gave me the opportunity to present my research regularly to many revenue management experts, from whom I have learned immensely. For the members' continued support of me, the Consortium and MIT, I am very thankful. A special thank you to Thomas Fiig and Mike Wittman at Amadeus for hosting me for a research visit. Your ideas and advice have propelled my progress as a researcher.

Matt Berge has been instrumental to the development and implementation of the models in this thesis in the Passenger Origin-Destination Simulator (PODS). I am very privileged to have worked with an excellent programmer like him. I would like to thank current and recent PODS student researchers, Adam, Akash, Alexander, Bazyli, Nicholas and Yuxuan, for collaborating with me. You made my time at MIT very enjoyable.

Along my journey at the Institvte, I am supported by a group of wonderful people, who have shaped my life like no other. My real and imaginary roommates. My occasionally rational and my sometimes wonderfully irrational friends. My mutually selected adventurers include clairet, juanmoo, larat, mau, maxli, pozhu, serenado, stellal, zyteka and many more.

Finally, I would not be graduating today without the unabated support of my parents, whose lifelong struggles have enabled me to pursue my passions at the Institvte. You would have loved seeing me walk down Killian Court. Thank you.

Contents

1 Introduction **17**

1.1 Airline Pricing & Revenue Management 18

1.2 Role of Ancillary Services in the Airline Industry 19

1.3 IATA's New Distribution Capability 21

1.4 Research Motivation 23

1.5 Thesis Outline and Contributions 25

2 Literature Review **27**

2.1 Airline Revenue Management 28

2.1.1 Restricted Fare Structures 28

2.1.2 Less Restricted Fare Structures 31

2.1.3 Continuous Flight Pricing Methods 33

2.2 Airline Ancillary Services 34

2.2.1 Passenger Ancillary Choice Behaviors 34

2.2.2 Impacts of Ancillary Services on Revenue Management 35

2.2.3 Joint Management of Flight and Ancillary Revenues 36

2.3 Bundling of Flights and Ancillary Services 37

2.3.1 Economics of Bundling 37

2.3.2 Pricing and Assortment Optimization 39

2.3.3 Airline Offer Generation 40

3 Dynamic Offer Generation Model **43**

3.1 Overview and Problem Formulation 44

3.2 Offer Set Price Optimization 48
3.2.1 Myopic Price Optimization 49
3.2.2 Concurrent Choice Assumption 50
3.2.3 Sequential Approximation 55
3.2.4 μ-Heuristic 56
3.2.5 Sensitivity of Optimized Offer Prices 57
3.2.6 Extension to Two Ancillary Services 59
3.3 Offer Set Selection 62
3.4 Integration with Revenue Management Systems 65
3.5 Summary 69

4 Baseline Results of Dynamic Offer Generation in PODS 71

4.1 The Passenger Origin-Destination Simulator 72
4.1.1 Overview of PODS Architecture 72
4.1.2 PODS Network U10 74
4.1.3 PODS Baseline Simulation 76
4.2 Baseline Dynamic Offer Generation Simulation Results 81
4.2.1 Dynamic A La Carte Results 83
4.2.2 Dynamic Bundling Results 88
4.2.3 Dynamic Offer Generation Results 91
4.3 Summary 98

5 Sensitivity Tests of Dynamic Offer Generation in PODS 101

5.1 DOG with Different Flight-related Parameters 101
5.2 DOG with Different Revenue Management Optimizers 108
5.3 Dynamic Bundling with Different Ancillary Parameters 111
5.4 DOG with Different Ancillary Parameters 116
5.5 DOG in a More Realistic Environment 122
5.6 Summary 129

6 Conclusions **131**

6.1 Research Findings and Implications 132

6.2 Suggested Future Research Directions 135

List of Figures

1-1 Worldwide ancillary revenue estimate 20

1-2 US domestic airfare evolution 20

1-3 Potential benefits of NDC distribution 23

1-4 Example of ancillary offer generation 24

1-5 Schematic illustration of the dynamic offer generation problem 24

3-1 DOG problem overview 45

3-2 DOG with one ancillary service 52

3-3 DOG concurrent choice assumption 54

3-4 DOG myopic ancillary price sensitivity 58

3-5 DOG offer price sensitivity 59

3-6 DOG with two ancillary services 60

3-7 DOG expected net revenue sensitivity 63

3-8 Integration of dynamic offer generation with RM systems 66

3-9 DOG price bounding impact on expected net revenue 69

4-1 PODS network U10 route network 74

4-2 PODS network U10 average fares 76

4-3 PODS network U10 maximum flight WTP 78

4-4 Estimated and true flight WTP 79

4-5 Dynamic a la carte experimental setup 83

4-6 DALC baseline revenue and load factor impact 84

4-7 DALC prices by timeframe 85

4-8 DALC bookings by timeframe 86

4-9 DALC change in booking mix by timeframe 87
4-10 Dynamic bundling experimental setup 89
4-11 DB baseline net revenue and load factor 90
4-12 DB prices by timeframe and bookings breakdown 91
4-13 No DOG baseline net revenue 92
4-14 No DOG baseline load factor and net yield 93
4-15 DOG net revenue impact 94
4-16 DOG net yield impact 95
4-17 DOG load factor impact 95
4-18 DOG prices by timeframe 96
4-19 DOG price segmentation by timeframe 97
4-20 DOG bookings breakdown by offer purchased 98
4-21 DOG bookings by passenger type, offer set and timeframe 99

5-1 DOG price sensitivity to Q multipliers 102
5-2 DOG revenue sensitivity to business Q multiplier 103
5-3 DOG revenue sensitivity to leisure Q multiplier 104
5-4 DOG revenue gain with identification accuracy 105
5-5 DOG price segmentation with identification accuracy 106
5-6 DOG price sensitivity to bounds 106
5-7 DOG bookings sensitivity to price bounds 108
5-8 DOG price sensitivity to RM optimizer 110
5-9 DOG revenue and load factor sensitivity to RM optimizer 111
5-10 DB revenue sensitivity to ancillary WTP 112
5-11 DB bookings breakdown with ancillary WTP 113
5-12 DB ancillary sales with ancillary WTP 114
5-13 DB revenue breakdown with ancillary WTP 115
5-14 DB revenue sensitivity to ancillary cost 116
5-15 DB ancillary sales with ancillary cost 117
5-16 DB revenue breakdown to ancillary cost 117

5-17 DOG bookings breakdown with ancillary WTP 118
5-18 DOG bookings breakdown with ancillary cost 119
5-19 DOG revenue sensitivity to ancillary WTP 120
5-20 DOG revenue sensitivity to ancillary cost 120
5-21 No DOG baseline revenue sensitivity to ancillary WTP misestimation 122
5-22 DOG revenue sensitivity to ancillary WTP misestimation 122
5-23 Re-calibrated DOG revenue and load factor impacts 126
5-24 Re-calibrated DOG average fare by timeframe 126
5-25 Re-calibrated DOG bookings by timeframe 127
5-26 Re-calibrated DOG bookings breakdown by offer purchased 128
5-27 Re-calibrated DOG revenue gain with identification accuracy 129

List of Tables

4.1 Summary Table of PODS Baseline Simulation Parameters 82
4.2 DALC baseline results 84
4.3 DB baseline results 89
4.4 DOG baseline results 94

5.1 DB sensitivity test on ancillary WTP 112
5.2 DB sensitivity test on ancillary cost 114
5.3 DOG sensitivity test on ancillary WTP misestimation 121
5.4 Re-calibrated DOG baseline results 125

Chapter 1

Introduction

The airline industry is diversifying. Driven by a period of record profitability from 2010-2019, many different airline business models have found success. Global network carriers and short-haul low-cost carriers alike have found their target customers with different service levels and fares. Airlines are pushing new retailing initiatives to better distinguish themselves from their competitors. More recently, airlines have used branded fares such as Basic Economy, Main Cabin or Economy Plus to deliver differentiated travel experiences within the same cabin. In the future, the New Distribution Capability (NDC) standard is expected to streamline the ticket distribution process and enable airlines to receive more information about their customers' preferences and in return generate more relevant offers, which combine flight itineraries and ancillary services to form new customized travel products.

For revenue management, this represents a new opportunity to improve flight and ancillary revenues. However, new revenue management models are needed to dynamically generate and price airline offers in response to customer requests. Our research explores this new problem space we call dynamic offer generation by developing new optimization models and heuristics to price flight and ancillary offers. In this chapter, we contextualize the recent trends in the airline industry by discussing the evolution of airline revenue management (1.1), the growth of ancillary services (1.2) and the development of NDC (1.3). We then motivate our research (1.4) and outline the remaining chapters of this thesis (1.5).

1.1 Airline Pricing & Revenue Management

Airlines have long aimed to price seats in a way that maximizes revenue. To account for the large variation in willingness-to-pay of different customer segments, airlines use restrictive fare rules to differentiate their fares such that each customer segment prefers a different fare class. Examples of such rules are advance purchase and point of commencement restrictions, change and cancellation fees, as well as round-trip and minimum stay requirements. Revenue management systems limit the availability of low fares on high-demand flights in order to reserve seats for higher revenue bookings in the future. Revenue management teams set seat inventory levels to maximize total flight revenue while balancing the trade-off between departing with unsold seats (spoilage of inventory) and not having enough seats to serve high-revenue, last-minute demand (spill of demand).

The introduction of the first automated revenue management system by American Airlines in the 1980s has been credited by its then-Chairman and CEO Robert Crandall as "the single most important technical development in transportation management". The airline estimated a sustained $500 million annual revenue benefit from its investment in yield management and overbooking systems (Smith et al., 1992). Initial revenue management systems optimized the seat inventory for each flight leg individually using heuristics such as the expected marginal seat revenue method introduced by Belobaba (1987). Today's advanced optimizers consider the demand across the entire flight network as a whole. This is especially important for the hub-and-spoke networks of many major airlines. As each flight acts as a source of onward connections for other flights, the revenue value of a booking extends beyond the flight itself to the entire network. Network revenue management models increased revenues by 1-2% over leg-based revenue management in simulations (Belobaba, 2002).

The effectiveness of differential pricing with fare restrictions reduces when low-cost carriers enter a market with simplified fare structures and almost no restrictions. Their one-way pricing typically offers a single fare for each flight at any one time and is easier to comprehend for consumers, but dilutes the pricing power of airlines with restricted fare structures. Airline revenue management systems had to adapt to unrestricted fare structures (Belobaba, 2011), where the core assumption of independent demand in each fare class was no longer valid: passengers would always choose the lowest available price and could no longer be segmented into different fare classes using restrictions.

The long legacy of airline distribution also imposes constraints on an airline's ability to price flights. These include limitations on the number of reservation booking designators (RBD) and thus price points that can be offered in each market at any one time, typically a maximum of 26. In addition, price points can only be updated in fixed intervals when airlines "publish" their fares publicly through organizations such as the Airline Tariff Publishing Company (ATPCO). Currently, airlines are preparing for a future with continuous prices, where the fare could be set at any price point without limitations. This would allow a revenue management system to freely set prices and potentially dynamically adjust them in real-time to match a competitor's fare. As airline revenue management continues to evolve, it is plausible that the concept of fare classes could be eliminated altogether.

1.2 Role of Ancillary Services in the Airline Industry

In the airline context, ancillary revenue can be defined as "revenue generated by activities and services that yield cashflow for airlines beyond the simple transportation of customers from A to B" (IdeaWorks, 2019). These come mainly from a la carte sales of services such as checked luggage, seat assignments, meals, onboard WiFi, or cancellation and rebooking fees. Another ancillary revenue stream is from commission-based sales and frequent flyer activity, for example credit card spend or rental car and hotel bookings made through the airline.

Ancillary services have been a fast-growing revenue stream in the last decade, expected to hit an average of $24 per passenger and 12.2% of total airline industry revenue in 2019 (IdeaWorks, 2019). Of this total, we focus our research on a la carte sales, which represent an estimated 70% of global ancillary revenues. The

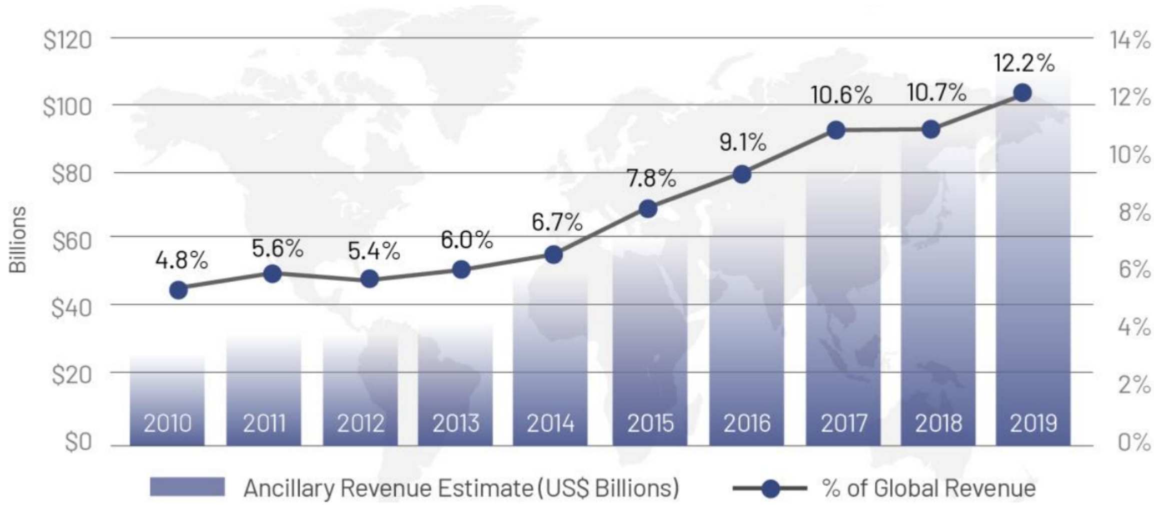

Figure 1-1: CarTrawler worldwide estimate of ancillary revenue (IdeaWorks, 2019)

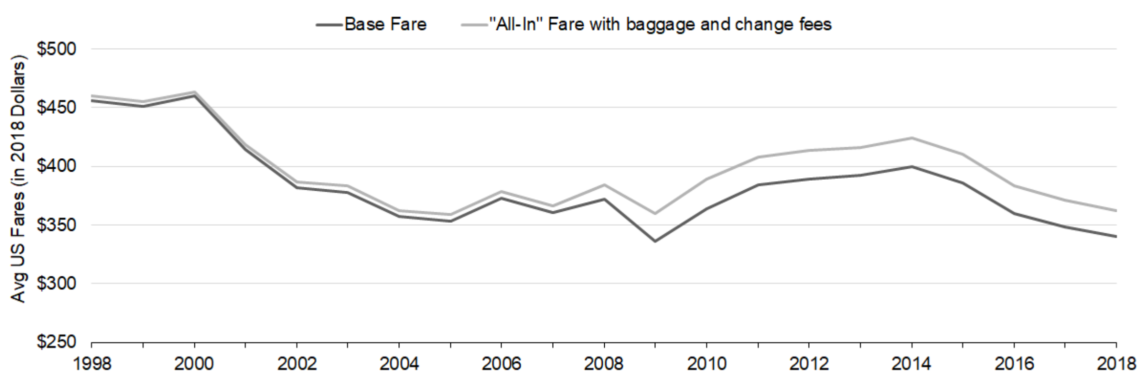

Figure 1-2: Average US domestic round-trip airfares (Airlines for America, 2018)

biggest source of global ancillary revenue are baggage fees, which represent 60% of total ancillary revenue for low-cost carriers (IdeaWorks, 2018). Airlines have achieved this growth in ancillary revenues by *unbundling* their fares, removing services such as inflight meals and checked luggage that were previously offered for free from the cheapest fares. This allows them to counter a trend of generally declining ticket prices (Figure 1-2).

Unlike the pricing of the flight itself, airlines have more flexibility to set ancillary prices, as most ancillaries are sold through the airline's direct distribution channels. As a result, ancillaries are a popular testbed for experiments with dynamic and continuous pricing (IdeaWorks, 2018). For example, Spirit Airlines, an ultra low-cost carrier in the United States, varies baggage fees based on the search request, travel date, route, and time of purchase (CAPA Centre for Aviation, 2019). In another experiment, the use of machine learning to individually price ancillary services based on passenger characteristics increased ancillary sales (Shukla et al., 2019).

As ancillary revenues increase in importance, revenue management systems have to take those potential revenue sources into account when they calculate seat protection levels. New research has shown that maximizing ancillary revenue can have a negative impact on flight and total revenues if revenue management systems are not adapted (Lu, 2019). Strategies and new optimization models to extend revenue management systems to account for ancillary sales were proposed by Bockelie (2019).

1.3 IATA's New Distribution Capability

Airlines primarily sell their flights and services through either the direct or indirect channels. When airlines sell tickets through the direct channel, they have full control over the booking flow: the order in which itineraries are shown, which fare classes are displayed and which ancillary services are offered. The indirect channel comprises bookings made through third parties such as travel agents, corporate travel providers, or metasearch engines. These retailers typically rely on global distribution systems (GDSs) to aggregate flight offers from multiple airlines and make bookings with the airlines. In 2015, the GDSs processed nearly half of global flight bookings (Taubmann, 2016).

In the indirect channel, airlines have much less control over how their services are priced and sold. The EDIFACT standard, which is used to communicate between airlines and the GDS, can only transmit basic information about flight schedule and fare class inventory availability (Wittman, 2018). In each available fare class, a price can be computed by looking up the corresponding published fares and verifying that the chosen itinerary satisfies all the rules and restrictions of the fare. As a result of the technological limitations of the indirect channel, GDSs and flight comparison websites rank itineraries almost exclusively based on price and schedule. The lack of detailed product information in the indirect channel makes it difficult for airlines to differentiate themselves with meal services, free checked baggage or more comfortable seating.

To overcome the limitations of the indirect channel, the International Air Transport Association (IATA) has launched the New Distribution Capability (NDC) program (Westermann, 2013). Its purpose is to drive the adoption of a new XML-based data transmission standard, which allows better communication between airlines, content aggregators and travel agencies. Under the NDC standard, the airline could receive more information about each search request, for example the geographical location, frequent flyer profile, or preferences for ancillary services. Instead of only providing fare class availability, the airline can respond with a set of customized offers. Beyond the itinerary and the flight price, these could include information about associated fare rules, included onboard services, or prices of additional ancillary services. NDC could give airlines more influence over how their services are presented and sold. The airline could sell its services as branded fares and enrich results with multimedia content about the onboard experience.

In particular, NDC could remove legacy restrictions and transform pricing and revenue management. For example, the elimination of fare classes and the limit on 26 fixed price points would pave the way for continuous pricing. Besides, the airline would receive more data about the customer making the request, which could improve an airline's understanding of customer willingness-to-pay. It is conceivable that more detailed price segmentation will become possible. Furthermore, NDC presents a large opportunity for airlines to improve the distribution of ancillary services. Currently, many ancillary services cannot be distributed through the indirect channel and have to be purchased in a second step on the airline's website or at the airport. With NDC, the ancillaries offered and their prices could even vary from one request to another.

In the limit, NDC enables dynamic offer generation (DOG), which is the focus of this thesis. In response to a search request, a customer would receive a set of dynamically generated offers that contain flights and ancillary services, the nature of which could vary based on a customer's profile. Some customers may receive *bundled* offers, where the ancillary services are already included in the flight price, especially

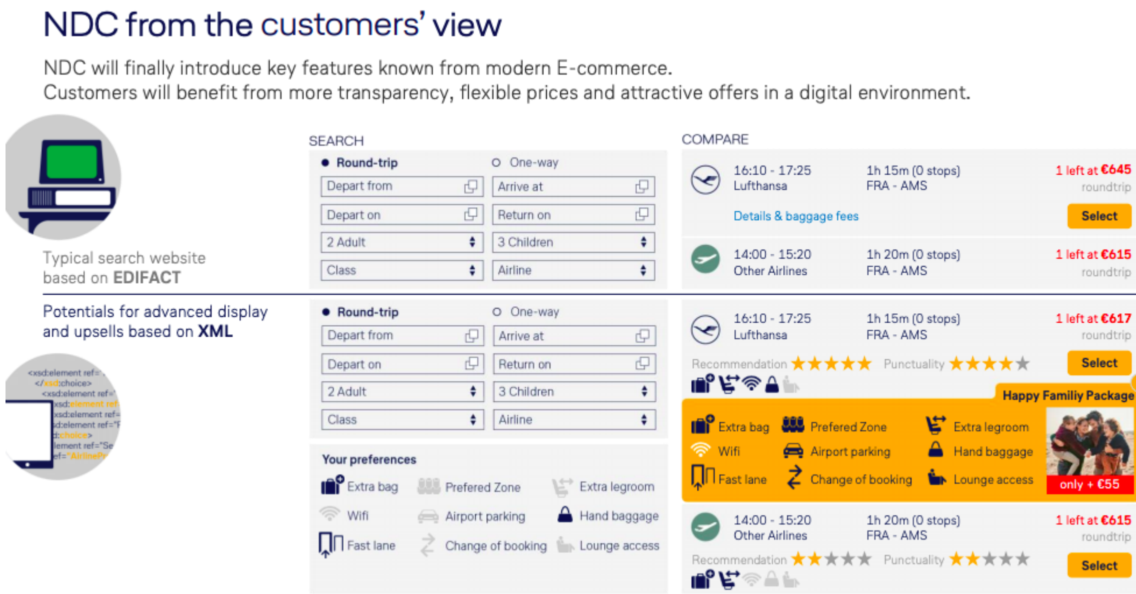

Figure 1-3: Illustration of potential benefits of NDC distribution (Lufthansa, 2018)

if they expressed a preference for such in the search request. Other customers may receive *a la carte or unbundled* offers, that allow them to purchase ancillary services as add-ons to the flight for a fee. Multiple ancillary services could be combined in bundles at a discount. With NDC, the flights, ancillary services and bundles could all be priced continuously without pre-determined price points.

1.4 Research Motivation

The focus of this thesis is to develop new revenue management models that enable dynamic offer generation in the airline industry. With dynamic offer generation and the new distribution capability, the scope of airline pricing & revenue management teams could expand. In addition to their existing responsibility to set the most appropriate flight and ancillary prices at any given point in time, revenue management would be able to control the subset of products displayed to the customer (assortment optimization) and vary this from one customer segment to another. The overall objective of the airline is to optimally price both flight and ancillary services and make the revenue-maximizing subset of services available for purchase.

This represents a fundamental departure from traditional pricing and revenue

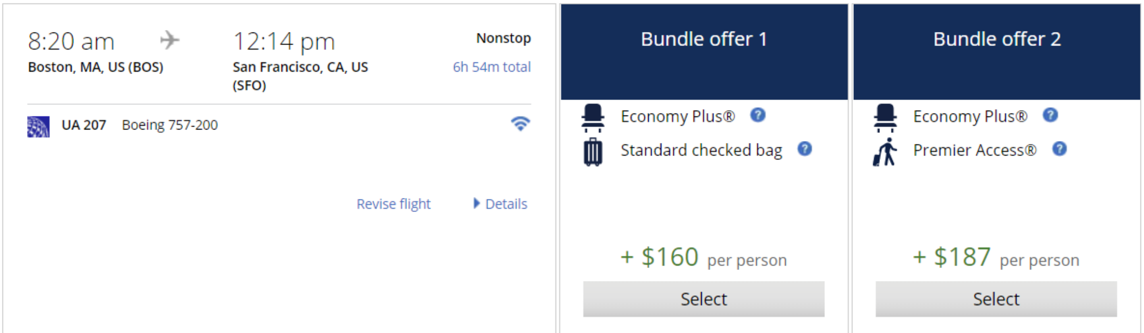

Figure 1-4: Example of ancillary offer generation (Source: united.com on 2019-12-04)

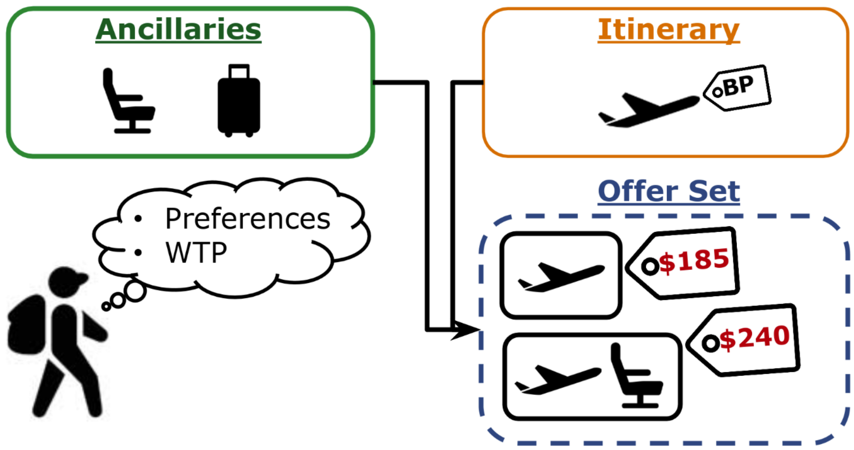

Figure 1-5: Schematic illustration of the dynamic offer generation problem

management functions. Typically, prices of ancillary services are predetermined and independent of the flight price, although some airlines include more ancillary services free of charge with more expensive fares. While studies have shown that ancillary prices can impact flight revenues (Lu, 2019), there exists little research on how ancillary services could be priced optimally alongside the flight.

For example, United Airlines generates two ancillary bundle offers during the flight purchase process in their direct distribution channel (Figure 1-4). It states that "the bundle offers you receive are customized to your trip." The goal of revenue management is to select two combinations of ancillary services and determine the best prices that will generate the highest expected revenue from each customer. Theoretically, the offers could change based on trip-specific parameters (distance, flight price, etc.), but also customer-specific information (trip purpose, loyalty status, etc.).

In our research, we explore the potential for dynamic offer generation to go one step further beyond ancillary offer generation (Figure 1-5). Instead of generating bundles of ancillary offers after the flight and its price has been determined, we are interested in optimizing flight and ancillary prices together. We would also like to evaluate bundles that combine the flight with ancillary services at a single price point. Optimizing flight and ancillary prices together raises previously unexplored questions in airline revenue management, such as:

- How should the prices for flight and ancillary *bundles* be determined?
- For which customers and when is it beneficial to offer services *a la carte* instead of as a *bundle*?
- When is it better to only offer *bundles* and no option to purchase the standalone flight?
- Can it be profitable to offer customers both *bundled* and *a la carte* options?
- How does the optimal price for a service depend on the alternative offers that are displayed alongside it?
- Does the optimal price for the bundle change if the *a la carte* option was available for purchase as well?

Making these pricing and assortment optimization decisions dynamically in response to each customer request calls for new optimization models, which is the motivation for our research.

1.5 Thesis Outline and Contributions

In this thesis, we provide a formulation for the general dynamic offer generation problem in the context of airline revenue management. We propose a solution that can be integrated with traditional revenue management systems and study its performance in a competitive airline network revenue management simulation. Our research shows when bundling ancillary services with the flight can improve revenue and explores the competitive implications of such a pricing strategy.

Chapter 2 provides a literature review on new developments in both airline flight revenue management and ancillary revenue management. We also review the literature on the topic of bundle pricing of two or more products and assortment optimization.

In Chapter 3, we formulate the joint price and assortment optimization problem and apply it to airline revenue management, in the context of dynamically pricing bundles of flights with ancillary services. We propose a dynamic offer generation algorithm to solve the problem for Gaussian distributed demand. We explore the robustness of the algorithm to its input parameters and derive conclusions about the profitability of bundling ancillary services. We illustrate how airlines can apply the solution and integrate it with existing revenue management systems.

In Chapter 4, we introduce the Passenger Origin-Destination Simulator (PODS) and define a set of baseline parameters to study the performance of the model in conjunction with a traditional flight revenue management system under competition. We explain the magnitude and source of observed revenue performance in detail by separating the benefits of dynamic offer generation into its two components: dynamic flight pricing and dynamic offer set selection (ancillary bundling). We also illustrate the competitive effects exhibited in the simulation.

In Chapter 5, we use further sensitivity tests to confirm the performance of dynamic offer generation under a variety of input parameters. We identify conditions under which dynamic offer generation performs better than traditional revenue management as well as the additional risks introduced by the strategy. We incorporate the insights in a simulation that reflects more realistic market conditions, where the ancillary service represents a checked bag.

Finally, we conclude the thesis with Chapter 6, summarizing the key findings and suggesting future research directions in the area of airline dynamic offer generation.

Chapter 2

Literature Review

This chapter surveys the revenue management literature with a focus on airline industry applications and offer generation problems. In Section 2.1, we discuss the origins of revenue management in the airline industry. Beginning with traditional optimal solutions and heuristics for the revenue management problem with independent demand, which is applicable in markets with restricted fare structures (2.1.1), we discuss the evolution of airline pricing to less restricted and unrestricted fare structures and the implications for revenue management algorithms (2.1.2) and look ahead to a future reality with continuous prices and dynamic price adjustment (2.1.3). We then discuss the impact of ancillary services on airline revenue management (2.2) and describe models that incorporate ancillary choice behavior into revenue management to maximize total revenue. Finally, we discuss the implications of bundling flights and ancillary services into offers for revenue management (2.3). We review the economic rationale behind bundling and existing results from the economics literature (2.3.1). We discuss how a set of offers can be generated (assortment optimization) and priced (pricing optimization) (2.3.2). We conclude this chapter by reviewing prior contributions to the airline dynamic offer generation problem (2.3.3), on which the contributions in this thesis are based.

2.1 Airline Revenue Management

Revenue management is the discipline of selling the right product to the right consumer at the right time and the right price through the use of analytics. The idea is especially applicable to airlines because they are faced with perishable inventory (unsold seats at the time of departure have no salvage value), variable and elastic demand (seasonal variability in demand, trip purpose and willingness-to-pay) and varying lead times (travel booked both far in advance and close to departure). As a result, the airline industry led the development of revenue management systems before the idea spread to other industries like hotels, rental cars and retail. Initial revenue management research made many assumptions specific to the airline industry and led to the rapid advancement of the field. Over time, as the business environment changed for the airlines, revenue management research evolved with it. Here, we provide a brief overview of the history of airline revenue management and the current research directions.

2.1.1 Restricted Fare Structures

McGill and Van Ryzin (1999) provide a seminal overview of the state of revenue management research at the time. They trace the origins of revenue management to the airline overbooking problem, in which airlines oversell flights to maximize revenue based on their forecast of cancellations and no-shows. The earliest seat protection mechanism is credited to Littlewood (1972), whose rule applies to an airline with two fare levels (high and low). Based on the newsvendor model of inventory theory, it limits the number of seats available at the lower fare in order to "protect" the remaining seat capacity for future high fare bookings. Littlewood's rule balances the revenue of the lower fare with the *expected* revenue of potentially selling at a higher fare in the future. Belobaba (1987) extended the rule to multiple fare classes using the *expected marginal seat revenue* (EMSR) heuristic. A refined version of EMSR, called EMSRb, was widely adopted in the airline industry due to its simplicity and intuitiveness (Van Ryzin and McGill, 2000). Though EMSR is not optimal (Robinson, 1995), the heuristic is faster to compute and produces near-optimal results with demand distributions commonly observed in the airline industry (Belobaba, 1992). Algorithms that determine *optimal booking limits* (OBL) for the single flight leg RM problem were published by Curry (1990) and Brumelle and McGill (1993). These are computationally more expensive, as they require convolutions of multiple probability density functions to compute the nested purchase probabilities for each class.

Both the OBL and EMSR methods are static optimization models that compute the seat allocations based on strict assumptions about the demand. For example, the methods assume that demand across each fare class is independent. This is not the case in reality, where passengers choose to sell-up and buy-down from one fare class to another. Furthermore, the demand arrivals are assumed to be ordered strictly in increasing order, with the demand for a cheaper fare class arriving entirely before the demand for a more expensive fare class.

To relax this assumption and better account for the fact that low-fare demand exists close to departure (and vice versa), an alternative revenue management optimization approach models demand arrivals as a Markov decision process. In this formulation, every single unit of demand is forecast to arrive in its own time slice t, the state is described by the number of accepted bookings x_t at time t. The two control actions are to either accept the new demand and transition to a state $x_{t+1} = x_t + 1$ or reject it and transition to state $x_{t+1} = x_t$. The seat allocations for a single flight leg are then solved using *dynamic programming* (DP) (Lautenbacher and Stidham, 1999). The optimal control policy is generated by taking the action that maximizes the value function at any point in time, which describes the maximum expected future revenue from state x_t at time t. This mathematically attractive formulation has not been widely adopted in modern airline RM systems, due to the computational complexity of the algorithm. DP has been shown to outperform EMSRb and OBL when demand variance is low, but underperform the static methods when true demand variance is higher than the underlying model assumption of Poisson arrivals (Diwan, 2010).

Most early optimization methods were based on key simplifications, such as optimizing for a single flight only without considering connecting traffic and assuming independent demand in each booking class (McGill and Van Ryzin, 1999). The second assumption was reasonable at the time, because airlines used highly restricted fare structures to segment demand into different fare products, whose availability was controlled with booking classes. For example, a business traveler desiring a round-trip ticket departing on a Monday and returning on a Thursday would not be able to choose a lower fare because the discounted fare products required a Saturday night stay at the destination. However, the first assumption was unsatisfactory since airlines with hub-and-spoke systems carried large amounts of connecting traffic. Accepting a booking on a connecting itinerary could displace higher revenue nonstop passengers, or vice versa. Further research described an approach that set seat protection levels on an origin-and-destination level (O-D control) (Smith and Penn, 1988).

Most commonly, network RM solutions rely on a linear program to produce an optimal allocation of deterministic (zero variance) O-D demand forecasts across the network (a minimum cost flow problem). Despite the deterministic demand assumption, the shadow (dual) prices of the capacity constraints can be used as estimates of the opportunity cost of displacing a single passenger on that flight segment. O-D control methods use these displacement costs to account for the network effects when calculating the seat protection limits. One approach called *Displacement Adjusted Virtual Nesting* (DAVN) maps the net revenue value of a connecting fare after adjusting for displacement costs into a corresponding local booking class on each flight segment (Smith and Penn, 1988). Thus, all O-D pairs are clustered and tied to the availability of local booking classes. An alternative approach uses *bid prices* (BP) to determine the minimum fare required to accept a booking. The bid price is the total displacement or opportunity cost of accepting a booking summed across all traversed flight legs. A booking class is available for purchase if its fare is higher than the corresponding bid price. One algorithm that uses bid price control is called *Probabilistic Bid Price* (ProBP), which nests and prorates connecting fares to each traversed leg based on their displacement costs (Bratu, 1998). Another widely adopted bid price control method is *Unbucketed Dynamic Programming* (UDP), a network extension of the single-leg DP (Lautenbacher and Stidham, 1999). Because solving the Markov decision process for an entire flight network is computationally infeasible, the DP is solved for each flight leg individually, after applying using a deterministic network LP to perform displacement adjustment. This reduces the total revenue value of accepting a connecting booking on the leg to be optimized by the displacement cost on the other legs. The bid price for each flight leg at time t is determined as the change of the DP value function (expected future revenue) if one unit of demand was accepted at time t.

2.1.2 Less Restricted Fare Structures

With the advent of low-cost carriers (LCCs), revenue management systems had to be fundamentally revised (Belobaba, 2011). To compete against legacy airlines, LCCs often use unrestricted fare structures, where every flight has a single price at any given time, no matter how many nights passengers stays at the destination or whether it is the outbound or the return flight. These simplified fares are popular with consumers, but weaken the airline's ability to segment demand into different fare products. Legacy airlines often respond by matching the less restricted fare structures, which causes an issue for revenue management systems. Less restricted fare structures fundamentally violate the second assumption in Section 2.1.1, namely that demand is independent in each fare class. Instead, customers mostly choose the lowest available fare, with few incentives to purchase a higher fare. When traditional revenue management systems are used with unrestricted fares, the well-documented *spiral down* effect occurs (Cooper et al., 2006). As passengers that were originally willing to pay higher fares buy down into the lower booking classes, the independent demand forecast expects fewer future bookings in higher classes and the revenue management system protects fewer seats, which allows for more even buy down to occur.

The notions of customers buying down to a lower fare class when it is available for purchase and selling up to a higher fare class when the lower one is closed were studied by Belobaba and Weatherford (1996), who proposed heuristics to account for these effects. New revenue management models had to account for customer choice behavior. The groundwork was laid in a general model proposed by Talluri and van

Ryzin (2004). Strauss et al. (2018) provide a good review of the developments since 2004 and popular choice models for revenue management.

In the airline context, *Q Forecasting* is a forecasting method developed by Belobaba and Hopperstad (2004) that accounts for sell-up effects. Instead of forecasting demand for all booking classes independently, it generates a single forecast of total demand at the lowest (Q) class. This forecast is then re-partitioned into higher class demand using a sell-up rate described by a *FRAT5* value (fare ratio at which 50% of demand is willing buy up from the lowest fare). For a detailed description of the method, see Cléaz-Savoyen (2005).

Q forecasting for unrestricted fare structures assumes all demand will buy down to the lowest available fare (*priceable* demand). With less restricted fare structures in reality there is a proportion of *yieldable* demand that seeks a specific booking class even though a lower fare is available (Boyd and Kallesen, 2004). This is especially true when airlines employ a mix of restricted and unrestricted fare structures across their network, or when airlines group booking classes into so-called *fare families* (Walczak and Kambour, 2014). Fare families have become a popular way to market airline fares, in which all fares within a family share the same set of restrictions and characteristics and thus demand is priceable within, but each family differs from another significantly. As a result, demand within a fare family can be considered priceable, whereas it is yieldable across different families. Hybrid forecasting methods were tested by Reyes (2006), in which yieldable and priceable demand are forecast independently using traditional class-independent forecasts and Q forecasting methods, respectively.

Beyond the effects on the demand forecast, priceable demand has to be accounted for in the optimization. This is especially important for O-D control systems, where some O-D markets may have restricted and others unrestricted fare structures. In both markets, the booking class availability on a shared flight leg is determined by comparing the displacement adjusted fare against a common bid price. However, in markets with unrestricted fare structures, the risk of buy down (or revenue dilution) is much higher if a lower booking class is made available. To account for this difference, Fiig et al. (2010) developed a marginal revenue transformation, which converts an unrestricted fare structure to an equivalent set of fare classes with independent demand. The transformation reduces the revenue management optimizer's perception of the revenue value of opening a lower fare class and accounts for the risk of buy down. The unrestricted fare is thus nested in a lower class than the restricted fare when using O-D control. As a result, a network RM system would reduce seat availability in markets with unrestricted fare structures. This method is called *Fare Adjustment* and has enabled the continued use of traditional revenue management systems in markets with unrestricted fare structures and prevented the "spiral down" effect, despite an assumption of independent demand (Fiig et al., 2005) (Fiig et al., 2010) (Walczak et al., 2010). This concept was later applied to the optimization of fare families (Fiig et al., 2012).

2.1.3 Continuous Flight Pricing Methods

Motivated by the advancements in revenue management and dynamic pricing for other industries (Golrezaei et al., 2014) (Gallego et al., 2016), airlines began to explore the possibilities of offering flights at more than 26 pre-defined price points. Wittman and Belobaba (2019) provide a definitional framework describing two different approaches to so-called "dynamic pricing", as compared to traditional revenue management, which is described as an assortment optimization problem of choosing which pre-defined fare classes to make available for purchase. The first, called *Dynamic Price Adjustment*, uses a traditional revenue management system to perform the assortment optimization process, but then adjusts the resulting fares away from the pre-defined published fares (Wittman and Belobaba, 2018). The adjustment could be based on an estimate of the individual passenger's willingness-to-pay or trip purpose, or on the current competitor fares available (Fiig et al., 2016). Wittman (2018) discusses the potential regulatory implications of segmented pricing, in which different customer segments could receive different fare quotes at the same time for the same itinerary, a possibility with dynamic price adjustment and IATA's New Distribution Capability (Westermann, 2013).

The second approach, called *Continuous Pricing*, lets a revenue management system freely choose an appropriate price from a continuous and unlimited set of possible price points. Liotta (2019) describes and tests two different continuous pricing methods in the Passenger Origin-Destination Simulator (PODS). The *class-based continuous* method uses the bid price calculated by a traditional class-based revenue management system as the continuous fare. On the other hand, the *classless* method removes the concept of fare classes altogether and generates demand forecasts by time period for the revenue management optimizer instead. Both methods result in positive revenue gains over traditional class-based pricing in both symmetric and asymmetric scenarios with unrestricted fare structures. The results also show that similar revenue gains can be achieved by increasing the number of pre-defined fare classes sold by an airline using traditional revenue management systems.

2.2 Airline Ancillary Services

As ancillary revenues have become more important to airlines (see Chapter 1.2), researchers and airline revenue management teams have begun studying the ancillary revenue management problem in more detail. In this section, we review existing literature on how ancillary purchase behavior of passengers can be modeled and how ancillary revenues impact the optimality of existing revenue management systems. Finally, we review existing literature on how revenue management can account for ancillary revenues with the goal of total revenue optimization.

2.2.1 Passenger Ancillary Choice Behaviors

For any model that attempts to optimize ancillary revenues, it is important to consider how and when passengers choose which ancillary services to purchase. In their research on ancillary price optimization, Odegaard and Wilson (2016) modeled a mix of three types of choice behaviors: passengers who never purchase any ancillaries, passengers who purchase ancillaries if the price is below their reservation price and passengers who always purchase the ancillary service with the flight. For the evaluation of price against reservation price, Bockelie and Belobaba (2017) propose an integrated flight and ancillary choice model and describe two distinct types of consumers: *Simultaneous* consumers select a flight "itinerary, a fare class and a set of ancillary services at the same time" based on their flight and ancillary reservation prices and a combined overall budget (willingness-to-pay). They choose rationally with full knowledge of all ancillary prices across all airlines at the time of booking. On the other hand, *sequential* consumers first choose a flight itinerary and fare class without any knowledge of ancillary offers or prices at all, before proceeding to evaluate the offered ancillary services against their reservation prices. This boundedly rational behavior represents consumers who choose the cheapest flight on a low-cost carrier and then encounter unexpected ancillary fees, while a competing full-service carrier airline offers a lower total price.

2.2.2 Impacts of Ancillary Services on Revenue Management

Bockelie and Belobaba (2017) performed initial studies of both simultaneous and sequential choice behaviors in PODS. They show that sequential passengers evaluate flight and ancillary service offers independently and thus an airline's flight revenue is not impacted by ancillary pricing. On the other hand, the ancillary purchase decision of simultaneous passengers can impact flight revenues: If higher booking classes include complimentary ancillary services, they may choose to buy-up and pay a higher fare instead of the ancillary fees. Simultaneous passengers can also buy down to a lower, more restricted fare class, and instead purchase ancillary services with their remaining budget. This has implications for the forecasting and optimization systems of airline revenue management systems. Lu (2019) extended the studies to include the competitive effects of ancillary price segmentation (if one or more airlines begin to charge a subset of passengers more or less for the same ancillary service) and ancillary differentiation (if a subset of passengers have a higher or lower reservation price for the ancillary service). These suggest that an airline's market share can increase when it offers a cheaper ancillary service than its competitors, because simultaneous passengers will switch their choice of airline. Thus, a minor difference in ancillary pricing can have large implications on an airline's ticket and overall revenue.

2.2.3 Joint Management of Flight and Ancillary Revenues

These results raise the question of how airline revenue management systems can account for the impacts of ancillary revenues. It also highlights the importance of collaboration within an airline's commercial department. If ancillary prices were set myopically to maximize ancillary revenues, an airline's overall and flight revenues would decline (Lu, 2019).

One approach called the *Optimizer Increment (OI)* involves informing the revenue management optimizer of the ancillary revenue value of a customer. This is done by increasing the fare values used in the revenue management optimizer by the expected ancillary revenue generated by an incremental passenger. As a result, it becomes more preferable to sell a seat for \$100 than to protect it for a potential \$200 booking in the future that materializes with a 50% probability. This is because each seat sold has a potential to generate ancillary revenues after booking. Hao (2014) first simulated Optimizer Increment in PODS and found that it results in overall revenue losses. These stem primarily from higher availability of lower fare classes and resulting buy-down of demand with higher willingness-to-pay. However, revenue management methods that incorporate passenger sell-up models (hybrid forecasting with fare adjustment) showed overall revenue increases with OI methods.

Bockelie (2019) proves that OI is an optimal control strategy under limited conditions. However, simulations of OI with the ancillary choice behaviors described in Section 2.2.1, show that revenues decrease with OI once feedback effects are considered: Buy down from higher classes lead to lower demand forecasts and thus more opportunity for buy down. Instead, Bockelie (2019) proposes the *Ancillary Choice Dynamic Program (ACDP)* for joint flight and ancillary revenue management, which is an extension to the dynamic program in Talluri and van Ryzin (2004) with an ancillary revenue term. He also develops the *Ancillary Marginal Demand (AMD)* and *Ancillary Marginal Revenue (AMR)* transformations, which are fare adjustment heuristics. Analogous to Fiig et al. (2010), these transformations incorporate ancillary-awareness (buy-up and buy-down) into traditional static revenue management optimizers that assume independent demand. He shows that in competitive PODS simulations, the AMD/AMR heuristics can outperform OI by 3.5% in total revenue in a hypothetical small market scenario.

The AMD/AMR heuristics, the ACDP and OI are all assortment optimization methods that determine the availability of flight and ancillary combinations from a pre-defined set of fares and ancillary prices. They do not optimize the flight and ancillary prices themselves. On the other hand, Odegaard and Wilson (2016) develop a multi-period dynamic program that solves for the optimal flight and ancillary prices in each period, but do not consider the nature of airline distribution systems that currently rely on pre-defined flight and ancillary price points.

2.3 Bundling of Flights and Ancillary Services

Previous research has explored the interactions between the flight and ancillary pricing, but to our knowledge no prior research has developed an optimization model to price a bundle of flights with ancillary services, as would be required for dynamic offer generation. To motivate the rationale behind bundling flights with ancillary services, this section reviews the economics literature on bundling and its profitability. We then review methods to optimize prices in an assortment of products. We conclude the literature review with existing research on the topic of dynamic offer generation in the airline industry.

2.3.1 Economics of Bundling

The first studies into the economics of bundling are attributed to Stigler (1963) and Adams and Yellen (1976). Early research generally showed that it can be profitable to offer bundles of two distinct products at a cheaper price than purchasing both items separately. These bundles can either be sold alongside the individual products (*mixed bundling*) or exclusively, such that the products are not available for purchase individually (*pure bundling*). Kobayashi (2005) provides a good overview of the different types of bundling and the relevant literature.

The seminal model of Adams and Yellen (1976) makes a number of key assumptions about bundling:

- A seller's marginal cost of selling one bundle is the sum of the marginal costs of its components.

- A consumer's reservation price (maximum willingness-to-pay) for a bundle is the sum of the reservation prices of its components.

- The reservation prices for the individual products are independent. In particular, there is no limit on total willingness-to-pay for all products and the reservation prices do not depend on the selling prices.

- A consumer's choice behavior is fundamentally rational, in that they purchase the products that maximizes their consumer surplus (reservation price - selling price). In *mixed bundling*, they will never purchase two products individually if a cheaper bundle is available.

Schmalensee (1984) explored this model where the reservation prices for two products were subject to a bivariate normal distribution and concluded that both *pure bundling* and *mixed bundling* could be more profitable than unbundled (or *a la carte*) sales. Under this model, bundling appears most effective when the reservation prices are negatively correlated and as a result the optimal price for the bundle is relatively close to the individual prices (in mixed bundling) (Adams and Yellen, 1976) (Schmalensee, 1984) (McAfee et al., 1989). This way, mixed bundling can be used as a price discrimination tool: Customers with a high willingness-to-pay for one product would be able to purchase a second product in a bundle at a small incremental cost. Only a relatively small willingness-to-pay for the second product is thus required for customers to be able to afford the incremental cost.

Bundling has been especially popular when the marginal cost of selling another item of the product is small and inventories are unlimited (Bakos and Brynjolfsson, 1999). This is especially the case with information goods, such as television channels (cable subscriptions usually bundle many channels at a single price) and software products (such as the Microsoft Office or Adobe Creative Cloud bundles).

All research above assumes the rationality of the consumer and the additivity of their reservation prices. These assumptions fail to consider the psychological influences of customer choice. For example, the "center-stage effect" states that humans are biased towards options in the middle when presented with a set of choices (Raghubir and Valenzuela, 2006) (Rodway et al., 2012). Yadav (1994) explores the psychology of bundling and suggests a model, in which consumers first evaluate the most relevant offer (anchoring) and then compare alternatives relative to it (adjustment). Further studies on how consumers evaluate bundles and partitioned prices were conducted by Morwitz et al. (1998), Johnson et al. (1999) and Chakravarti and Paul (2002). Ben-Akiva and Gershenfeld (1998) provide an example of how more realistic discrete choice models could be calibrated using stated preference data of customers choosing among different bundle options.

2.3.2 Pricing and Assortment Optimization

From the perspective of revenue management, the design and pricing of bundles represents a joint pricing and assortment optimization problem. When bundles of different products are offered, it is reasonable to assume that the optimal price for one product depends on the other products and bundles available for purchase. Early on, Adams and Yellen (1976) showed that the optimal price for a bundle could vary across the *pure bundling* and *mixed bundling* strategies.

The question of how bundles of n products could be priced optimally was explored by Hanson and Martin (1990), who propose a general linear program to generate prices for bundles of any possible subset of the n products. The optimization assumes that customers will rationally choose the bundle that maximizes their consumer surplus among all available options. More recently, Bulut et al. (2009) and Gürler et al. (2009) explored optimal bundle pricing and assortment optimization problems subject to inventory constraints.

As a field of research distinct from price optimization, assortment optimization originated in the retail industry and was used to optimize the selection of products on a store shelf (Kök et al., 2008). However, in the context of bundling, the joint optimization of price and assortment becomes highly relevant. The chosen assortment defines the bundling strategy: Which products are bundled together? Is the bundling strategy mixed or pure? One example of an integrated model is the product planning model by Ferreira and Wu (2011).

2.3.3 Airline Offer Generation

The topic of designing and optimizing the assortment of travel products in the airline industry grew in relevance with the introduction of fare families and branded fares. Branded fares are a product bundle marketed by airlines that combine an airfare with a set of ancillary services or fare restrictions ("Basic Economy" or "Economy Plus"), which make it easier for customers to understand the product they are purchasing. A customer choice framework to design and evaluate branded fare products, as well as their pricing, was initially proposed by Ratliff and Gallego (2013).

Madireddy et al. (2017) and Vinod et al. (2018) proposed solutions that allows an airline to transition from branded fares to offer generation, where the offered travel products could vary from one customer to another. They explore how customers can be clustered into different segments using trip characteristics, how appropriate offers can be designed for each segment and how price experimentation methods like Thompson sampling in the multi-armed bandit problem can be used to price the offers. Fiig et al. (2018) outline the unsolved scientific challenges of dynamically pricing airline offers, but also highlight the opportunities presented by the New Distribution Capability, the improved availability of shopping data and advancements in statistical data analysis and machine learning.

The dynamic offer generation algorithm in this thesis is based on unpublished work by Bockelie and Wittman (2017). They propose a joint pricing and assortment optimization algorithm for bundling flights with ancillaries, in which the airline decides dynamically whether to sell the ancillaries separately (unbundled) or whether to include them with the flight (pure bundling). Unlike the existing literature, this model optimizes prices by assuming an underlying *concurrent customer choice model* (Bockelie and Wittman, 2018) (Bockelie, 2018) and does not perform price experimentation or reinforcement learning. Instead, machine learning could be used to estimate the choice model parameters from real data and to segment customers based on booking characteristics. However, the authors do not provide any formulations for these components and the algorithm assumes that these tasks can be performed externally to generate the required inputs. Similarly, the assortment optimization process selects the offer set that maximizes expected revenue in the concurrent choice model, which is a variant of the ancillary choice behaviors described in Section 2.2.1. It shares many characteristics with the rational choice model described by Adams and Yellen (1976) and Schmalensee (1984), but incorporates the dependency that ancillary services cannot be purchased without the flight. In Chapter 3, this thesis presents the first detailed description of this dynamic offer generation algorithm.

Chapter 3

Dynamic Offer Generation Model

This chapter defines the dynamic offer generation problem that is the focus of this thesis and proposes a solution heuristic that can be incorporated as an extension to a traditional airline revenue management system. This thesis represents the first formal description of this dynamic offer generation algorithm, which is based on an initial formulation by Bockelie and Wittman (2017).

Section 3.1 defines key terms and separates the original dynamic offer generation problem into two independent subproblems: offer set price optimization and offer set selection (assortment optimization). In Section 3.2, we describe in detail how our proposed heuristic solves the offer set price optimization problem. We show how the resulting prices depend intuitively on willingness-to-pay, the flight bid price and the cost of providing ancillary services. In Section 3.3, we then explain how the heuristic selects the set of offers that are distributed to the customer in a way that maximizes expected revenue per customer (assortment optimization). We also show how the chosen set of offers varies with the willingness-to-pay and cost of the offers.

We conclude this chapter with a discussion on how the dynamic offer generation model can be integrated with traditional airline revenue management systems in Section 3.4. Here, we present an improved flight price bounding heuristic that better isolates the offer set selection process from the flight price, compared to the original implementation of Bockelie and Wittman (2017).

3.1 Overview and Problem Formulation

The International Air Transport Association (IATA) defines dynamic offer generation (or dynamic offer creation) as the "construction by an airline of an offer for a defined set of products and services, with a defined set of conditions." The offer is called *dynamic*, because it is "provided in real-time on a one-time basis and in response to a request" (Touraine and Coles, 2018). This process is enabled by the IATA New Distribution Capability and represents a fundamental shift in airline distribution, as airlines would directly handle all search requests and gain more control over the set of offers that are made available to each customer.

While dynamic offer generation (DOG) does not require dynamic pricing, the New Distribution Capability represents an opportunity for airlines to transform their approach to revenue management. Because offers can be customized to each search request, these travel products can be better targeted to a customer segment, both in terms of the included services and the pricing. Airlines would have the capability to, for example discount the price for leisure passengers and bundle popular ancillary services such as checked luggage in the total offer price. On the other hand, business passengers may receive a more flexible, refundable offer at a higher price that includes free onboard internet. By combining the pricing of ancillary services with the flight itself, this enables next-generation revenue management algorithms to perform *total offer management* and price offers at any price point (*continuous pricing*) without legacy limitations on the number of price points that can be offered. The purpose of our research is to propose such a revenue management algorithm that leverages the full potential enabled by dynamic offer generation, including total offer management and continuous pricing.

We begin the formulation of our dynamic offer generation (DOG) algorithm by narrowing the problem definition and introducing key terms used in the rest of the thesis. As illustrated in Figure 3-1, the purpose of the algorithm is to respond to each flight search request with a customized offer set. Each offer set includes one or more offers that contain a flight and zero or more ancillary services. Each offer has

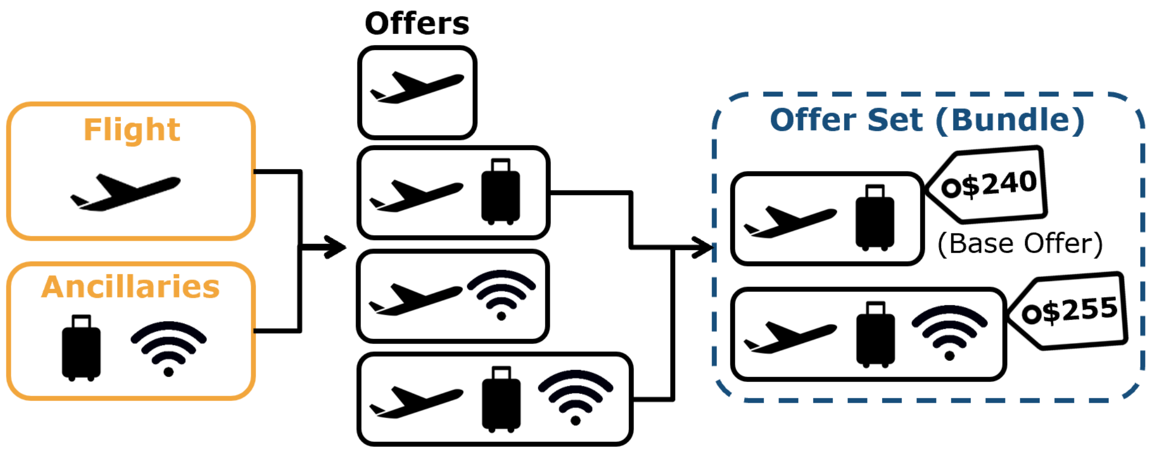

Figure 3-1: Example of two ancillary services in a bundled offer set with two offers

a continuous and dynamically generated price. The offer set contains a subset (an assortment) of all possible offers that the airline can sell based on its range of flights and ancillary services it offers. In the figure, an example airline offers one flight and two ancillary services. With these components, four offers can be created. In response to a request, the DOG algorithm chooses to make two of the four offers available for purchase. Both are priced based on the characteristics of the request, such as trip purpose, trip duration and estimated willingness-to-pay. In this particular case, the generated offer set is *bundled* (Definition 7), as all offers include the checked bag, which is bundled with the flight. On the other hand, the onboard internet service is optional and a customer may choose not to purchase it.

To reduce the scope of the problem, we limit ourselves to offer sets with a single flight itinerary. As a result, we treat each flight independently and neglect any alternative flights when performing revenue management, a common assumption in existing RM algorithms. We also neglect the restrictions and conditions of the fare, such that every flight will have only one price point available at any given time. This is comparable to unrestricted fare structures commonly used by low-cost carriers, where the fare rules are identical across different fare classes. To extend the algorithm to restricted fare structures, where the cheaper fares have more restrictions such as change/cancel fees, one could treat the attribute of "no change fees" as an ancillary service that the algorithm can choose to add to the offer.

Definition 1 (**Flight**). A flight f is a seat supplied by the airline on a specific itinerary and path for travel from an origin to a destination. In this definition, one *flight* can involve boarding more than one airplane if it includes a connecting itinerary.

Definition 2 (**Ancillary**). Ancillaries $a_1, a_2, ..., a_m$ are services provided by the airline in conjunction with a flight that enhances the travel experience (i.e. checked baggage or internet).

Definition 3 (**Offer**). An offer $O_i = \{f, a_k, ...\}$ can be sold by an airline at its corresponding price p_i. Every non-empty offer $O \neq \{\emptyset\}$ includes exactly one flight and zero or more distinct ancillary services.

Definition 4 (**Offer Set**). An offer set $S = \{O_\emptyset, O_1, O_2, ..., O_n\}$ is a collection of one or more offers with the same flight. For any flight f, the offer set encompasses all offers that the airline makes available for purchase. Every offer set includes an empty offer $O_\emptyset = \{\emptyset\}$ with price $p_\emptyset = \$0$.

Definition 5 (**Base Offer**). When it exists, we define an offer in the offer set that is also a subset of all other non-empty offers in the offer set as the base offer O_1. For example, in the following offer set S_1, the base offer is bolded ($O_1 = \{f, a_1\}$). On the other hand, no base offer exists in S_2.

$$S_1 = \{\{\emptyset\}, \{\mathbf{f}, \mathbf{a_1}\}, \{f, a_1, a_2\}\}$$

$$S_2 = \{\{\emptyset\}, \{f, a_1\}, \{f, a_2\}, \{f, a_1, a_2\}\}$$

Definition 6 (**A La Carte Offer Set**). An offer set is called *a la carte* when its base offer includes only the flight and no ancillary services. Since the flight can be purchased without any ancillary services, all ancillaries are sold optionally for an additional charge. S_3 is an example of an a la carte offer set.

$$O_1 = \{f\}$$

$$S_3 = \{\{\emptyset\}, \{f\}, \{f, a_1\}, \{f, a_1, a_2\}\}$$

Definition 7 (**Bundled Offer Set**)**.** An offer set is called *bundled* when its base offer includes the flight and at least one ancillary service. The ancillary services in the base offer are bundled with the flight and must be purchased in order to travel. S_1 is an example of a bundled offer set, where ancillary a_1 is bundled in the base offer.

$$O_1 \supsetneq \{f\}$$

$$S_1 = \{\{\emptyset\}, \{f, a_1\}, \{f, a_1, a_2\}\}$$

For the DOG algorithm, the objective is to choose the optimal offer set and correspondingly the optimal offer prices for each booking request. This objective can be expressed through an offer set's expected net revenue as defined in Equation 3.1. The algorithm needs to find the offer set S and the prices p_i that maximize this expression.

Definition 8 (**Offer Set Expected Net Revenue**)**.** The expected net revenue $\mathbb{E}(S)$ of an offer set S is the probability that a customer will purchase an offer O_i multiplied by each offer's net revenue, summed across all offers $O_i \in S$. The net revenue is defined as the difference between the offer price p_i and the cost of providing all the services in the offer c_i.

$$\mathbb{E}(S) = \sum_{O_i \in S} (p_i - c_i) P(O_i | \vec{p}(S)) \tag{3.1}$$

The dynamic offer generation algorithm can be separated into two distinct optimization problems, which we formulate in this chapter:

- **Offer Set Price Optimization:** How should each offer in an offer set be priced to maximize expected net revenue? In particular, how should the optimal price p^* for an offer change with the offer set it is included in.

- **Offer Set Selection:** Which offer set S should be shown to the customer to maximize expected net revenue? In particular, when should the airline show an *a la carte* offer set and when should the airline show a *bundled* offer set?

3.2 Offer Set Price Optimization

We begin by discussing the offer set price optimization problem. The goal is to generate a price for every offer in every offer set, such that each offer set's total expected net revenue (cf. Equation 3.1) is maximized. This is in the spirit of total revenue maximization, where the airline tries to maximize the total revenue (net of costs) received from a customer. Every customer always purchases one offer from the offer set they are presented with. By optimizing a vector of prices $\vec{p^*}(S)$ for the whole offer set S at once, the dependence of one offer's purchase probability on other offers' prices can be accounted for. For example, the algorithm may choose to set a higher price for one offer in order to incentivize buy-up to another higher revenue offer.

Definition 9 (**Offer Set Price Optimization**)**.** Find the prices $\vec{p^*}$ for each offer O_i in the offer set S that maximize the offer set's total expected net revenue $\mathbb{E}(S)$:

$$\vec{p^*}(S) = \arg\max_{\vec{p}} \mathbb{E}(S) = \arg\max_{\vec{p}} \sum_{O_i \in S} (p_i - c_i) P(O_i | \vec{p}(S)) \tag{3.2}$$

The price optimization uses the following input parameters:

- The list of offers O_i in the offer set S
- For each offer, a function that links price to purchase probability, typically using a probability distribution of customer willingness-to-pay W_i
- For each offer, the cost of providing the services included in the offer c_i

We define the willingness-to-pay distribution for an offer as the sum of the willingness-to-pay distributions for its components $W_i = W_f + \sum_{a_k \in O_i} W_{a_k}$ (the flight and any included ancillary services). While our formulation is general for any probability distribution W, our implementation will assume that all W_f and W_{a_k} are normally distributed. The (scalar) cost of an offer is the sum of the costs of the flight and the ancillaries $c_i = c_f + \sum_{a_k \in O_i} c_{a_k}$.

Both the willingness-to-pay estimates and the costs could vary in the model based on the customer segment, time of booking and revenue management system availability. For instance, the mean willingness-to-pay for the flight could increase toward departure, while the mean willingness-to-pay for a checked bag could decrease, as the demand shifts from advance-booking leisure passengers to close-in bookings by business passengers. The airline could also use different parameters for business and leisure passengers and attempt to identify a booking request as either business or leisure. The resulting optimized offer prices would vary for each customer segment.

The cost of providing the flight c_f is variable and represents the opportunity cost of selling the seat. This is often called the *bid price* and can be computed with a traditional revenue management system. The cost of providing an ancillary service c_{a_k} is assumed to be a fixed amount per ancillary service in our implementation. However, it could be variable when ancillary services are capacity-constrained, such as extra-legroom seating.

3.2.1 Myopic Price Optimization

To quantify and compare the benefits of dynamic offer generation, offer set price optimization and ancillary bundling, we define the baseline pricing strategy as one of myopic price optimization. In this simple approach to pricing, each offer's price is calculated myopically to maximize that offer's expected revenue. As a result, an offer's price is independent of the offer set it is included in. The potential to cannibalize the revenue of the other offers is not considered. In the context of the airline industry, this situation could arise when ancillary pricing is set myopically to maximize ancillary revenue and as a result reduces flight revenue. In our experiments in Chapters 4 and 5, we compare dynamic offer generation to an a la carte pricing strategy, where the ancillary service is priced myopically and independently of the flight.

Definition 10 (**Myopic Price Optimization**)**.** Myopically find the price p^* for an offer O that maximizes the offer's total expected revenue $\mathbb{E}(O)$ net of unit cost c, without consideration of the overall offer set's expected revenue:

$$p^* = \arg\max_p \mathbb{E}(O) = \arg\max_p (p - c)P(O|p) \tag{3.3}$$

A simple numerical example illustrates myopic price optimization: Let the cost of an offer O_1 be $c_1 = \$0$ and the probability of purchase decreases linearly with the price from $P(O_1|\$0) = 1$ to $P(O_1|\$10) = 0$. Then, the myopic optimal price is $p_1^* = \$5$ with an expected net revenue of $\mathbb{E}(O_1) = \$2.5$. Now consider the case where a second offer O_2 is available for purchase in the offer set ($c_2 = \$0$). Its probability of purchase depends on the price of both offer as follows:

$$P(O_2|p_1, p_2) = \begin{cases} 0, if p_1 < \$8 \vee p_2 > \$100 \\ 1, if p_1 \geq \$8 \wedge p_2 \leq \$100 \end{cases} \tag{3.4}$$

Now, if both offers are priced myopically ($p_1 = \$5$ and $p_2 = \$100$), then the customer would not purchase offer O_2. Here myopic price optimization does not maximize total offer set revenue, as the airline could receive \$100 from selling O_2 instead of O_1, if O_1 were priced at $p_1 \geq \$8$.

This simple example illustrates how offer set price optimization differs from myopic price optimization: It pursues to maximize total offer set revenue. In general, an offer's purchase probability depends on the prices of all offers in the offer set: $P(O_i|\vec{p}(S))$. This allows the optimization to consider the effects of sell-up and buy-down, where customers choose between different offers depending on their relative prices. This makes it a more difficult optimization problem than myopic price optimization.

3.2.2 Concurrent Choice Assumption

The dynamic offer generation algorithm uses a concurrent choice assumption to translate customer willingness-to-pay distributions and prices into the purchase probabilities required for the expected net revenue calculation. It assumes customers choose among offers in an offer set according to the concurrent choice model (Bockelie and Wittman, 2018) (Bockelie, 2018). Customers are modeled as rational and fully informed. They attempt to maximize their consumer surplus by balancing the utility of receiving a service with its price.

Under the concurrent choice assumption, passengers are modeled as having a utility (or willingness-to-pay (WTP)) for each flight w_f and ancillary service w_{a_k}, which is drawn from underlying random variables (W_f, W_{a_k}). The WTP for an offer is taken as the sum of its components' WTP: $w_i = \sum_{k \in O_i} w_k$. Passengers then choose among all offers in an offer set based on the difference between their individual WTP w_i and the price p_i.

Definition 11 (**Concurrent Choice Assumption**)**.** Consumers will choose the offer O_i^* within an offer set S with the highest consumer surplus, the difference between their total willingness-to-pay for the offer w_i and its price p_i.

$$O_i^* = \underset{O_i \in S}{\arg\max}\, w_i - p_i \tag{3.5}$$

This choice assumption represents rational consumers who are aware of all offers in the offer set. They select the best offer by weighing their WTP against its price. It is equivalent to the assumptions used in the economic studies of bundling by Adams and Yellen (1976) and Schmalensee (1984). In our adaptation of the choice assumption for the dynamic offer generation algorithm, we assume that the offer set S has the following two properties, which were included in Definition 4:

1. The offer set S always includes a no-purchase option with zero surplus: $S \supseteq O_\emptyset = \{\varnothing\}$ with $w_\emptyset - p_\emptyset = 0$. This ensures that consumers do not choose offers with a negative surplus, which they are not willing to pay for.

2. The flight f is included in all non-empty offers $\forall O_i \in S \setminus O_\emptyset : f \in O_i$. This ensures that no ancillary services can be purchased without the flight.

We will illustrate the concurrent choice assumption first for one flight f and one ancillary service a, before extending it to two ancillary services in Section 3.2.6. As shown in Fig. 3-2, two non-empty offers can be created: $O_f = \{f\}$ and $O_{fa} = \{f, a\}$. The two possible offer sets are: $S_1 = \{O_\emptyset, O_f, O_{fa}\}$ and $S_2 = \{O_\emptyset, O_{fa}\}$. Following Definitions 6 and 7 in Section 3.1, S_1 is an a la carte offer set where O_f is the base offer and S_2 is a bundled offer set where O_{fa} is the base offer. Note that while the

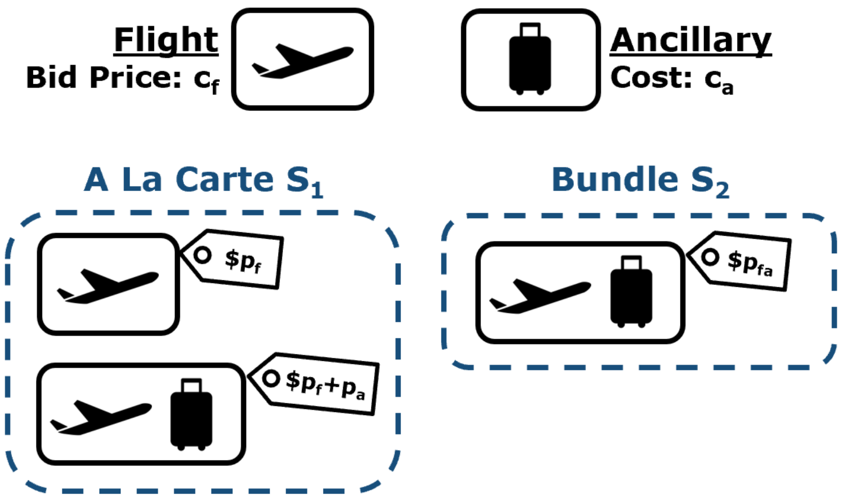

Figure 3-2: Illustration of dynamic offer generation with one ancillary service

offer O_{fa} is present in both offer sets, its price can differ across the offer sets. In S_1, we can split the price of O_{fa} into a flight component p_f determined by the price of offer O_f and an ancillary component p_a. We are interested in the purchase probabilities $P(O_f|p_f, p_a)$ and $P(O_{fa}|p_f, p_a)$. Similarly, in S_2 we would like to find the purchase probability $P(O_{fa}|p_{fa})$ of the offer O_{fa} given its price p_{fa}.

We derive these purchase probabilities with the concurrent choice assumption, assuming the willingness-to-pay distributions for the flight and ancillary service (W_f, W_a) are independent. For a customer presented with S_1, they will choose O_f, if their surplus for O_f is higher than for O_{fa} and positive:

$$\begin{aligned} P(O_f|p_f, p_a) &= P((W_f - p_f > W_f + W_a - p_f - p_a) \wedge (W_f - p_f > 0)) \\ &= P((0 > W_a - p_a) \wedge (W_f > p_f)) \\ &= P(W_a < p_a) \cdot P(W_f > p_f) \end{aligned} \tag{3.6}$$

Conversely, they will choose O_{fa}, if their surplus for O_{fa} is higher than for O_f and positive:

$$\begin{aligned} P(O_{fa}|p_f, p_a) &= P((W_f + W_a - p_f - p_a > W_f - p_f) \wedge (W_f + W_a - p_f - p_a > 0)) \\ &= P((W_a > p_a) \wedge (W_f + W_a > p_f + p_a)) \\ &= P(W_a > p_a) \cdot P(W_f + W_a > p_f + p_a | W_a > p_a)) \end{aligned} \tag{3.7}$$

Since S_2 only contains one offer, passengers will purchase it as long as their surplus for O_{fa} is positive:

$$\begin{aligned} P(O_{fa}|p_{fa}) &= P(W_f + W_a - p_{fa} > 0) \\ &= P(W_f + W_a > p_{fa}) \end{aligned} \tag{3.8}$$

In summary, the purchase probabilities for each offer under the concurrent choice assumption and independently distributed W are:

$$S_1 : \begin{cases} P(O_f|p_f, p_a) & = P(W_a < p_a) \cdot P(W_f > p_f)) \\ P(O_{fa}|p_f, p_a) & = P(W_a > p_a) \cdot P(W_f + W_a > p_f + p_a | W_a > p_a)) \end{cases} \tag{3.9}$$

$$S_2 : \begin{cases} P(O_{fa}|p_{fa}) & = P(W_f + W_a > p_{fa}) \end{cases} \tag{3.10}$$

These expressions link the probability of purchase within an offer set with input WTP distributions and the offer price. By evaluating the probabilities, the offer set price optimization problem can be solved.

We visualize these purchase probabilities in Figure 3-3, which marks the purchase decision as an area on a graph of the two drawn willingness-to-pay variables w_f and w_a. The flight and ancillary prices are marked as lines on the axes, while the combined offer price p_{fa} is marked as a diagonal line. For example, a consumer with low flight and ancillary WTP (compared to the prices) would fall on the lower left-hand corner of the chart and correspondingly choose the no-purchase option. On the other hand, a consumer with both high flight and ancillary WTP would fall on the upper right-hand corner and purchase the combined offer O_{fa}.

If the flight and ancillary WTP were independent and *uniformly* distributed, the

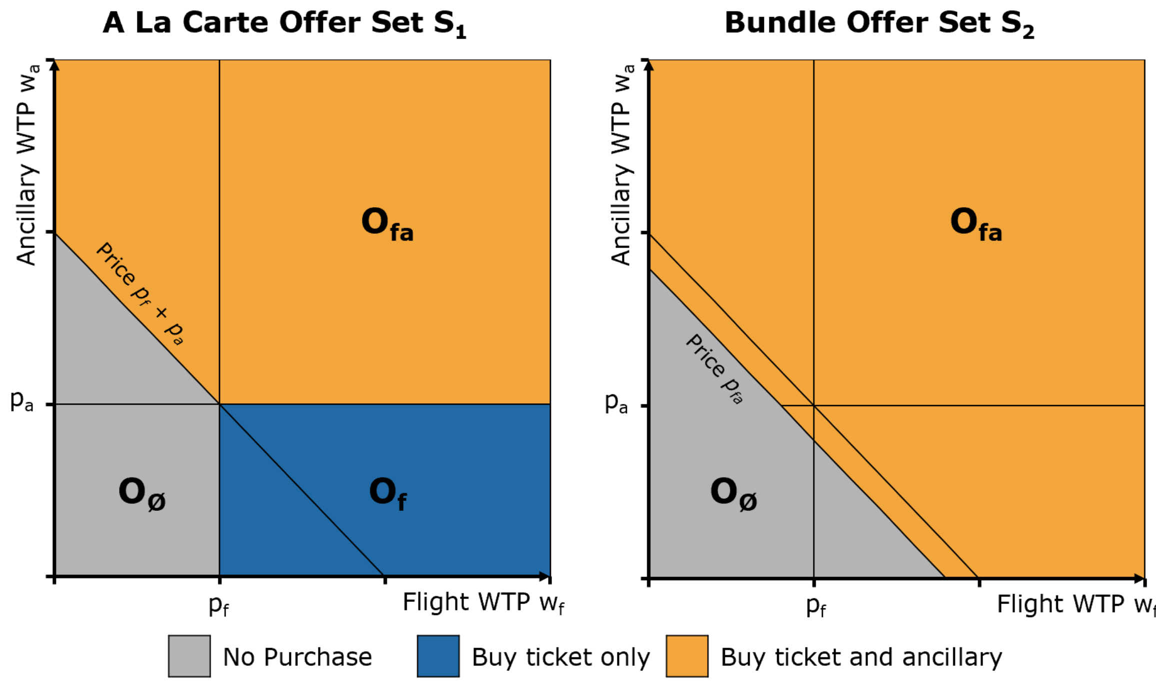

Figure 3-3: Visualization of the concurrent choice assumption

purchase probabilities of each offer in an offer set would correspond to the *area* of the colored sections on the graph. The graph also illustrates that a consumer's purchase decision depends on which offer set they are shown, even if the WTP's are the same. A consumer with $w_f + w_a > p_f + p_a$ and $w_a < p_a$ (in the bottom right-hand corner) purchases only the flight (O_f) when shown S_1, but buys the ancillary too (O_{fa}) when shown a bundle S_2. The graph also hints that, in general, the optimized price of O_{fa} in offer set S_2 (p_{fa}) is lower than the optimized price of the same offer in offer set S_1 ($p_f + p_a$) once the offer set price optimization problem is solved.

While the concurrent choice model is based on common assumptions made in rational choice theory and other existing literature on the economics of bundling, it has limitations. Customer choices cannot always be modeled with a single utility (WTP) function as in rational choice theory. In particular in pricing, psychological effects can influence purchase behavior, which are not reflected in the model. Furthermore, the concurrent choice assumption used in the algorithm has very few input parameters, which can make it difficult to fit the model to real purchase data. The general formulation of the dynamic offer generation problem allows the use of any customer choice model, as long as a conditional probability of purchase $P(O_i|\vec{p}(S))$ can be calculated. As a result, other customer choice models could be explored in further research.

3.2.3 Sequential Approximation

Based on the formulation of the offer set price optimization in Equation 3.2, any offer set with more than one offer requires a multivariate optimization that simultaneously determines the prices of all offers. In initial tests, this led to instabilities where the optimal a la carte ancillary price p_a^* of S_1 could drop to \$0 when bundling was the revenue-maximizing strategy. In effect, the pricing algorithm would make S_1 equivalent to the bundle S_2 by setting $p_f^* = p_{fa}^*$ and thus incentivizing all customers to purchase O_{fa} with the ancillary service. Computational complexity is also a consideration, if the algorithm is to generate real-time offer sets in response to individual booking requests. To simplify and speed up the computations, we approximate the problem using a series of sequential and univariate price optimizations.

The sequential method first determines optional ancillary prices $p_{a_i}^*$ in an offer set using myopic price optimization (Section 3.2.1). The $p_{a_i}^*$ are then used as fixed inputs when performing the offer set price optimization to determine the price of the base offer O_1 of an offer set. The remaining offers in the offer set are priced by summing the price of the base offer $p_{O_1}^*$ and the prices of any additional ancillary prices $p_{a_i}^*$ included in the offer.

In the single ancillary example of Figure 3-2, the sequential approximation is only used to compute prices in offer set S_1: First, p_a is determined myopically as $p_a^* = \arg\max_{p_a}(p_a - c_a)P(W_a > p_a)$. Then, p_f of the flight-only base offer is solved by taking p_a^* as given:

$$p_f^*(S) = \arg\max_{p_f}(p_f - c_f)P(O_f|p_f, p_a^*) + (p_f + p_a^* - c_f - c_a)P(O_{fa}|p_f, p_a^*) \tag{3.11}$$

The remaining offer O_{fa} will then be priced as the sum of $p_f^* + p_a^*$. With this approach, all numerical optimizations are univariate and the global maximum can be found using common optimization methods.

3.2.4 μ-Heuristic

Computing the purchase probabilities under the concurrent choice assumption in Equations 3.9 and 3.10 can present a challenge depending on the probability distributions used for W_f and W_a. In our implementation, we used independently and normally distributed W_f and W_a similar to the work of Schmalensee (1984), for which the following expression from equation 3.9 is difficult to evaluate:

$$P(W_f + W_a > p_f + p_a | W_a > p_a) \tag{3.12}$$

As far as we are aware, there is no closed-form expression for this conditional probability and computing it requires numerical integration. To reduce computational complexity, Berge and Bockelie (2018) developed a heuristic approximation called the μ-Heuristic.

Definition 12 (**μ-Heuristic**)**.** Equation 3.12 can be approximated by replacing a random variable W with a scalar that represents the mean of the truncated normal distribution $\mu_{TR} \coloneqq \mathbb{E}[W_a | W_a > p_a]$.

$$P(W_f + W_a > p_f + p_a | W_a > p_a) \approx P(W_f + \mu_{TR} > p_f + p_a)$$

For any normally distributed $W = N(\mu, \sigma^2)$, there exists a closed-form expression for μ_{TR} based on the standard normal PDF $\phi(z)$ and standard normal CDF $\Phi(z)$:

$$\mu_{TR} \coloneqq \mathbb{E}[W | W > p] = \mu + \sigma \cdot \frac{\phi(\frac{p-\mu}{\sigma})}{1 - \Phi(\frac{p-\mu}{\sigma})}$$

For normally distributed W, the μ-Heuristic enables all purchase probabilities to be calculated through standard normal functions without requiring numerical integration.

3.2.5 Sensitivity of Optimized Offer Prices

To gain an initial understanding of the behavior of the model, we perform a sensitivity test on the input parameters of the DOG algorithm with one ancillary service. We study the optimized, unbounded offer prices that are generated by the price optimizer with the sequential approximation and μ-Heuristic. As shown in Figure 3-2, three prices are computed by the algorithm:

p_f^***:** The flight price in the a la carte offer set S_1

p_a^***:** The price for the optional ancillary service in S_1

p_{fa}^***:** The combined bundle price for both flight and ancillary service in offer set S_2

DOG generates these prices using four input variables, two of which are related to the flight and two of which are related to the ancillary service:

W_f**:** The estimated flight WTP distribution for the flight

c_f**:** The opportunity cost of using the flight capacity (bid price)

W_a**:** The estimated ancillary WTP distribution

c_a**:** The cost of providing one unit of the ancillary service

We optimize the three prices for different levels of input parameters to observe how stable the optimization process is. We also gain insights into how the offer prices vary with input parameters. First, we study the sensitivity of p_f^*, p_a^* and p_{fa}^* to W_f and c_f. We test pairs of flight-related input parameters on a grid, where the mean flight WTP $E[W_f]$ ranges from \$10 to \$400 and the bid price c_a ranges from \$0 to \$200. In all cases, we hold the ancillary-related parameters constant at $c_a = \$20$ and $W_f \sim \mathcal{N}(\$25, \$7.5^2)$, which will form the baseline for our simulations in Chapter 4. Then, we optimize p_f^* and p_{fa}^* for pairs of ancillary-related parameters, where both $E[W_a]$ and c_a range from \$0 to \$70. Here, we hold the flight-related parameters constant at $c_f = \$50$ and $W_f \sim \mathcal{N}(\$200, \$60^2)$. In all cases, W_f and W_a are normally

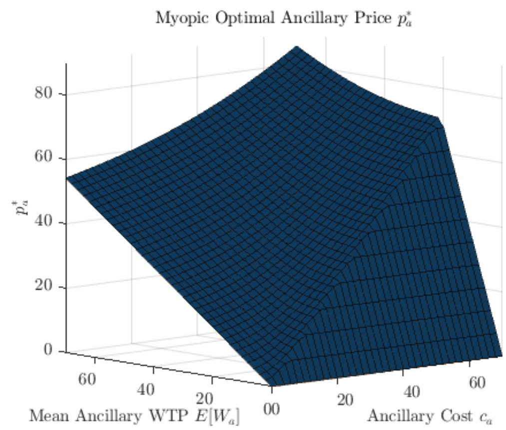

Figure 3-4: Sensitivity of myopic optimal ancillary prices p_a^* to inputs W_a,c_a

distributed and as the mean changes, its standard deviation scales as 30% of the mean.

The first price computed according to the sequential approximation is the ancillary price p_a^*, which is myopically chosen to maximize ancillary net revenue. Figure 3-4 shows how the optimal ancillary price varies with both the ancillary cost c_a and the ancillary WTP distribution W_a. p_a^* does not depend on the flight-related parameters. We observe that at zero ancillary cost, p_a^* scales linearly with $E[W_a]$. As c_a increases, the ancillary price also increases non-linearly along a mostly smooth surface. There is a discontinuity and a sharp drop in p_a^* in the region of very high c_a and low $E[W_a]$. In general, the myopic price optimization is well-behaved and intuitive, with the price increasing with both WTP and cost.

Next, we observe how the flight price p_f^* and bundle price p_{fa}^* respond to both flight-related input parameters (Figure 3-5a) and ancillary-related input parameters (Figure 3-5b). Both p_f^* and p_{fa}^* increase with both the flight WTP $E[W_f]$ and the bid price c_f. The increase is nearly linear in regions where the WTP much higher than the bid price. Given that passengers are expected to be willing to pay more for the bundle that includes the ancillary service, p_{fa}^* is consistently more expensive than the flight price p_f^*. This is also evident when considering the sensitivity to ancillary-related inputs, where the flight price is nearly independent of W_a and c_a, but the bundle price p_{fa}^* shows a strong dependence on both parameters. Intuitively, when

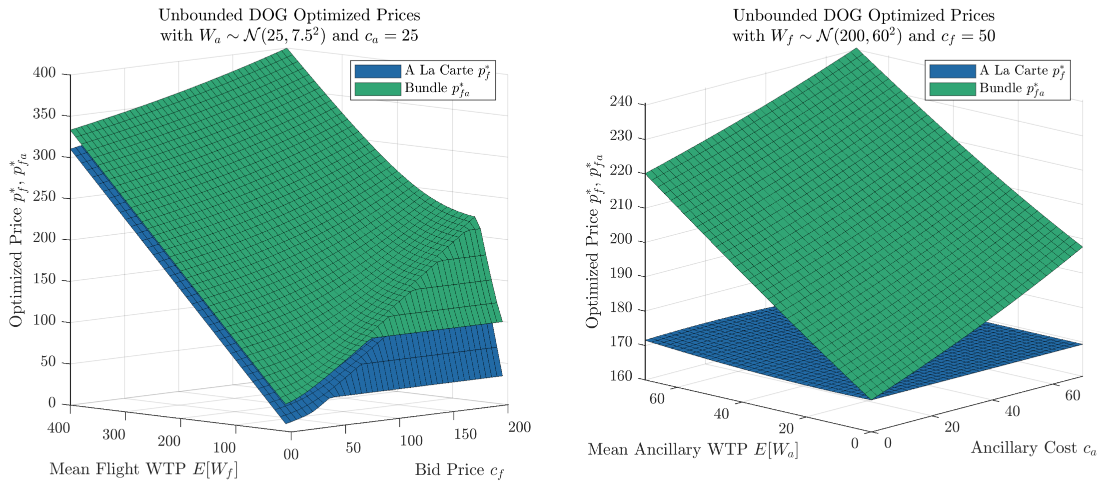

(a) Flight-related parameters W_f, c_f (b) Ancillary-related parameters W_a, c_a

Figure 3-5: Sensitivity of DOG optimized offer prices p_f^* and p_{fa}^*

$W_a = c_a = \$0$, both prices are identical.

These tests have given us insight into how the algorithm prices offers across a variety of input parameters. The results are intuitive and reasonable, boosting confidence in the model. Furthermore, the algorithm can find optima across a wide range of input parameters.

3.2.6 Extension to Two Ancillary Services

The offer set price optimization with concurrent choice assumption, sequential approximation and μ-Heuristic, can be applied to optimize prices for offer sets with more than one ancillary services using the same principles. This is an important property, as airlines have expanded their ancillary service offerings and require an algorithm that can price a multitude of ancillary offers. In this section, we provide an example of how the algorithm would be extended to two ancillary services.

In Figure 3-6, we show four offer sets that can be generated with two ancillary services. While other offer sets are possible, we deem these the most commercially relevant, as they each have a base offer that is either only the flight (*a la carte*) or also includes bundled ancillaries. We call the four different offers present in the offer sets $O_f, O_{fa1}, O_{fa2}, O_{fa12}$ according to their components.

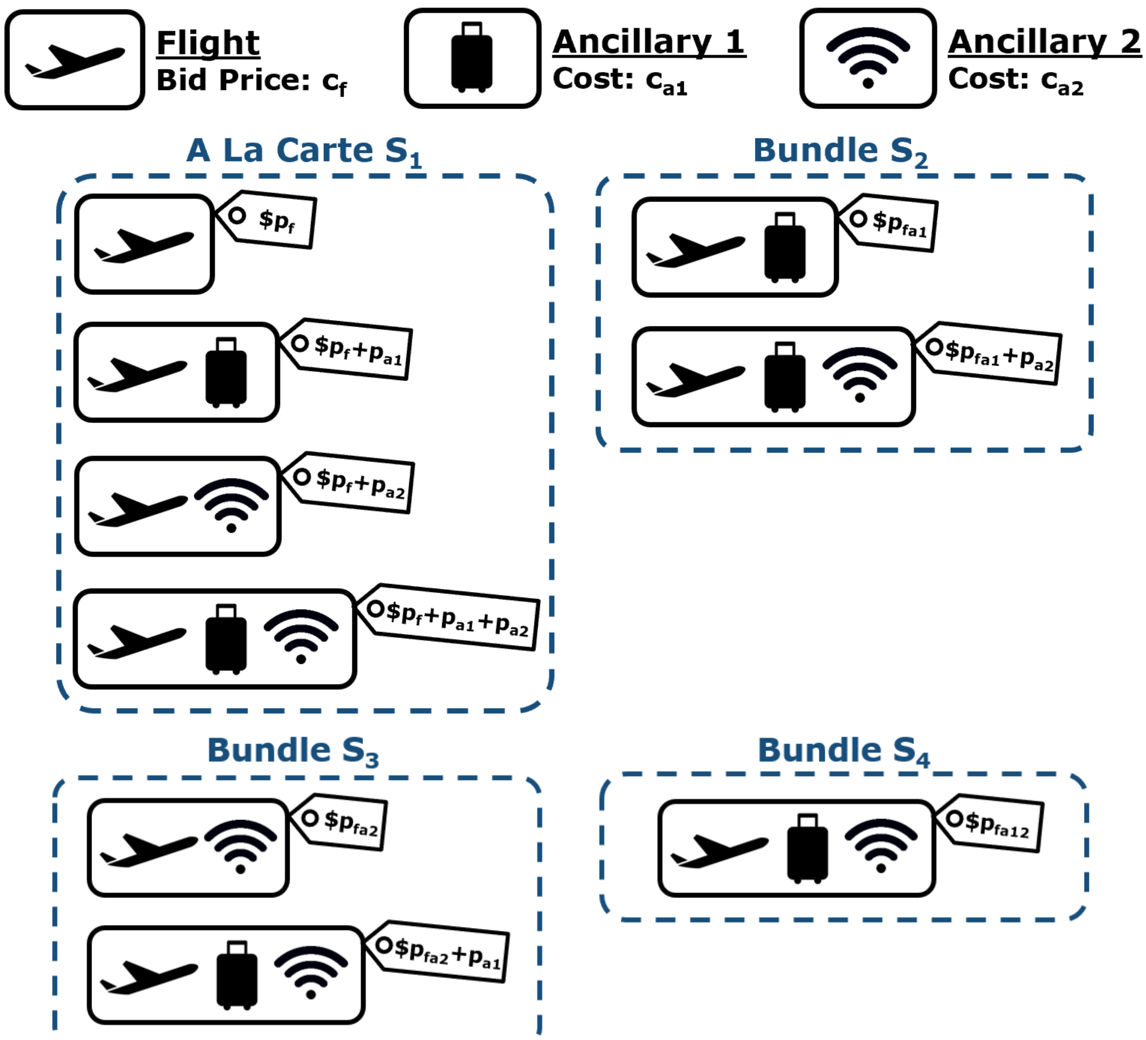

Figure 3-6: Example of dynamic offer generation with two ancillary services showing four possible offer sets with different base offers

In each of the offer sets, one of the offers acts as the base offer O_1, and all other offers are a superset of the base offer. These additional offers can be viewed as ancillaries offered as an optional add-on alongside the base offer at an incremental price p_{a_i}. Note that it will always be possible to break down two offers with two different prices p_1, p_{12} into the price of a base offer p_1 and an incremental price $p_{12} = p_1 + p_2$. However, we have artificially imposed an additional constraint in the *a la carte* offer set S_1 that the incremental price of offer O_{fa12} over the base offer is the sum of the incremental prices of offers O_{fa1} and O_{fa2}. As a result, no discount is given to customers for purchasing both ancillary services.

Applying the concurrent choice assumption to these offer sets, we yield the following purchase probabilities:

$$S_1 \begin{cases} P(O_f|p_f, p_{a1}, p_{a2}) & = P((W_{a_1} < p_{a1}) \wedge (W_{a_2} < p_{a2}) \wedge (W_f > p_f)) \\ P(O_{fa1}|p_f, p_{a1}, p_{a2}) & = P((W_{a_1} > p_{a1}) \wedge (W_{a_2} < p_{a2}) \wedge (W_f + W_{a_1} > p_f + p_{a1})) \\ P(O_{fa2}|p_f, p_{a1}, p_{a2}) & = P((W_{a_1} < p_{a1}) \wedge (W_{a_2} > p_{a2}) \wedge (W_f + W_{a_2} > p_f + p_{a2})) \\ P(O_{fa12}|p_f, p_{a1}, p_{a2}) & = P((W_{a_1} > p_{a1}) \wedge (W_{a_2} > p_{a2}) \wedge (W_f + W_{a_1} + W_{a_2} > p_f + p_{a1} + p_{a2})) \end{cases} \tag{3.13}$$

$$S_2 \begin{cases} P(O_{fa1}|p_{fa1}, p_{a2}) & = P((W_{a_2} < p_{a2}) \wedge (W_f + W_{a_1} > p_{fa})) \\ P(O_{fa12}|p_{fa1}, p_{a2}) & = P((W_{a_2} > p_{a2}) \wedge (W_f + W_{a_1} + W_{a_2} > p_{fa1} + p_{a2})) \end{cases} \tag{3.14}$$

$$S_3 \begin{cases} P(O_{fa2}|p_{fa2}, p_{a1}) & = P((W_{a_1} < p_{a1}) \wedge (W_f + W_{a_2} > p_{fa2})) \\ P(O_{fa12}|p_{fa2}, p_{a1}) & = P((W_{a_1} > p_{a1}) \wedge (W_f + W_{a_1} + W_{a_2} > p_{fa2} + p_{a1})) \end{cases} \tag{3.15}$$

$$S_4 \begin{cases} P(O_{fa12}|p_{fa12}) & = P(W_f + W_{a_1} + W_{a_2} > p_{fa12}) \end{cases} \tag{3.16}$$

According to the sequential approximation, our implementation would calculate the incremental prices p_{a1}, p_{a2} first for each offer set using myopic price optimization (Section 3.2.1): $p_a^* = \arg\max_{p_a}(p_a - c_a)P(W_a > p_a)$. With these optimized p_a^*, the offer set price optimization becomes univariate for each offer set and the μ-Heuristic (Section 3.2.4) for independent and normally distributed W is used to simplify the purchase probability calculations.

The most complex expression, $P(O_{fa12}|p_f, p^*_{a1}, p^*_{a2})$, is approximated with two truncated means $\mu_{TR1} = \mathbb{E}[W_{a_1}|W_{a_1} > p^*_{a1}]$ and $\mu_{TR2} = \mathbb{E}[W_{a_2}|W_{a_2} > p^*_{a2}]$ as follows:

$$\begin{aligned}
P((W_{a_1} > p^*_{a1}) \wedge (W_{a_2} > p^*_{a2}) \wedge (W_f + W_{a_1} + W_{a_2} > p_f + p^*_{a1} + p^*_{a2}))& \\
= P(W_{a_1} > p^*_{a1}) \cdot P(W_{a_2} > p^*_{a2})& \\
\cdot P(W_f + W_{a_1} + W_{a_2} > p_f + p^*_{a1} + p^*_{a2}|(W_{a_1} > p^*_{a1}) \wedge (W_{a_2} > p^*_{a2}))& \\
\approx P(W_{a_1} > p^*_{a1}) \cdot P(W_{a_2} > p^*_{a2}) \cdot P(W_f + \mu_{TR1} + \mu_{TR2} > p_f + p^*_{a1} + p^*_{a2})& \quad (3.17)
\end{aligned}$$

The number of possible offer sets expands quickly with the number of ancillaries in the problem. While the general formulations of the offer set price optimization and offer set selection problem can be extended to multiple ancillaries, the scale of the problem becomes evident. In particular, the use of the concurrent choice assumption when calculating the purchase probabilities increases the computational complexity. We have introduced approximations to keep the problem tractable, yet the implementation becomes increasingly complex for a large number of ancillary services. This is an area where further research is required to simplify the algorithm for more than one ancillary service and rationalize the number of offer sets evaluated.

3.3 Offer Set Selection

Once all the offer prices are calculated using the algorithm in Section 3.2, the offer set selection problem can be solved to complete the dynamic offer generation algorithm. In offer set selection, the combination of offers that are expected to generate the highest net revenue is selected from all possible offer sets. The chosen set of offers can then be distributed by the airline to the final customer. In this section, we first formulate the offer set selection algorithm. Using our example of DOG with one ancillary service, we then show when each offer set (S_1 or S_2) is selected by the algorithm.

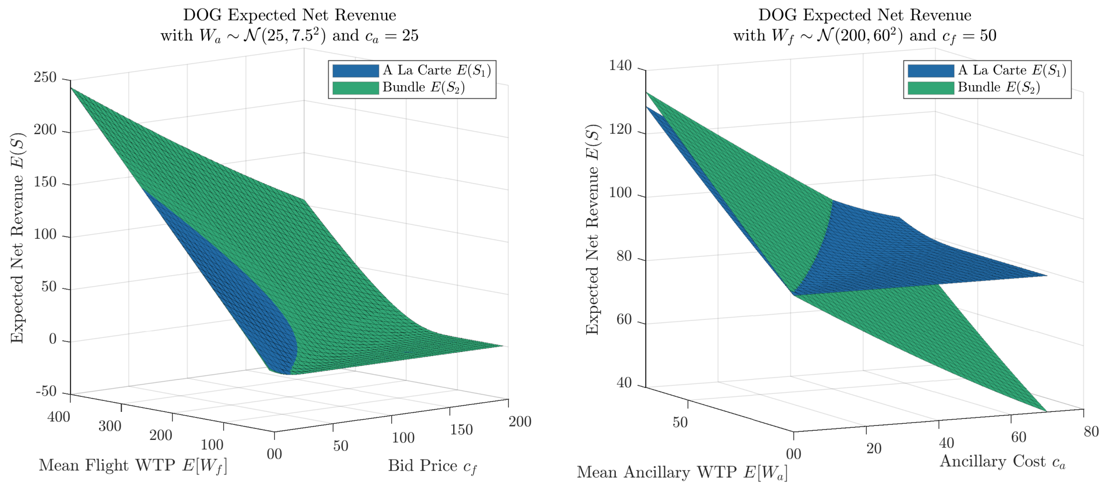

(a) Flight-related parameters W_f, c_f (b) Ancillary-related parameters W_a, c_a

Figure 3-7: Sensitivity of DOG offer set expected net revenue $\mathbb{E}(S)$

Definition 13 (Offer Set Selection). Find the offer set S with the highest total expected net revenue $\mathbb{E}(S)$ given each offer set's optimized prices $\vec{p^*}(S)$:

$$S^* = \arg\max_S \mathbb{E}(S) = \arg\max_S \sum_{O_i \in S} (p_i^* - c_i) P(O_i | \vec{p^*}(S)) \tag{3.18}$$

Given each offer set's optimal prices $\vec{p^*}(S)$ from the offer set price optimization problem, the expression above is to be evaluated for all offer sets. In this step, we use the same purchase probabilities based on a concurrent choice assumption and μ-heuristic. The offer set with the highest expected revenue is then shown to the customer.

We examine the offer set selection process by continuing the sensitivity tests on the one-ancillary case from Section 3.2.5. We use the optimized prices p_f^*, p_a^*, p_{fa}^* calculated there and compute the expected net revenue of each offer set according to equations 3.9 and 3.10. DOG will choose to offer either an a la carte offer set S_1 or the bundled offer set S_2, depending on which has the highest expected net revenue under the concurrent choice assumption. For the same range of flight-related and ancillary-related input parameters as before, the expected net revenue per passenger is represented by a colored surface in Figure 3-7.

In Figure 3-7a, we observe that the flight-related input parameters only have a very small effect on the offer set selection decision. While the expected net revenue depends heavily on the flight-related inputs, it is very similar for the two offer sets S_1 and S_2. The difference between the two surfaces is very small. Nonetheless, we observe that when both W_f and c_f are low, there is a region when the algorithm would offer the a la carte offer set. In the remaining regions, the bundle is the preferred offer set given the ancillary parameters of this experiment.

The ancillary-related parameters have a sizable influence on the offer set selection decision (Figure 3-7b). When the ancillary WTP is higher than the cost, bundling generates higher net revenue. On the other hand, when the ancillary service is expensive to provide but does not have a broad appeal (high WTP), it is better to offer it a la carte. This illustrates the potential risk/reward trade-off of bundling, where the incremental net revenue gain from bundling is small compared to the potential revenue loss of bundling the wrong ancillary service.

The offer set selection process does not currently incorporate this trade-off. The offer set with the highest expected net revenue may not always be the best choice in reality, especially if there are uncertainties and estimation errors in the input parameters used in the model. Instead, an offer set with a lower expected net revenue could be a more attractive choice, provided its expected net revenue is less susceptible to uncertainties in the input parameters. By incorporating a notion of robustness, one could account for uncertainties in the estimated input parameters. This could be achieved by establishing a range of input parameters and selecting the offer set based on the average expected net revenue within the uncertainty range.

The figure shows that there is a nonlinear decision boundary, where both offer sets have the same $E(S)$. This boundary is relatively consistent across a variety of tested W_f and c_f parameters and can be linearly approximated by the following rule of thumb:

$$\text{If } W_a \sim \mathcal{N}(E[W_a], (0.3E[W_a])^2)\text{, then: } \mathbb{E}(S_2) > \mathbb{E}(S_1) \iff E[W_a] \geq 1.25 \cdot c_a \tag{3.19}$$

The high dependence of the offer set selection process on ancillary-related variables compared to flight-related variables has implications for dynamic offer generation: For a segment of passengers with static ancillary-related parameters, the offer set selection decision is not very *dynamic*, meaning that they are generally offered the same offer set throughout the booking window, independent of the flight price. As a result, it is even more important for airlines to segment their passengers into groups with different ancillary WTP. They also need to be able to classify a booking request into the correct segment. This allows dynamic offer generation to show more relevant offer sets and achieve the revenue benefits of dynamic bundling.

3.4 Integration with Revenue Management Systems

Our experimental results in Chapters 4 and 5 come from simulations of dynamic offer generation working alongside a traditional airline revenue management system in a competitive market. Here, we illustrate how the dynamic offer generation model might be integrated with traditional revenue management systems, before elaborating on the mechanics of the Passenger Origin-Destination Simulator (PODS) used to test the model in Chapter 4.1.

The dynamic offer generation model cannot directly replace existing revenue management optimizers, as it lacks an important component: RM optimizers use network optimization to calculate the opportunity cost of selling a seat on a flight leg (*bid price*) for a given forecast of future demand and willingness-to-pay, which the dynamic offer generation model assumes as a given input (cost of providing the flight c_f).

Because DOG can generate prices on a continuous spectrum and do not select from a limited set of possible price points (fare classes), it represents a departure from traditional airline pricing and RM systems. In an ideal setting, DOG could be integrated directly with an RM optimizer that generates continuous flight prices and does not rely on fare classes. However, such classless RM algorithms are still under active development and have not been implemented yet. We now show that DOG can also be used alongside a traditional class-based revenue management system, where

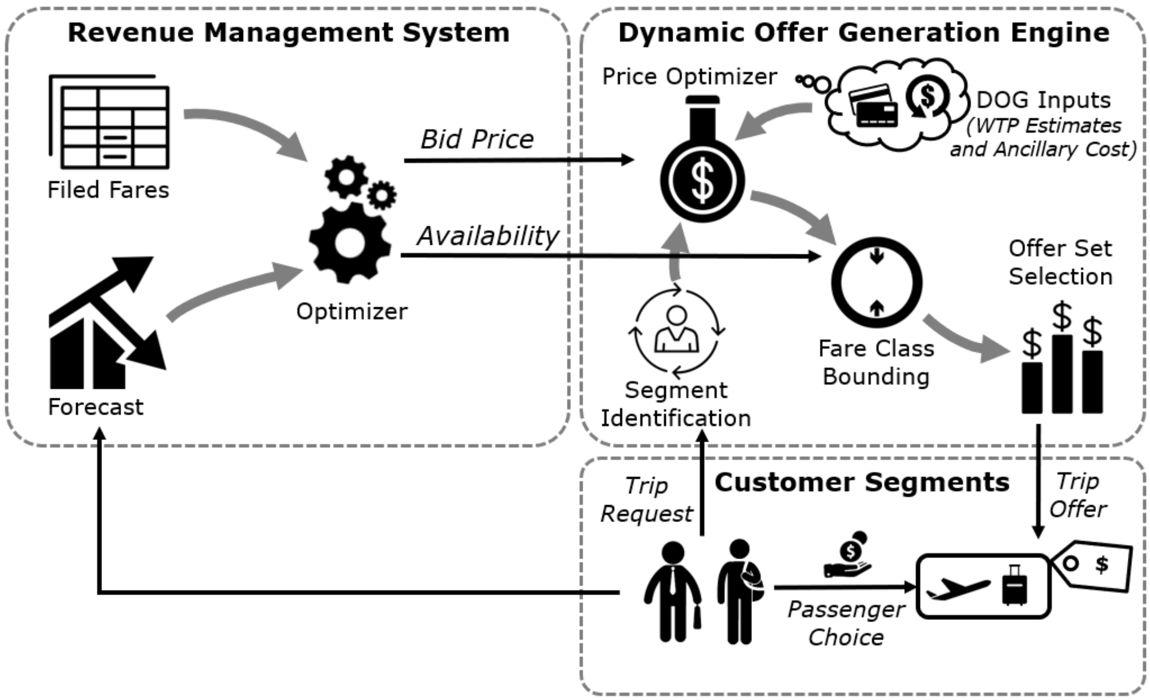

Figure 3-8: Integration of dynamic offer generation with RM systems

airlines can leverage existing RM infrastructure to deliver dynamic offers.

We use the traditional revenue management system to calculate a bid price, select a pre-defined flight price and set a range of permitted price deviations from the pre-defined price point. Dynamic offer generation then uses the bid price as c_f to determine the final flight and ancillary prices within the established price range, as well as the offer set shown to each customer segment. In this capacity, dynamic offer generation acts as a dynamic price adjustment mechanism for the flight price as defined by Wittman and Belobaba (2019), which has been extended to include ancillary pricing and offer set selection.

An illustration of this integration is shown in Figure 3-8. At every re-optimization point, the revenue management system recalculates the *bid price* and the number of seats available at each filed fare based on the current demand forecast. This could occur for an individual flight segment with a leg-based RM system, or an entire network of flights with a network RM system. In the dynamic offer generation engine, the remaining inputs (estimated flight and ancillary WTP distributions, as well as ancillary unit costs) are specified for each point in time and customer segment. This allows us to show different offers and/or prices to different customer segments. In our simulations, each booking request's customer segment is correctly identified with a certain probability (identification accuracy) and the offer set price optimization is solved using the identified segment's inputs and the current RM bid price. In practice, a separate customer segmentation model would have to be trained to classify the booking request and deliver the appropriate WTP inputs.

Currently, the DOG algorithm only generates a single flight price p_f, and as such it is most compatible with an unrestricted fare structure and corresponding RM system, where all passengers are assumed to purchase the lowest available fare. Because DOG computes different prices than the traditional RM system, airlines will likely want to reconcile the two systems and limit the deviation of DOG prices from the RM system's prices. In our implementation, we use the existing RM system to determine a range of allowed prices that the DOG prices are bounded to. These bounded DOG offer prices $\overline{p_i}$ are marked with an overbar.

Let us call F_0 the filed fare of the lowest available fare class in a traditional class-based RM system, F_{+1} the filed fare of the next higher (more expensive) fare class and F_{-1} the filed fare of the next lower (unavailable) fare class. If there F_0 is the lowest (cheapest) fare class of the airline, we set $F_{-1} = F_0$ and if F_0 is the highest fare class, then $F_{+1} = F_0$. In general, it always holds that: $F_{-1} \leq F_0 \leq F_{+1}$.

We propose and implement a bounding rule, where the flight price in the *a la carte* offer set p_f^* is bounded first to be within a range of F_0 determined by the bounding parameter b. When $b = 0$, no deviation from the RM system's price is permitted $\overline{p_f} = F_0$. When $b = 1$, the flight price is allowed to be adjusted by up to one fare class away from F_0.

$$\overline{p_f} = \min[\max[p_f^*, F_0 - b(F_0 - F_{-1})], F_0 + b(F_{+1} - F_0)] \tag{3.20}$$

The price of all other offers p_{fa_i} that include one or more ancillary services in addition to the flight, including those in *bundle* offer sets, will be bounded based on

$\overline{p_f}$ to maintain the relative price difference $p^*_{fa_i} - p^*_f$ of the original unbounded prices.

$$\overline{p_{fa_i}} = \overline{p_f} + (p^*_{fa_i} - p^*_f) \tag{3.21}$$

This bounding rule aims to preserve the pricing structure and minimize the impact on the offer set selection decision, i.e. the difference between the expected net revenues of the offer sets, while acknowledging that the bounded prices no longer maximize each offer set's expected net revenue. To illustrate this, we employ a numerical example with one ancillary service, where the DOG input parameters are $W_f \sim \mathcal{N}(\$200, \$60^2)$, $c_f = \$50$, $W_a \sim \mathcal{N}(\$25, \$7.5^2)$ and $c_a = \$20$. The corresponding optimized prices before bounding from the offer set price optimization are $p^*_f = \$169.72$, $p^*_a = \$27.41$ and $p^*_{fa} = \$192.78$. Now with $F_{-1} = \$200, F_0 = \$250, F_{+1} = \$300$ and $b = 1$, the flight price will be bounded to $\overline{p_f} = \$200$. The flight and ancillary bundle price will be bounded to $\overline{p_{fa}} = \$200 + \$23.06 = \$223.06$.

Figure 3-9 shows the offer set expected net revenue $\mathbb{E}(S_1)$ and $\mathbb{E}(S_2)$ at various bounded flight prices $\overline{p_f}$. At all times, the bundle price is bounded to be \$23.06 above $\overline{p_f}$. The offer set expected net revenue for both S_1 and S_2 are very similar, indicating that this numerical example lies on the decision boundary where both offer sets perform equally well. As the prices deviate from their optimum values, the expected net revenue for both decreases simultaneously. However, the difference between the two $\Delta\mathbb{E}(S)$ changes very little, by at most \$0.40. It is this difference that determines whether the a la carte offer set ($\Delta\mathbb{E}(S) \leq 0$) or the bundle ($\Delta\mathbb{E}(S) > 0$) is selected. $\Delta\mathbb{E}(S)$ is much more sensitive to the main DOG input parameters, as we have shown in Section 3.3. This suggests that the impact of bounding on the offer set selection decision is small, unless the chosen parameters are close to the decision boundary $\Delta\mathbb{E}(S) = 0$, as is the case in this numerical example.

As shown in Figure 3-8, the offer set selection is performed with the *bounded* prices to determine which offer set is shown to the customer. The bounded prices $\overline{p}$ are used as p^* in Equation 3.18. In the final step of the interactions between DOG and the traditional RM system, all offer purchases are reported back to the RM

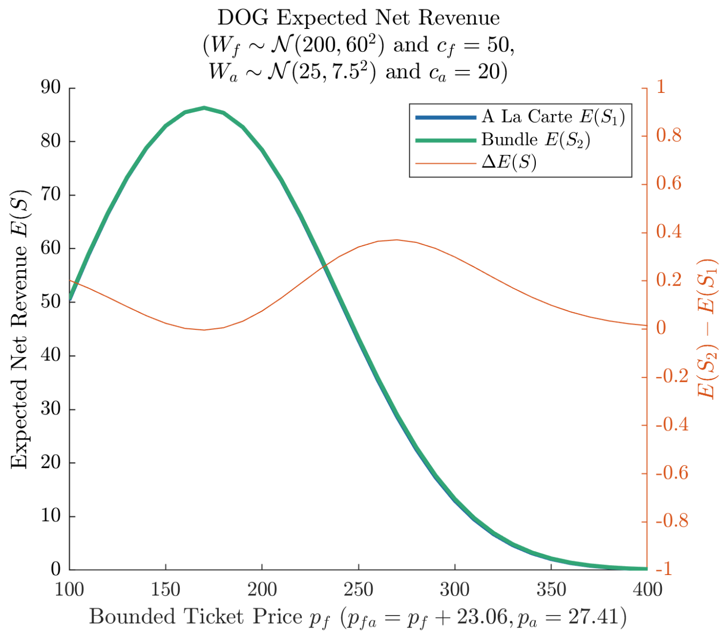

Figure 3-9: Impact of price bounding on offer set expected net revenue

system and recorded in the lowest available fare class at its filed fare F_0. As opposed to directly recording the flight price paid $\overline{p_f}$, this solution is more compatible and easier to integrate with traditional class-based forecasters and optimizers (we deem this simplification reasonable for $b \leq 1$).

3.5 Summary

In this chapter, we introduced a new heuristic to solve the dynamic offer generation problem. For any offer set with one flight and up to two ancillary services, we provided formulations that allow the algorithm to compute optimized offer prices under the concurrent choice assumption. We simplified the computations required using the μ-Heuristic and sequential approximation. We then explained how the assortment of offers can be optimized by choosing the set of offers that result in the highest expected revenue.

Using numerical tests, we showed that the algorithm produces optimized offer prices reliably across a wide spectrum of input parameters. The results were intuitive to interpret. Similarly, the outcomes of the offer set selection process are intuitive as well and heavily depend on the balance between ancillary WTP and ancillary cost. We discussed that bundling can improve revenue under the assumptions of the concurrent choice model, but that it carries a higher risk of revenue loss especially when the input parameters used carry estimation errors.

Finally, we outlined how our algorithm could be integrated by airlines alongside a traditional airline revenue management system. The algorithm relies on the bid prices generated by the RM system to accurately estimate the opportunity cost of selling a seat. It uses a relatively simple bounding heuristic to limit the deviation of DOG prices from the original RM system fares. DOG bookings are recorded back in the historical bookings database and used to generate the demand forecast for the RM system. However, in this approach the conditional willingness-to-pay estimates used by DOG are static inputs. We did not study how these would be estimated from historical bookings. We also did not discuss in detail how booking requests can be accurately segmented based on their search request or what the potential impacts of adjusting the prices after RM optimization are in terms of the potential feedback effects this may generate in the RM system.

In the following Chapters 4 and 5, we apply the dynamic offer generation algorithm in a competitive airline revenue management simulation and examine its impact on airline bookings, fares and revenues under hypothetical but more realistic market conditions.

Chapter 4

Baseline Results of Dynamic Offer Generation in PODS

In this chapter, we present experimental results for the dynamic offer generation algorithm introduced in Chapter 3. We introduce the Passenger Origin-Destination Simulator (PODS) in Section 4.1, which we use to simulate the combination of DOG with a traditional revenue management system in a competitive airline network.

We define a set of baseline parameters for a large network with four airlines using traditional RM systems (4.1.3). This sets up the subsequent experiments that compare the performance of DOG against the traditional RM baseline.

Section 4.2 presents the results from the simulations. We first test the two components of DOG individually: dynamic flight price optimization (4.2.1) and dynamic ancillary bundling (4.2.2). In Section 4.2.3, we combine these components and show the first PODS simulation results of a full implementation of dynamic offer generation using baseline parameters.

We close this chapter with a summary of results and initial conclusions about the potential benefits of using DOG and ancillary bundling under the baseline parameters (4.3). The experimental studies continue in Chapter 5 with extensive sensitivity testing of the algorithm, to determine whether the observed results are applicable across a broad spectrum of input parameters.

4.1 The Passenger Origin-Destination Simulator

The Passenger Origin-Destination Simulator (PODS) is an advanced revenue management simulator, originally developed in the 1990s at Boeing Commercial Airplanes and more recently by PODS Research LLC. Today, it is used by the MIT-PODS Revenue Management Consortium to test new airline revenue management strategies. Dynamic offer generation as presented in Chapter 3 has been implemented in PODS as part of the PODS Consortium's research program.

4.1.1 Overview of PODS Architecture

In this section, we provide a summary of the basic architecture of PODS. For a more detailed description of the PODS simulator and its models, see works by Gorin (2000), Cléaz-Savoyen (2005) and Carrier (2008).

The simulator consists of two primary components:

- **Airlines** schedule flights between airports and sell tickets across their network in each market (origin-destination pair) that they serve. Each airline uses a revenue management system to determine fare availability on its flights, which is based on a demand forecast generated from the airline's historical booking observations. Each airline operates independently from its competitors and uses a separate revenue management system.

- **Passengers** are individually generated and have a demand for travel from an origin to a destination in a specific time window. They evaluate all available itineraries across all airlines and book a flight that best satisfies their willingness-to-pay, fare option preference and schedule sensitivity. Their sensitivity parameters, called *disutility* in PODS, are individually drawn for each passenger from an underlying probability distribution, which can vary by passenger type (business or leisure, for example). Similarly, their maximum willingness-to-pay for the flight is also individually drawn and determines the threshold at which passengers decide not to travel (*no-go*).

The functionality of PODS has been extended beyond its core itinerary and fare choice model over the years. For dynamic offer generation, two extensions are particularly relevant:

- **Dynamic Price Adjustment**: Traditional revenue management optimizers in PODS set the number of seats available for sale in each fare class. Each fare class has a pre-determined (*filed* or *published*) price. PODS has a Dynamic Price Adjustment capability based on research on continuous pricing by Wittman and Belobaba (2018), which allows an airline to dynamically adjust the flight price away from the pre-determined price determined by the revenue management optimizer to any price within a continuous range determined by bounds. This adjustment could vary for each customer segment with a different willingness-to-pay (i.e. business or leisure). Probabilistic Fare-Based Dynamic Availability (PFDynA) is one example of a heuristic to optimize the adjusted continuous price for a customer segment, which is implemented in PODS (Wittman, 2018).

- **Ancillary Choice**: In PODS, airlines can offer ancillary services in addition to the flight. The passengers choose among itineraries, fare classes and ancillary services according to an integrated passenger choice model first introduced by Bockelie and Belobaba (2017). Besides the sequential and simultaneous choice behaviors described in the paper, passengers can also behave according to concurrent choice behavior, in which they compare the flight and ancillary prices rationally against their separately drawn flight and ancillary willingness-to-pay. This behavior matches the concurrent choice assumption (cf. Chapter 3.2.2) used in the DOG pricing algorithm exactly. At the end of the simulation, PODS reports ancillary purchases, ancillary revenues and costs for each airline separately.

These two PODS extensions are the foundation for our implementation of dynamic offer generation, as illustrated in Figure 3-8 (Chapter 3.4). The airlines use dynamic price adjustment to adjust the pre-determined (filed) fares to the DOG offer prices based on the identified customer segment. The passengers use concurrent ancillary

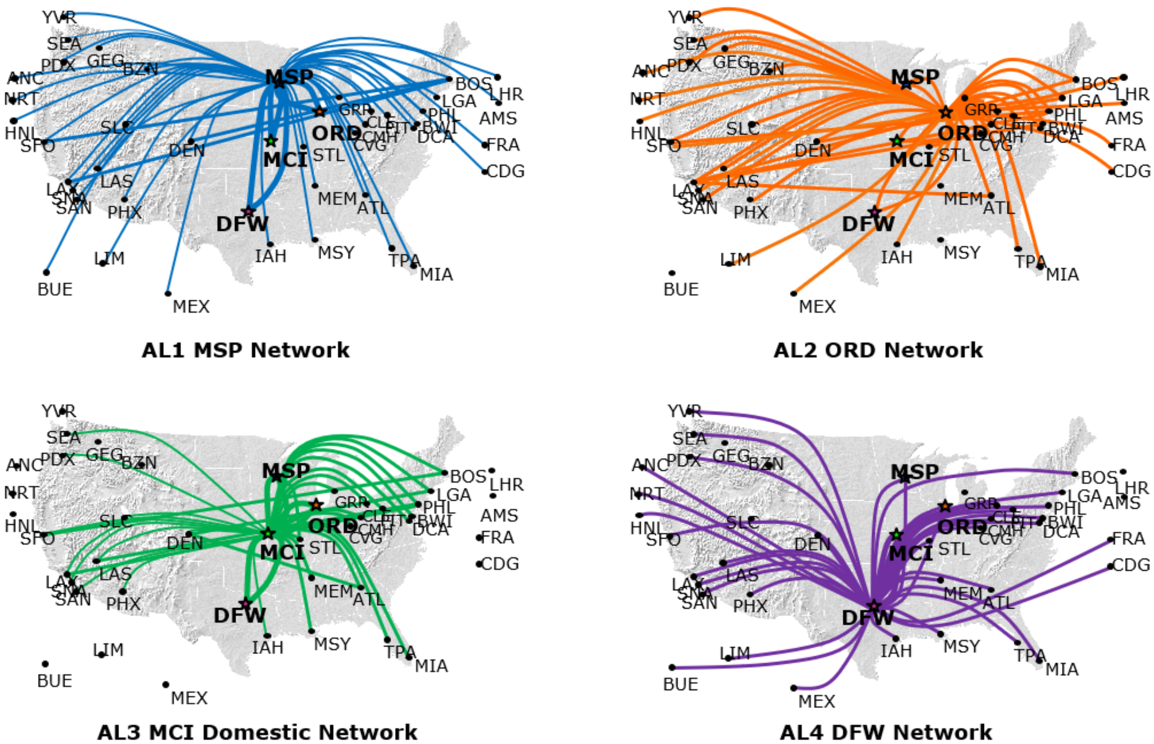

Figure 4-1: Route network of the four airlines in PODS Network U10

choice behavior to choose which itinerary and which offer to purchase.

4.1.2 PODS Network U10

All of our simulations are conducted in Network U10, the largest and most realistic network currently simulated with PODS, with four competing airlines operating in both short-haul domestic and long-haul international markets. Each airline primarily operates a hub-and-spoke network, offering connections over a single hub airport. In aggregate, the four airlines offer 442 daily flights across 572 origin-destination market pairs, with multiple frequencies on each route spread across the day. The route network is illustrated in Figure 4-1.

The demand and supply in the network are both one-directional, with all flights flowing from origins west of the hub to destinations east of the hubs. In total, there are 4 hubs and 40 spoke cities served. Each airline offers 10 different fare classes (and corresponding price points) in each O-D market pair. The price of each fare class (the filed fare) for a market pair is identical across all airlines, representing a market in which competitors match each other's prices. These fares were derived from published fares provided by an airline when the network was first designed. The average fare ratio F_i/F_Q between the fare in a fare class i and the lowest fare in the market F_Q is shown in Figure 4-2. The airlines' revenue management systems allocate the number of seats to sell in each fare class. This leads to differences in quoted prices across airlines, as one airline may close a low fare class in response to a high demand forecast, while another airline still sells seats in that fare class.

In all of our studies for this thesis, all four airlines use a fully unrestricted fare structure in all markets. This means that the fare rules are identical for all fare classes 1-10, with the only difference being the price. None of the fares have any advance purchase requirements either: they can be purchased at any point before departure as long as there are seats available in that class. As a result, all passengers naturally desire and purchase the lowest available fare class. The airlines are effectively offering only a single price point in every market at any point in time. This is an attractive case for continuous pricing applications, as it allows us to use the DOG algorithm from Chapter 3 to determine a single optimized flight price $\overline{p_f}$ for each passenger segment.

Demand is generated for each O-D market pair in the form of individual passengers. These passengers can come from one of two segments: business or leisure. Each O-D market pair has a different mix of business and leisure demand, with the proportion of business demand generated in a market averaging 37% with a minimum of 15% and a maximum of 74%. The passengers are generated continuously during the booking window: In short-haul markets, 35% of business demand and 78% of leisure demand is generated more than three weeks before departure. In long-haul markets, the demand arrives slightly earlier and 54% of business demand and 85% of leisure demand arrives more than three weeks before departure.

The passengers from each segment draw their schedule and price sensitivity from a probability distribution. Most relevant to our study of dynamic offer generation are the maximum flight and ancillary willingness-to-pay (WTP), which are described by the random variables X_f and X_a. While in practice, passenger's *conditional WTP*

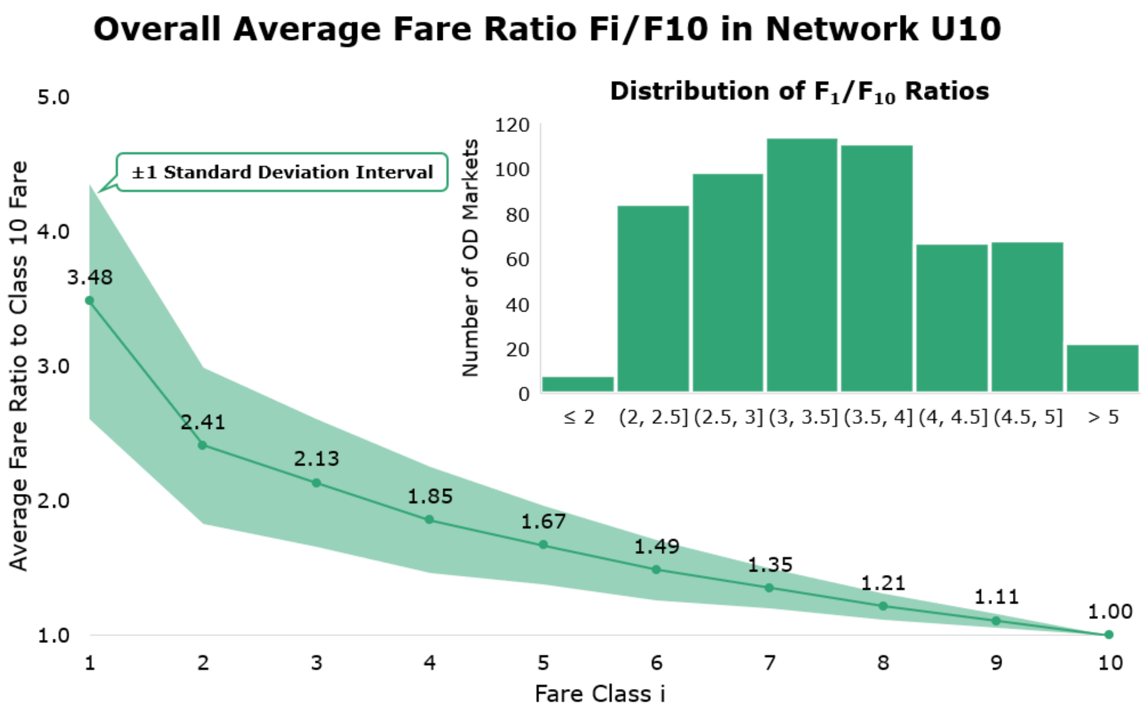

Figure 4-2: Comparison of average fares across fare classes

for an offer is dependent on competitor's offers and prices, the maximum WTP is the limit at which passengers would rather choose not to travel (*no-go*) or not to purchase an ancillary service. During the simulation, the underlying distributions of X_f and X_a are unknown to the airlines in PODS.

4.1.3 PODS Baseline Simulation

In this section, we outline the baseline settings used in the PODS simulator for our experimental results on dynamic offer generation. A summary of all simulation settings we use in our baseline is provided at the end of this section in Table 4.1.

In all experiments, we restrict ourselves to the dynamic offer generation problem with one ancillary service as illustrated in Figure 3-2. Each airline incurs a variable cost of $20 for every ancillary service sold in the baseline. To explore a wider range of parameters, we also conduct sensitivity tests on individual input parameters around these baseline settings.

In our experiments in Network U10, the simulated passengers draw X_f from an exponential distribution $Exp(\lambda)$. The distribution is scaled for each O-D market pair based on the lowest fare F_Q in the market, which is the filed fare of the tenth fare class:

$$Business : X_f/(2.5 \cdot F_Q) - 1 \sim Exponential(ln(2)/0.6); X_f \in [2.5 \cdot F_Q, \infty) \quad (4.1)$$

$$Leisure : X_f/F_Q - 1 \sim Exponential(ln(2)/0.5); X_f \in [F_Q, \infty) \quad (4.2)$$

This distribution has the following properties:

- 100% of all business passengers are willing to pay at least 2.5 times the lowest fare in the market F_Q
- 100% of all leisure passengers are willing to pay at least F_Q
- 50% of all business passengers have $X_f \geq 4.0 \cdot F_Q$. $[X_f/(2.5 \cdot F_Q) - 1 \geq 0.6]$
- 50% of all leisure passengers have $X_f \geq 1.5 \cdot F_Q$. $[X_f/F_Q - 1 \geq 0.5]$

The resulting probability that a random passenger can afford a certain flight price p_f is shown in Figure 4-3 for both passenger segments. The average network-wide fares in U10 for each fare class are overlaid. The spacing between fare classes corresponds to the fare ratios from Figure 4-2. The graph shows that the maximum WTP of business passengers is high compared to the fares, such that they can purchase most fare classes. On the other hand, leisure passengers are very price sensitive and the airlines can stimulate additional demand by making lower fare classes available for purchase.

The maximum ancillary WTP X_a is drawn from a normal distribution that is identical for both business and leisure passenger segments in the baseline case. The mean $E[X_a]$ is \$25 and the coefficient of variation is 30%.

$$X_a \sim \mathcal{N}(\$25, (0.3 \cdot \$25)^2) \quad (4.3)$$

In the simulation, the passengers will use their drawn WTP X_f and X_a to decide which

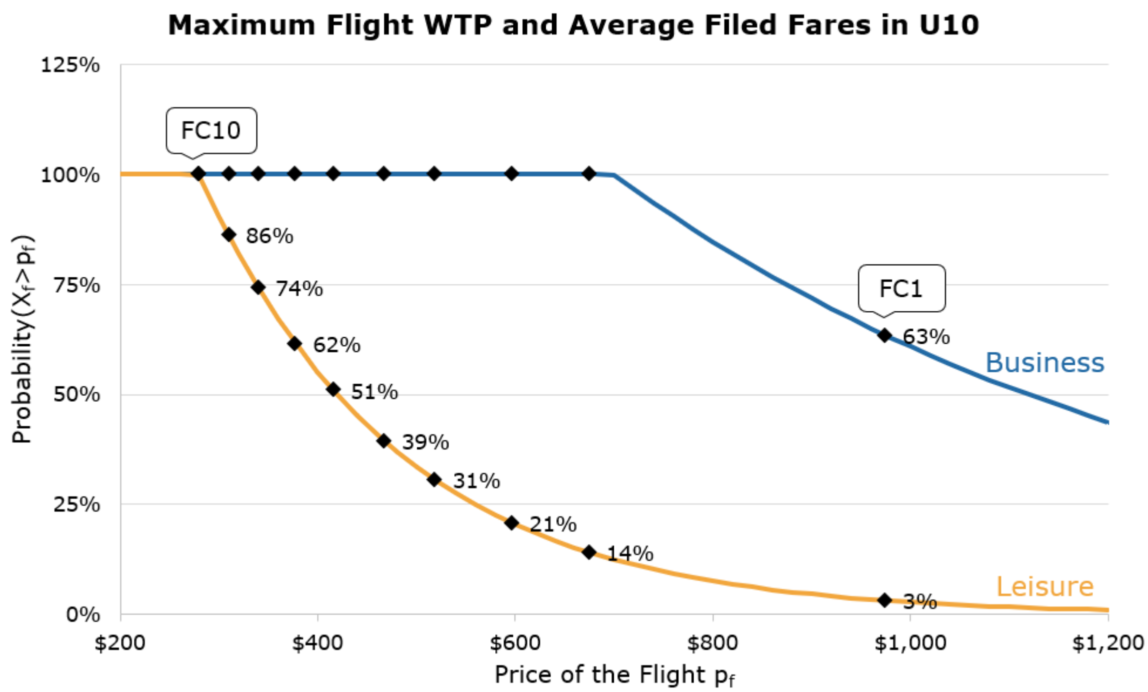

Figure 4-3: Comparison of cumulative flight WTP curves with average fares

offer to purchase according to concurrent choice behavior as described in Chapter 3.2.2.

Since the airlines do not know the underlying maximum WTP distributions X_f and X_a, they use estimated conditional WTP distributions W_f and W_a in Equations 3.9 and 3.10 to perform the DOG price optimization under competition. We use normally distributed W_f and W_a in our tests, so that the μ-Heuristic can be applied (Chapter 3.2.4). In our baseline, we use the following distributions with a coefficient of variation of 30% of the mean:

$$Business : W_f \sim \mathcal{N}(3.0F_Q, (0.3 \cdot 3.0 \cdot F_Q)^2) \tag{4.4}$$

$$Leisure : W_f \sim \mathcal{N}(1.2F_Q, (0.3 \cdot 1.2 \cdot F_Q)^2) \tag{4.5}$$

$$W_a \sim \mathcal{N}(\$25, (0.3 \cdot \$25)^2) \tag{4.6}$$

We compare the estimated conditional flight WTP W_f with the true maximum flight WTP X_f in Figure 4-4 using the U10 network-wide average $F_Q = \$280$. The so-

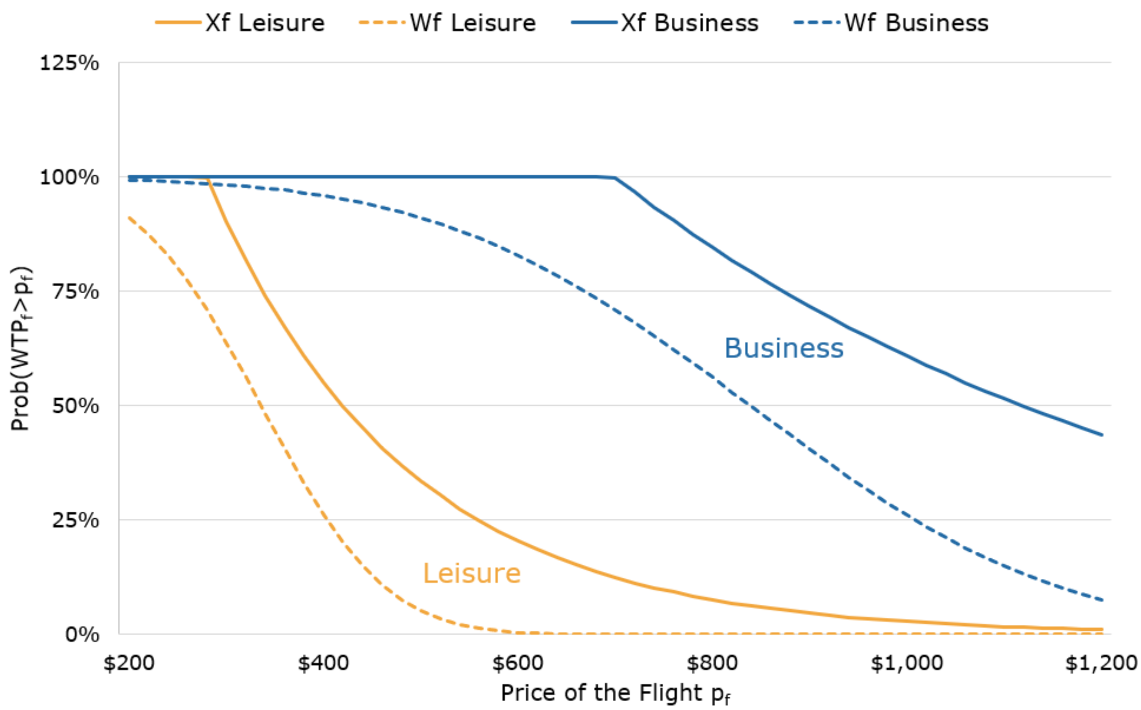

Figure 4-4: Comparison of estimated (W_f) and true (X_f) cumulative flight WTP

called *Q multipliers* of 3.0 and 1.2 used to determine W_f were chosen because they maximized revenue among a set of PODS simulations with the following Q multipliers tested: {2.0, 2.5, 3.0, 3.5} for business and {1.0, 1.2, 1.4} for leisure. In reality, an airline's choice of W_f would likely be dictated by their pricing power in the market and their competitors' prices. We observe that the estimated conditional WTP is lower than the true maximum flight WTP for both passenger segments, which reflects a competitive market where none of the airlines can charge fares corresponding to the passengers' maximum flight WTP.

Note that $W_a = X_a$, meaning that all airlines accurately estimate the true ancillary WTP distribution of the passengers. Whenever an airline does not use DOG in our tests, they are in the *No DOG* baseline state: They always present the *a la carte* offer set, in which the flight is priced by a traditional RM system at the lowest available filed fare without dynamic price adjustment and the ancillary service is priced at the myopic optimized price. Using W_a from Equation 4.6 to calculate $P(O|p) = P(W_a \geq p_a)$ in Equation 3.3, we obtain a myopic optimal ancillary price of

$p_a^* = \$27.41$ with an ancillary cost of $c_a = \$20$.

In the baseline, the passenger segments are identified with perfect accuracy (100%). This means that the airlines always use the appropriate estimated WTP W_f to optimize the DOG prices for business and leisure passengers in response to a booking request. We also perform sensitivity tests with lower than 100% identification accuracy to reflect more realistic customer segmentation. In our tests, the optimized DOG prices are bounded as described in Equations 3.20 and 3.21 with $b = 1$, meaning that the flight price can be dynamically adjusted up or down one fare class.

Since the focus of our studies is the performance of the DOG algorithm, we will not change the settings of the traditional RM system used by the airlines. All four airlines use a network RM system for unrestricted fare structures using concepts described in Chapter 2.1.2. Each airline uses O-D path-based sell-up *Q forecasting* based on an exponential conditional WTP described by a *FRAT5* curve, which is appropriate for the fully unrestricted fare structure used. Rejected demand when a flight has sold out completely is unconstrained in the forecast using booking curve detruncation.

Seat inventory control is performed using the Displacement Adjusted Virtual Nesting (DAVN) heuristic (Smith and Penn, 1988) with the marginal revenue transformation (Fiig et al., 2010). On each flight leg, all of the connecting O-D fares are displacement adjusted by deducting the bid price of the other legs from the fare, which is obtained using a deterministic linear program on the network. The marginal revenue is also deducted from the displacement adjusted fare to account for buy-down in unrestricted fare structures. The fully adjusted fares are then nested into virtual classes by sorting them by revenue value. Then, the leg-based EMSRb heuristic computes seat protection levels on each leg for the virtual classes. The optimization is re-run at every time frame during the booking window. The time frames are longer at the beginning of the booking window and progressively shorten towards departure, leading to more frequent RM re-optimizations.

Table 4.1 summarizes our chosen experimental baseline. In our experiments, we primarily focus on the impacts of airlines switching from a traditional RM methodology to dynamic offer generation in an unrestricted fare structure. We study the following scenarios:

- **No DOG:** All four airlines use traditional RM with filed fares and an a la carte offer set where the ancillary is always optional and priced at the myopic optimal ancillary price p_a^*.

- **AL1 Only DOG:** Airline 1 switches to dynamic offer generation and segmented pricing. DOG chooses to offer either the a la carte or bundle offer set. Airlines 2-4 retain their traditional RM system and a la carte offer set. This scenario is used to study the benefit of DOG with asymmetric competition and its first-mover advantage.

- **All ALs DOG:** All airlines use dynamic offer generation. This scenario is used to study the benefit of DOG with symmetric competition.

4.2 Baseline Dynamic Offer Generation Simulation Results

In this section, we present our first results of dynamic offer generation in PODS simulations. In 4.2.1, we first study the impact of the DOG pricing algorithm for the a la carte offer set, a mode we call *dynamic a la carte* (DALC). All airlines sell the ancillary service optionally but begin to perform segmented continuous pricing for the flight. These experiments are similar to studies of Probabilistic Fare-based Dynamic Adjustment by Wittman (2018). In 4.2.2, we then study the impact of *dynamic bundling* (DB) without continuous pricing. All offers are priced such that the flight price remains the filed fare without dynamic price adjustment. The DOG algorithm compares the expected revenue of the bundle and a la carte offer sets and decides when to offer dynamic bundles instead of the a la carte offer set. These experiments provide insights about the revenue impact of bundling alone without the competitive effects of changing the flight prices. Finally, in 4.2.3 we present the first results of full *dynamic offer generation* (DOG) with the baseline settings. We use the insights from

Category	Input parameter	Baseline setting
Simulation	Network	PODS network U10 with four airlines
	Fare structure	No fare restrictions & No advance purchase requirements
Passengers	Segments	Two passenger types (business and leisure)
	Max flight WTP	$X_f/(2.5F_Q) - 1 \sim Exp(ln(2)/0.6)$ for business $X_f/F_Q - 1 \sim Exp(ln(2)/0.5)$ for leisure
	Max ancillary WTP	$X_a \sim \mathcal{N}(\$25, (0.3 \cdot \$25)^2)$ for both segments
	Choice behavior	Concurrent choice behavior
Forecaster	Method	Q forecasting for each O-D path
	Detruncation	Booking curve detruncation
	Sell-up	Sell-up based on FRAT5c curve
Optimizer	Method	DAVN O-D Network RM (Displacement Adjusted Virtual Nesting)
	Frequency	Reoptimized every simulated timeframe
	Fare adjustment	Marginal revenue transformation (based on FRAT5c sell-up)
DOG	Estim. flight WTP	$W_f \sim \mathcal{N}(3.0F_Q, (0.3 \cdot 3.0 \cdot F_Q)^2)$ for business $W_f \sim \mathcal{N}(1.2F_Q, (0.3 \cdot 1.2 \cdot F_Q)^2)$ for leisure
	Estim. ancillary WTP	$W_a \sim \mathcal{N}(\$25, (0.3 \cdot \$25)^2)$ for both segments
	Passenger segment identification accuracy	100% correct segment identification
	Bounding	$b = 1$ (± 1 fare class from lowest available)
Ancillary	Unit Cost of Provision	\$20

Table 4.1: Summary Table of PODS Baseline Simulation Parameters

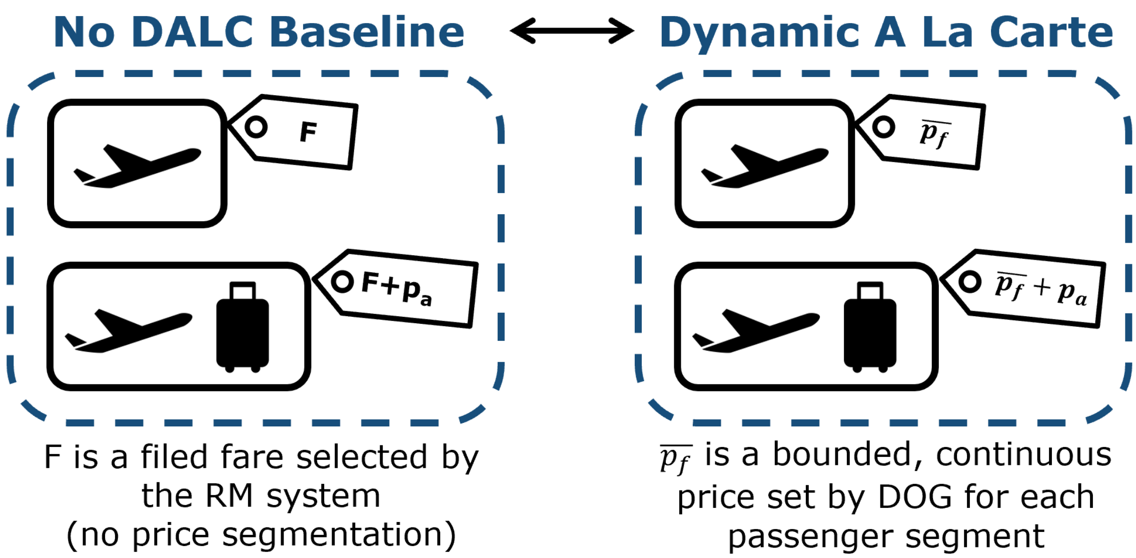

Figure 4-5: Comparison of the offer set in the baseline (No DALC or No DOG) and dynamic a la carte (DALC)

DALC and DB to explain the effects of DOG on total revenue, load factor, prices and bookings.

4.2.1 Dynamic A La Carte Results

In *dynamic a la carte* (DALC), we test dynamic offer generation without offer set selection. All customers are always shown the a la carte offer set, in which the flight is priced using the DOG algorithm at $\overline{p_f}$, and the ancillary service is available optionally for an additional p_a. Just like with DOG, p_a is set to be the myopic optimal price according to the sequential approximation (cf. Chapter 3.2.3). We compare DALC experiments to the standard baseline without dynamic offer generation, where all airlines sell the ancillary services optionally at the myopic optimal price. As a result, no changes are made to the ancillary price in these experiments and we are focused on generating revenue through gains by dynamically adjusting the flight price $\overline{p_f}$. Figure 4-5 summarizes the differences between the baseline and DALC. We show results from both an asymmetric case where only Airline 1 uses DALC and a symmetric case where all four airlines use DALC.

Figure 4-6 and Table 4.2 summarize the overall results. With our baseline set-

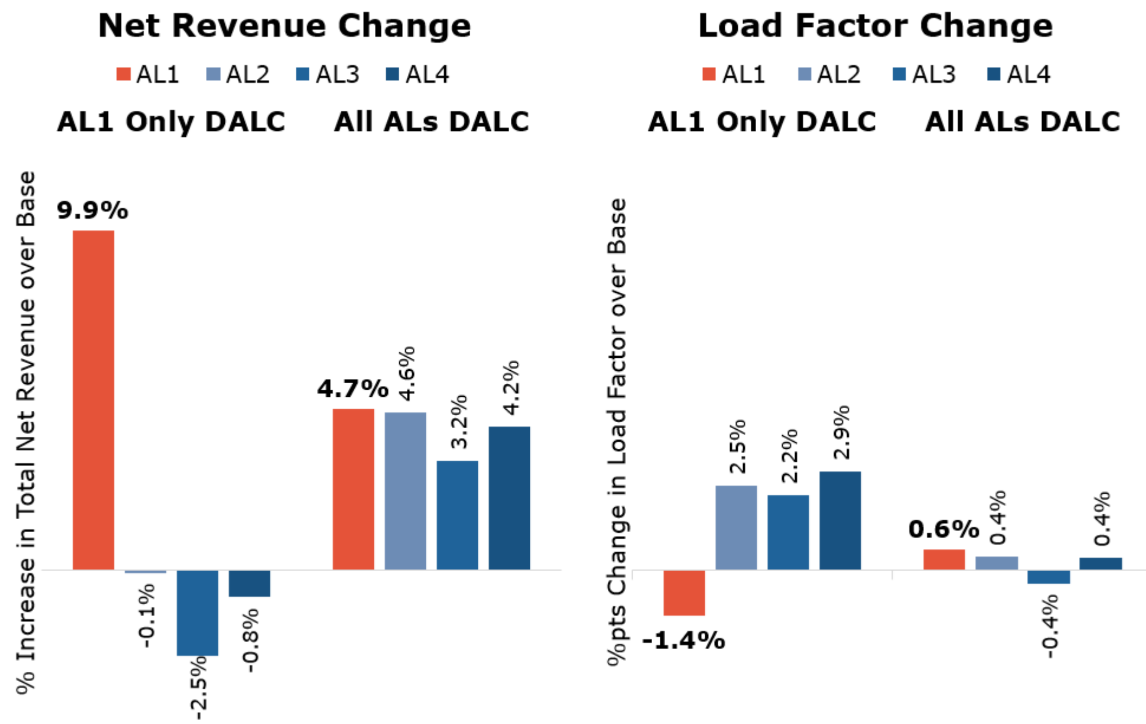

Figure 4-6: Overview of dynamic a la carte (DALC) net revenue gains and load factor changes in symmetric and asymmetric tests

Airline 1	**Base: No DALC**	**AL1 Only DALC**	**All ALs DALC**
Flight Revenue	$2,920,160	$3,211,120	$3,058,163
Ancillary Revenue	$90,937	$90,322	$92,544
Ancillary Cost	$66,354	$65,704	$67,555
Total Net Revenue	$2,944,743	(+**9.9%**) $3,235,738	(+**4.7%**) $3,083,152
RPM	22,889,814	22,506,11	23,059,985
ASM	28,588,660	28,588,660	28,588,660
Net Yield	$0.1286	$0.1438	$0.1337
Load Factor	80.07%	(**-1.4pts**) 78.72%	(+**0.6pts**) 80.66%

Table 4.2: Overall simulated results of dynamic a la carte (DALC) for Airline 1

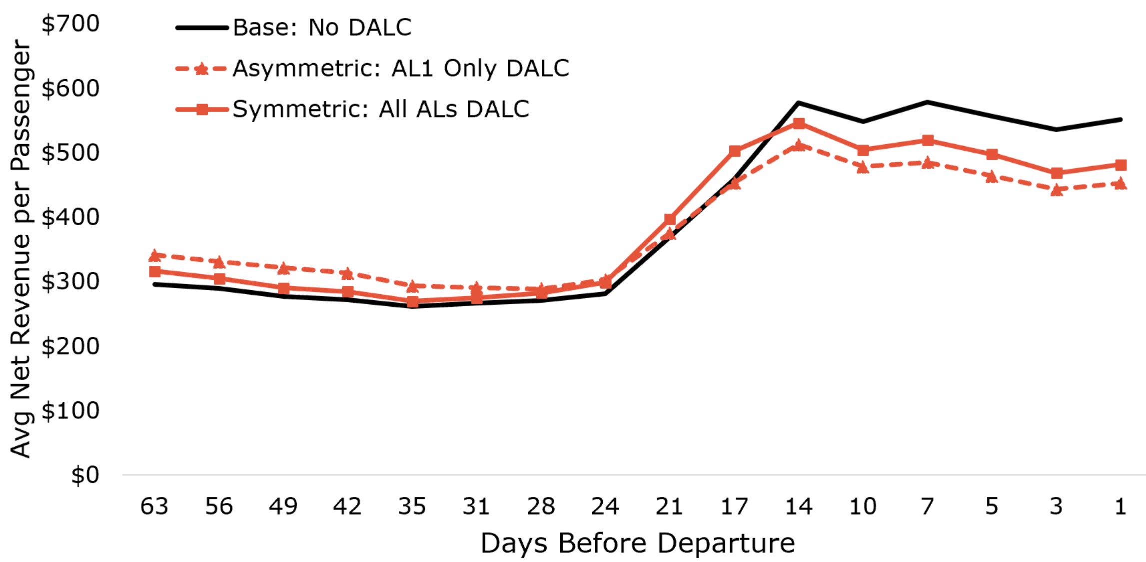

Figure 4-7: Airline 1 average revenue per passenger by timeframe before departure

tings, we observe strong revenue gains for airlines that employ the DOG algorithm to dynamically price the flight (only). When Airline 1 exclusively uses DALC, it observes an up to 10% net revenue increase with a -1.4pts load factor decrease. In the symmetric experiment when all airlines use DALC, all airlines observe a 3-5% increase in net revenue. We define net revenue as total flight and ancillary revenue net of ancillary cost of provision. This matches the objective of the DOG algorithm, which attempts to find offer prices that maximize an offer set's expected net revenue (Equation 3.1).

These net revenue gains are very large and specific to the chosen baseline test parameters and the simulation environment. To explain the source of the revenue gains, we first examine the average fares paid by passengers in Figure 4-7. We observe that the adjusted fares are different than the baseline: The DOG prices are adjusted downward (discount) in the last two weeks before departure and adjusted higher (increment) early on in the booking process. This is a result of the chosen normal WTP distribution W_f, which is constant in time for each passenger segment. If the RM bid price remains fixed, the resulting DOG optimized price will also be constant in time. As the RM system closes lower fare classes towards departure, the lowest available filed fare increases in line with the baseline curve. Whenever the lowest

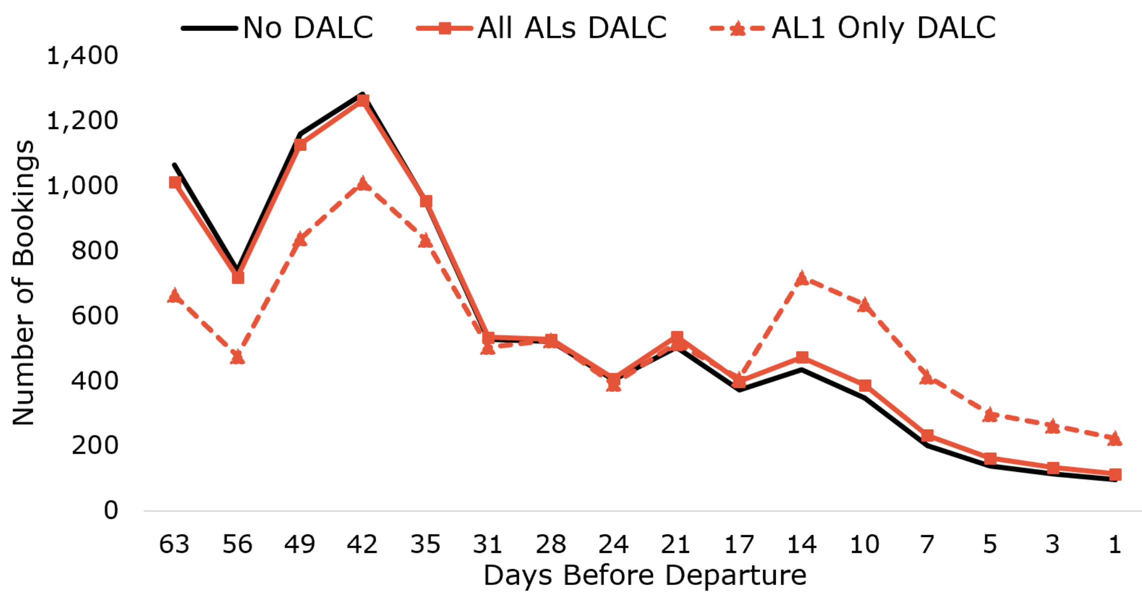

Figure 4-8: Number of bookings observed by Airline 1 in each timeframe before departure

available filed fare is below the DOG optimized price (early in the booking process), DOG will adjust the fare upward closer to the optimized price, subject to bounds on how far the fare can be adjusted. When the lowest available filed fare is higher than the DOG optimized price (in the last two weeks of the booking process), DOG will adjust the fare downward.

The observed change in average fare (or net revenue per passenger) is stronger when Airline 1 asymmetrically uses DALC than when all airlines symmetrically use DALC. This is because in the asymmetric case, Airline 1 is consistently undercutting its competitors, who are offering baseline filed fares, in the last two weeks before departure and thus receiving a lot more bookings during this period, lowering the average fare per passenger. Because of the additional close-in demand observed, the RM system makes fewer seats available in lower classes early on in the booking process, which increases Airline 1's average fares in those periods. This is a first indication of the competitive effects and revenue management feedback that can be observed in the PODS simulator. However, it is unlikely that revenue gains of this magnitude can be achieved with DALC in reality, since competitors could respond by matching the lower close-in prices.

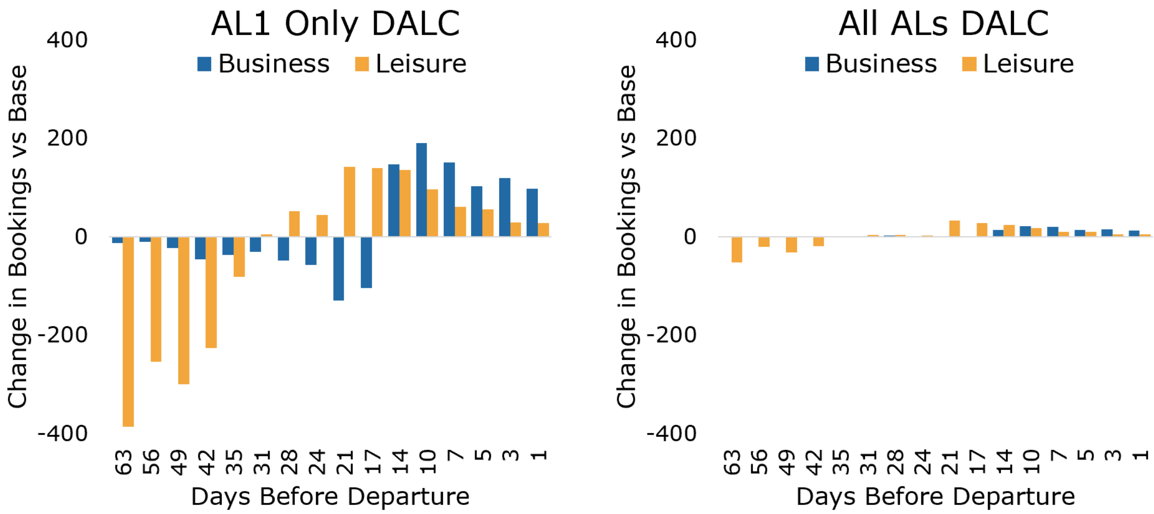

Figure 4-9: Change in bookings by segment observed by Airline 1 in each timeframe before departure

The consistently lower fares offered by DALC in the last 14 days doubles the number of Airline 1 bookings when it is the only airline using DALC (Figure 4-8). These bookings are gained by undercutting the competitors' prices and consists of passengers that would have previously flown with airlines 2-4. As Airline 1 sells more seats close to departure, it compensates and spills more passengers early in the bookings process to protect seats for the higher revenue demand. In the symmetric case, when all airlines use DALC, this competitive share shift is not observed. Airline 1 is only able to gain 13% more bookings in the last two weeks by stimulating new demand that had previously not been able to afford the high close-in prices. The overall effect seen on the booking curve is much smaller in the symmetric case.

Figure 4-9 shows the change in the booking curve by passenger type. It shows that DALC gains both business and leisure bookings close to departure, while primarily spilling leisure demand early in the booking process. As usual, the effect is smaller in a symmetric test.

The experiments show that DALC generally increases bid prices, as the airline observes more high revenue demand close-in, leading to higher initial bid prices and fewer seats available at low fares. This effect is magnified in the asymmetric test, as even more close-in bookings were observed in Figure 4-8. These bid prices are also used in the DOG price optimization, so they influence the DOG optimized prices offered to the customers.

Since DALC always offers the ancillary service optionally in the a la carte offer set, a consistent set of 38% of passengers choose to purchase it, same as in the traditional No DALC baseline. This is because DALC introduces no changes to the ancillary price or the offer set.

In conclusion, through these tests of DALC we gain an initial understanding of the dynamic price adjustment component of dynamic offer generation. Compared to the traditional RM baseline, the adjusted prices are higher in advance, but lower close-in. Combined with the segmented pricing for business and leisure passengers, DALC introduces a large shift in the booking curve. Airlines attract more high revenue close-in bookings, but fewer advance bookings in the lowest fare class. Competitive feedback effects follow, where the DALC airline attracts new close-in traffic, while traditional airlines capture more of the early demand at lower fares. As a result, Airline 1's total net revenue increases by 9.9% in asymmetric tests and 4.7% in symmetric tests. DALC does not influence ancillary revenue, as no bundles are offered and the same proportion of ancillaries are sold.

4.2.2 Dynamic Bundling Results

In *dynamic bundling* (DB), we focus on offer set selection and disable the dynamic price adjustment component of DOG, as illustrated in Figure 4-10. We let the DOG algorithm choose whether to show offer set S_1 (a la carte) or S_2 (bundle) to customers based on the expected revenue. Unlike dynamic offer generation, we bound the flight price in S_1 to be the filed fare of the lowest available fare class $\overline{p_f} = F$ with bounding parameter $b = 0$ according to Equation 3.20. Similarly, the bundle price is bounded according to Equation 3.21 as $\overline{p_{fa}} = F + (p^*_{fa} - p^*_f)$, the filed fare F plus the unbounded price difference between the bundle and the standalone flight. This choice makes the DB prices very similar to the baseline so that we can isolate the impacts of offering bundles, where the flight cannot be purchased without the ancillary service. As before, we test asymmetric cases with Airline 1 using DB and symmetric cases with all four

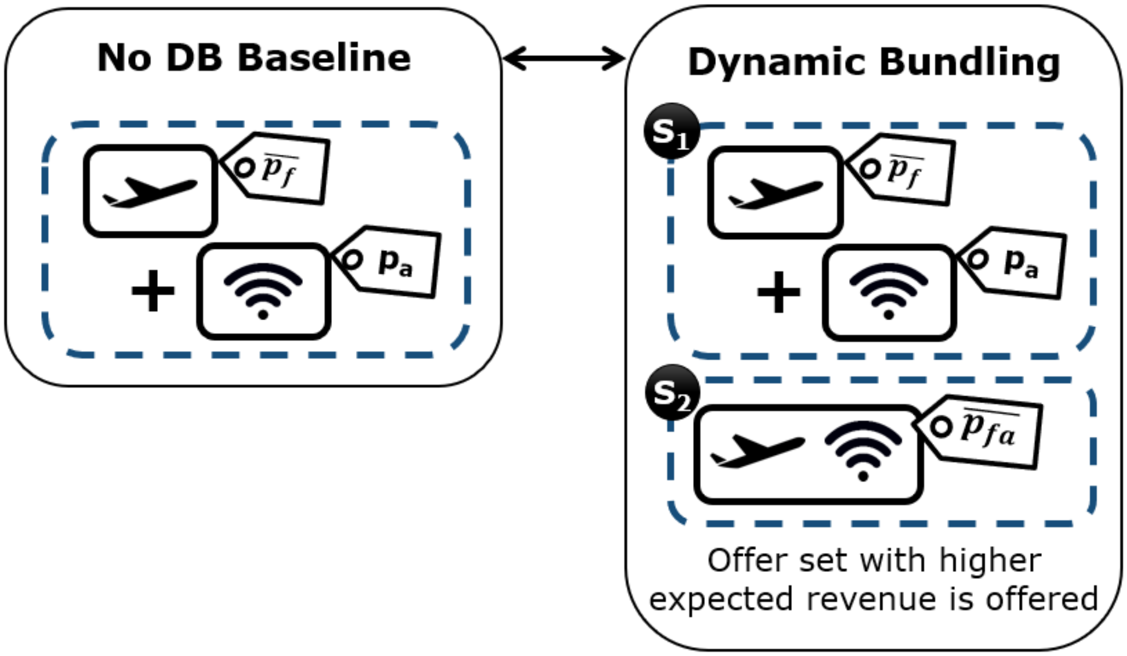

Figure 4-10: Comparison of the baseline (No DB) and dynamic bundling (DB)

Airline 1	**Base: No DB**	**AL1 Only DB**	**All ALs DB**
Flight Revenue	\$2,920,029	\$2,956,929	\$2,924,935
Ancillary Revenue	\$91,137	\$181,172	\$176,257
Ancillary Cost	\$66,383	\$154,506	\$150,427
Total Net Revenue	\$2,944,783	(+**1.3%**) \$2,983,595	(+**0.2%**) \$2,950,765
RPM	22,889,372	23,043,665	22,867,947
ASM	28,588,660	28,588,660	28,588,660
Net Yield	\$0.1286	\$0.1295	\$0.1290
Load Factor	80.06%	(+**0.54pts**) 80.60%	(**-0.07pts**) 79.99%

Table 4.3: Overall simulated results of dynamic bundling (DB) for Airline 1

airlines using DB.

Figure 4-11 and Table 4.3 show the overall revenue and load factor results of the test cases. In asymmetric tests, Airline 1 can achieve a 1.3% total net revenue gain by changing its ancillary pricing strategy, even though ancillary revenues represent only 3% of the airline's total revenue in the baseline scenario. In symmetric gains, dynamic bundling changes the four airlines' net revenue by -0.2% to +0.2%.

The larger asymmetric revenue gains are explained by competitive effects: Since passengers are modeled using concurrent choice behavior as well-informed about an-

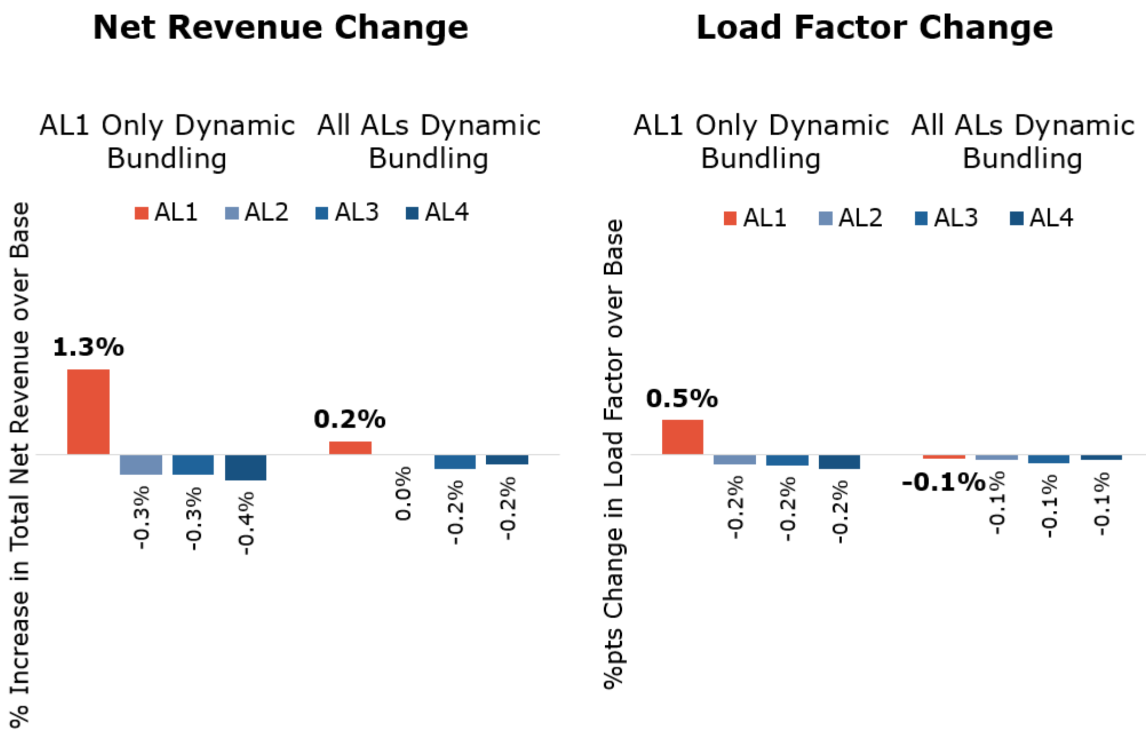

Figure 4-11: Overview of dynamic bundling (DB) net revenue gains and load factor changes in symmetric and asymmetric tests

cillary fees, Airline 1 is able to gain market share from its competitors offering a cheaper bundled ancillary service and gain flight revenue. In fact, ancillary net revenue increases by +7.7% in the asymmetric case (only +\$1,912 in absolute terms), but flight revenue increases by +1.3% (+\$36,900). Compare this to the symmetric case, where Airline 1's ancillary net revenue increases by +\$1,076, but flight revenue increases only by +\$4,907. Further sensitivity tests on this effect are performed in Chapter 5.3.

In Figure 4-12, we verify that we the flight price is constant across the experiments and that there are no significant changes to the average net revenue between the baseline and dynamic bundling. However, dynamic bundling has a large effect on the offers being purchased. Unlike in the baseline, many customers are only offered the bundle and as a result, the number of ancillary services sold increases. In this experiment, we have gained an initial insight into the potential benefits of a bundled pricing strategy.

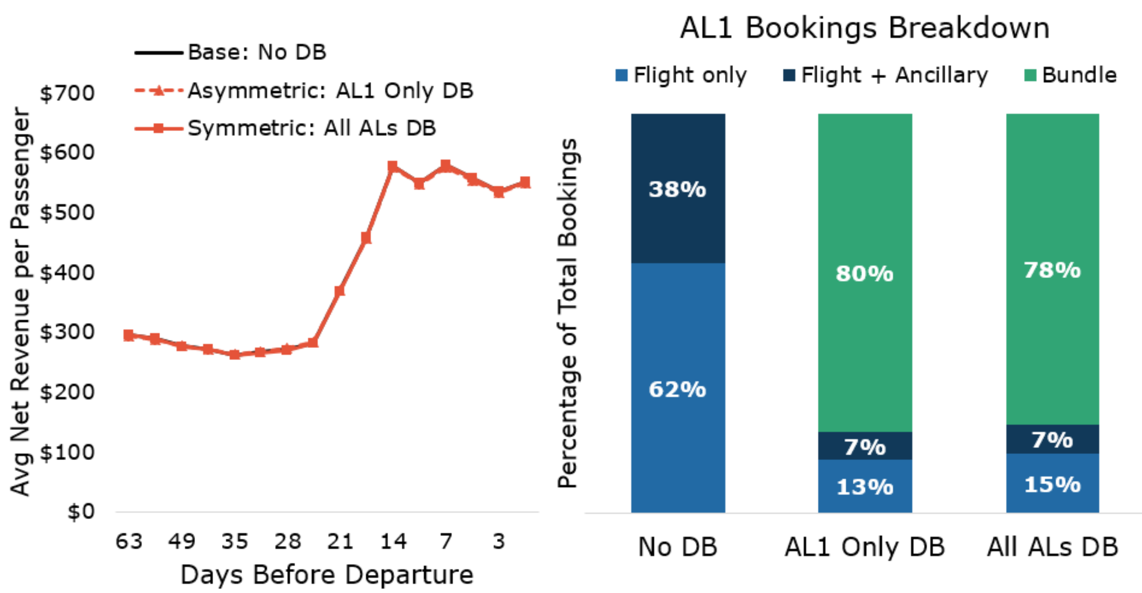

Figure 4-12: Airline 1 average net revenue by timeframe and breakdown of purchases by offer

4.2.3 Dynamic Offer Generation Results

Here we present the first results of *dynamic offer generation* (DOG) by combining the dynamic price adjustment heuristic tested in DALC with the offer set selection heuristic tested in DB. As introduced in Chapter 3, DOG uses the bid price of the RM system to compute segmented, continuous prices for both passenger segments (business and leisure) and both offer sets (S_1 and S_2). Like DB, the offer set with the higher expected net revenue under the concurrent choice assumption is then shown to the customer. The simulation baseline parameters were summarized in Table 4.1.

Since this represents the complete implementation of DOG and our main focus of the thesis, we expand our testing and report three sets of test results to study the impact of DOG in simulations with different demand levels. The standard (medium) demand level achieves an average load factor of 80% in the fully unrestricted baseline without DOG. It was used in the previous DALC and DB tests. We also test a new low demand scenario with 10% less demand generated across all passenger types, OD markets and simulated days before departure. The baseline average load factor is 73%. Similarly, a high demand scenario has 10% more generated demand and an average load factor of 83%. Figure 4-13 shows the total net revenue of the traditional

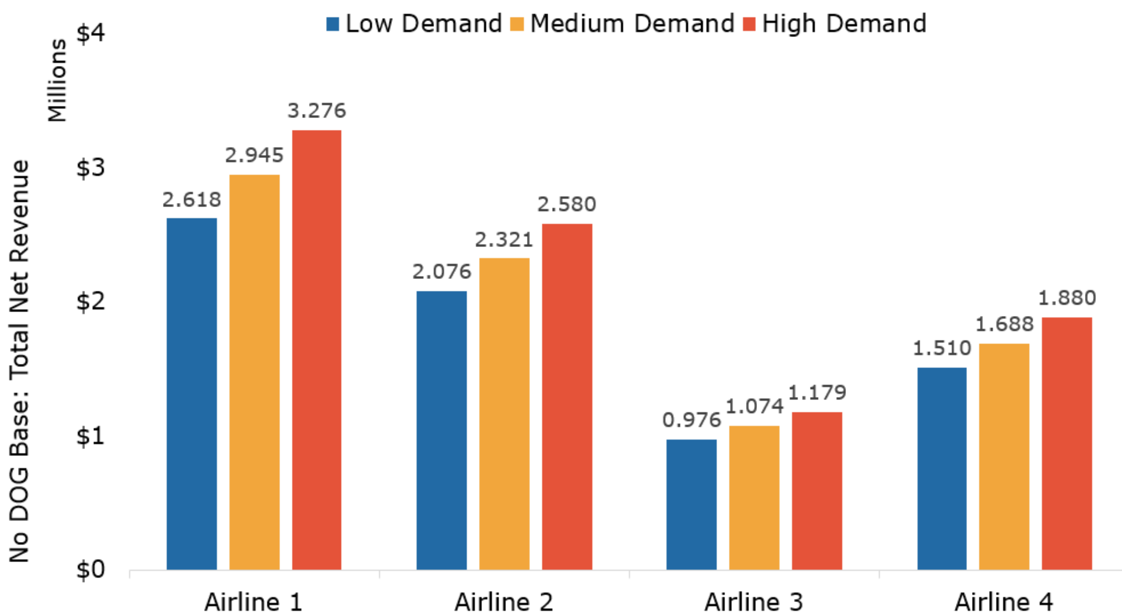

Figure 4-13: Total net revenue for the baseline scenarios without DOG

RM baseline at the three demand levels (no airline uses DOG). Figure 4-14 shows the load factor for each airline, as well as the net yield. Increasing demand increases both the load factor and the yield.

Next, we compare DOG against these baseline results. The asymmetric test as Airline 1 using DOG, while the other airlines continue with a traditional RM system and a la carte ancillary service. In the symmetric case, all airlines use DOG with the same input parameters.

Figure 4-15 shows the net revenue increases observed for airlines that implement DOG. In the medium demand case, Airline 1 gains 10.5% in net revenue asymmetrically and 5.1% symmetrically. This is higher than what we observed in DALC (+9.9% and +4.7%, Figure 4-6), reflecting the additional revenue gains from dynamic bundling. In the asymmetric case, DOG achieves the highest revenue gains in the low demand case. In the symmetric case, the high demand case produces the highest DOG revenue gains for all airlines. This is a result of the balance between changes in net yield (Figure 4-16) and load factor (Figure 4-17).

The net yield increase from DOG is highest in the high demand case. This implies that the DOG prices deviate most from the traditional RM system under high

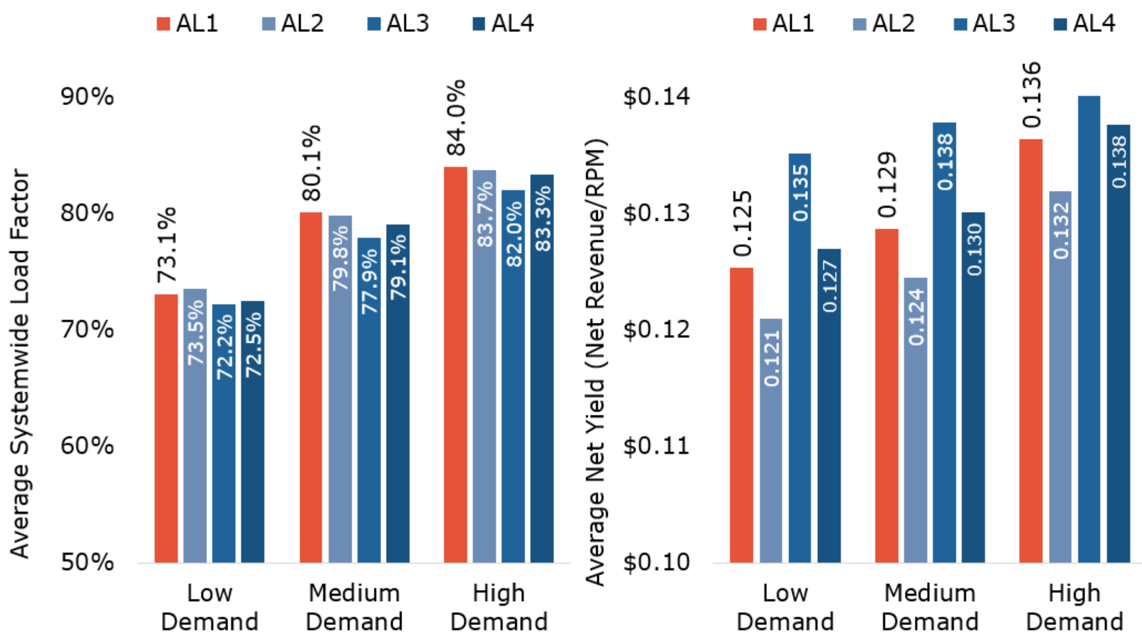

Figure 4-14: Load Factor and Net Yield for the baseline scenarios without DOG

demand. In the asymmetric tests, this large asymmetry in pricing leads to a large market share shift and the Airline 1 sees a 3.9% drop in load factor with DOG. The large drop in passenger numbers reduces the total net revenue benefit of DOG to 9.7% in the high demand case. On the other hand, in symmetric tests a change in price levels has a smaller impact on load factor as all other airlines use the same pricing and the market remains in equilibrium. This allows the net revenue gain from DOG to increase to 6.2% in the high demand case.

Conversely, in the low demand case, DOG acts to increase load factor while forgoing the high increase in yield observed at higher demand levels. DOG sells more seats early on, competing with the traditional airlines for leisure traffic that it would otherwise spill to protect seats for close-in demand. The asymmetric revenue gain is an extremely high +14.4%, because DOG is able to reap the rewards of undercutting traditional airlines for high-yielding demand without spilling the lower yielding demand to its competitors. In the symmetric case, this advantage disappears and revenue gains are reduced to +3.8%.

We can observe how the DOG pricing changes with demand in Figure 4-18. It shows the average net revenue per passenger in the symmetric DOG case and the

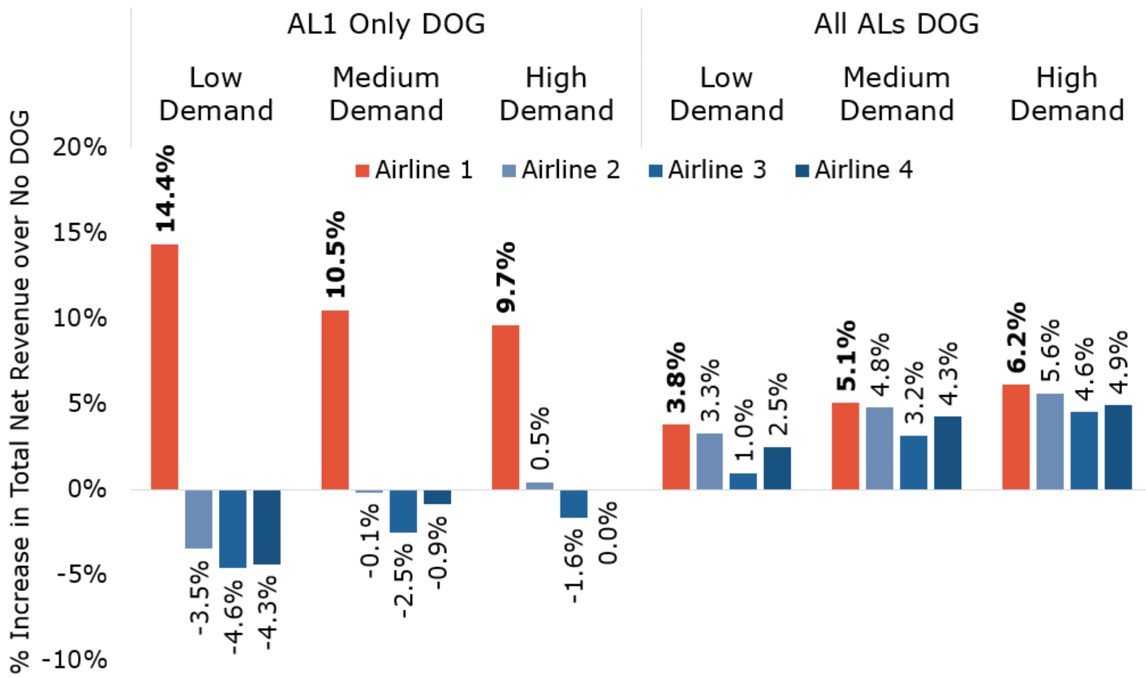

Figure 4-15: Total net revenue gains with the asymmetric and symmetric use of DOG

Airline 1	**No DOG**	**AL1 Only DOG**	**All ALs DOG**
Total Net Revenue	\$2,944,743	(+**10.5%**) \$3,253,192	(+**5.1%**) \$3,094,238
RPM	22,889,814	22,514,074	23,058,657
ASM	28,588,660	28,588,660	28,588,660
Net Yield	\$0.1286	\$0.1445	\$0.1342
Load Factor	80.07%	(**-1.3pts**) 78.75%	(+**0.6pts**) 80.66%

Table 4.4: Overall simulated results for Airline 1 of dynamic offer generation (DOG) with medium demand and baseline parameters

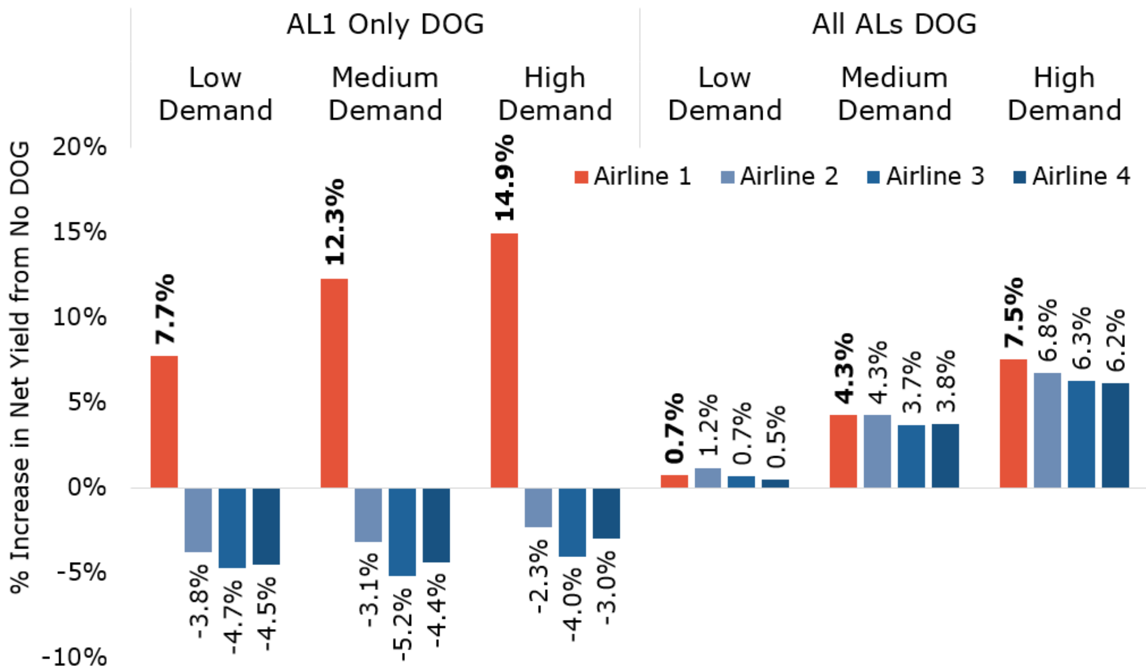

Figure 4-16: Increase in net yield with dynamic offer generation

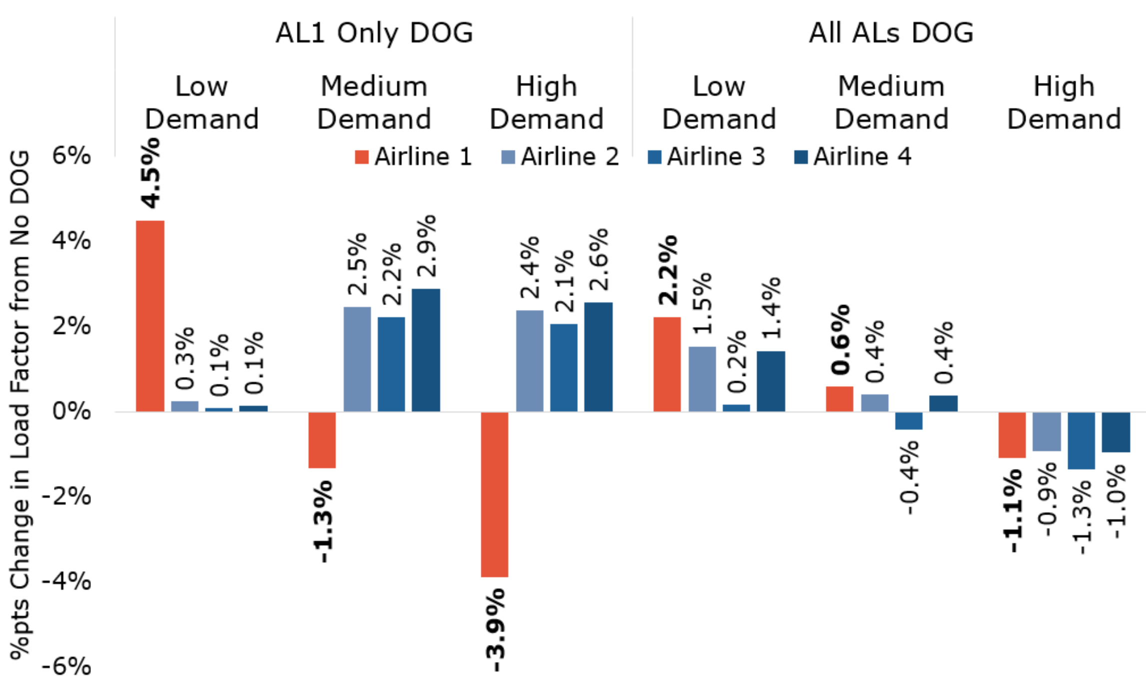

Figure 4-17: Change in load factor with the asymmetric and symmetric use of DOG

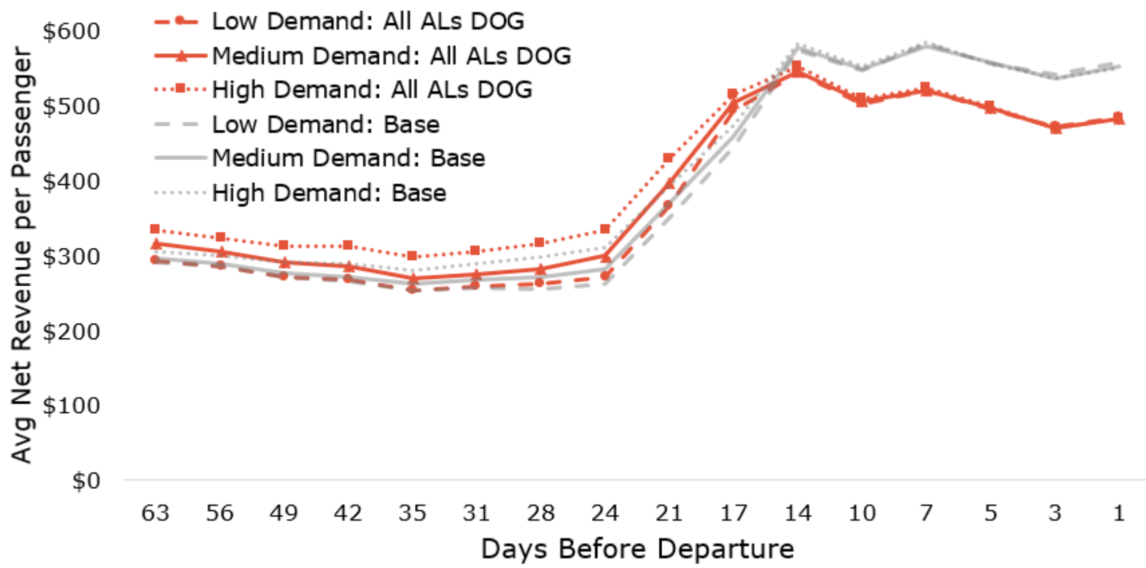

Figure 4-18: Airline 1 average net revenue per passenger by timeframe

No DOG baseline. The overall trends are similar to what was observed in DALC (Figure 4-7): Within 14 days of departure, DOG prices are consistently lower than No DOG, irrespective of demand level. This allows DOG airlines to capture more close-in demand, boosting their average net yield. More than 14 days before departure, we observe that DOG consistently has higher pricing than the baseline. Here, the higher demand levels also translate directly into higher pricing. This is an effect of the revenue management system protecting more seats for higher revenue passengers close to departure, which leads to fewer seats available at low fares early in the booking process and a higher bid price.

DOG includes the same price segmentation between business and leisure passengers that was applied in DALC. Figure 4-19 shows the difference between the average prices paid by business and leisure passengers in each timeframe. The domestic-only Airline 3 clearly exhibits the price segmentation in DOG, with business passengers consistently paying higher fares than leisure passengers in DOG. There is no price segmentation in the traditional baseline without DOG. For the larger Airline 1, the same effect exists, but it also shows that in the No DOG baseline, business passengers pay higher prices than leisure passengers. This is not because of active price segmentation by the airline, but an effect of the market mix. Business passengers that book

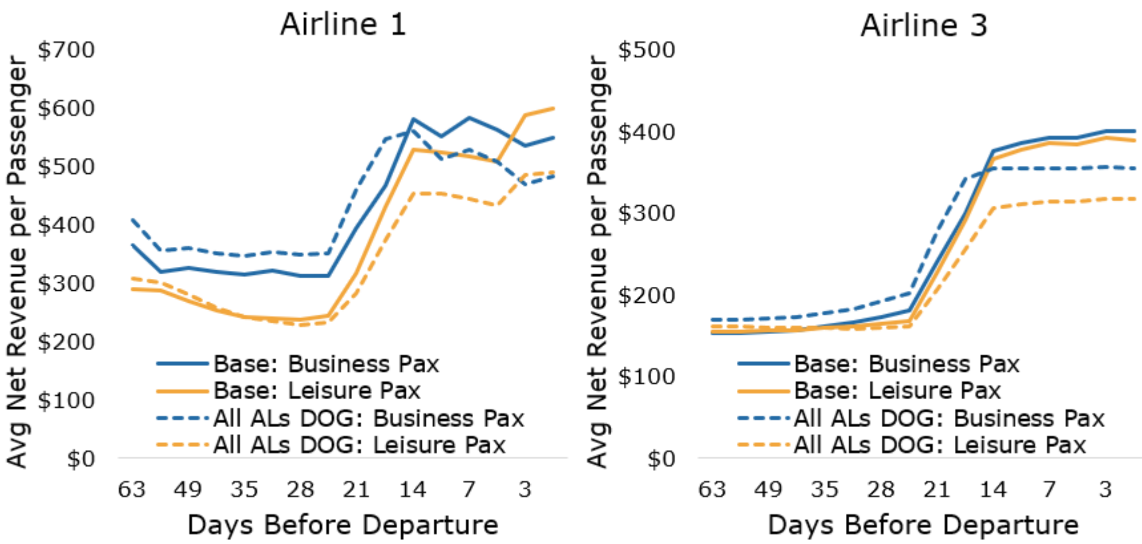

Figure 4-19: Airline 1 and 3 average net revenue by timeframe and passenger type showing price segmentation by DOG compared to the base

early are mostly long-haul international passengers, whereas many short-haul leisure passengers also book far in advance. As a result, the average net revenue for business passengers more than 14 days before departure is distorted by a high proportion of international bookings at high fares. In the same timeframe, leisure bookings include a higher proportion of domestic markets at lower fares, which explains the apparent price segmentation for Airline 1.

DOG chooses the offer set dynamically at each timeframe for each passenger type. Figure 4-20 shows that the demand level does not change an individual passenger's ancillary WTP since the percentage of passengers choosing to purchase the ancillary optionally (*Flight + Ancillary*) is nearly the same across all No DOG tests. However, DOG offers more bundles when demand is high in both asymmetric and symmetric cases. To explain this, Figure 4-21 further separates the offer set shown to customers by passenger type and timeframe for the symmetric (All ALs DOG) and medium demand case, where 76% of overall AL1 passengers booked a bundle. We observe that only some leisure passengers booking more than 21 days in advance received the a la carte offer set and almost all business passengers received the bundle offer set. This can be explained by our sensitivity test in Figure 3-7a, showing that the flight WTP and bid price have a secondary effect on offer set selection. When the flight

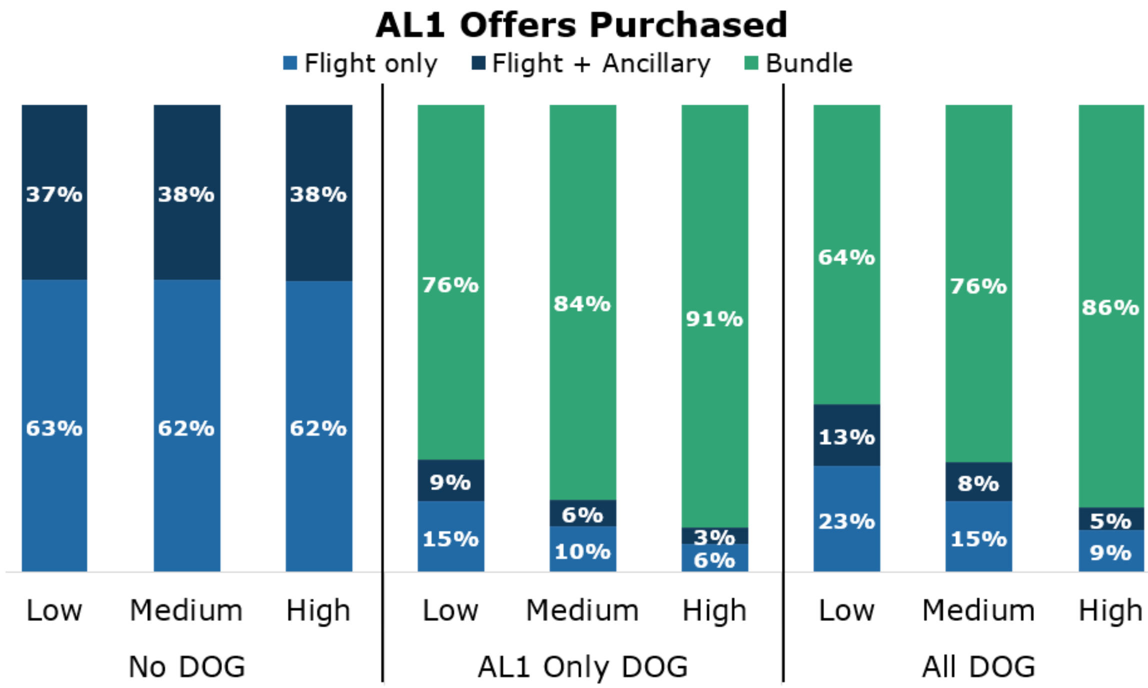

Figure 4-20: Airline 1 bookings by offer type, with and without DOG in asymmetric and symmetric tests

prices are low, there is a region when the a la carte offer set has higher expected net revenue than the bundle. This occurs especially for leisure passengers more than 21 days before departure. Fewer a la carte bookings are observed when demand is high because the higher bid prices make DOG offer more bundles.

4.3 Summary

In this chapter, we presented simulations of dynamic offer generation algorithm with one ancillary service in the Passenger Origin-Destination Simulator. The DOG algorithm was used in conjunction with a traditional network revenue management system. We studied the competitive implications of applying DOG for large network airlines.

With the baseline settings, DOG increased total net revenue in both asymmetric and symmetric cases: We showed that dynamic a la carte alone can lead to large asymmetric revenue increases of +9.9%, but that these revenue gains reduce to +4.7%

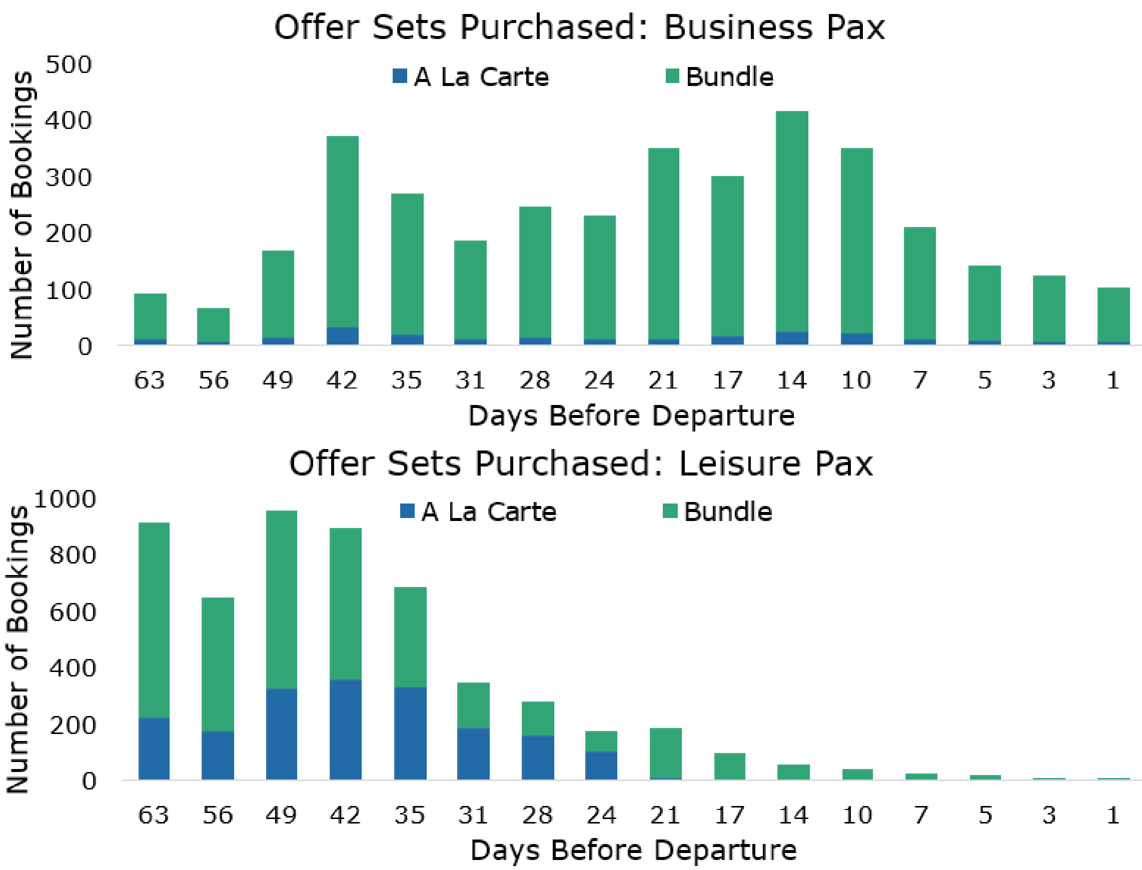

Figure 4-21: Bookings received by AL1 by passenger type, offer set and timeframe (symmetric DOG, medium demand)

once competitors respond and also implement DOG. These results are highly sensitive to the exact parameters assumed in the simulation and are the result of DOG offering lower close-in fares than the traditional baseline, as well as accurately segmenting prices for business and leisure passengers. We also assume that customers choose their itinerary, fare and ancillary services in a single step according to the concurrent choice model.

Similarly, dynamic bundling can lead to revenue gains of +1.3% asymmetrically under the same assumptions. Then, one airline can attract a higher market share by offering a cheaper ancillary service than its competitors and increase its flight revenue. In symmetric environments, the impact of bundling on total revenue was relatively small ranging from -0.2% to +0.2% across the four airlines.

Overall, dynamic offer generation performed better than both dynamic a la carte and dynamic bundling alone, as it combines the benefits of both. The revenue gains in both asymmetric and symmetric cases were consistently positive across three demand levels. In the low demand scenario, Airline 1 gained +14.4% in total net revenue asymmetrically and +3.8% symmetrically. For medium demand level, the gains were +10.5% asymmetrically and +5.1% symmetrically with baseline simulation parameters. At high demand, DOG increased revenue by +9.7% asymmetrically and +6.2% symmetrically.

In the following Chapter 5, we show sensitivity tests of DOG outside of the baseline scenario studied in this chapter. This shows how our simulated revenue gains depend on the input parameters chosen and under which circumstances DOG or dynamic bundling can achieve the highest revenue gains.

Chapter 5

Sensitivity Tests of Dynamic Offer Generation in PODS

In this chapter, we present additional sensitivity analyses on the input parameters for dynamic offer generation under the medium demand scenario in PODS. We begin in Section 5.1 with the flight-related parameters: we test different estimated flight WTP levels, the passenger type segmentation accuracy and flight price bounding parameters. We further study the impact of different traditional RM optimizers and their bid price computation methods on DOG prices and simulated revenue in Section 5.2. Next, we test the ancillary-related parameters of ancillary WTP and ancillary cost to study the impact of bundling on DB (5.3) and DOG (5.4). We use the insights gained to close this chapter by calibrating a new set of DOG results in Section 5.5, which better reflect the benefits in a realistic market conditions.

5.1 DOG with Different Flight-related Parameters

First, we perform a sensitivity test on the *Q multipliers* used in our dynamic flight price adjustment heuristic for DOG. The *Q multipliers* act as estimates of the mean conditional flight WTP of a passenger segment. It is a multiplicative factor on the lowest published fare in each market (Wittman, 2018). The lowest booking class is called Q class in some airlines' revenue management systems. In PODS network U10,

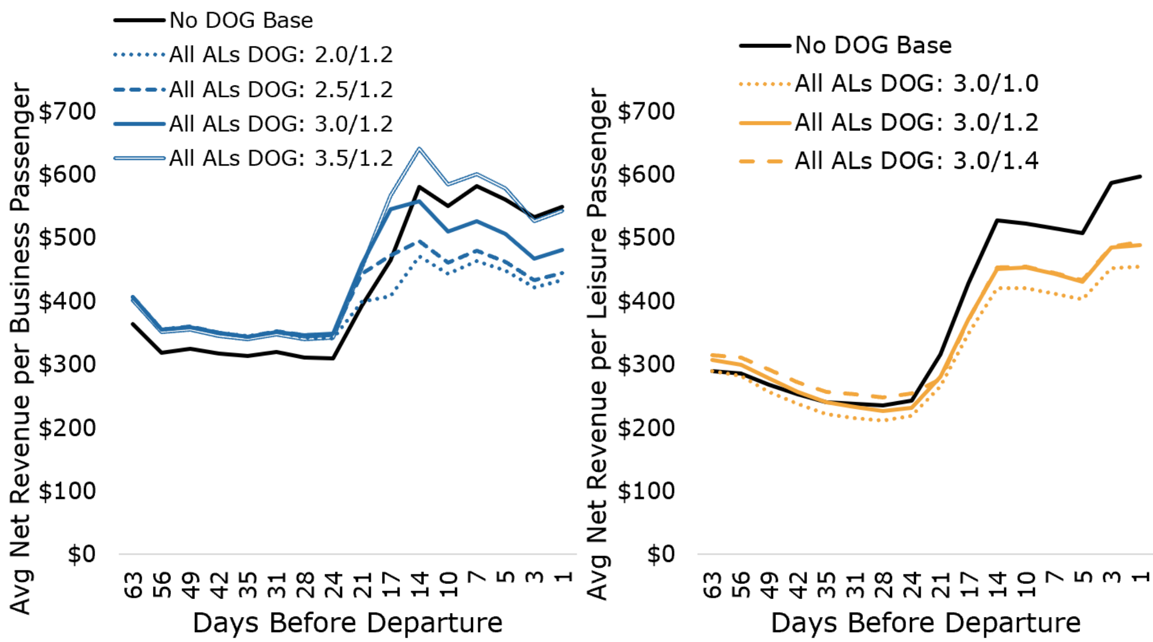

Figure 5-1: Average net revenue per business passenger with different business Q multipliers (left) and net revenue per leisure passenger with different leisure Q multipliers (right)

the lowest booking class is 10 and we call its fare F_Q. DOG receives the Q multiplier as an input for the normally distributed conditional flight WTP with mean $\mu_{W_f} = Q \cdot F_Q$ and standard deviation $\sigma_{W_f} = 0.3\mu_{W_f}$, separately for each passenger segment. In our baseline case with constant Q multipliers of 3.0/1.2 (business Q/leisure Q), the airline assumes that on average business passengers are willing to pay \$300 and leisure passengers are willing to pay \$120 at all times in a market where the lowest fare F_Q is \$100.

In a first set of tests, we test four different business Q multipliers {2.0,2.5,3.0,3.5} while holding the leisure Q multiplier constant at 1.2. In a second test, we hold the business Q multiplier constant at 3.0 and vary the leisure Q multiplier {1.0,1.2,1.4}. In all cases, we maintain the same baseline bounding on the offer prices. The Q multipliers determine the optimal DOG offer prices before the prices are bounded to the lowest open fare class of the RM system. As a result, changing the Q multiplier affects the average net revenue paid by passengers in timeframes when the optimized DOG price is within the range allowed by the fare class bounds.

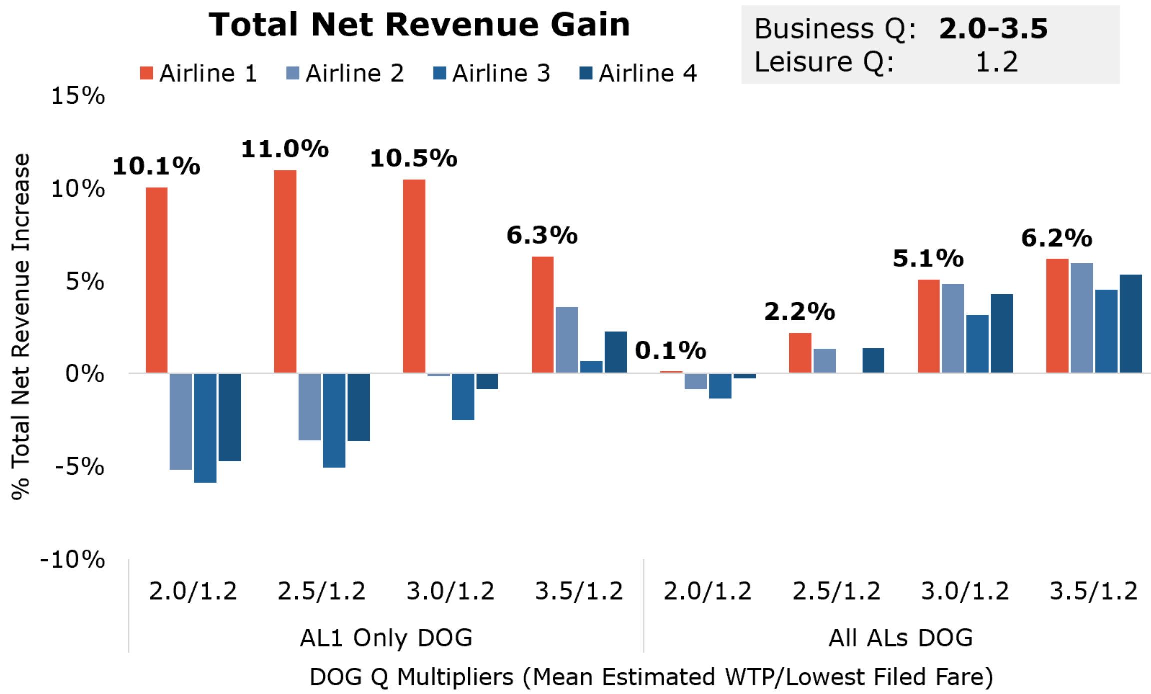

Figure 5-2: DOG net revenue increase from base with different business flight WTP estimates (Q multipliers)

The impact of Q multipliers on fares is shown in Figure 5-1. It shows how a lower Q multiplier decreases the close-in business fares but has nearly no impact on advance business fares. In early timeframes, the optimal DOG price adjustment is bigger than the upper bound allows, even at a lower business Q multiplier of 2.0. As a result, the upper bound fare is charged irrespective of Q multiplier. Similarly, when the leisure Q multiplier is increased, the average net revenue paid by leisure passengers increases in early timeframes. However, close to departure, the optimized DOG adjusted price is much lower than the lower bound permitted by the RM system, so the lower bound fare is charged irrespective of leisure Q multiplier.

Figures 5-2 and 5-3 show the simulated revenue gains as the business and leisure Q multipliers are varied independently. The revenue impact is largely driven by the differences in pricing shown in Figure 5-1. In the symmetric cases, a higher business Q multiplier of 3.5 increases revenue gains as all airlines would increase their close-in business fares. However, in the asymmetric case, a 3.5 business Q multiplier increases close-in business fares above the traditional airlines' prices. This makes the DOG

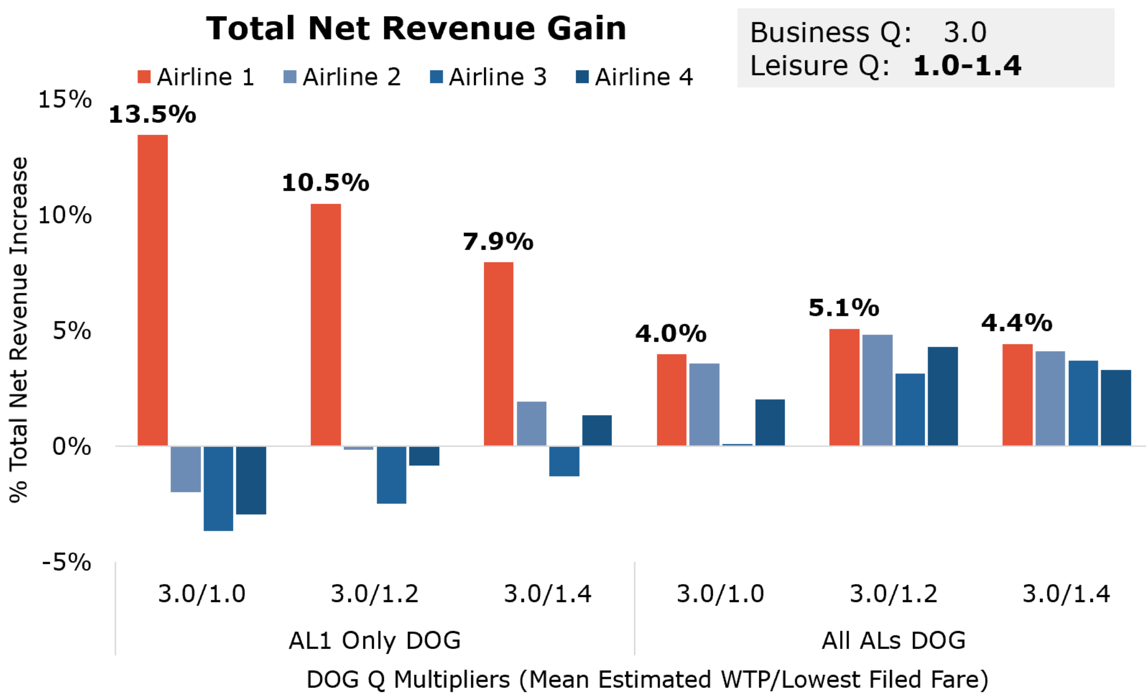

Figure 5-3: DOG net revenue increase from base with different leisure flight WTP estimates (Q multipliers)

airline less competitive and leads to a drop in revenue. A lower Q multiplier of 2.5 produces the highest DOG revenue gains in asymmetric tests. Our chosen baseline Q multipliers of 3.0 and 1.2 represent a compromise that produces good revenue gains in both asymmetric and symmetric cases.

With the Q multipliers, the airlines control the degree of price segmentation between the passenger segments. The more aggressive the price segmentation, the higher the potential risk of misidentifying passengers becomes. In all tests so far, the identification accuracy was 100%, but in reality a classification algorithm would have lower performance. The experiment in Figure 5-4 shows the robustness of DOG at lower identification accuracy with 3.0/1.2 Q multipliers and standard b=1 bounding. There is a nearly linear decline of the percentage net revenue gain as the identification accuracy decreases. However, even a random guess (50% identification accuracy) can produce asymmetric revenue gains in DOG. In the symmetric case where all airlines independently misidentify at the same accuracy, greater than 60% accuracy is needed with this set of Q multipliers to achieve a positive revenue gain.

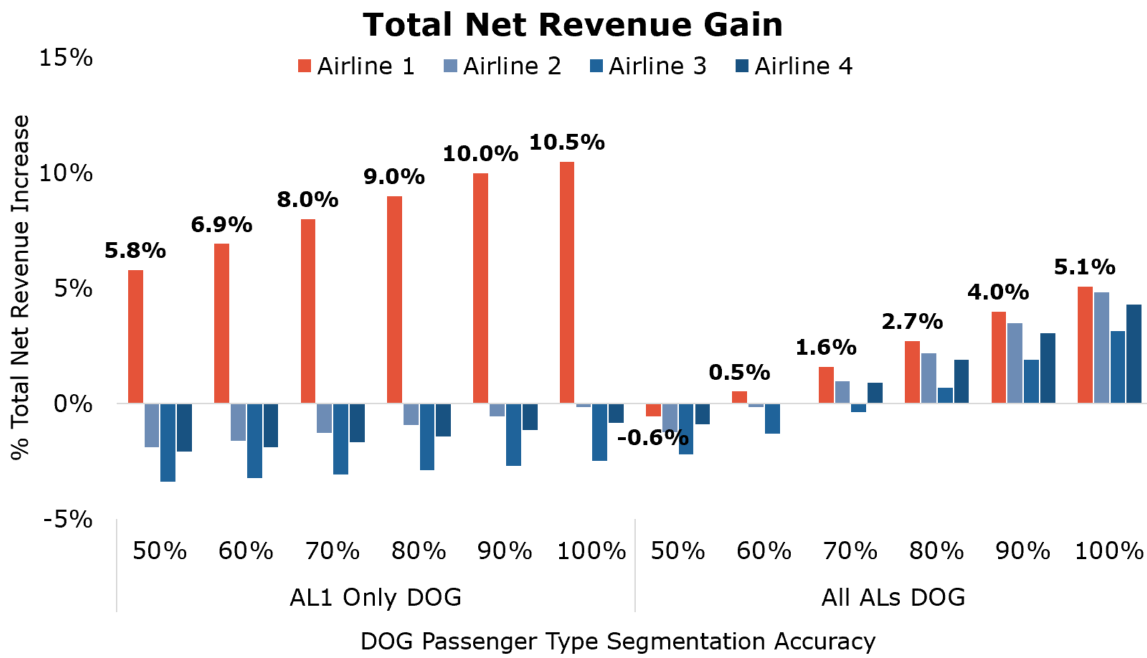

Figure 5-4: DOG net revenue increase from base with different passenger type identification accuracy (price segmentation)

Two effects are driving this behavior at lower identification accuracy: First, more business passengers are misidentified as leisure by at least one DOG airline. As a result, they receive cheaper offers and buy-down, causing revenue losses in the business segment. Second, more leisure passengers are misidentified and receive more expensive business offers. These are often not affordable for leisure passengers and they choose to book with a competitor or not travel at all.

These effects can be observed in the average net revenue that the DOG airline receives from passengers in Figure 5-5. The average net revenue from actual business passengers decreases with lower identification accuracy on the left (buy down). However, only little change is observed in the average net revenue paid by leisure passengers, which suggests that the airline does not get a high purchase rate on its more expensive offers for misidentified passengers.

Finally, we test the sensitivity of DOG to the strictness of the bounds on the offer prices imposed by the RM system. The bounding parameter b as defined in Equation 3.20 refers to a multiplier on the price difference from the current lowest available filed fare to the fare of the next higher and lower fare class. Our standard tests use

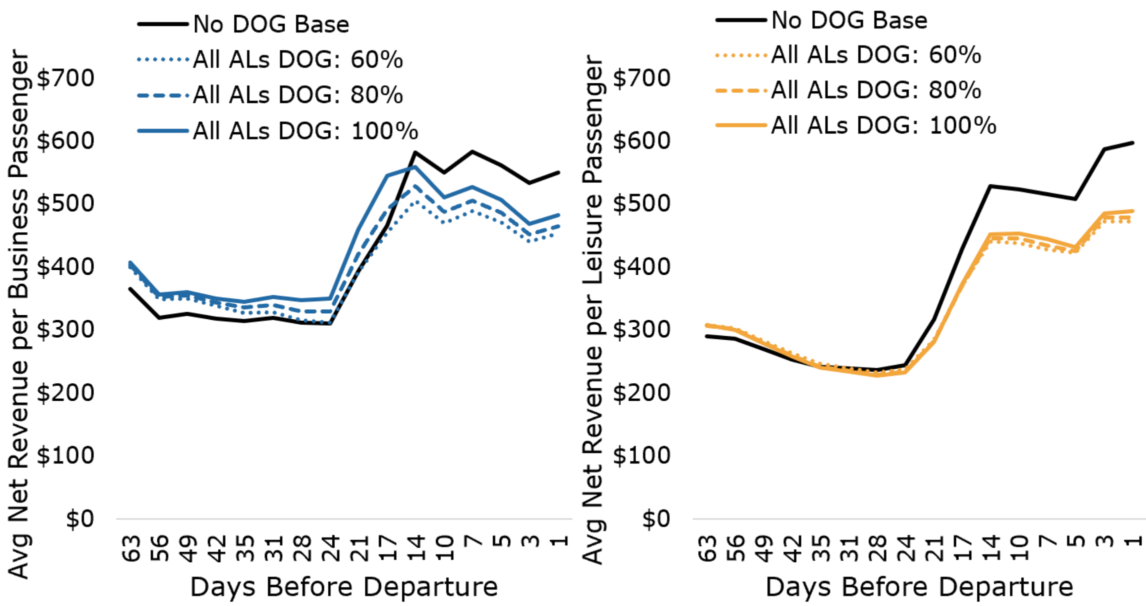

Figure 5-5: Airline 1 average net revenue paid by business and leisure passengers in each timeframe with varying identification accuracy

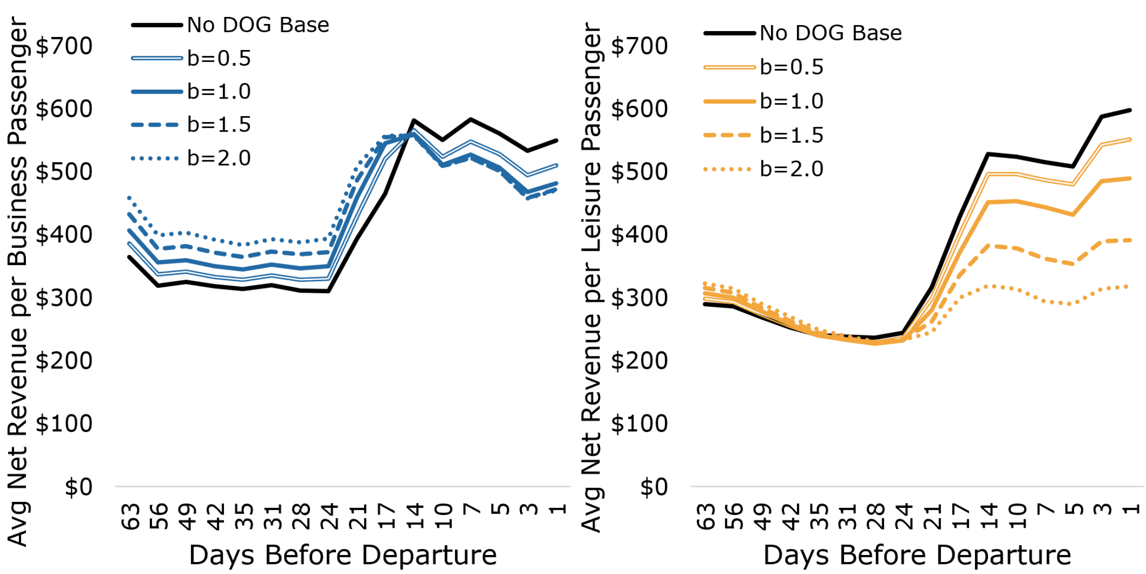

Figure 5-6: Airline 1 average net revenue per passenger with different bounding widths b (medium demand)

b=1, meaning that DOG can adjust the lowest available filed fare determined by the RM system up to ± 1 fare class. When b=0, no deviation from the flight price is permitted, which represents the dynamic bundling scenario tested in Chapter 4.2.2.

In Figure 5-6, we test bounds of b={0.0,0.5,1.0,1.5,2.0}. As the bounding parameter b increases, both business and leisure fares deviate more from the traditional RM baseline and also from each other (increasing price segmentation). Especially close-in leisure bookings benefit from large discounts and advance business bookings are charged higher prices. This leads to a shift in the booking curve shown in Figure 5-7: In both the symmetric and asymmetric cases, a more relaxed bounding leads to DOG airlines accepting more bookings close to departure and fewer bookings in advance. This effect is strongest in the asymmetric case, as the DOG airline undercuts other traditional airlines. This allows it to attract a large number of close-in bookings from its competitors. In these tests, the traditional airlines do not respond to the shift in pricing introduced by DOG and thus lose high-revenue bookings. This allows DOG to sustain its high asymmetric revenue gains. The traditional airlines could respond by changing their pricing to match the DOG airline and recover some of the revenue lost to DOG. Possible competitive responses of traditional airlines to continuous pricing algorithms were tested in PODS by Papen (2020).

Through these tests, we have seen that each DOG parameter can have a large impact on the offer prices. The DOG price adjustment heuristic gives an airline many levers to control the level of price segmentation and booking curve they desire, while also applying optimization to find an optimal price point given the constraints. On the flip side, the additional parameters create a calibration and estimation challenge for airlines. It can be difficult to accurately estimate the flight WTP of each passenger segment, as well as the segment that an individual booking request belongs to. We showed that the reported DOG net revenue gains are reduced at lower passenger type identification accuracy. The revenue gain observed in PODS is also highly dependent on the competitive situation, with DOG performing better in asymmetric scenarios where its prices undercut the traditional airlines that do not adapt their pricing. In a real-world implementation, it is unlikely that the same magnitude of revenue increases

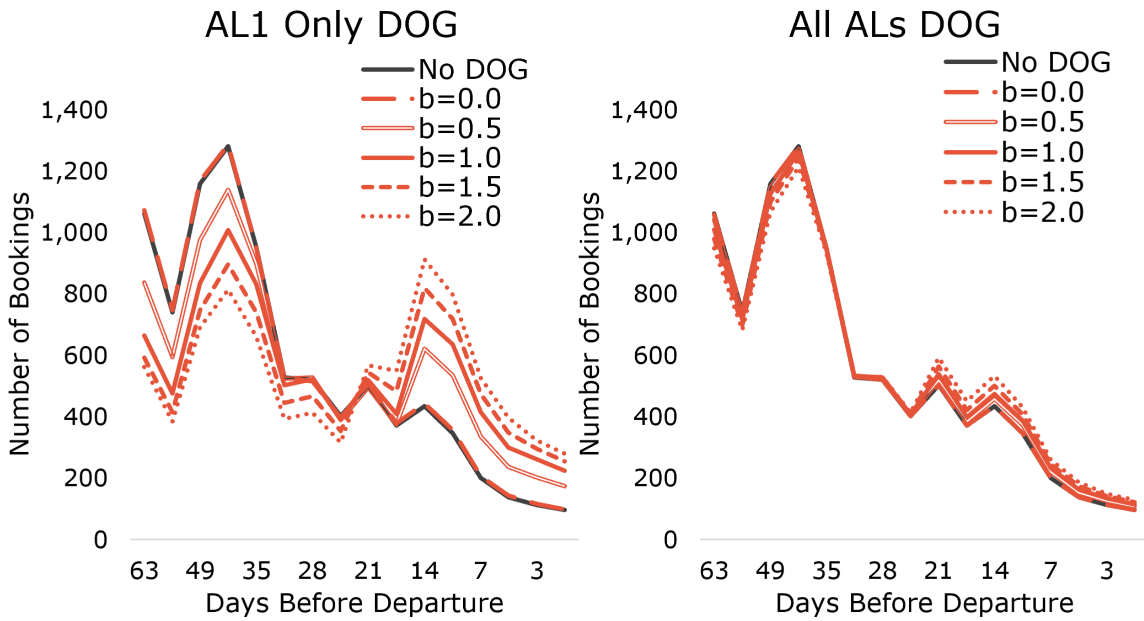

Figure 5-7: Airline 1 bookings by timeframe for asymmetric and symmetric cases of DOG with different bounding widths b (medium demand)

can be achieved, as traditional airlines can respond by adapting their pricing.

5.2 DOG with Different Revenue Management Optimizers

The revenue management forecaster and optimizer remain the backbone of the airline's RM system, even with dynamic offer generation. As presented in Chapter 3.4, DOG relies on the RM optimizer's bid price and its computed seat availability across fare classes to compute bounded offer prices. We are interested in studying how sensitive the DOG offer prices are to the RM optimizer used by the airline. In our standard baseline, all airlines use the displacement adjusted virtual nesting (DAVN) network RM algorithm.

In this section, we test DOG with two additional bid price control methods for network RM: Probabilistic Bid Price (ProBP) and Unbucketed Dynamic Programming (UDP). ProBP was developed by Bratu (1998) and UDP is an extension of the leg-based dynamic program by Lautenbacher and Stidham (1999) for network O-D control. We hold all baseline RM and DOG parameters constant, but change the RM optimizer used by the DOG airline. We study an asymmetric scenario, where only Airline 1 uses DOG with either DAVN, ProBP or UDP as the RM optimizer. We compare the performance of DOG with ProBP (or UDP) at Airline 1 against baselines where Airline 1 asymmetrically uses ProBP (or UDP) without DOG. All other airlines always retain a traditional RM system with DAVN as the optimizer.

The RM optimizer controls the DOG prices through the bid price in two ways: First, a higher average bid price closes more fare classes, thus increasing the corresponding DOG bounds and the final prices. Second, the bid price is used in the DOG price optimization. A higher bid price leads to higher optimized flight and bundle prices in DOG (cf. Figure 3-5a). There are differences in the bid prices produced by each method: The ProBP bid price is generally lower than DAVN throughout the booking window. UDP bid prices are even lower than ProBP early on, but rapidly increase in the last two weeks before departure. The final UDP bid prices are higher than both ProBP and DAVN. This effect has been attributed to the inherent Poisson demand arrival assumption in the dynamic programming approach, which underestimates the true variability in the simulated demand. This makes UDP overly certain of the future demand and causes it to accept more bookings early on.

These bid price differences are reflected in the DOG prices shown in Figure 5-8. Comparing the three DOG cases, we observe that the different RM optimizers directly impact the prices paid by customers: DOG with ProBP charges slightly lower prices than DOG with DAVN, especially more than three weeks before departure. DOG with UDP has significantly lower prices than both DAVN and ProBP up to 14 days before departure, when the surge in UDP bid price leads to higher close-in prices than DOG with DAVN. For reference, the average net revenue with DAVN, but without DOG, is shown in black. This is representative of the competing fares of the other airlines 2-4 in this asymmetric test. With DAVN, we concluded that the DOG heuristic gains revenue by primarily lowering close-in fares. The DOG airline compensates by increasing advance fares more than two weeks before departure. This trend still holds for DOG with ProBP, but not with UDP. When DOG is used with

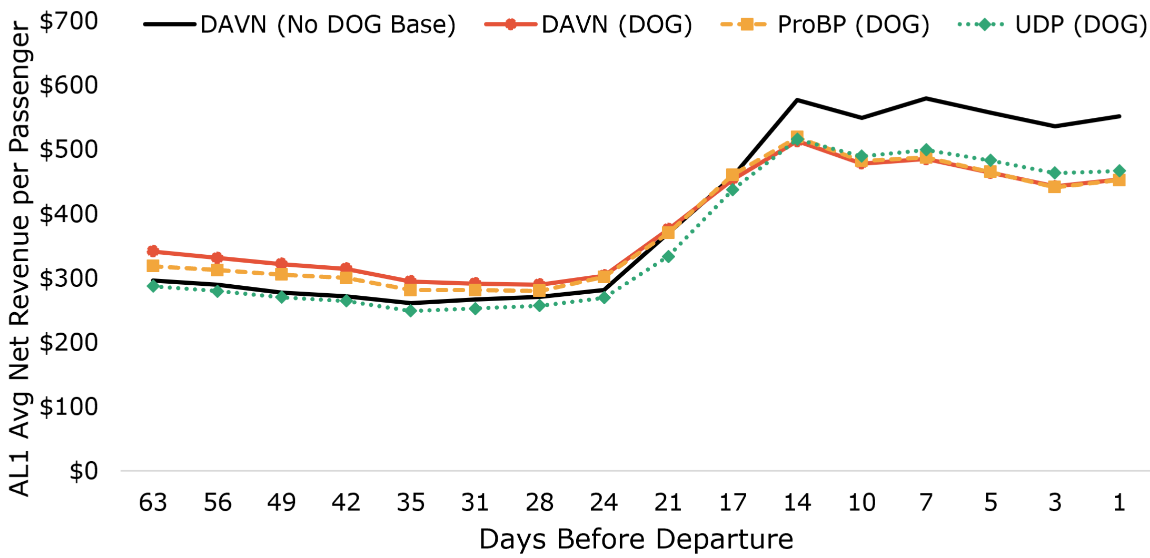

Figure 5-8: Airline 1's average net revenue per passenger generated by DOG airline with different RM optimizers and bid price inputs, compared to the symmetric DAVN No DOG baseline

UDP, the simulation shows that it lowers fares compared to the traditional baseline throughout the booking process.

By undercutting the other airlines throughout, Airline 1 with UDP was able to gain the most revenue and load factor in this asymmetric PODS simulation, as shown in Figure 5-9. It compares the incremental revenue and load factor gains achieved by using DOG compared to a No DOG baseline when Airline 1 asymmetrically uses different RM optimizers. DAVN and ProBP generate additional revenue with DOG by increasing the average yield and accepting a higher share of close-in high-revenue bookings. UDP is notably different and attracts additional passengers with DOG, increasing load factor by +7.5pts to a record 91.7%. While it performed best in our simulations, such high load factors have not been achieved in reality and come with operational risks for the airline. We also did not simulate interventions by RM analysts, which could close down lower fare classes earlier and lower overall load factors in an attempt to increase yield.

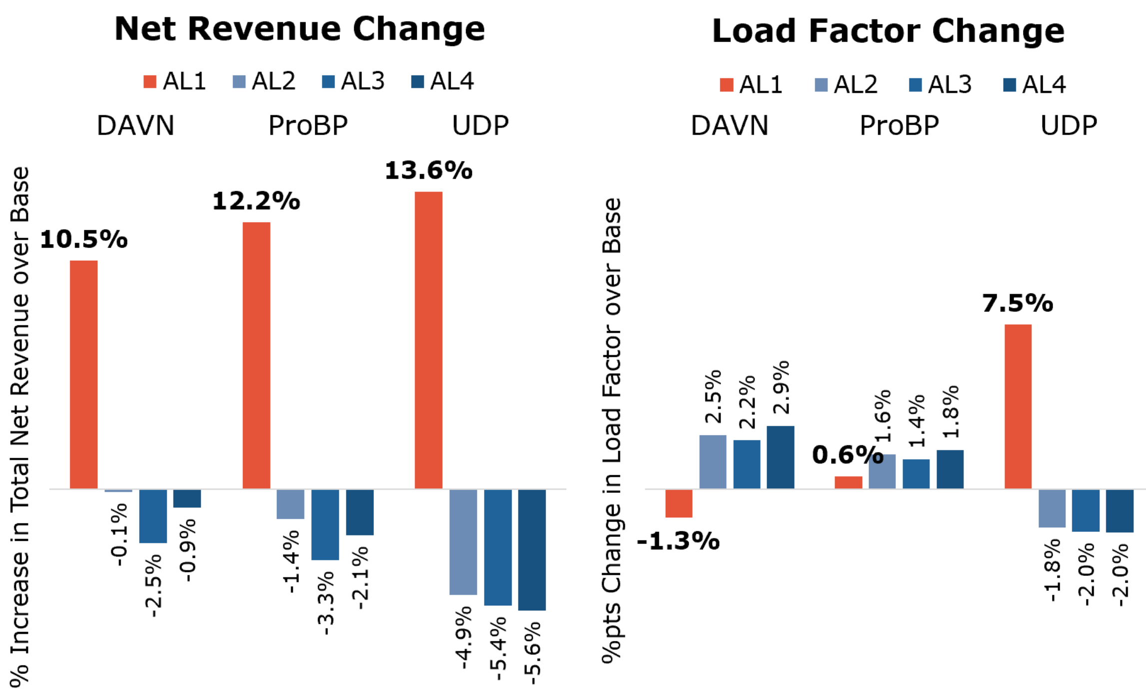

Figure 5-9: Average net revenue and load factor change when Airline 1 switches from base to DOG. Airline 1 asymmetrically uses a different RM optimizer in both base and DOG.

5.3 Dynamic Bundling with Different Ancillary Parameters

To further explore the impact of ancillary-related parameters and bundling on DOG, we first perform two sensitivity tests on *dynamic bundling* (DB) as introduced in Chapter 4.2.2. DB allows us to better isolate the impact of bundling on the results and separate total revenue into flight and ancillary components. This gives us additional insight into the revenue impacts. In the next section (5.4), we repeat these tests for DOG.

The first study will vary both the actual ancillary WTP distribution of the passengers X_a and the airline's estimated ancillary WTP distribution W_a as shown in Table 5.1. The true mean ancillary WTP ranges from \$15 to \$30 and the standard deviation scales as 30% of the mean. All airlines accurately estimate this WTP distribution and their ancillary cost remains constant at \$20. As a result, the price of

WTP Sensitivity	Case 1	Case 2	Case 3 (Base)	Case 4
Actual WTP X_a	$\mathcal{N}(\$15, \$4.5^2)$	$\mathcal{N}(\$20, \$6^2)$	$\mathcal{N}(\$25, \$7.5^2)$	$\mathcal{N}(\$30, \$9^2)$
Estimated WTP W_a	$\mathcal{N}(\$15, \$4.5^2)$	$\mathcal{N}(\$20, \$6^2)$	$\mathcal{N}(\$25, \$7.5^2)$	$\mathcal{N}(\$30, \$9^2)$
Ancillary Cost c_a	\$20	\$20	\$20	\$20
Myopic optimal p_a	\$22.22	\$24.51	\$27.41	\$30.65

Table 5.1: Sensitivity test of dynamic bundling (DB) on ancillary WTP X_a, W_a for all airlines

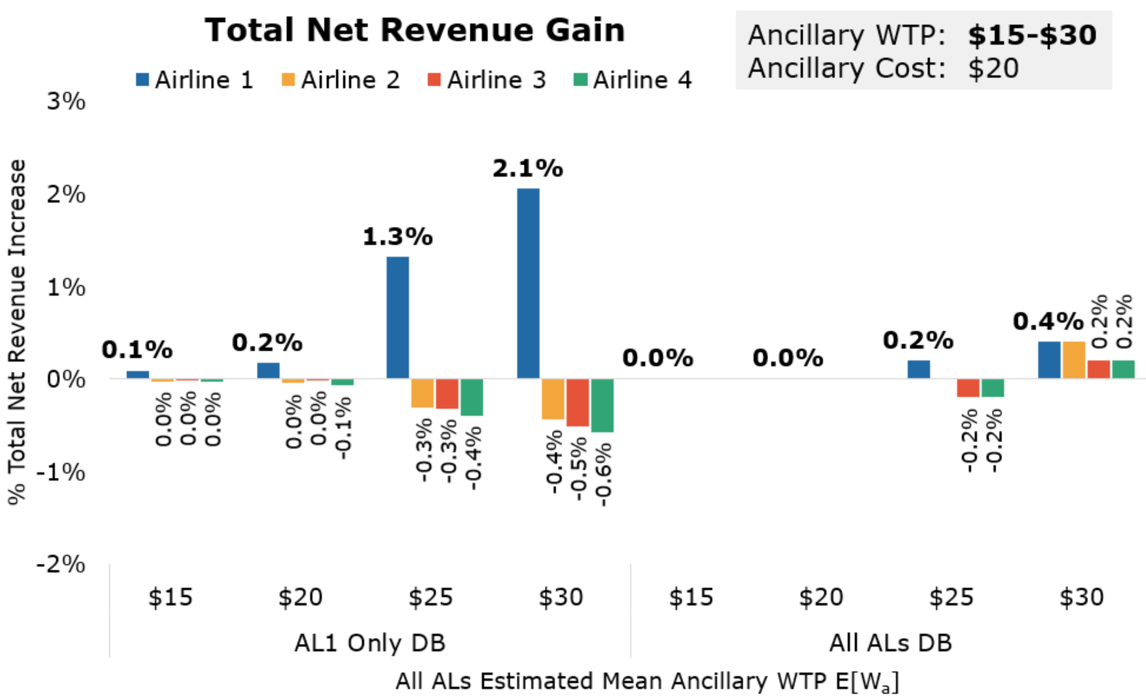

Figure 5-10: Variation of dynamic bundling (DB) net revenue gains with ancillary WTP at a constant ancillary cost of \$20

the ancillary service varies in each case for both the DB test case and the No DB baseline. When airlines do not use DB, they charge the listed myopic optimal price.

Figure 5-10 shows the overall net revenue benefit of using DB compared to always selling the a la carte offer set (No DB). In the asymmetric tests, the net revenue gain is only +0.1% when the WTP is low but rises to +2.1% when the ancillary WTP is high. The same trend is seen in the symmetric cases, but the net revenue gain ranges from +0.2% to +0.4% for all airlines at a \$30 mean WTP.

Figure 5-11 shows the source of the revenue gains when WTP W_a is high relative to the cost c_a: Dynamic bundling begins to offer the bundled offer set when the mean

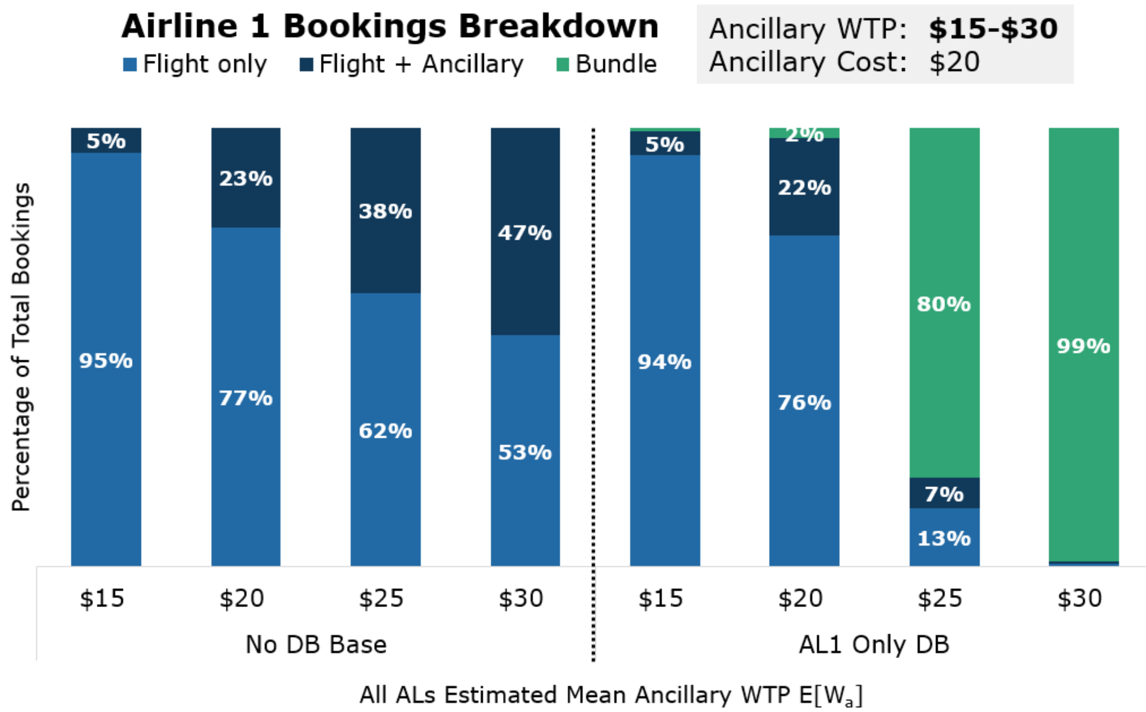

Figure 5-11: Variation of offer purchases with ancillary WTP at a constant ancillary cost of $20

WTP exceeds $25. In low WTP cases, dynamic bundling very rarely bundles and uses the same a la carte pricing strategy as the No DB baseline. This is consistent with our observations in Figure 3-7b, indicating that the bundled offer set has higher expected net revenue when $E[W_a]/c_a > 1.25$. It is notable that even when the a la carte offer set is offered (No DB) the ancillary purchase rate increases with ancillary WTP, even though the myopic optimal price increases as well.

Figure 5-12 explains why DB begins to offer bundles at higher WTP. The left graph shows the average selling price of the ancillary and the cost the airlines incur for selling an ancillary. When the WTP is high, the airline can achieve a sufficient margin (price - cost) to allow it to discount the ancillary service in a bundle. By selling the ancillary service in a bundle, the purchase rate increases significantly, since passengers are now required to purchase the ancillary service when buying the flight and the discounted price is sufficiently attractive to a large proportion of customers.

Figure 5-13 summarizes the revenue impact of dynamic bundling. In both asymmetric and symmetric cases, dynamic bundling increases ancillary net revenue (an-

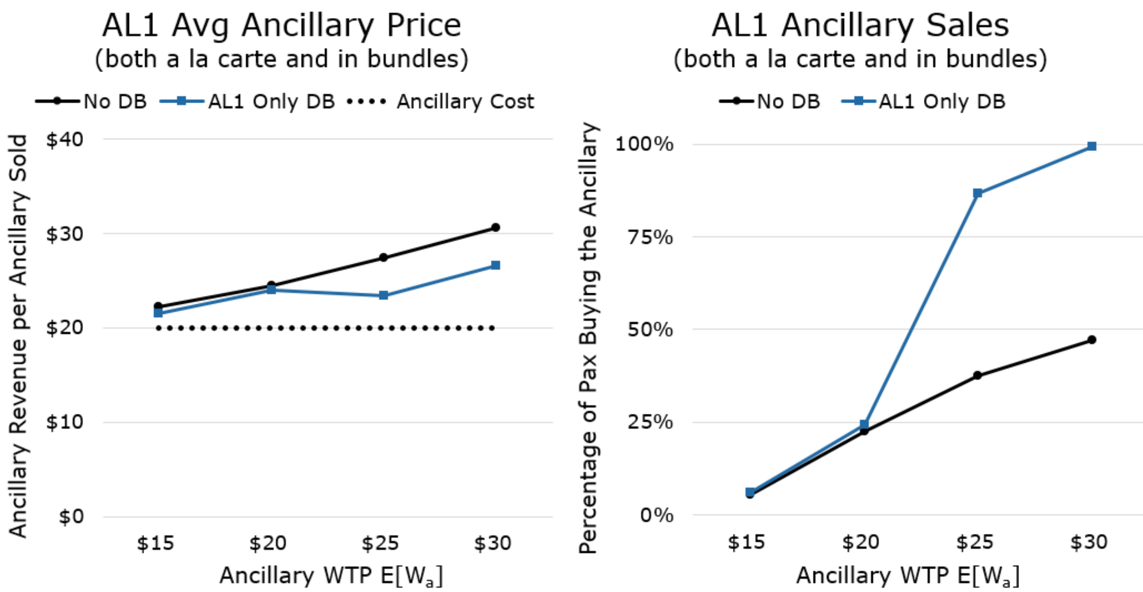

Figure 5-12: Variation of ancillary prices and purchases with ancillary WTP at a constant ancillary cost of $20

Cost Sensitivity	**Case 1**	**Case 2**	**Case 3 (Base)**	**Case 4**
Actual WTP X_a	$\mathcal{N}(\$25, \$7.5^2)$	$\mathcal{N}(\$25, \$7.5^2)$	$\mathcal{N}(\$25, \$7.5^2)$	$\mathcal{N}(\$25, \$7.5^2)$
Estim. WTP W_a	$\mathcal{N}(\$25, \$7.5^2)$	$\mathcal{N}(\$25, \$7.5^2)$	$\mathcal{N}(\$25, \$7.5^2)$	$\mathcal{N}(\$25, \$7.5^2)$
Ancillary Cost c_a	\$0	\$10	\$20	\$30
Myopic optimal p_a	\$19.35	\$22.51	\$27.41	\$34.35

Table 5.2: Sensitivity test of dynamic bundling (DB) on ancillary cost c_a for all airlines

cillary revenue - ancillary cost) by a similar amount. However, the majority of asymmetric revenue gains come from increased flight revenue, indicating that Airline 1 can attract passengers from its competitors by offering a bundle with a cheaper ancillary service. This competitive effect is not observed in the symmetric cases when all airlines offer bundles. Notably, a small change in flight revenue can outweigh a large gain in ancillary revenue.

The second sensitivity test is on the ancillary cost incurred by airlines for every ancillary sold. The parameters of the test cases are outlined in Table 5.2. Holding the ancillary WTP X_a and W_a constant, we change the cost from \$0 to \$30 for all airlines. As a result, the pricing of the ancillary changes in both the No DB baseline and the DB cases.

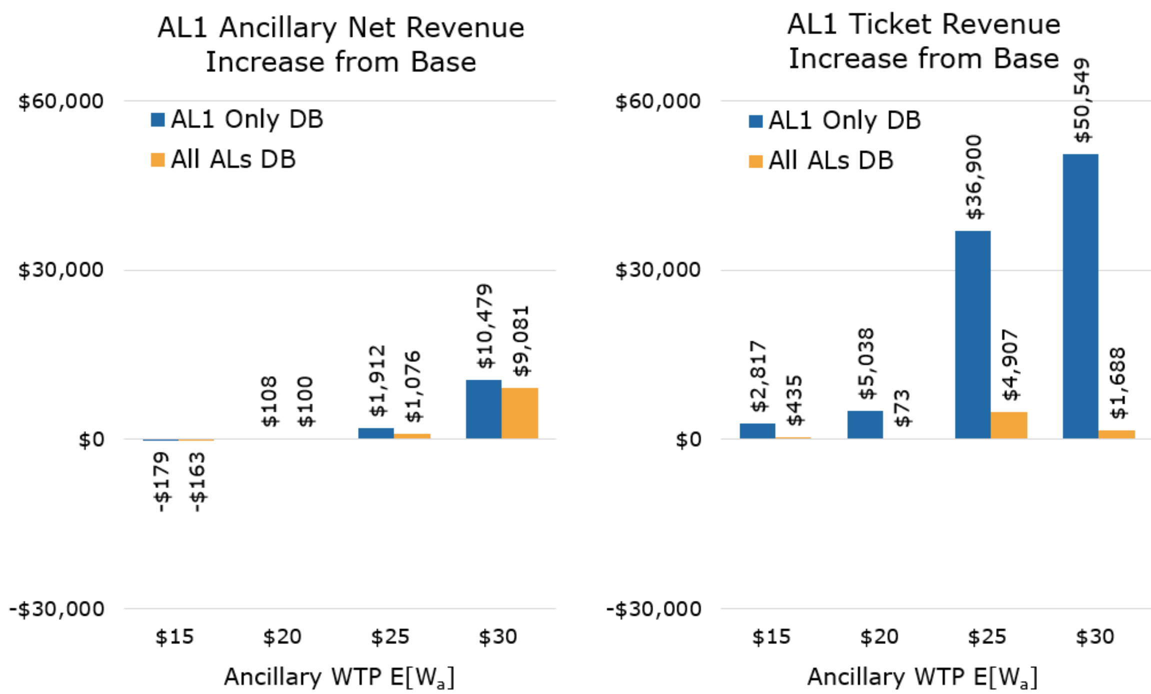

Figure 5-13: Increase of net ancillary and flight revenue with DB under a variety of ancillary WTP distributions

Figure 5-14 shows that dynamic bundling can increase total revenue by +0.9% in symmetric cases when the ancillary service is free to provide and by +2.6% in asymmetric cases. The asymmetric revenue gains for Airline 1 lead to losses at airlines 2-4, as passengers are attracted by the cheaper bundle offer at Airline 1.

The revenue gains are explained by Figure 5-15. The myopic optimal selling price of the ancillary service in the No DB case reduces with ancillary cost. This makes the ancillary service more popular and affordable, while at the same time the reduction in cost increases the margin for the airline. The ancillary service is in high demand at $0 cost: 78% of passengers choose to purchase it when given the choice. As a result, the bundles offered by DB are very attractive as it discounts the ancillary service even further. At low cost, dynamic bundling chooses to show the bundle offer set to >99% of all customers.

The impact on revenue is shown in Figure 5-16. As the cost decreases, the ancillary net revenue gains of dynamic bundling increase in both symmetric and asymmetric cases. Moreover, in the asymmetric case, Airline 1 can increase its flight revenue by

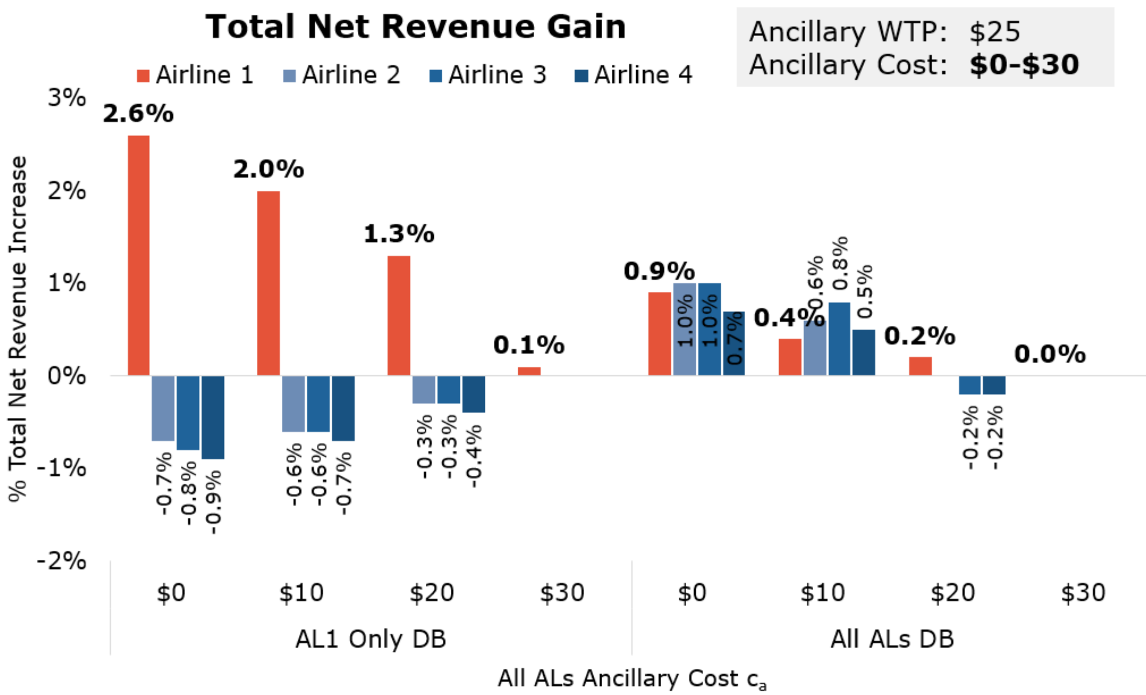

Figure 5-14: Variation of dynamic bundling (DB) net revenue gains with ancillary cost

undercutting its a la carte competitors with the bundle offers.

By testing DB, we observed how the balance between ancillary WTP and ancillary cost strongly influences when bundles are offered by the DOG offer set selection algorithm. Bundling can increase revenue when ancillary WTP is relatively high compared to the ancillary cost. Especially in asymmetric cases dynamic bundling is a competitive advantage and attracts new customers from a la carte competitors, thus increasing the airline's flight revenue. However, in symmetric cases the revenue gains are much smaller, coming primarily from increased ancillary revenue alone. No bundles were offered and no revenue benefits realized when ancillary WTP was relatively low compared to the cost.

5.4 DOG with Different Ancillary Parameters

With the insights gained from dynamic bundling, we now repeat these sensitivity tests on the ancillary WTP distributions X_a/W_a and the ancillary cost c_a with DOG.

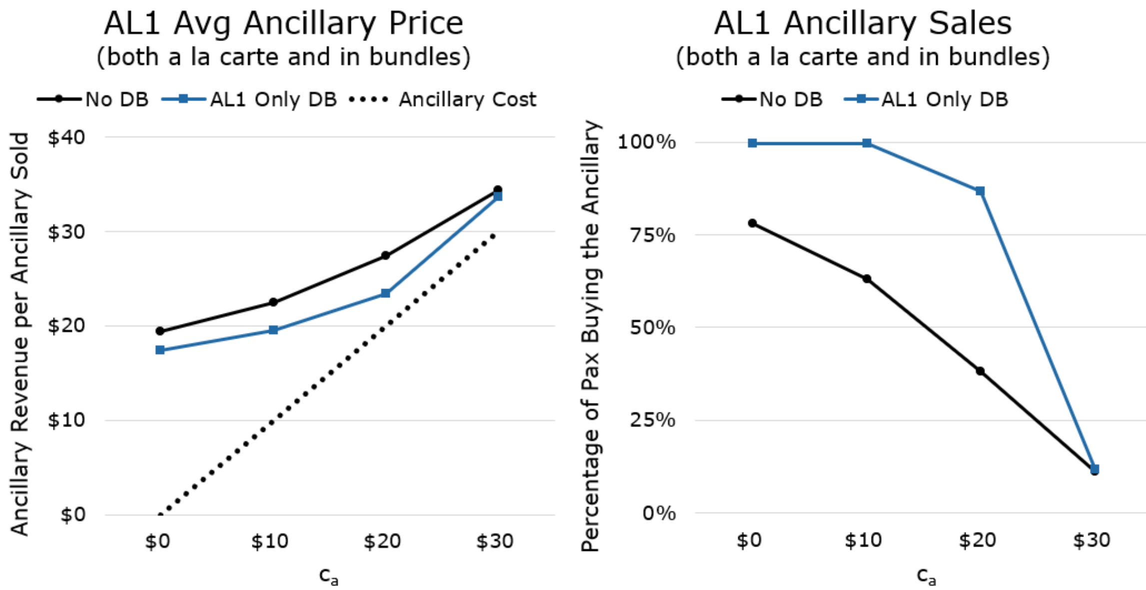

Figure 5-15: Variation of ancillary prices and purchases with ancillary cost

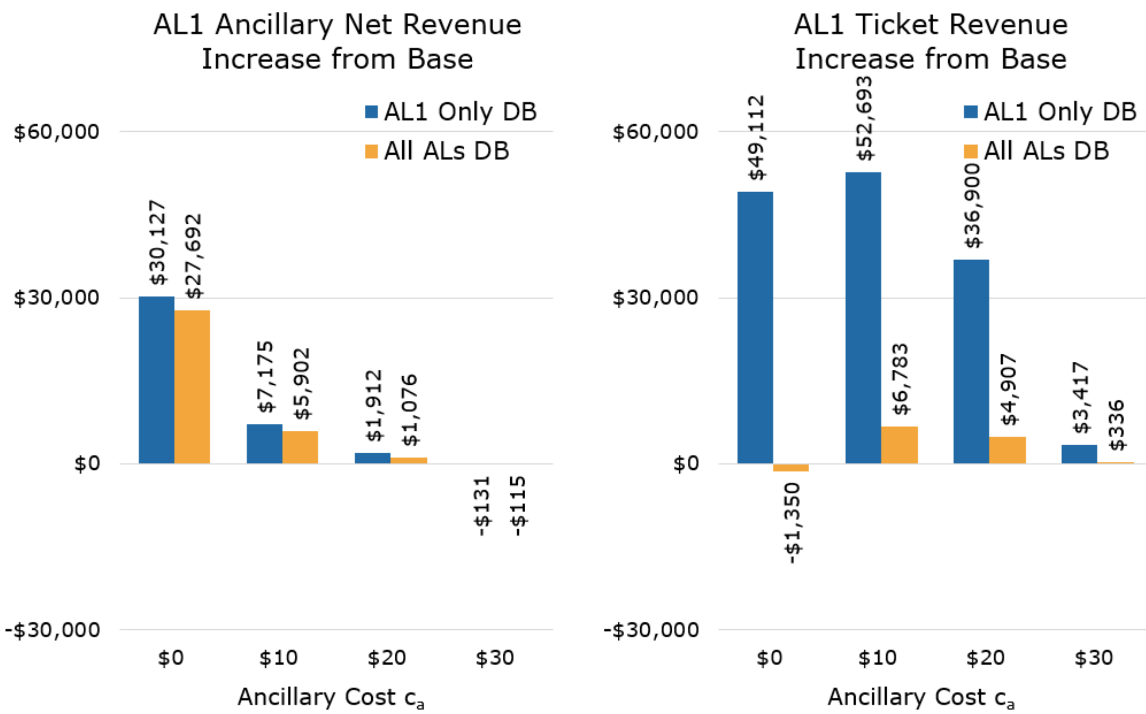

Figure 5-16: Increase of net ancillary and flight revenue with DB at a variety of ancillary costs

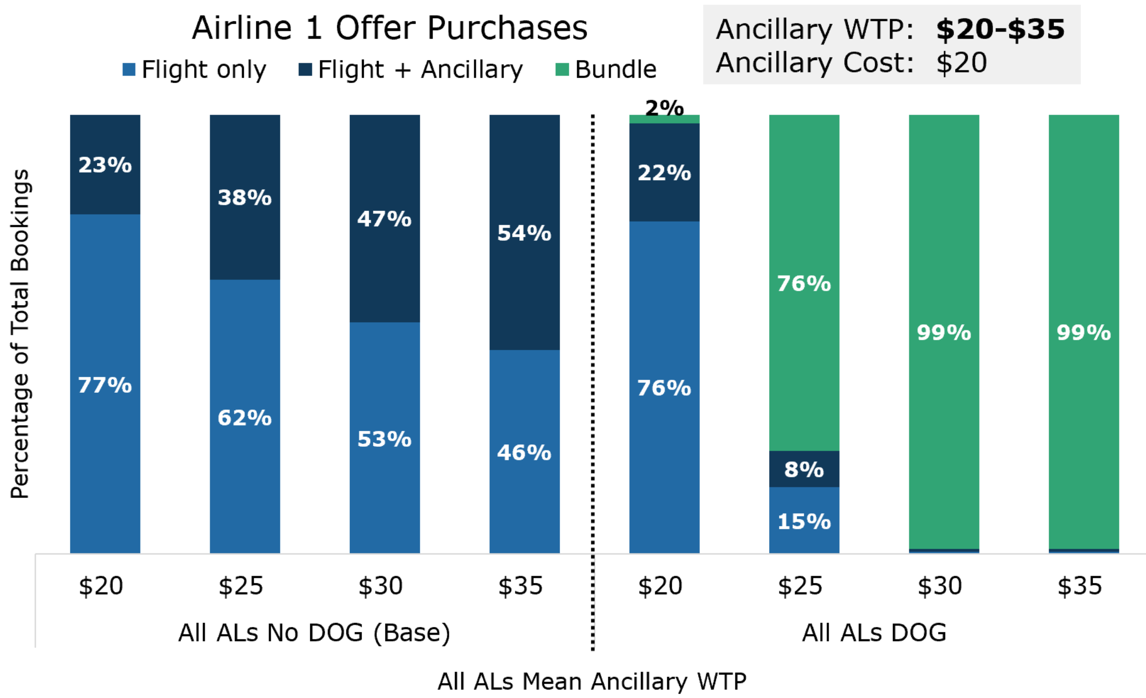

Figure 5-17: Airline 1 bookings by purchased offer for symmetric DOG and No DOG with different ancillary WTP W_a

In the first test, we vary the mean and standard deviation of the normally distributed ancillary WTP distribution while holding the cost constant. Both the underlying passenger ancillary WTP X_a and the airline's estimated WTP distribution W_a change from a mean of \$20 to \$35. Refer to Table 5.1 in the previous section (5.3) for the detailed description of this test. In the second test, we hold the ancillary WTP constant at the baseline setting while varying the ancillary cost c_a. As shown in Table 5.2, the airlines adjust their ancillary prices p_a accordingly to optimize their margin on the ancillary service in both the No DOG and DOG cases.

In terms of offer set selection, full DOG behaves the same as dynamic bundling alone. The bundled offer set has higher expected revenue and is offered the majority of the time when the mean ancillary WTP is higher than the ancillary cost. On the other hand, DOG offers the a la carte offer set when the mean ancillary WTP is smaller or equal to the ancillary cost. This trend is seen in the offer sets purchased by passengers of Airline 1 at various ancillary WTP (Figure 5-17) and ancillary cost (Figure 5-18). These results are very similar to those seen in dynamic bundling alone

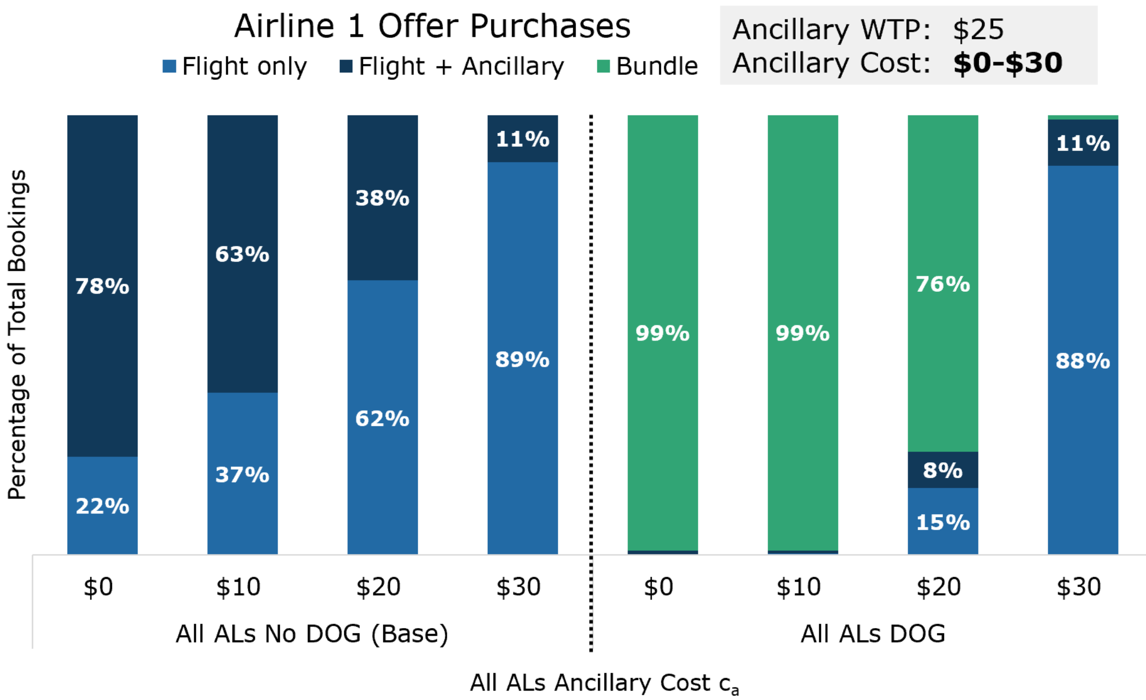

Figure 5-18: Airline 1 bookings by purchased offer for symmetric DOG and No DOG with different ancillary cost c_a

(Figure 5-11).

We observe that when the bundled offer set is offered by DOG, total net revenue is improved by an incremental 1-2% over dynamic a la carte in both asymmetric and symmetric tests. This leads to a total net revenue gain over the No DOG base of up to 12% in asymmetric and 7% in symmetric cases. In particular, the highest gains were achieved when WTP is much higher than the cost: Either with \$0 ancillary cost and \$25 mean WTP, or with \$35 mean WTP and \$20 cost.

Finally, we study how sensitive DOG is to over- and underestimation of ancillary WTP. We would like to understand the behavior of the algorithm if the airline introduces a bias in their estimate W_a of the actual underlying passenger ancillary WTP X_a in the simulation. In all tests, we maintain $X_a \sim \mathcal{N}(\$25, \$7.5^2)$ as normally distributed with a mean of \$25 and the ancillary cost to the airline at $c_a = \$20$. All four airlines misestimate the mean ancillary WTP by up to \$5 as indicated by Table 5.3. In the No DOG baseline, they sell the ancillary service at the listed myopic optimal price p_a, which is based on the misestimated W_a. In DOG, the algorithm computes

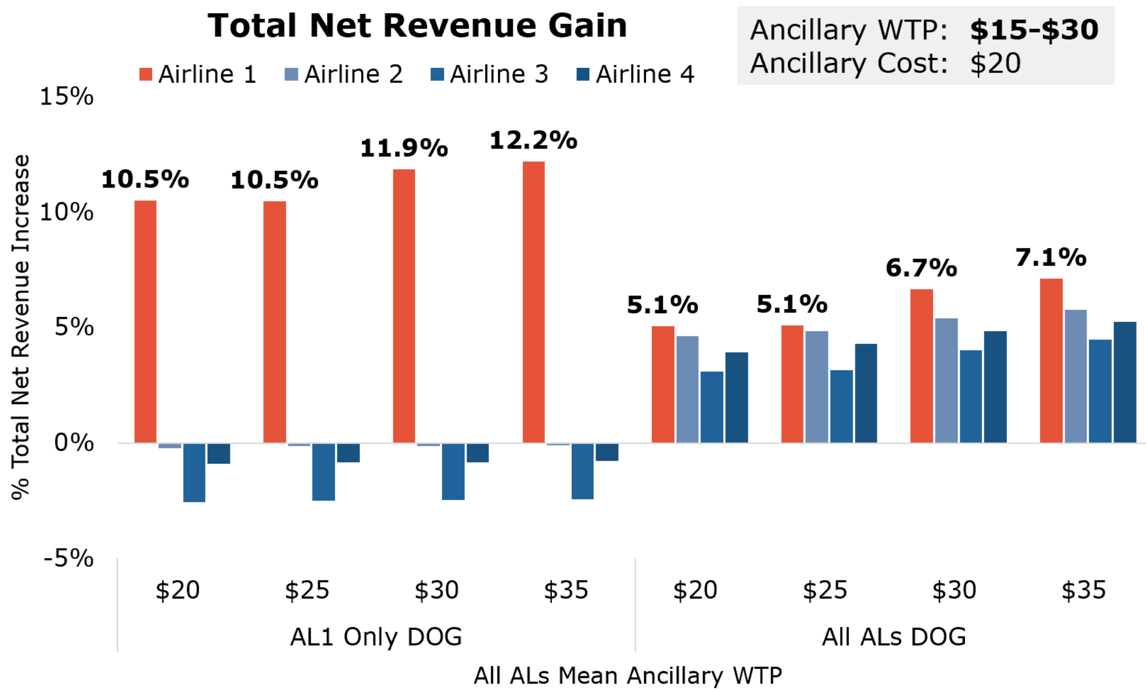

Figure 5-19: DOG total net revenue gains with varying ancillary WTP W_a at constant ancillary cost $c_a = \$20$

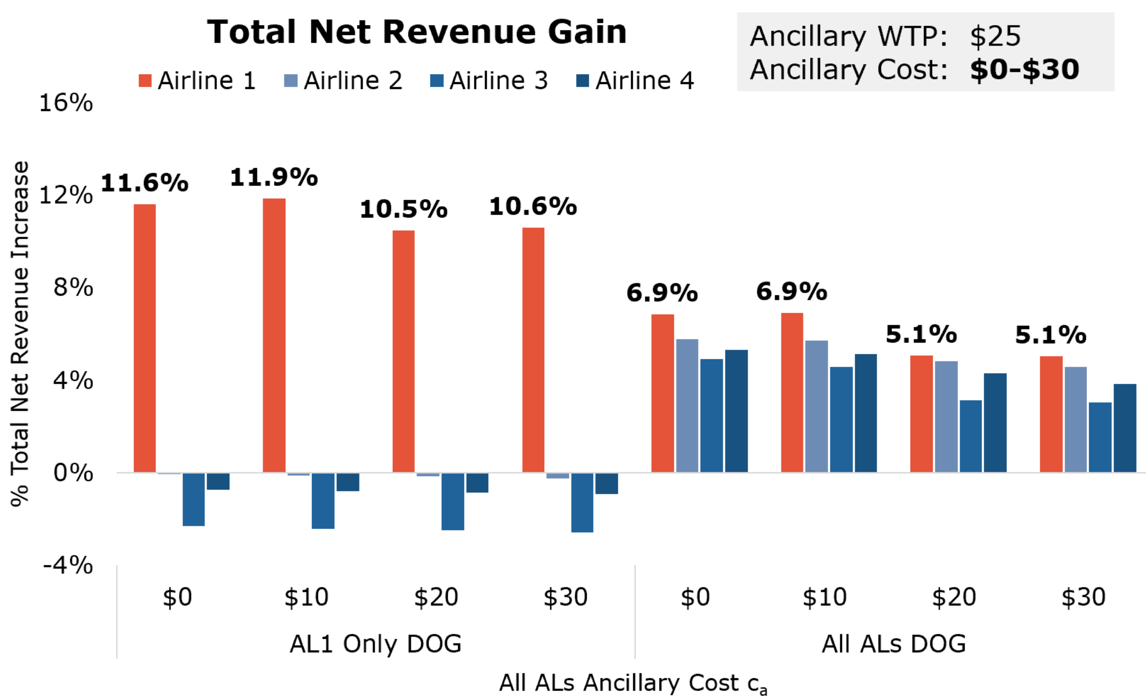

Figure 5-20: DOG total net revenue gains with varying ancillary cost c_a at constant ancillary WTP $E[W_a] = \$25$

WTP Sensitivity	Case 1	Case 2	Case 3 (Base)	Case 4	Case 5
Actual $E[X_a]$	\$25	\$25	\$25	\$25	\$25
Mean $E[W_a]$	\$20	\$22.5	\$25	\$27.5	\$30
Stdev $\sigma[W_a]$	\$6	\$6.75	\$7.5	\$8.25	\$9
Ancillary Cost c_a	\$20	\$20	\$20	\$20	\$20
Myopic optimal p_a^*	\$24.51	\$25.90	\$27.41	\$29.00	\$30.65

Table 5.3: Sensitivity of DOG and No DOG to misestimation of ancillary WTP $W_a \sim \mathcal{N}(E[W_a], \sigma[W_a]^2) \neq X_a$ at all airlines

the offer prices dynamically based on W_a.

First, we study the sensitivity of the baseline without DOG, where all airlines use traditional RM and sell the ancillary service a la carte at the myopic optimal price. In the baseline, Figure 5-21 shows that myopic ancillary prices indeed maximize ancillary net revenue when the ancillary WTP is accurately estimated $W_a = X_a$. This is an expected result since myopic prices were optimized to maximize expected net ancillary revenue $p_a - c_a$. However, in the PODS simulation we also observe that myopic ALC prices do not maximize total net revenue for the airline when passengers choose between flights, itineraries and ancillaries using concurrent choice behavior. Instead, a lower ancillary price can lead to an increase in flight revenue for the airline. While the increase in flight revenue is relatively small (+0.2%), it outweighs the small decrease in ancillary net revenue as the ancillary service accounts for only a small proportion of total revenue in this simulation. From this test, we learned that in an a la carte setting with concurrent choice behavior, total net revenue is maximized when the ancillary service is priced slightly below the myopic optimal price.

In dynamic offer generation, W_a also controls the offer set selection. When the airline overestimates W_a, we have seen in Figure 5-17 that DOG offers more bundles. When it underestimates W_a, DOG chooses the a la carte offer set more frequently. This additional effect makes DOG more sensitive to misestimation than the (a la carte) No DOG baseline. Without DOG, the difference in Airline 1's total net revenue at $E[W_a] = E[X_a] = \$25$ and the highest observed point at $E[W_a] = \$22.5$ was \$548, an increase of 0.02% as shown by the black bars in the center graph in Figure 5-22. However, when all airlines use DOG symmetrically, the total net revenue is

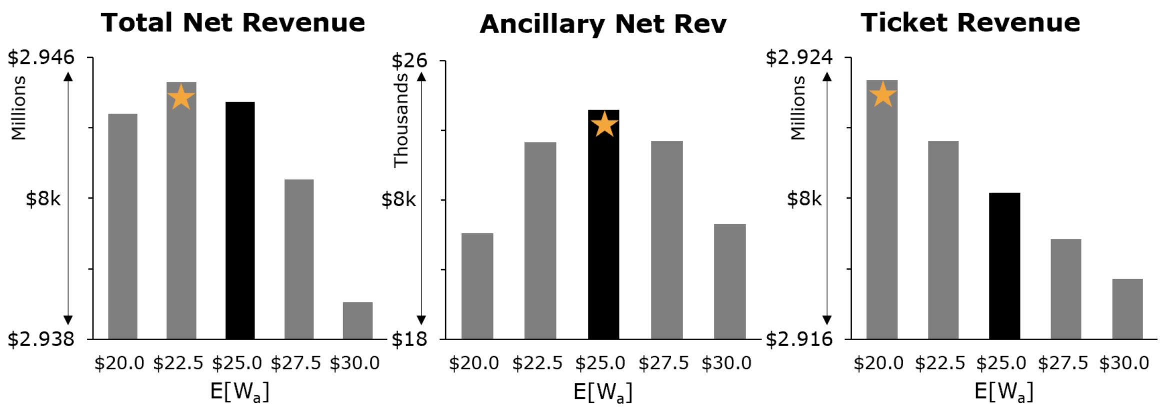

Figure 5-21: Airline 1 total net revenue, ancillary net revenue and flight revenue at different estimated ancillary WTP W_a

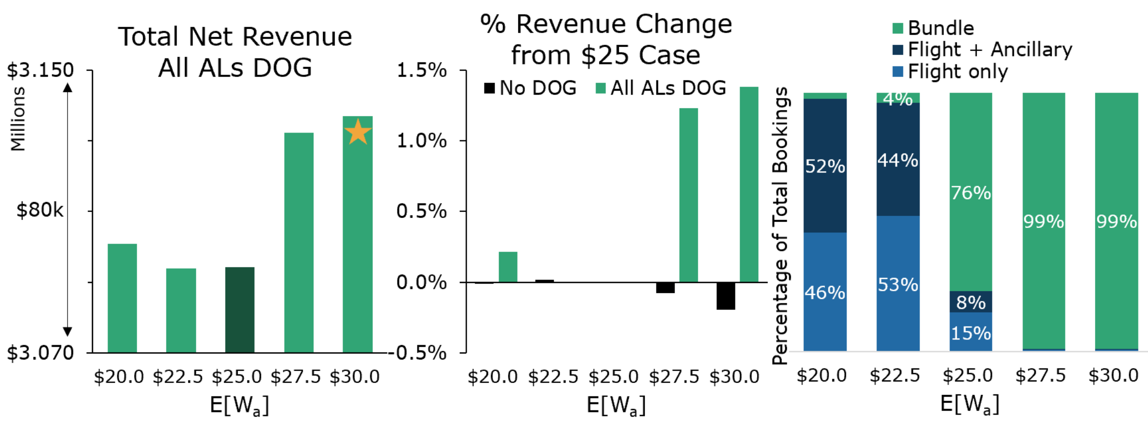

Figure 5-22: Impact of misestimated ancillary WTP W_a when all airlines use DOG symmetrically on Airline 1's total net revenue and offer purchases

maximized when W_a is overestimated. The difference in total net revenue at $E[W_a] = E[X_a] = \$25$ and the highest observed point at $E[W_a] = \$30$ is \$42,733 (+1.38%). This shows that DOG revenues are significantly more sensitive to the ancillary WTP parameter. In this simulated scenario, total net revenue is maximized when all airlines overestimate W_a and offer the bundled offer set 99% of the time.

5.5 DOG in a More Realistic Environment

Compared to an airline that does not use revenue management systems and performed manual seat inventory management, a leg-based revenue management system can increase passenger revenues by 4-6% (Belobaba, 1989) (Smith et al., 1992). In simulations, an origin-destination network RM system can add an incremental 1-2% in revenue (Belobaba, 2002). With this context in mind, it is clear that the total net revenue increases observed in our DOG baseline (Figure 4-15) might not be achievable in reality. While the baseline 5-10% revenue gains we previously observed were a good foundation to study the source of DOG revenue gains and the sensitivity of the algorithm, we strive to calibrate a more realistic case in this section to assess the potential impacts of DOG in a realistic market environment.

In this section, we simulate dynamic offer generation with one ancillary service in the context of the largest current source of airline ancillary revenue: checked baggage fees (IdeaWorks, 2018). All four airlines in this PODS network U10 charge a fee for the first checked bag, which is the ancillary service. We exclude other ancillary services from the simulation, including the second or more checked bags. An airline with a dynamic offer generation system can now decide for each passenger request, whether to offer the checked bag *a la carte* for an additional fee, or whether to include it in a *bundled* offer set. As before, the offer set selection can vary dynamically based on the flight price and the identified passenger type (business or leisure).

As before, we use DOG in a network with fully unrestricted fares. All airlines use sell-up WTP forecasting with a network RM optimizer (DAVN). However, steps were taken to reduce the pricing asymmetry introduced by DOG, which was able to generate large revenue gains by charging lower fares in the last two weeks before departure than the traditional RM system. The new simulation differs from the baseline settings introduced in Chapter 4.1.3 in the following four ways:

1. **Lower sell-up rate** used in the forecast and fare adjustment for both No DOG and DOG cases. The airlines use a FRAT5d sell-up curve instead of FRAT5c, which reduces their estimate of customer conditional flight WTP. The underlying maximum WTP of the simulated passengers in unchanged. The RM system believes that fewer passengers can afford to purchase the higher fare classes. As a result, the more seats in lower fare classes become available and thus the average flight price drops.

2. **Tighter bounds on DOG offer prices.** The width parameter is reduced from $b = 1.0$ to $b = 0.5$. The ability of DOG to adjust its prices away from the lowest available fare determined by the RM system is reduced by half. As a result, the average flight prices in DOG are going to be more similar to the prices in the No DOG baseline. This better reflects the limited ability of an airline to change its pricing without eliciting a competitive response in a realistic market.

3. **Lower DOG Q multipliers**, which represent the airline's estimate of conditional flight WTP for each passenger segment. In the previous tests, the airlines estimated the mean conditional flight WTP using a *Q multiplier* of 3.0 times the lowest filed fare in the market for passengers and 1.2 for leisure passengers. We use new values of 2.7 and 1.1, respectively. As the sell-up rate in the RM system is reduced, the Q multipliers in DOG are adjusted accordingly to lower the DOG prices quoted to passengers. The new values were chosen to match the DOG airline's average fares as closely as possible to the No DOG airline's average fares, simulating a competitive equilibrium.

4. **Differentiated ancillary parameters** between business and leisure passenger segments. As we model the ancillary service as a checked baggage fee, leisure passengers are willing to pay more, while some business passengers travel without checked baggage at all. The parameters are chosen such that the ancillary revenue represents around 3% of the airlines' total passenger revenue, which is in line with American Airlines' 2018 revenue from baggage fees.

Specifically, the following parameters were used for the ancillary service: Each leisure passenger's ancillary WTP is normally distributed with a mean of \$31 ($X_a \sim \mathcal{N}(\$31, (0.3 \cdot \$31)^2)$). Half of all generated business passengers have a normally distributed ancillary WTP with a mean of \$25 ($X_a \sim \mathcal{N}(\$25, (0.3 \cdot \$25)^2)$), while the other half does not travel with checked baggage and has an ancillary WTP $X_a = 0$. In the initial test, all DOG airlines can distinguish between business and leisure passenger types with 100% accuracy. They also accurately estimate both passenger types'

Airline 1	No DOG	AL1 Only DOG	All ALs DOG
Total Net Revenue	$2,858,341	(+**4.6%**) $2,990,873	(+**1.7%**) $2,905,600
RPM	23,460,878	23,621,755	23,525,016
ASM	28,588,660	28,588,660	28,588,660
Net Yield	$0.1218	$0.1266	$0.1235
Load Factor	82.06%	(+**0.6pts**) 82.63%	(+**0.2pts**) 82.29%

Table 5.4: Overall simulated results for Airline 1 of dynamic offer generation (DOG) with medium demand and new parameters

ancillary WTP distributions $W_a = X_a$. All airlines incur a cost of $c_a = \$25$ per ancillary service sold. When an airline does not use DOG, they are unable to segment the ancillary price between business and leisure passengers. As a result, they sell the ancillary service a la carte at a unified price of $p_a = \$33.59$, which is the myopic optimal price that maximizes ancillary net revenue under the 39%/61% mix of business and leisure bookings seen in the medium demand PODS simulation.

Compared to the original DOG baseline (Figure 4-15), the revised parameters lead to smaller total net revenue increases as airlines switch from traditional RM to DOG. As seen in Figure 5-23 and Table 5.4, when Airline 1 asymmetrically implements DOG under medium demand levels, their net revenue increases by +4.6% (previously: +10.5%) at the expense of the remaining airlines with traditional RM. This increase is achieved with a load factor increase of +0.6pts, whereas previously asymmetric DOG decreased the airline's load factor by -1.3pts. As all four airlines symmetrically use DOG, Airline 1's net revenue gain reduces to +1.7% (previously: +5.1%) with a load factor increase of +0.2pts (previously: +0.6pts). These more conservative revenue increases are more achievable under real competitive market conditions.

Figure 5-24 explains how the new parameters reduced the revenue gains from DOG. In the previous baseline (indicated with transparent lines), we observed that DOG undercut the traditional RM system's prices in the last two weeks before departure, which allowed the airline to gain a lot of high-revenue business passengers. To protect additional seats for close-in demand, DOG charged higher advance fares than the No DOG airlines. This introduced a strong shift in the booking curve compared to the No DOG case (Figure 4-8), allowing large revenue gains to occur. With the

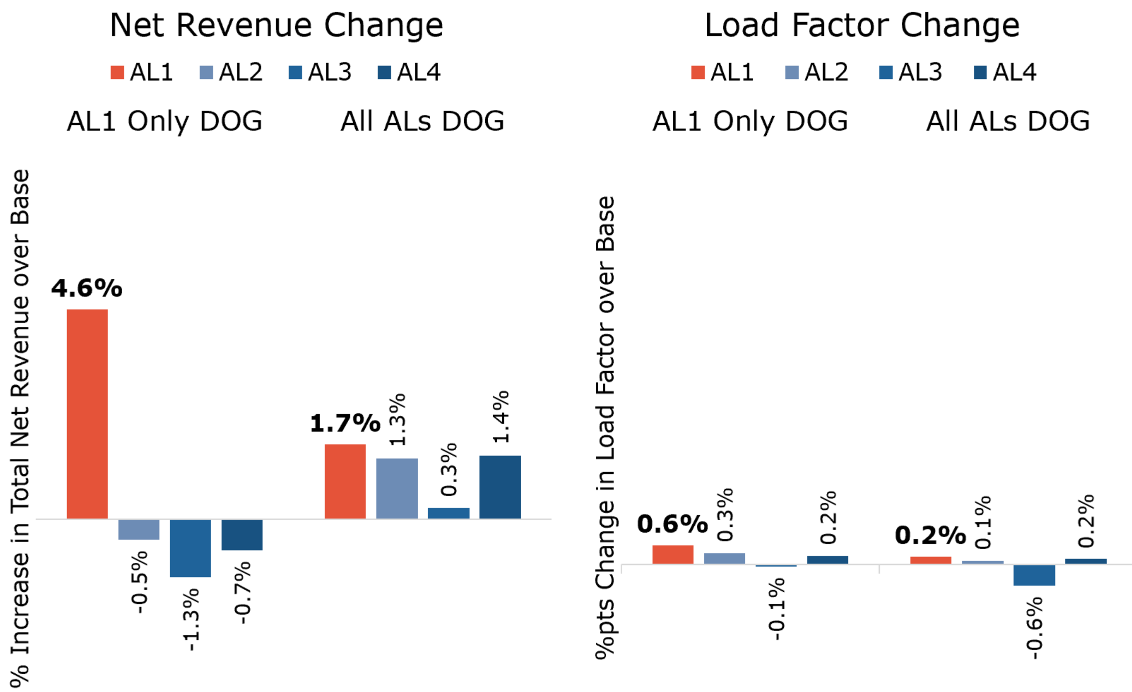

Figure 5-23: Overview of DOG net revenue gains and load factor changes in symmetric and asymmetric tests (medium demand level, ancillary service modeled as checked baggage fee)

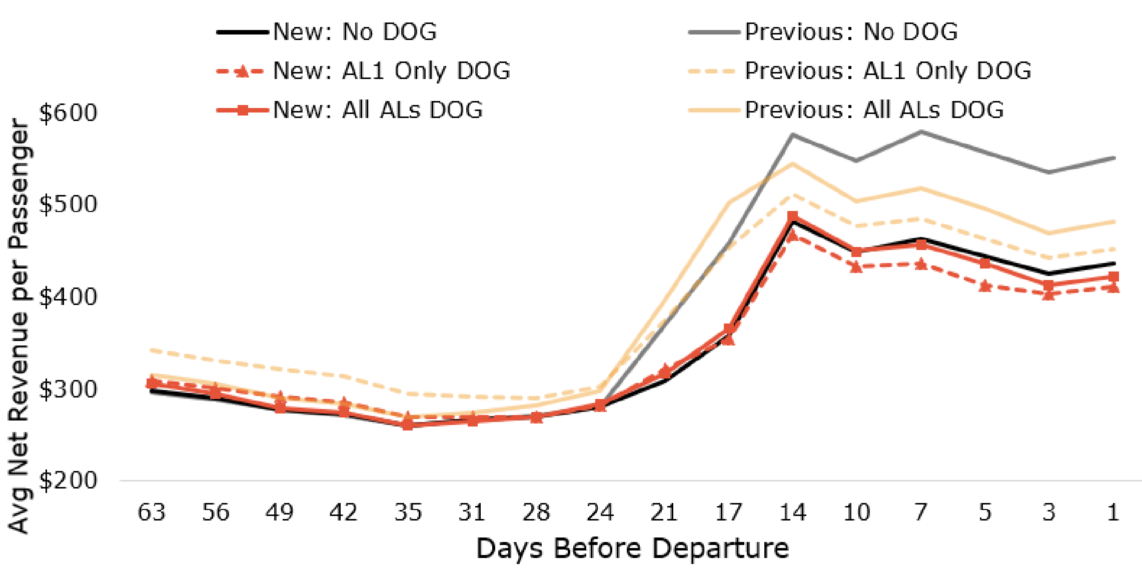

Figure 5-24: Airline 1 average net revenue per passenger for both old (medium demand, FRAT5c, Q=3.0/1.2, b=1.0) and new parameters (medium demand, FRAT5d, Q=2.7/1.1, b=0.5): Symmetric DOG fares are very similar to No DOG baseline fares with new parameters

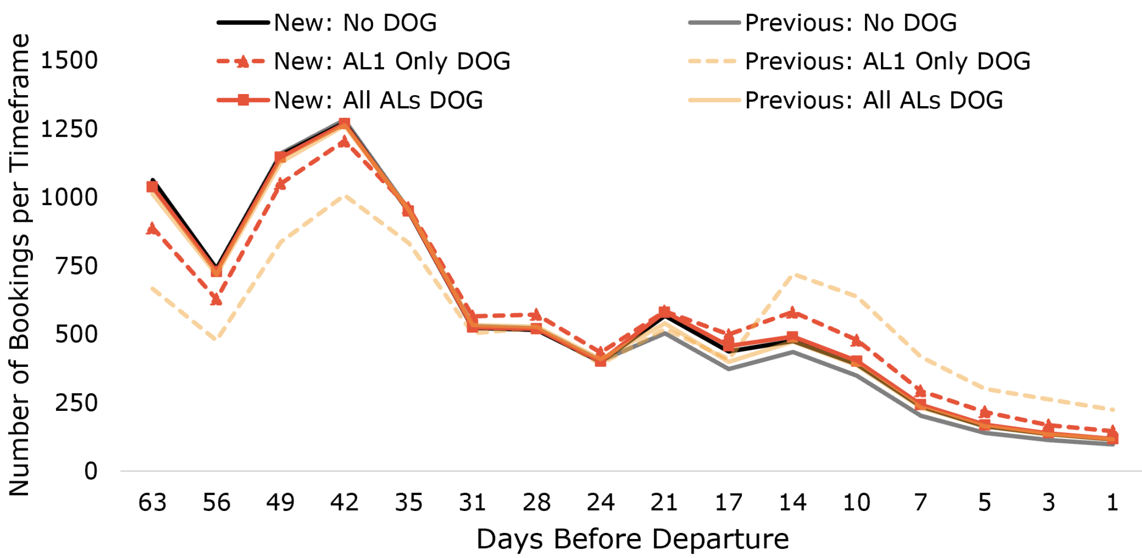

Figure 5-25: Airline 1 bookings by timeframe for both old (medium demand, FRAT5c, Q=3.0/1.2, b=1.0) and new parameters (medium demand, FRAT5d, Q=2.7/1.1, b=0.5): DOG introduces a smaller booking curve shift with new parameters

new parameters, the symmetric DOG prices are much closer to the No DOG prices. In particular, the switch to a less aggressive sell-up curve (FRAT5d) reduced No DOG prices close-in. Close-in DOG prices are lower than before and better matched with the new No DOG baseline throughout the booking window using Q multipliers of 2.7/1.1. This reduces the effect of DOG undercutting the traditional RM's pricing, lowering the revenue gains observed as well as the booking curve and market share shift (Figure 5-25). Most importantly, DOG remains revenue positive in both asymmetric and symmetric tests, even after reducing these competitive effects.

Dynamic offer generation continues to choose between offering the a la carte and bundled offer set for each booking request. Across all passengers on Airline 1, Figure 5-26 shows the breakdown of offers purchased with the revised ancillary parameters, which are modeled as a checked baggage fee. In the No DOG case, when the ancillary service is optional for all passengers, 26% of passengers purchase it at the myopic optimal price of \$33.59. 39.0% of leisure passengers chose to purchase the service, while only 6.3% of business passengers purchased the service, given their lower willingness-to-pay. In the symmetric case when all airlines use DOG (All DOG), 36% of all passengers were offered no choice and the ancillary service was already bundled

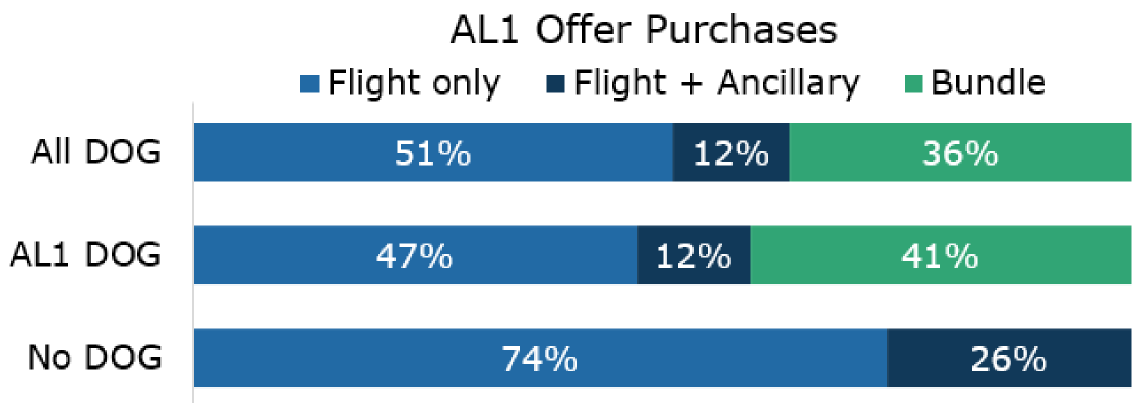

Figure 5-26: Airline 1 bookings by offer purchased: Proportion of overall bundle purchases reduced from 76% to 36% in symmetric DOG with new differentiated ancillary parameters

in their fare. The remaining 63% of passengers could purchase the service at a segmented myopic optimal price of \$30.64 for business passengers or \$34.10 for leisure passengers. Since DOG was able to distinguish passenger types with 100% accuracy, all business passengers were offered the a la carte offer set, which maximized expected net revenue given the low WTP for the ancillary service. On the other hand, 60% of leisure passengers were offered the bundle, whenever its expected net revenue was higher than the a la carte offer set. This generally occurred at higher flight prices.

In reality, the passenger type identification cannot be 100% accurate and airlines sometimes misestimate the true passenger WTP. We explore this in a sensitivity test, in which we reduce the probability of correct identification of each passenger from 100% to 50% for all airlines simultaneously. Whenever passengers are incorrectly identified, they are shown the offer set and prices that correspond to the other passenger segment (business or leisure). Airline 1's total net revenue gain in both asymmetric and symmetric cases compared to the traditional RM baseline, where no passenger type segmentation is performed, is shown in Figure 5-27. The previous DOG baseline was shown to be sensitive to identification accuracy in Figure 5-4. However, the new parameters show a strong resilience of DOG revenue to changes in identification accuracy, with asymmetric gains between 4.1%-4.6% (previous: 5.8%-10.5%) and symmetric gains between 1.2%-1.7% (previous: -0.6%-5.1%). This is largely

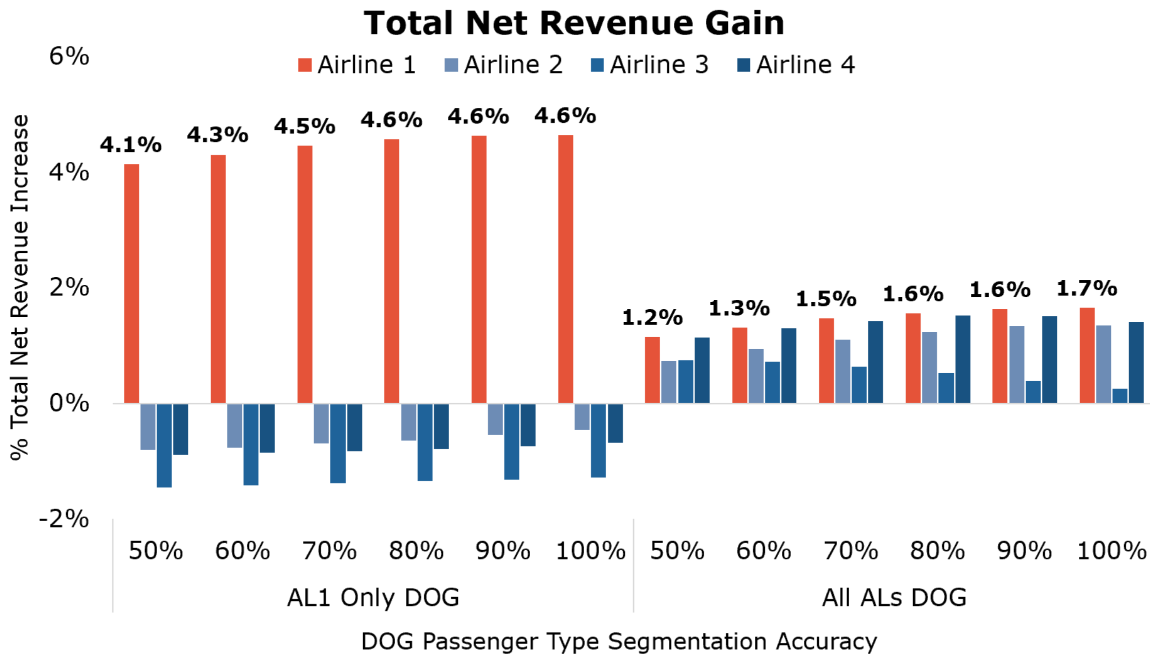

Figure 5-27: Airline 1 total net revenue gain with the asymmetric and symmetric use of DOG at varying passenger type identification accuracy (medium demand, FRAT5d, Q=2.7/1.1, b=0.5)

attributed to the reduced bounding parameter b=0.5, which reduces the degree of dynamic price adjustment performed for both passenger segments. The prices for business and leisure passengers are closer to each other, which reduces the impact of misidentification.

5.6 Summary

In this chapter, we performed extensive sensitivity tests on the performance of DOG measured by simulated revenue gain. We showed that competitive effects dominate the simulated revenue impacts of DOG, as they would in reality. The flight-related Q multipliers that are used by DOG to estimate flight WTP and associated passenger type identification accuracy were shown to have the biggest impact on DOG revenue gains, as they control the flight prices. The revenue gains were highest when DOG parameters were set to generate prices that undercut traditional airlines, since there was no competitive response in our simulations.

Comparatively, ancillary services generate only a small proportion of total airline revenue and thus have a smaller impact on total revenue gains. Nonetheless, there is a revenue benefit of 0-2% from asymmetrically bundling the ancillary service with the flight. This is maximized when bundles are offered to passengers with high ancillary WTP and/or for ancillary services that have a low cost of provision. There exists value in offering specific ancillary bundles to passenger segments that have a high WTP for those services, but it can be challenging to accurately segment passengers and estimate their ancillary WTP.

Finally, we tested DOG with revised parameters that more accurately reflect real market conditions. Using our insights from the sensitivity tests, DOG was calibrated to reduce the degree of flight price segmentation between business and leisure passengers. As DOG flight prices are then more similar to the No DOG baseline, simulated revenue gains from DOG are more modest, up to +1.7% in the symmetric case. The ancillary service was sometimes bundled for leisure passengers, representing a checked baggage fee that many business passengers have a lower willingness-to-pay for. We showed that dynamic offer generation can lead to revenue gains when applied to a single ancillary service. We also showed that by placing bounds on the price segmentation, the algorithm can be made more resilient to passenger type classification errors.

Chapter 6

Conclusions

Our research is guided by two major trends in the airline industry: the transformation of the airline distribution process with the New Distribution Capability (NDC) and the diversification of airline ancillary revenue streams. NDC has sparked interest in both *continuous pricing* and *segmented pricing*, for which airlines are seeking new revenue management (RM) models. Following the mission of RM, offering the right product to the right customer at the right price at the right time, NDC could allow airlines to better identify a customer's valuations and shape their products and prices accordingly. At the same time, the RM focus has shifted from one of flight ticket revenue maximization to *total revenue maximization* and *offer generation*.

As ancillary revenues become more important to airlines, new RM models are required to optimize ancillary pricing. Given the variety of ancillary services an airline sells, a new opportunity arises for *bundling* and the generation of different product offerings for different passenger segments. While a family on vacation may be most interested in checked baggage and seat assignment services, business travelers might prefer different services such as priority boarding or onboard internet.

In this thesis, we presented a new RM heuristic that enables *dynamic offer generation*, which utilizes the New Distribution Capability to generate customized offers and bundles. Together with a traditional RM system, it jointly optimizes prices for both flights and ancillaries. We used this model to study the benefits of dynamic offer generation and its implications for the airline industry.

6.1 Research Findings and Implications

Our literature review (Chapter 2) explored the extensive body of literature on airline flight revenue management. The prospect of continuous pricing has only heightened interest in the area, as the flight RM process evolves from one of assortment optimization (the allocation of seat inventory to a pre-defined set of products and price points) to one of price optimization (the determination of the optimal price point for a given seat). On the other hand, the area of ancillary RM is relatively nascent as researchers begin to explore how ancillary pricing can impact flight revenue and how passengers might choose their airline and itinerary based on the provided ancillary services. In particular, the area of bundle price optimization contains only limited literature, with initial economic studies proving that the joint bundle pricing of products can increase revenue for the retailer under the assumption of rational choice.

Our *dynamic offer generation* (DOG) model (Chapter 3) combines research in all these areas and integrates with existing airline revenue management systems to deliver new capabilities for airline distribution. We show how relatively simple heuristics can be used to optimize the prices of offers that combine flights with ancillary services. They can also determine the revenue-maximizing offer set and present it to the customer. We illustrate how the optimal price of an offer depends on the prices of all other offers in the offer set, as consumers make their purchase decision among all offers in the offer set. We show how the generated prices depend on the model's input variables in intuitive ways: The price generally increases with the cost of providing the service, as well as the passenger's willingness-to-pay (WTP). The price of a bundle of a flight and an ancillary service always costs more than the flight itself, but not more than purchasing both services a la carte. We also show that expected revenue can be increased by offering only the bundle and no option to purchase the standalone flight when the ancillary WTP sufficiently exceeds the cost of provision.

While our problem statement is very general, the solution heuristic provided relies on several strong assumptions that limit its optimality. In line with the PFDynA heuristic for dynamic flight price adjustment (Wittman, 2018), it assumes a normally distributed conditional flight willingness-to-pay. It also fundamentally assumes concurrent choice behavior, where passengers rationally weigh all offers against their willingness-to-pay. In reality, some passengers could be less well-informed about ancillary services and may choose differently. In addition, rational choice behavior has natural limitations in pricing, which is explored in the field of pricing psychology. The heuristic also relies on airlines to calibrate the model to their purchase data and accurately segment the booking requests, processes which we do not explore in this thesis.

We showed tests of DOG in the Passenger Origin-Destination Simulator (Chapters 4 and 5). Under baseline settings with fully unrestricted fares, willingness-to-pay forecasting, concurrent passenger choice behavior, and a single ancillary service, DOG delivers strong total net revenue gains of +10.5% in asymmetric tests, when one of the four competing airlines implements it. We show that this comes at the expense of the other airlines using traditional pricing and RM, which are being undercut by the DOG airline. When all airlines implement DOG symmetrically, the net revenue gains are reduced to 4-5%.

Further analysis showed that the majority of this increase comes from dynamic price adjustment and segmented pricing of the flight, which leads to a large increase in revenue as the algorithm discounts close-in fares to attract more bookings. In the traditional pricing structure that we used as the baseline, the price gap between two fare classes is largest in the highest booking classes sold close to departure. The continuous prices produced by the algorithm close the gap and sell the flights at intermediate price points. It is important to note that the simulations did not include competitive responses to the new pricing structure. In the real world, it is likely that traditional airlines would have lowered their fares accordingly, reducing the revenue gain observed by the DOG airline.

Our tests identify a smaller revenue benefit of offer set selection and bundling, when isolated from the impacts of flight pricing. Under concurrent choice behavior, it can be the revenue-maximizing strategy to only offer the bundled product, requiring the customer to purchase the ancillary service with the flight. This strategy alone can lead to simulated net revenue gains between +0.1% and +2.6% asymmetrically and between -0.2% and +0.9% symmetrically, depending on the ancillary WTP and cost assumptions. In the asymmetric case, the airline that offers bundles can attract customers from its a la carte competitors by charging a lower total price, further increasing its flight revenue. In the symmetric case, this competitive advantage disappears and the simulated net revenue gain reduces. Since the ancillary revenue represents only 3-5% of an airline's total revenue, any ancillary revenue gains can be outweighed by slight decreases in flight revenue.

In a final test, we sought to reduce the potentially exaggerated competitive effects observed in the simulation. After adjusting the pricing, limiting the dynamic price adjustment range, and recalibrating the ancillary service parameters to be primarily attractive to leisure passengers, we observe lower DOG net revenue gains of +4.6% asymmetrically and +1.7% symmetrically. These gains are less sensitive to passenger segmentation accuracy and offer a deeper insight into how DOG may actually perform under realistic market conditions in an unrestricted fare environment.

Our research shows that there exists a revenue benefit attributable to price segmentation for both flights and ancillary services. Airlines historically practiced flight price segmentation by offering a set of different price points for purchase with different restrictions that appeal to different customer segments. For instance, a flexible and refundable fare might be offered alongside a discounted, nonrefundable fare with round-trip restrictions. The effectiveness of such restrictions has reduced as competition by low-cost carriers has led to the introduction of less-restricted and unrestricted fare structures. Some low-cost carriers only offer one price point on a flight at any given time. Our simulations show that revenue can be gained by dynamically adjusting this price point for business and leisure passengers on a continuous spectrum, which could be enabled by the New Distribution Capability.

Furthermore, our results show that airlines can gain additional revenue by selectively bundling ancillary services with the flight. When such a bundle is offered to customers that have a high WTP for the included services, the airline can increase its revenue. In a competitive simulation, this revenue gain depends on the ancillary pricing of the competitors and can only be achieved when the bundle is cheaper than the competitors' offers. However, in all of our tests with DOG, the majority of revenue gains are attributed to the dynamic pricing of the flight, which undercuts the traditional flight prices, and not the bundling of ancillaries.

These observations open the door to further research on dynamic offer generation: With the New Distribution Capability, airlines will be able to adjust their pricing in real-time, potentially also in response to competitor pricing. In the future, RM algorithms could incorporate information about competitor's prices and their impacts on customer conditional WTP.

6.2 Suggested Future Research Directions

The dynamic offer generation model introduced in this thesis can be used as an initial step towards a world of offer generation and continuous pricing. However, its assumptions cause limitations that could be resolved with further research, namely:

- Intentionally, the heuristic uses a traditional class-based RM system as the backbone to perform demand forecasting and bid price calculations. While this is desirable to enable airlines with existing RM systems to adapt to the New Distribution Capability, it introduces additional complexity to the solution. For example, DOG bookings are currently assigned to a fare class, whose filed fare (instead of the continuous DOG offer price) is used in the RM forecast. In the future, DOG may be directly integrated into a classless RM system (Liotta, 2019) (Papen, 2020) that directly generates a demand forecast from the DOG prices paid. The conditional WTP estimates used in the RM system and the DOG heuristic could also be unified, alleviating the need to reconcile and bound the DOG-optimized prices to those generated by the RM system.

- We assumed that the airline only offers one price point at any given time on a flight. Airlines commonly sell flights at multiple price points that differ in their attributes, such as change or cancellation penalties. Further research could extend DOG to optimize multiple flight price points with different attributes.

- In our tests, the opportunity cost of selling a seat varied with the bid price provided by the traditional RM system, yet the cost of selling an ancillary service was fixed and known. To extend the model to capacity-constrained ancillary services, such as extra-legroom seating, a variable cost of the ancillary service could be used. At the limit, the ancillary service could have its own bid price that represents the opportunity of selling the last unit of ancillary capacity, which could be generated from a separate ancillary demand forecast.

- In this thesis, we primarily studied DOG with one ancillary service, where the algorithm decides to either offer the ancillary service a la carte (a *pure components* strategy) or bundled with the flight (*pure bundling*). When extended to two or more ancillary services, an additional pricing strategy can be explored: *mixed bundling*. In mixed bundling, an airline offers the individual ancillary services at an (a la carte) price, yet in the same offer set also offers a discounted price when more than one ancillary service is purchased. Such an offer set that combines individual component prices with bundle discounts can generate higher revenue under certain conditions (Adams and Yellen, 1976).

- In our model, we assume that all customers' conditional WTP for all flights and ancillary services are drawn from independent probability distributions. However, further research could be done to relax this assumption. According to the economics literature, bundling products that have a negatively correlated WTP can be a revenue-maximizing strategy (McAfee et al., 1989). Accounting for correlations in the DOG calculations may further increase the revenue benefit of dynamic bundling.

- Currently, the DOG price optimization is challenging to solve optimally, especially as the number of ancillary services increases. This is because the concurrent choice assumption used requires the convolution of different probability distributions, which is especially complex for the Gaussian conditional WTP distribution, for which we developed the μ-Heuristic (Chapter 3.2.4). Future research could explore different choice models, where the solution to the price optimization problem scales better to a large number of ancillary services and which can be more easily calibrated to real-world purchase behavior.

- The concurrent choice assumption prescribes a specific rational choice behavior that might not be able to model the real-world purchases of airline customers. As a result, it is not a given that the optimized DOG prices and selected offer sets would maximize expected revenue in a real implementation. For example, the ability of airlines to distribute information about offered ancillary services and their pricing is currently limited in indirect channels such as the Global Distribution Systems commonly used by travel agents. This reduces the price transparency and the ability of customers to make rational and fully-informed purchase decisions. A more flexible choice model may be able to better capture these diverse purchase behaviors observed in reality, as well as some aspects of pricing psychology such as the effect of anchoring and adjustment.

- In this thesis, we do not discuss the problem of model calibration and parameter estimation. Estimating passenger's conditional flight and ancillary WTP is a challenging problem on its own, and our results assume that these parameters are provided to the algorithm. Further research in this area could determine the achievable estimation accuracy and incorporate parameter uncertainty into the DOG algorithm, such that its solutions are robust for a range of parameters. This is particularly relevant for offer set selection, where offering a bundled offer set can lead to revenue losses, if the provided input parameters were inaccurate.

- We do not study the challenges of passenger segmentation either, we instead assume a given identification accuracy can be achieved. Further research in the area could develop new classification models that assign a booking request to a given customer segment, based on the parameters in the search request (destination, day-of-week, time-of-day, advance booking period, trip length, number of passengers, etc.).

The New Distribution Capability has opened the possibility of fundamentally rethinking airline distribution and removing decades of legacy distribution system restrictions. But the technological difficulties of integrating a completely new revenue management system or even the dynamic offer generation heuristic should not be underestimated. The human aspect of this change alone presents challenges: how can existing pricing and inventory teams at airlines work with an algorithm to optimize offer prices, but also respond to competitor action? As our results show, airlines that adapt and successfully employ the new capabilities enabled by NDC such as price segmentation, continuous pricing and ancillary bundling may reap the rewards of being the first mover: In all of our tests, the revenue gains were higher for the DOG airline when its competitors were still using traditional pricing and RM than in the fully symmetric tests where all airlines used DOG.

There also remains future work for revenue management researchers in this area, as existing RM systems expand their reach into ancillary services and total revenue optimization. In particular, the offer set selection problem adds a new dimension to the revenue management problem, significantly expanding the solution space. As we have shown, the opportunity for revenue gains exists for airlines that optimize the offers shown to customers. As such, the mission of revenue management may have to expand from *selling the right product* to the right consumer at the right time and at the right price. Instead, the challenge now is to *sell the right set of products* to the right consumer at the right time and at the right prices.

Bibliography

Adams, W. J. and Yellen, J. L. (1976). Commodity bundling and the burden of monopoly. *The Quarterly Journal of Economics*, 90(3):475–498.

Airlines for America (2018). Domestic round-trip fares and fees. `https://www.airlines.org/dataset/annual-round-trip-fares-and-fees-domestic/`. Accessed: 2019-12-04.

Bakos, Y. and Brynjolfsson, E. (1999). Bundling information goods: Pricing, profits, and efficiency. *Management Science*, 45(12):1613–1630.

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

Belobaba, P. P. (1989). Application of a probabilistic decision model to airline seat inventory control. *Operations Research*, 37(2):183–197.

Belobaba, P. P. (1992). Optimal vs. heuristic methods for nested seat allocation. In *Presentation at ORSA/TIMS Joint National Meeting*.

Belobaba, P. P. (2002). Airline network revenue management: recent developments and state of the practice. In Jenkins, D., editor, *Handbook of Airline Economics*, chapter 10. McGraw-Hill, 2nd edition.

Belobaba, P. P. (2011). Did LCCs save airline revenue management? *Journal of Revenue and Pricing Management*, 10(1):19–22.

Belobaba, P. P. and Hopperstad, C. (2004). Algorithms for revenue management in unrestricted fare markets. In *Presentation at AGIFORS Reservations and Yield Management Study Group*.

Belobaba, P. P. and Weatherford, L. R. (1996). Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations. *Decision Science*, 27(2):343–363.

Ben-Akiva, M. and Gershenfeld, S. (1998). Multi-featured products and services: Analysing pricing and bundling strategies. *Journal of Forecasting*, 17(34):175–196.

Bockelie, A. (2018). Tests of dynamic offer generation with modified concurrent choice behavior. In *Presentation at PODS Consortium Meeting (Nice, FR)*.

Bockelie, A. (2019). *Ancillary services in the airline industry: Passenger choice and revenue management optimization*. PhD thesis, Massachusetts Institute of Technology.

Bockelie, A. and Belobaba, P. (2017). Incorporating ancillary services in airline passenger choice models. *Journal of Revenue and Pricing Management*, 16(6):553–568.

Bockelie, A. and Wittman, M. D. (2017). Some first thoughts on dynamic offer generation. In *Presentation at PODS Consortium Meeting (Montreal, QC)*.

Bockelie, A. and Wittman, M. D. (2018). Proposal: Modified concurrent choice behavior with IIA property. In *Presentation at PODS Consortium Meeting (Seattle, WA)*.

Boyd, A. E. and Kallesen, R. (2004). Practice Papers: The science of revenue management when passengers purchase the lowest available fare. *Journal of Revenue and Pricing Management*, 3(2):171–177.

Bratu, S. (1998). Network value concept in airline revenue management. Master's thesis, Massachusetts Institute of Technology.

Brumelle, S. L. and McGill, J. I. (1993). Airline seat allocation with multiple nested fare classes. *Operations Research*, 41(1):127–137.

Bulut, Z., Gürler, U., and Şen, A. (2009). Bundle pricing of inventories with stochastic demand. *European Journal of Operational Research*, 197(3):897–911.

CAPA Centre for Aviation (2019). Spirit Airlines pushes up ancillaries, offsetting other challenges. `https://centreforaviation.com/analysis/reports/spirit-airlines-pushes-up-ancillaries-offsetting-other-challenges-501102`. Accessed: 2020-03-23.

Carrier, E. (2008). *Modeling the choice of an airline itinerary and fare product using booking and seat availability data*. PhD thesis, Massachusetts Institute of Technology.

Chakravarti, D. and Paul, P. (2002). Partitioned presentation of multicomponent bundle prices: evaluation, choice and underlying processing effects. *Journal of Consumer Psychology*, 12(3):215–229.

Cléaz-Savoyen, R. (2005). Airline revenue management methods for less restricted fare strucutres. Master's thesis, Massachusetts Institute of Technology.

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

Curry, R. E. (1990). Optimal airline seat allocation with fare classes nested by origins and destinations. *Transportation Science*, 24(3):193–204.

Diwan, S. (2010). *Performance of dynamic programming methods in airline revenue management*. PhD thesis, Massachusetts Institute of Technology.

Ferreira, K. D. and Wu, D. D. (2011). An integrated product planning model for pricing and bundle selection using Markov decision processes and data envelope analysis. *International Journal of Production Economics*, 134(1):95–107.

Fiig, T., Goyons, O., Adelving, R., and Smith, B. (2016). Dynamic pricing - the next revolution in RM? *Journal of Revenue and Pricing Management*, 15(5):360–379.

Fiig, T., Isler, K., Hopperstad, C., and Belobaba, P. (2010). Optimization of mixed fare structures: Theory and applications. *Journal of Revenue and Pricing Management*, 9(1-2):152–170.

Fiig, T., Isler, K., Hopperstad, C., and Cleaz-Savoyen, R. (2005). DAVN-MR: A unified theory of O&D optimization in a mixed network with restricted and unrestricted fare products. In *AGIFORS Reservations and Yield Management Study Group Meeting, Cape Town, South Africa, May*.

Fiig, T., Isler, K., Hopperstad, C., and Olsen, S. S. (2012). Forecasting and optimization of fare families. *Journal of Revenue and Pricing Management*, 11(3):322–342.

Fiig, T., Le Guen, R., and Gauchet, M. (2018). Dynamic pricing of airline offers. *Journal of Revenue and Pricing Management*, 17(6):381–393.

Gallego, G., Li, A., Truong, V.-A., and Wang, X. (2016). Online resource allocation with customer choice.

Golrezaei, N., Nazerzadeh, H., and Rusmevichientong, P. (2014). Real-time optimization of personalized assortments. *Management Science*, 60(6):1532–1551.

Gorin, T. (2000). *Airline revenue management: Sell-up and forecasting algorithms*. PhD thesis, Massachusetts Institute of Technology.

Gürler, U., Öztop, S., and Şen, A. (2009). Optimal bundle formation and pricing of two products with limited stock. *International Journal of Production Economics*, 118(2):442–462.

Hanson, W. and Martin, R. K. (1990). Optimal bundle pricing. *Management Science*, 36(2):155–174.

Hao, E. (2014). Ancillary revenues in the airline industry: Impacts on revenue management and distribution systems. Master's thesis, Massachusetts Institute of Technology.

IdeaWorks (2018). Airline ancillary revenue projected to be $92.9 billion worldwide in 2018. `https://www.ideaworkscompany.com/wp-content/uploads/2018/11/Press-Release-133-Global-Estimate-2018.pdf`. Accessed: 2019-11-21.

IdeaWorks (2019). Airline ancillary revenue projected to leap to $109.5 billion worldwide in 2019. https://www.ideaworkscompany.com/wp-content/uploads/2019/11/Press-Release-142-Global-Estimate-2019.pdf. Accessed: 2019-11-21.

Johnson, M. D., Herrmann, A., and Bauer, H. H. (1999). The effects of price bundling on consumer evaluations of product offerings. *International Journal of Research in Marketing*, 16(2):129–142.

Kobayashi, B. H. (2005). Does economics provide a reliable guide to regulating commodity bundling by firms? A survey of the economic literature. *Journal of Competition Law & Economics*, 1(4):707–746.

Kök, A. G., Fisher, M. L., and Vaidyanathan, R. (2008). Assortment planning: Review of literature and industry practice. In *Retail supply chain management*, pages 99–153. Springer.

Lautenbacher, C. J. and Stidham, S. (1999). The underlying Markov decision process in the single-leg airline yield-management problem. *Transportation Science*, 33(2):136–146.

Liotta, N. J. (2019). Airline revenue management for continuous pricing: Class-based and classless methods. Master's thesis, Massachusetts Institute of Technology.

Littlewood, K. (1972). Forecasting and control of passenger bookings. In *Proceedings of the 12th AGIFORS Symposium*.

Lu, Y. (2019). Impacts of ancillary services on airline revenue management. Master's thesis, Massachusetts Institute of Technology.

Lufthansa (2018). Conference & roadshow presentation. https://investor-relations.lufthansagroup.com/fileadmin/downloads/en/charts-speeches/presentations/RS_Presentation_Q3_CS_Conference_29Nov.pdf. Accessed: 2019-12-03.

Madireddy, M., Sundararajan, R., Doreswamy, G., Nia, M. H., and Mital, A. (2017). Constructing bundled offers for airline customers. *Journal of Revenue and Pricing Management*, 16(6):532–552.

McAfee, R. P., McMillan, J., and Whinston, M. D. (1989). Multiproduct monopoly, commodity bundling, and correlation of values. *The Quarterly Journal of Economics*, 104(2):371–383.

McGill, J. I. and Van Ryzin, G. J. (1999). Revenue management: research overview and prospects. *Transportation Science*, 33(2):233–256.

Morwitz, V. G., Greenleaf, E. A., and Johnson, E. J. (1998). Divide and prosper: Consumers' reactions to partitioned prices. *Journal of Marketing Research*, 35(4):453–463.

Odegaard, F. and Wilson, J. G. (2016). Dynamic pricing of primary products and ancillary services. *European Journal of Operational Research*, 251(2):586–599.

Papen, A. R. R. (2020). Competitive effects of continuous pricing mechanisms in airline revenue management. Master's thesis, Massachusetts Institute of Technology.

Raghubir, P. and Valenzuela, A. (2006). Center-of-inattention: Position biases in decision-making. *Organizational Behavior and Human Decision Processes*, 99(1):66–80.

Ratliff, R. and Gallego, G. (2013). Estimating sales and profitability impacts of airline branded-fares product design and pricing decisions using customer choice models. *Journal of Revenue and Pricing Management*, 12(6):509–523.

Reyes, M. H. (2006). Hybrid forecasting for airline revenue management in semi-restricted fare structures. Master's thesis, Massachusetts Institute of Technology.

Robinson, L. W. (1995). Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes. *Operations Research*, 43(2):252–263.

Rodway, P., Schepman, A., and Lambert, J. (2012). Preferring the one in the middle: Further evidence for the centre-stage effect. *Applied Cognitive Psychology*, 26(2):215–222.

Schmalensee, R. (1984). Gaussian demand and commodity bundling. *Journal of Business*, 57:S211–S230.

Shukla, N., Kolbeinsson, A., Otwell, K., Marla, L., and Yellepeddi, K. (2019). Dynamic pricing for airline ancillaries with customer context. *arXiv preprint arXiv:1902.02236*.

Smith, B. C., Leimkuhler, J. F., and Darrow, R. M. (1992). Yield management at American Airlines. *Interfaces*, 22(1):8–31.

Smith, B. C. and Penn, C. W. (1988). Analysis of alternate origin-destination control strategies. In *Proceedings of the Twenty Eigth Annual AGIFORS Symposium, New Seabury, Massachusetts*.

Stigler, G. J. (1963). United States v. Loew's Inc.: A note on block-booking. *The Supreme Court Review*, 1(1):152–157.

Strauss, A. K., Klein, R., and Steinhardt, C. (2018). A review of choice-based revenue management: Theory and methods. *European Journal of Operational Research*, 271(2):375–387.

Talluri, K. and van Ryzin, G. (2004). Revenue management under a general discrete choice model of consumer behavior. *Management Science*, 50(1):15–33.

Taubmann, H. (2016). One year since Lufthansa's distribution cost charge – what we've learned. `https://amadeus.com/en/insights/blog/one-year-since-lufthansas-distribution-cost-charge-weve-learned`. Accessed: 2019-11-26.

Touraine, S. and Coles, H. (2018). Dynamic offer creation white paper. Technical Report October, IATA.

Van Ryzin, G. and McGill, J. (2000). Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels. *Management Science*, 46(6):760–775.

Vinod, B., Ratliff, R., and Jayaram, V. (2018). An approach to offer management: Maximizing sales with fare products and ancillaries. *Journal of Revenue and Pricing Management*, 17(2):91–101.

Walczak, D. and Kambour, E. (2014). Revenue management for fare families with price-sensitive demand. *Journal of Revenue and Pricing Management*, 13(4):273–290.

Walczak, D., Mardan, S., and Kallesen, R. (2010). Customer choice, fare adjustments and the marginal expected revenue data transformation: A note on using old yield management techniques in the brave new world of pricing. *Journal of Revenue and Pricing Management*, 9(1-2):94–109.

Westermann, D. (2013). The potential impact of IATA's New Distribution Capability (NDC) on revenue management and pricing. *Journal of Revenue and Pricing Management*, 12(6):565–568.

Wittman, M. (2018). *Dynamic pricing mechanisms for airline revenue management: Theory, heuristics, and implications*. PhD thesis, Massachusetts Institute of Technology.

Wittman, M. D. and Belobaba, P. P. (2018). Customized dynamic pricing of airline fare products. *Journal of Revenue and Pricing Management*, 17(2):78–90.

Wittman, M. D. and Belobaba, P. P. (2019). Dynamic pricing mechanisms for the airline industry: a definitional framework. *Journal of Revenue and Pricing Management*, 18(2):100–106.

Yadav, M. S. (1994). How buyers evaluate product bundles: A model of anchoring and adjustment. *Journal of Consumer Research*, 21(2):342–353.