Ancillary Services in the Airline Industry: Passenger Choice and Revenue Management Optimization

by

Adam Bockelie

Bachelor of Science, Massachusetts Institute of Technology (2011)
Master of Engineering, Massachusetts Institute of Technology (2012)

Submitted to the Department of Civil and Environmental Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Transportation

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2019

Signature redacted

Author ....
Department of Civil and Environmental Engineering
January 16, 2019

Signature redacted

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

Signature redacted

Accepted by ....
Heidi Nepf
Donald and Martha Harleman Professor of Civil and Environmental Engineering
Chair, Graduate Program Committee

Ancillary Services in the Airline Industry: Passenger Choice and Revenue Management Optimization

by

Adam Bockelie

Submitted to the Department of Civil and Environmental Engineering
on January 16, 2019, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Transportation

Abstract

The recent proliferation of ancillary services means that airline passengers can face substantially different ancillary service prices and offerings based on their itinerary and fare class selection. At the same time, airlines have become interested in accounting for this supplementary revenue stream in their revenue management (RM) systems to maximize total, not just ticket, revenue. This thesis develops models for both of these issues, with a goal of providing a better understanding of how ancillary services affect the airline industry.

We develop the Ancillary Choice Model (ACM) to describe how passengers make purchase decisions about ancillary services in conjunction with the selection of a fare class. We model two extremes of passenger knowledge and awareness of ancillary services, which we term *simultaneous* and *sequential*. We show that under the simultaneous model, the presence and price of ancillary services can affect the fare class selection of a passenger, even when all fare classes have the same ancillary prices.

The second part of this thesis studies total revenue optimization. We provide a detailed assessment of a prior total revenue maximization approach, the Optimizer Increment (OI), proving that it can be an optimal revenue management strategy under limited conditions, but also showing through the Passenger Origin-Destination Simulator (PODS) that it decreases revenue in more realistic environments.

We then develop a new revenue management optimization model, the Ancillary Choice Dynamic Program (ACDP), which maximizes total revenue by explicitly including the revenue and fare class choice impacts of ancillary services. We describe an Ancillary Marginal Demand (AMD) and Ancillary Marginal Revenue (AMR) transformation that can be used as heuristics to provide the ancillary and choice awareness benefits of ACDP to existing RM optimization models.

We test the revenue performance of our new AMD and AMR heuristics using PODS in a wide range of scenarios. In a network with competing airlines and hundreds of flights, our heuristics can increase total revenue by 2–3%. A consistent trend throughout our simulations is that the forecasting and optimization model that maximizes total revenue is often *not* the model that maximizes ancillary revenue, because models that maximize ancillary revenue often do so to the detriment of ticket revenue.

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

Acknowledgments

I would like to thank my advisor, Peter Belobaba, for the wealth of knowledge he has passed to me. He has always provided very helpful feedback and support, allowing me to explore my interests. His mentorship has made me a better researcher.

My committee members, Carolina Osorio and Moshe Ben-Akiva, have also contributed to my learning and their suggestions have improved this thesis.

The feedback and funding from the members of MIT PODS Research Consortium has made this work possible. In particular, I'd like to thank Jonas Rauch, Eric Ruhlin, Dan Zawack, Thomas Fiig, Richard Cleaz-Savoyen, and Emmanuel Carrier for their very helpful discussions, as well as Amadeus for hosting me for a summer research visit. I would also like to thank the airline that supplied data for the choice model calibration in Appendix A.

Testing the models developed in this thesis required modifying nearly every aspect of the PODS simulator, an extensive task accomplished by Craig Hopperstad and Matt Berge. Thank you for programming my wild ideas and for finding many unanswered implementation questions.

My sincerest thanks go to my friend and colleague Mike Wittman. You have been a constant source of much appreciated advice. Our extensive conversations on esoteric revenue management topics have helped me shape my views, improve my research, and enjoy my time in ICAT.

Of course there is more to life than convex hulls and ancillary services, and I'm grateful to have friends in Boston and beyond who remind me to take a break to ski, play games, or enjoy the beach.

None of this would have been possible without the lifelong support of my parents, who have always encouraged me to pursue my dreams. Thank you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Since the advent of discounted leisure fares in the 1970s, airlines have invested in revenue management systems to maximize the proceeds from ticket sales. In the mid-2000s, however, airlines began developing a secondary, ancillary, revenue stream by both unbundling their fares and offering new products, services, and amenities for sale. Because these ancillary services have been traditionally less important to overall profitability, or not offered at all, airlines do not have a good understanding of how passenger decisions to purchase ancillary services are related to decisions about itineraries and fare classes. As the number and price of ancillary services grow, the availability and price of ancillary services may alter the way in which passengers select itineraries and fare classes. In addition, airlines have little understanding of how to account for this new and growing revenue stream within revenue management systems to maximize total revenue, not just ticket revenue. This thesis explores both topics.

This chapter is organized as follows. Section 1.1 provides the motivation and context for the research, describing the development of ancillary services and their associated revenues, and describing the airline systems involved in selling tickets and ancillary services. Section 1.2 describes the research objectives and contributions of this thesis. Section 1.3 outlines

Table 1.1: Ancillary service categorization

Itinerary	Trip	Relationship
Baggage	Hotels	Loyalty
Seating	Rental cars	Co-branded credit card
Meals	Destination activities	Subscription clubs
Priority boarding		
Lounge access		

the remaining chapters of the thesis.

1.1 Motivation and Context

Passenger airlines sell more than just basic transportation of passengers. In the broadest sense, all revenue from sources other than passenger tickets can be considered ancillary revenue: revenue from performing maintenance for other carriers; revenue from transporting mail, cargo, and freight; and non-ticket revenue from passengers, for example. This thesis focuses on non-ticket revenue from passengers, or revenue from selling passenger-related ancillary services.

Passenger-related ancillary services are optional services sold by airlines that are related, in some way, to passenger transportation. These services can be roughly divided into three groups: services related to a specific itinerary, services related to a particular trip, and services related to a passenger's relationship with the airline. Table 1.1 summarizes this categorization.

Itinerary-related services are directly related to a specific itinerary and are fulfilled as part of transporting the passenger. Examples include checked and carry-on baggage, seating upgrades and assignments, inflight meals and entertainment, priority boarding, and lounge access. Upgrade revenues for Delta's additional legroom Comfort+ seating section provided

$125 million in 4Q 2015, and the airline expected these revenues to grow.[1] Spirit Airlines earns more than $110 million, or 5% of total operating revenue, from seat assignment fees.[2] Checked and carry-on baggage fees provided nearly $1.5 billion for American Airlines in 2016, representing 2.8% of total operating revenue.[3]

Trip-related services are related to a specific trip, but are not fulfilled as part of transporting the passenger. Examples include hotel rooms, rental cars, and destination activities, which provided 3% of Allegiant Airline's total operating revenue in 2016.[4]

Relationship-related services are not linked to a specific trip or itinerary, but occur as part of a passenger's relationship with the airline. Examples include loyalty programs, co-branded credit cards, and subscription discount clubs. United Airlines reported more than $3 billion in revenue from frequent flyer mile sales (primarily related to a Chase Bank co-branded credit card).[5]

1.1.1 Ancillary Revenue Reporting and Data Sources

The revenue impact of ancillary services is difficult to quantify given limited data and reporting. Airlines in the United States are required to report financial information on Form 41 to the US Department of Transportation (DOT), which releases the data to the public. The US DOT typically defines ancillary revenue as three line items within Form 41, Schedule P-1.2: Reservation Cancellation Fees, Miscellaneous Operating Revenues, and Property - Passenger Baggage Fees. These categorizations, however, provide only a limited view of ancillary services: they contain several revenue streams that are not typically considered as ancillary products (like reservation cancellation fees and compensation for collecting airport Passenger Facility Charges). They also exclude several important ancillary revenue

[1] Delta Air Lines Earnings Call (4Q 2015)
[2] Spirit Airlines Form 10-K (2016)
[3] US DOT Form 41, Schedule P-1.2
[4] Allegiant Airlines Form 10-K (2016)
[5] United Airlines Form 10-K (2016)

streams, like onboard food/drink sales (categorized as Transport Related Revenue, which also includes codeshare ticket revenue and contracted maintenance service revenue), loyalty program income (categorized as ticket revenue) and seat assignment fees (categorized as ticket revenue). Outside of the United States, there are no comparable and consistent governmental financial reporting requirements.

Financial statements for publicly-traded airlines (such as 10-K and 20-F Securities and Exchange Commission (SEC) filings in the United States and equivalent filings in other countries) provide some additional data on ancillary services. However, each carrier has its own reporting methodologies and no two carriers use the same revenue categories. Low cost and ultra-low cost carriers (LCCs and ULCCs) that specifically target ancillary revenue tend to provide more detail about the revenue impacts of their ancillary services. For example, Allegiant reports three different categories of scheduled service revenue in its 10-K filings: ticket revenue, "ancillary air-related revenue," and "ancillary third-party revenue" (charter operations are reported separately as well). Delta reports four different categories of passenger revenue: mainline and regional ticket revenue; "loyalty programs;" administrative fees, club, and on-board sales;" and baggage fees in its 10-K filings—a very different system than used by Allegiant. Other carriers have other schemes, and an individual carrier may change reporting practices over time, making comparisons between carriers and years challenging.

In general, the only comprehensive global ancillary revenue reports are estimated and compiled by third-party consulting firms like IdeaWorks, by assessing 10-K and equivalent filings, Form 41 data, press releases, and executive interviews to assemble estimates. IdeaWorks therefore has a more holistic view of ancillary revenue than Form 41, but because of the variety of data sources, is potentially also more variable. In addition, IdeaWorks specifically attempts to include the impact of loyalty programs (e.g. sale of frequent flyer miles/points to banks issuing co-branded credit cards); such revenues are not otherwise considered in this thesis.

1.1.2 Ancillary Service History and Trends

IdeaWorks estimates that the major US carriers collected more than $18 billion in passenger-related ancillary revenue in 2015 (about 11% of total revenue), and that airlines around the world collected $59 billion in ancillary revenue in the same year,[6] far greater than the global airline profit of $33 billion.[7] Although ancillary revenue is clearly important to airlines today, it is a relatively new revenue stream that airlines have put substantial effort into developing over the last ten years. Airlines have increased ancillary revenues both by developing new products and by "unbundling," the practice of charging separately for products or services that were traditionally included as part of a plane ticket. Unbundling initiatives have typically been led by LCCs and ULCCs whose business models rely heavily on revenue from ancillary fees.

Although no longer operational, People Express introduced checked baggage fees and paid onboard meals for trans-Atlantic flights in the 1980s (Conrady, 2013). More recently, Ryanair in Europe and Allegiant Airlines and Spirit Airlines in the United States have transformed into LCCs and ULCCs, with substantial increases in ancillary revenue (as shown in Figure 1.1). Ryanair began the transition to low fares and no frills in the early 1990s, and then began to grow ancillary revenues. By 1997, Ryanair obtained 5% of its revenue from inflight sales, primarily from duty-free items.[8] Until the early 2000s, ancillary services (in the form of inflight sales, destination car rentals, and other non-flight ancillary services) provided about 9% of Ryanair's operating revenue.[9] Ancillary revenue grew through fiscal year 2016 to 24% as the airline implemented fees for checked baggage, airport check-in, flight notifications, seat assignments, and extra legroom seats.[10]

[6]IdeaWorks. (November 9, 2015). *Airline ancillary revenue projected to be $59.2 billion worldwide in 2015* [Press release]. Retrieved from `http://www.ideaworkscompany.com/wp-content/uploads/2015/11/Press-Release-103-Global-Estimate.pdf`
[7]IATA Economic Performance of the Airline Industry End-Year Report (2015)
[8]Ryanair Initial Public Offering Prospectus (1997)
[9]Ryanair Initial Public Offering Prospectus (1997); Ryanair Annual Report (1999–2001)
[10]Ryanair Form 20-F (2001–2016)

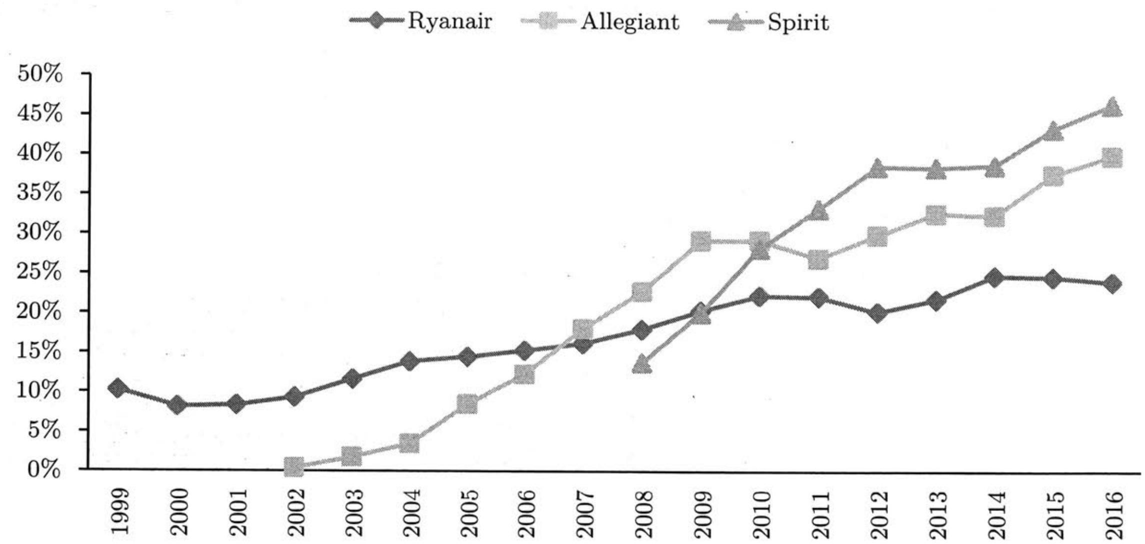

Figure 1.1: Portion of revenue from ancillary services for Ryanair, Allegiant, and Spirit. Source: Ryanair Annual Report (1999–2001), Ryanair Form 20-F (2001–2016), Allegiant Airlines Initial Public Offering Prospectus (2006), Allegiant Airlines Form 10-K (2006–2016), Spirit Airlines Initial Public Offering Prospectus (2012), Spirit Airlines Form 10-K (2010–2016).

Allegiant started its transformation to a ULCC, and began a concerted effort to increase ancillary revenue, when it emerged from bankruptcy in 2002. In 2003, only 2% of total operating revenue was derived from ancillary services.[11] In 2006, the airline implemented fees for checked baggage and for bookings made through the airline's website (as opposed to bookings made at the airport) and ancillary revenue rose to 12% of total operating revenue. The airline implemented fees for carry-on baggage in 2012; by 2016, 40% of Allegiant's revenue came from ancillary services, as shown in Table 1.2.

Spirit Airlines began its transition to a ULCC in 2006. It implemented seat selection fees in 2008, an online booking fee in 2009, and a call center booking fee in 2010. In 2010, Spirit was one of the first airlines to charge customers for carry-on baggage.[12] These and other charges increased Spirit's ancillary revenue from 14% of total operating revenue in 2008 to 46% in 2016.[13]

[11] Allegiant Airlines Initial Public Offering Prospectus (2006)
[12] Spirit Airlines Initial Public Offering Prospectus (2012)
[13] Spirit Airlines Initial Public Offering Prospectus (2012); Spirit Airlines Form 10-K (2011–2016)

Table 1.2: History of ancillary revenue at Allegiant Airlines

Year	New ancillary services	Portion of total revenue from ancillary services
2002	None	0%
2005	Seat assignment fees	8%
2006	Checked baggage fee, online booking fee	12%
2012	Carry-on baggage fee	30%
2016	Refinement of existing fees	40%

Source: Allegiant Airlines Initial Public Offering Prospectus (2006), Allegiant Airlines Form 10-K (2006–2016).

Legacy network carriers have also adopted the ancillary fee model. Most US network carriers implemented fees for checked bags in 2008 and 2009. By 2010, all US legacy carriers had discontinued free domestic economy meal service, in favor of onboard food sales. Lufthansa began charging for seat assignments on short-haul flights in 2013, and later expanded the fees to its long-haul network. In 2017, United Airlines and American Airlines launched basic economy tickets, which provide seating in the standard economy cabin but prohibit carry-on bags (passengers must pay to check all baggage).[14] Basic economy fares have since spread to other North American and European airlines.

In addition to unbundling, airlines have been developing new services and products. Lufthansa launched the first Wi-Fi equipped flights in 2004, and Wi-Fi is now sold on many flights worldwide. In 2010, Air New Zealand launched "SkyCouches," which are sets of three economy class seats that convert into a bed for two passengers. In 2015, American Airlines led US carriers in launching a premium economy cabin, which provides more legroom, wider seats, and improved service (non-US carriers had previously offered premium economy).

Although many large airlines utilize ancillary fees, not all have completely unbundled their product offerings. Some airlines have pursued a hybrid approach to bundling: offering a mix of bundled and unbundled fares. For example, Delta Air Lines provides complimentary upgrades to "preferred" seats to the highest value economy fare classes, while selling such seats

[14]Delta introduced Basic Economy tickets in 2016, but includes complimentary carry-on bags. In 2018, American Airlines dropped its basic economy carry-on prohibition.

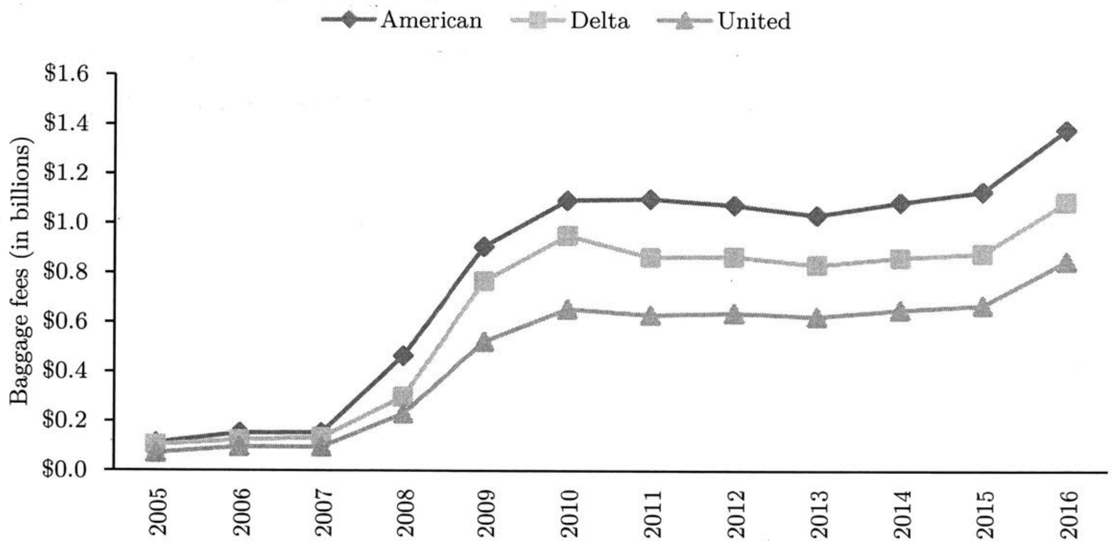

Figure 1.2: Baggage revenue for US network airlines
Source: US DOT Form 41, Schedule P-1.2.

to passengers in lower value fare classes. Other airlines, such as Qantas and Air Canada, offer "branded" fares or bundles whose various restrictions and ancillary fee structures are clearly marketed; the goal is to make clear to consumers the benefit of selling-up from the lowest offered fare (Vinod and Moore, 2009). For more details on the evolution of ancillary fees, see Garrow et al. (2012).

According to the US DOT definition of ancillary revenue, the impact of unbundling has been a substantial increase in ancillary revenue. As shown in Figure 1.2, baggage-related fees for major US network carriers (including first and second checked bags, overweight bags, oversize bags, and excess bags) increased by a factor of five between 2007 and 2010 as these airlines unbundled fares. Baggage fees now account for nearly $1.5 billion per year for American Airlines, and are equal to 3%–4% of ticket revenue for major US network airlines. Ancillary revenue as a whole, as reported to the US DOT, has grown in the last decade from 2% to 8% of total revenue, or from $2.2 billion to $15.0 billion per year, as shown in Figure 1.3. As discussed above, these figures likely underestimate total ancillary revenue production.

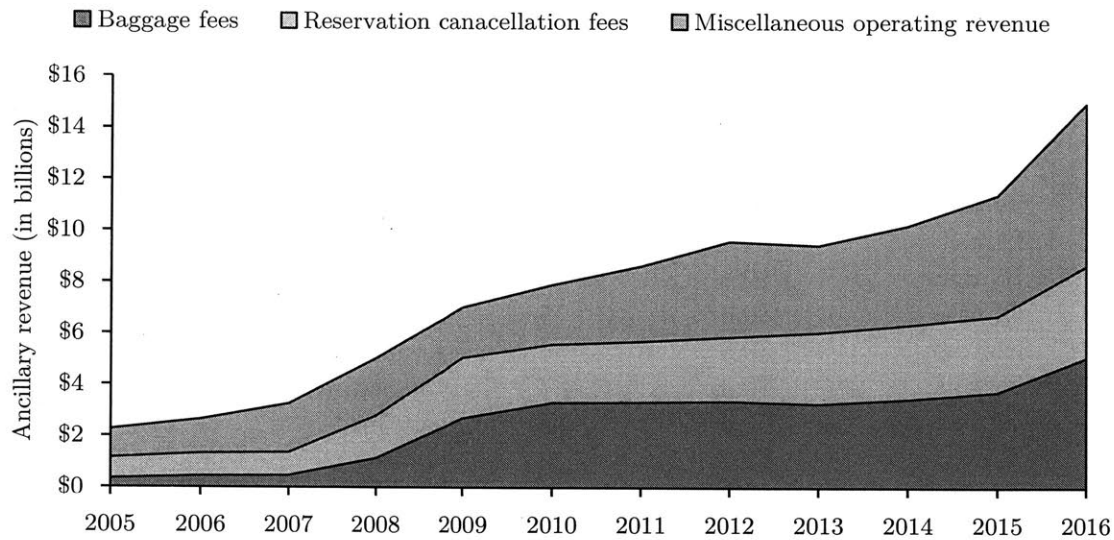

Figure 1.3: Total ancillary revenue for US airlines, as reported to US DOT.
Source: US DOT Form 41, Schedule P-1.2.

IdeaWorks ancillary revenue estimates also include the significant financial impact of loyalty programs, which are not otherwise considered in this thesis. Estimated total passenger-related ancillary revenue for United Airlines in 2015, for example, was $6.2 billion, of which $3.0 billion is attributed to the MileagePlus loyalty program (primarily through revenue from a co-branded credit card). Sales of frequent flyer miles are estimated to provide 47% of American's 2015 IdeaWorks ancillary revenue, and 64% of Delta's.[15] IdeaWorks estimates that ancillary revenues grew dramatically between 2008 and 2015, as shown in Table 1.3. For carriers that were in the top ten in ancillary revenue production in both 2008 and 2015, estimated ancillary revenues grew between 90% and 350%.

Although ancillary revenues have increased, empirical studies of add-on pricing suggest the increases have led to decreases to base (flight) prices. Ancarani et al. (2009) use transaction data for airlines, hotels, online retailers, and restaurants to demonstrate that increases in add-on prices typically lead to decreases in base prices. Scotti and Dresner (2015) focus on the airline industry and evaluate average fares, passenger traffic, and ancillary fees between

[15]IdeaWorks Yearbook of Ancillary Revenue (2015)

Table 1.3: IdeaWorks estimates of ten largest airlines by ancillary revenue (in billions).

2008			2015		
Rank	Airline	Ancillary revenue	Rank	Airline	Ancillary revenue
1	American	$2.2	1	United	$6.2
2	United	$1.6	2	American	$4.7
3	Delta	$1.5	3	Delta	$3.8
4	Ryanair	$0.8	4	Air France/KLM	$2.2
5	Qantas	$0.6	5	Southwest	$2.1
6	easyJet	$0.5	6	Ryanair	$1.7
7	JetBlue	$0.4	7	Lufthansa	$1.5
8	Emirates	$0.3	8	easyJet	$1.5
9	TAM	$0.2	9	Qantas	$1.2
10	Alaska	$0.2	10	Alaska	$1.1

Source: IdeaWorks Top 10 Ancillary Revenue Rankings (2015).

2007 and 2010. They find that increasing checked baggage prices results in lower average base fares and passenger counts. Brueckner et al. (2015) also show a reduction in average base fares when airlines implement baggage fees. Zou et al. (2017) find that, when a la carte airlines compete against airlines that bundle checked baggage, there is a correlation between the baggage price for the a la carte carriers and average fares for the bundled carrier.

Because ancillary services were previously not offered, or were only a minor component of an airline's revenue and a passenger's cost to travel, they have not been extensively analyzed. This thesis will examine the impact of ancillary services on passengers and on airline revenue management systems.

1.1.3 Overview of the Airline Industry

Airlines offer networks of flights (which consist of one take-off and one landing) that serve a variety of origin-destination (OD) markets (pairs of cities or airports where passengers begin or end an air travel journey) and passenger types (i.e. business, leisure, etc.). Each market has its own demand characteristics, and each passenger has their own budget, schedule, and quality preferences. Airlines leverage passenger heterogeneity to increase profits by price and product discrimination: on each flight, airlines may offer several distinct cabins of service (e.g. first class, business class, and economy class). Within each market and each cabin of service, airlines offer a variety of different fare classes, or price points with purchase and use restrictions. Typically, the most expensive fare class has no restrictions, while the least expensive fare class is highly restricted. These restrictions may include a round trip purchase requirement, a minimum and/or maximum stay requirement, a Saturday night stay requirement, or a cancellation penalty. For example, when Delta was the only non-stop carrier between Boston and Detroit, it offered the restricted fare structure shown in Table 1.4. Passengers wanting to purchase the lowest-value class V had to purchase round trip tickets three weeks in advance with a Saturday night at their destination, and pay a penalty to make changes or cancel the reservation. Airlines impose these restrictions on discount fares in an attempt segment demand: restrictions force restriction-averse (but typically high budget) business travelers to purchase more expensive, less-restricted tickets while still allowing restriction-tolerant (and typically low budget) leisure travelers to purchase less expensive discount tickets.

Not all airlines impose such extensive restrictions on their low-value fare classes, and airlines tend to offer similar fare structures as their non-stop competitors. In 2014, Delta modified its Boston–Detroit fare structure as JetBlue introduced new non-stop service. Delta eliminated the Saturday night stay requirement and dramatically lowered its lowest-value fares—to \$69 in V class and \$99 in X class from \$205 and \$215, respectively (Belobaba, 2015).

In an airline network, each physical seat on each flight is a perishable asset that can be sold to passengers flying in different markets and paying different fares. Typically, airlines establish schedules and the fares discussed above far in advance of operating a flight; see Belobaba et al. (2009) for an overview of the airline planning process. During the booking window, when travelers are shopping, the airline's revenue management (RM) system determines which fare classes to sell in each market at any given time. The optimization model inside

Table 1.4: Restricted Boston-Detroit fare structure for Delta Air Lines in September 2013.

Fare class	One way fare	Advance purchase	Refundable	Change fee	Round trip required	Minimum stay
Y	$936	None	Yes	None	No	None
B	$794	None	No	$200	No	None
M	$603	None	No	$200	No	None
H	$501	14 days	No	$200	No	None
K	$365	None	No	$200	Yes	Sat Night
T	$249	7 days	No	$200	Yes	Sat Night
X	$215	14 days	No	$200	Yes	Sat Night
V	$205	21 days	No	$200	Yes	Sat Night

Source: Belobaba (2015).

the RM system attempts to offer the revenue-maximizing fares, considering both the supply of seats on each flight leg and the demand within each OD market. The demand forecast is generated based on historical booking data. More details about RM and forecasting methods can be found in Chapter 4.

The booking window for a future flight departure begins up to 330 days before departure. Conceptually, consumers purchase travel via a travel retailer, which could be an airline website or call center, an online travel agent or meta-search (such as Expedia or Google Flights), or other service. Retailers provided by the airline are direct booking channels; those provided by others are indirect booking channels. A consumer provides the retailer with a booking request, which consists of an origin, destination, departure dates, and desired class of service. Direct booking channels send the booking request directly to the airline's reservations inventory system, which returns a set of available options for the passenger. As illustrated in Figure 1.4, indirect booking channels typically send the booking request to a Global Distribution System (GDS), which combines information from three sources to return a set of options for the passenger: schedules, from a third-party data source (typically OAG, or Official Airline Guide); availability, from the marketing airline for each flight; and fares, from another third-party data source (typically ATPCO, or Airline Tariff Publishing Company). The retailer presents the set of options (from either the GDS or the airline)

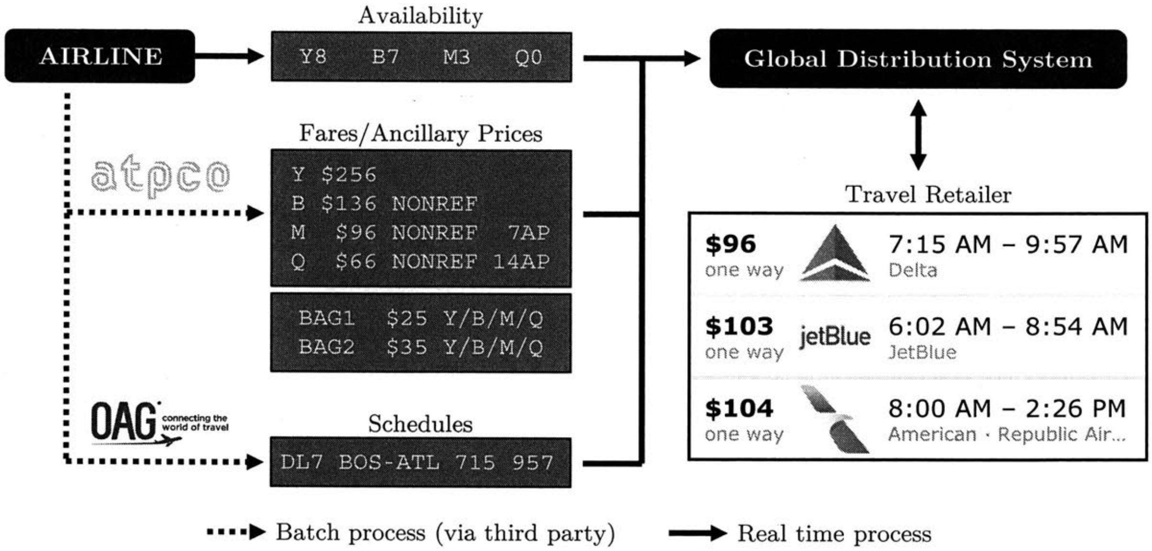

Figure 1.4: Airline distribution system schematic. Consumers request travel options from a travel retailer, who then passes the request to a Global Distribution System. While the airline supplies schedules and fares/prices to OAG and ATPCO, the only real time request to the airline is for fare class availability.

to the consumer, who then books one option, makes a new request, or leaves the system. If the consumer books, the retailer notifies the GDS or airline, which issues a ticket and a reservation.

Approximately 50% of bookings worldwide are made through a GDS, and therefore the structure of GDSs has a significant impact on how airlines sell travel (Taubmann, 2016). As GDSs were originally designed to sell tickets, they are anchored around fare classes—availability is controlled at the fare class level, fares are filed at the fare class level, and booking options are specific itinerary and fare class combinations. GDSs have limited capability to offer and sell ancillary services. For example, according to IdeaWorks, American Airlines sells paid seat assignments, but not checked baggage, though Sabre GDS. Delta Air Lines sells bundled fares and paid seat assignments through Amadeus GDS, only bundled fares through Sabre GDS, and only paid seat assignments through Travelport GDS.[16]

[16]IdeaWorks Yearbook of Ancillary Revenue (2016)

The industry is in the process of replacing this workflow with New Distribution Capability (NDC), which features several important changes. When NDC is fully implemented, GDSs will no longer need to aggregate schedules, availability, and fares to assemble sets of booking options—the GDS could request *offers* from airline offer management systems. These offers would no longer necessarily be centered on a fare class—each offer could consist of an itinerary, a set of zero or more ancillary services, various purchase/use restrictions, and a single price.

The shift from a traditional distribution environment to NDC has significant implications for airlines. Because NDC improve travel retailers' abilities to sell ancillary services, airlines expect ancillary revenues to rise with the implementation of NDC. In addition, because NDC could move away from fare class-centered availability control, airlines will be able to develop offer management systems. NDC also allows more detailed and personalized booking requests (including information such as frequent flyer number and marital status in the initial booking request), and allows airlines to respond with individualized sets of offers. At the limit, with NDC, each offer could be dynamically and personally constructed and priced for each consumer.

1.2 Research Objectives and Contributions

This thesis address two primary research questions surrounding airline ancillary services. First, it examines how the presence and/or price of ancillary services might affect the way passengers make choices about air travel. Second, it investigates mechanisms for airlines to effectively incorporate information about ancillary revenues into their revenue management systems, to maximize the total combination of ticket and ancillary revenue.

These questions reflect two significant gaps in the literature. Knowledge about the choice process(es) airline consumers use to select an itinerary and fare class in conjunction with ancillary services is lacking. Although authors have studied passenger selection of airline, itinerary, and fare class or have studied valuation of ancillary services, and other authors have studied consumer selection of bundled vs unbundled products, no research has integrated these questions. How and when consumers consider the price and presence of ancillary services in conjunction with other booking decisions has not been modeled.

Second, an effective optimization process incorporating ancillary revenues is lacking. Optimizing for total expected contribution (i.e. the optimizer increment, discussed in Chapter 5) leads to revenue losses in many airline simulations, but reportedly increases revenue for hotel casinos. No formal proof of optimality exists under any set of assumptions, and no research addresses which assumptions are most important in maintaining revenue performance.

No studies have connected ancillary-aware passenger choice with ancillary-aware optimization and no research explains how different choice processes may affect revenue generation or the optimality of any particular RM method.

This thesis addresses both knowledge gaps with three principal contributions:

1. Development of the Ancillary Choice Model framework for extending existing itinerary/fare class choice models to incorporate ancillary services, while allowing passengers to vary in their degree of knowledge about ancillary services (Chapter 3).
2. Development of a theoretical background for the existing Optimizer Increment and simulation results illustrating why the approach does not provide consistent revenue benefits (Chapter 5).
3. Development of the Ancillary Choice Dynamic Program (ACDP), a RM optimization model that is both ancillary-aware and passenger choice-aware; development of two associated heuristics, the Ancillary Marginal Demand (AMD) transformation and Ancillary Marginal Revenue (AMR) transformation, that allow the benefits of ACDP to be applied to existing RM optimization models; and development of additional approximations and processes required the operationalize ACDP and the AMD/AMR

heuristics (Chapter 6). Simulation results compare AMD and AMR against previous models and show revenue benefits in a variety of environment (Chapter 7).

Together, these contributions create a better understanding of the impacts ancillary services have on passengers and on airline pricing, revenue management, and distribution systems. They suggest methods for airlines to leverage ancillary services to increase revenues, especially with the imminent development of NDC.

1.3 Thesis Outline

The remainder of this thesis is organized into two parts. Part I focuses on the behavioral impacts of ancillary services and contains two chapters. Chapter 2 reviews the literature on consumer behavior, discrete choice modeling, and passenger choice modeling and identifies deficiencies in the manner in which previously published models handle ancillary services. Chapter 3 presents a new Ancillary Choice Model, which describes two methods by which passengers might integrate the presence and price of ancillary services into their itinerary and fare class selection. The chapter also presents simulated booking and revenue impacts of the proposed model.

Part II contains four chapters and focuses on revenue management methods that incorporate ancillary services. Chapter 4 reviews the literature on pricing theories, revenue management, forecasting, and distribution. Chapter 5 describes a previously-proposed ancillary-aware RM method, the optimizer increment, and develops a theoretical background for the approach in a limited setting. The chapter also explains why the optimizer increment is not beneficial in more general settings. Chapter 6 presents the new Ancillary Choice Dynamic Program (ACDP) for ancillary and choice-aware revenue management, as well as an Ancillary Marginal Demand (AMD) transformation and Ancillary Marginal Revenue (AMR) transformation that allow the benefits of new model to be incorporated into existing RM systems. Simulated performance for the heuristics are included in Chapter 7 and show a revenue benefit in a variety of environments and networks, including competitive scenarios.

Finally, Chapter 8 provides a conclusion to the thesis, summarizing the key findings and contributions, and suggesting future research directions.

Part I

Impacts of Ancillary Services on Passenger Behavior

Chapter 2

Literature Review: Passenger Behavior and Decision Making

As the number and value of ancillary services grows, so does the need to understand how they affect airline passengers. As discussed in Section 1.1.2, airlines are currently employing a variety of pricing and marketing schemes for ancillary services. A better understanding of how passengers purchase ancillary services in conjunction with itineraries and fare classes can further inform airlines on optimal pricing, bundling, and marketing strategies, as well as on strategies for total revenue optimization, since the design of those strategies depends on how ancillary valuations vary across different demand segments (Stremersch and Tellis, 2002). With the advent of New Distribution Capability and the promise of increased distribution flexibility, airlines will have more latitude to leverage their knowledge of the passenger choice process to develop more customized offers and increase revenue and market share, and the potential benefit of additional passenger choice knowledge will increase.

In this Part, we develop and assess an integrated model of passenger choice with ancillary services, fare classes, and itineraries that could be used to assess potential pricing, bundling, marketing, and total revenue optimization strategies. This chapter provides a review of the relevant literature, while Chapter 3 presents our Ancillary Choice Model. The remainder of this chapter is organized as follows. Section 2.1 describes key characteristics of the passenger choice problem. Section 2.2 reviews a framework for modeling discrete choice problems, including general mathematical formulations. Section 2.3 describes relevant previous approaches to modeling passenger choice and to modeling the purchase of ancillary services, both in the airline and the economics literature. Finally, Section 2.4 summarizes the key gaps in the literature that our Ancillary Choice Model seeks to address.

2.1 Characteristics of the Passenger Choice Problem

The airline passenger choice problem—the set of tasks associated with evaluating (and purchasing) an airline travel plan—is a discrete choice problem with consumer heterogeneity, alternatives with complicated attributes, and very large choice sets.

Consider the example of two passengers shopping for air travel. At one extreme may be a business executive traveling to an important away-from-home conference at an airport hotel. Her destination, date, and arrival time are likely dictated by the conference schedule and are not part of her choice process; she only needs to select a specific itinerary, booking class, and set of onboard amenities. At the other extreme may be a college student planning a summer vacation to Europe. His choice task will include the dimensions faced by the executive, as well as possibly the selection of destination country/city, travel date, and travel time. Figure 2.1 lists some of the dimensions that may be included in the problem.

The attributes of each alternative are potentially numerous and not necessarily easy to measure. Obvious attributes include departure and arrival times, duration, number of connections, and price, but airlines and itineraries may vary in other ways as well. In the winter, a connection in Chicago may be less desirable than a connection in Dallas because of Chicago's higher risk of disruption due to snow and ice. When destination (or origin) is included in the choice problem, as in the case of our college student, two airports may be

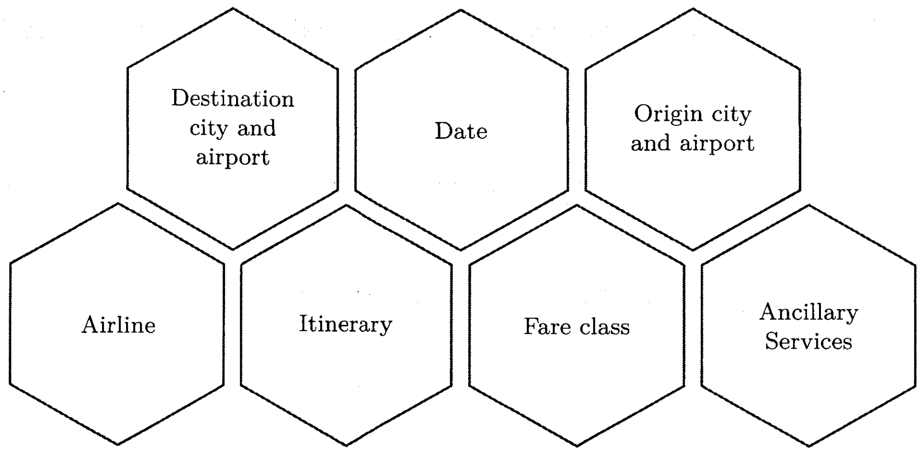

Figure 2.1: Possible dimensions in the airline passenger choice problem

perceived similarly if they are close geographically (i.e. London's Heathrow and Gatwick airports) or if they offer similar destination activities or amenities (e.g. Spanish beaches and Greek beaches).

Consider one aspect the passenger choice problem: the selection of an airline and a route (a sequence of connection airports, without regard to the specific flights between those airports) for a travel from San Francisco International Airport (SFO) to Boston Logan International Airport (BOS). Between April and June 2015, US Department of Transportation data indicates that passengers actually traveled on at least 260 different airline and route combinations, thus the number of alternatives is at least 260.[1]. When the problem is expanded to cover itineraries and not just routes (i.e. considering the specific flights between airports), the choice set grows, and when fare classes are also considered the choice set grows again. De Marcken (2003) reports that for a round trip from SFO to BOS and back, constrained to only American Airlines flights connecting through Chicago O'Hare International Airport (ORD) on the outbound and Dallas/Fort Worth International Airport (DFW) on the inbound, had more than 25 *million* valid flight and fare combinations in 2003.

[1]US DOT Airline Origin and Destination Survey (DB1B) which is a 10% sample of all domestic tickets sold in the United States.

Harvey (1987) proposes simplifying passenger choice models to a set of sequential decisions for both calibration and application. While studying the selection of airport and airport access mode in the San Francisco area, he finds that access mode choice does not appear to be made simultaneously with airport choice. This framework of simplifying the passenger choice problem to a few simultaneous dimensions is used implicitly in many models, including ours.

In addition, because alternative attributes are numerous and because airlines are continually adjusting their pricing and marketing strategies, the problem is characterized by limited knowledge on the part of decision makers. Evidence indicates that airline passengers are not always knowledgeable about ancillary services: a study of US adults who had flown at least once in the previous year showed that 55% of travelers were surprised by additional ancillary fees after their ticket purchase, and that 47% found it "very difficult or nearly impossible" to compare the total cost of travel (including desired ancillary fees) across different airlines and itineraries (Open Allies, 2014, Exhibit 1, pg. 1–2). In studies on the perceived fairness of ancillary services, data from travel agent customer surveys shows that baggage fees, particularly carry-on bag fees, generate a sense of "betrayal" amongst airline customers (Chung and Petrick, 2013; Tuzovic et al., 2014), and the Open Allies (2014, Exhibit 1, pg. 2) survey found that 81% of respondents believe that airline ancillary fees are "unfair or deceptive."

Limitations in knowledge or awareness are also found in a variety of empirical pricing studies outside of the airline context, with several studies providing evidence that consumers cannot accurately compare equivalent combined and partitioned/unbundled prices. Johnson et al. (1999) demonstrate that consumer perceptions of a two-product offer improve as price information is bundled and as discount information is unbundled, in accordance with prospect theory. Chakravarti et al. (2002) and Morwitz et al. (1998), however, show that consumers use may heuristics to compute the total cost for product offers, and that the heuristics tend to underestimate the total cost of a partitioned offer. Morwitz et al. (1998)

suggest that consumers may even ignore secondary pricing components completely to reduce their information processing effort. Xia and Monroe (2004) find that as the number and value of partitions increases, consumers may begin to overestimate the total cost of a partitioned offer. See Greenleaf et al. (2016) for other behavioral impacts of partitioned pricing.

Finally, as suggested above, consumers shopping for air travel have a variety of trip purposes and preferences. These different demand segments may be based on travel purpose or perceived value of transportation. Kothari et al. (2016) subdivide the business and leisure segments by class of service and as a visitor or resident. Schwieterman (1985) suggests three segments for "highly discretionary consumers" who travel for personal reasons and alter their travel plans to obtain a low fare, "moderately discretionary consumers" who cannot comply with minimum stay requirements, and "non-discretionary consumers" who do not compromise on schedule convenience. Belobaba (1987) proposes a four-segment model, with each segment corresponding to a combination of high/low price sensitivity and high/low time sensitivity. For example, a time-sensitive, price-sensitive consumer has a preferred departure time, but will make some schedule concessions to obtain lower fares. Carrier (2008) notes that as the number of segments increases, the ability to deterministically assign a travel to any individual segment decreases, and that a probabilistic approach may be more appropriate. Carrier (2008), Teichert et al. (2008), and Bruning et al. (2009) all develop probabilistic membership models with two to five price, quality, and schedule-oriented segments. A sophisticated passenger choice model must support variations in demand segments.

The high problem dimensionality, limitations in decision maker knowledge, and heterogeneity of preferences makes accurately modeling the passenger choice problem challenging.

2.2 Framework for Discrete Choice Modeling

Discrete choice models predict the decision of an individual faced with a choice task—such as the selection of an airline ticket. In this section we review key concepts in the general choice modeling literature, following the ideas of Ben-Akiva and Lerman (1985). They view a choice as the outcome of a multi-step process that includes with problem definition, generation of alternatives, evaluation of alternative attributes, and application of a decision rule.

Problem definition includes an assessment of the problem dimensions to include, which, in the passenger choice problem, could be any of the dimensions in Figure 2.1. Generation of alternatives refers to the process of determining possible options from which to choose. These options must be known to the decision maker and must be feasible for the decision maker to select. Each alternative has a variety of attributes (such as price, schedule, quality, etc. in the context of choosing an airline ticket); the decision maker must assess each of the attributes she intends to use when applying the decision rule. The set of alternatives for any individual decision maker is the choice set.

2.2.1 Decision Rules

Decision rules govern how the assessment of alternative attributes leads to a choice outcome. Many theoretical decision rules have been proposed. The simplest is the dominance rule, which states that the decision maker will choose the alternative that is best in all attributes. Relaxing the requirement for strict dominance leads to a threshold rule or lexicographic rule. Under a threshold policy, an individual eliminates any alternative with attributes worse than his or her threshold level (which may be deterministic or stochastic). Threshold rules do not typically result in a unique alternative, so other decision rules must also employed. Under a lexicographic policy, an individual selects the alternative(s) that are best according to one single (most important) attribute. If multiple alternatives are chosen, the decision maker considers the best according to the next most important attribute. This process continues until only one alternative is selected.

Dominance, threshold, and lexicographic rules are examples of non-compensatory decision rules. A compensatory rule, on the other hand, states that decision makers will weigh each attribute against all other attributes, making trade-offs between various attributes. Under a compensatory rule, the decision maker chooses the alternative with the greatest utility, where the utility of each alternative is some weighted function of the alternative's attributes.

In practice, these theoretical decision rules could be combined such that decision makers are non-compensatory in some aspects, but compensatory in others. Decision makers know alternative attributes either deterministically or stochastically. Finally, decision makers may consider all attributes at the same time (simultaneously) or in groups (sequentially, as in Harvey (1987)). In a simultaneous approach, attributes from all dimensions are considered and, assuming utility-maximization, the combination of all dimensions with the highest total utility will be selected. In a sequential approach, groups of attributes are evaluated independently; no attributes from other groups are included in the utility function. Sequential models are often used when the choice process can be decomposed into multiple different, somewhat independent, dimensions. Simultaneous and sequential decision processes result in different types of utility maximizing behavior: simultaneous decisions maximize utility globally (over all alternatives), while sequential decisions maximize utility locally (within each group/dimension).

2.2.2 Rationality

In their simplest form, the theoretical models of the preceding section assume that consumers behave rationally: they use "a consistent and calculated decision process" with "consistent and transitive preferences" (Ben-Akiva and Lerman, 1985, pg. 38). In classic economics, rational actors are assumed to assess alternatives carefully, correctly, and with full information.

Concerns about this perfectly intelligent model were raised early, with claims of "a complete lack of evidence" that actual human behavior in any sort of complex choice environment followed the classic economic model (Simon, 1955, pg. 104). Instead, Simon proposed that humans have limited knowledge and calculating abilities—a model referred to as *bounded rationality.* This more restricted view of human ability may include limitations in assessing risky outcomes and performing mental calculations (Kahneman and Tversky, 1979), in assessing time (Schiffman, 1990), in knowledge and awareness (Simon, 1989), or in accounting for the impact of emotional and physical conditions on preferences (Loewenstein, 1996).

Kahneman and Tversky (1979) propose prospect theory, which states that consumer decisions are made not by comparing absolute values of each option, but by comparing gains or losses relative to some reference point, and that gains are valued with diminishing margins of positive returns while losses are valued with diminishing margins of negative returns. Thaler (1985) utilizes prospect theory in developing the principle of mental accounting, which states that people "track" different types of spending or consumption in different "accounts," which are not generally allowed to cross-subsidize eachother. McFadden (1999) provides an extensive review of rationality and its limitations. As discussed above, there is substantial evidence of bounded rationality, in both the airline and general marketing/pricing literature.

2.2.3 Mathematical Formulations

Threshold and Dominance Models

Threshold and dominance models are common in the economics literature. In the classic work of Adams and Yellen (1976), for example, consumers choosing between two products (or a bundle of the two products) select the option that maximizes their surplus, defined as the difference between their reservation price (or valuation of the product(s)) and the price charged by the firm; this process is subject to the threshold that surplus must be positive.

These economic models are typically focused on more providing a platform for analyzing pricing strategies by the firm and their consequences on consumer welfare, rather than on providing a sophisticated view of how consumers make decisions.

Random Utility Models

A commonly employed mathematical model of choice behavior is a random utility model (RUM), in which a decision maker assigns a utility to each alternative and then selects the alternative with the highest utility. The utility is typically decomposed into a deterministic systematic utility, which can be observed and measured by an analyst and is a direct function of alternative attributes, and a stochastic error term, which accounts for unobserved attributes of alternatives, preference variations amongst consumers, and measurement errors by the analyst. The distributions and covariances of the error terms affects choice probabilities. The classic multinomial logit model has independent and identically distributed type-I extreme value error terms and rational decision makers. A typical use of a multinomial logit model decomposes systematic utility into a linear function of preferences and alternative attributes:

$$U_i = V_i + \varepsilon_i \qquad V_i = \beta X_i \qquad \Pr(i) = \frac{\exp(V_i)}{\sum_{j=1}^{n} \exp(V_j)}$$

where U_i is the utility of alternative i, V_i is the systematic utility, ε_i is the error term, β is a vector of preferences, X_i is a vector of attributes, and $\Pr(i)$ is the probability the decision maker selects alternative i (out of n possible alternatives). The variance of the error terms ε_i is $\pi^2/6\mu^2$, where μ is the *scale* and is typically normalized to 1. At the extreme, a model with infinite scale (zero variance) error terms is deterministic; a model with zero scale (infinite variance) error terms results in equal choice probabilities regardless of alternative attributes.

Multinomial logit models are convenient because they have closed form expressions for choice probabilities and therefore for maximum likelihood estimation of the (typically unknown) preferences β. However, the assumption of independent and identically distributed error terms limits their flexibility. In particular, multinomial logit models misspecify choice probabilities when some alternatives share unobserved characteristics (as in the "red bus, blue bus" problem). In addition, the models are inherently compensatory and rational.

Extensions to the multinomial logit specification alleviate some of these concerns. These extensions may introduce correlations amongst the error terms within groups of alternatives (nested logit, see Ben-Akiva (1973)) or allow decision makers to consider only a portion of the choice set (probabilistic choice multinomial logit, see Swait and Ben-Akiva (1987)).

The most flexible and general extension is a mixture of logit model, in which choice probabilities follow traditional multinomial logit model conditional on a vector of parameters β, which is distributed according to some distribution $G(\beta, \theta)$ (where θ is a vector of parameters controlling the distribution G):

$$V_i(\beta) = \beta X_i \qquad \Pr(i \mid \beta) = \frac{\exp(V_i(\beta))}{\sum_{j=1}^{n} \exp(V_j(\beta))} \qquad \Pr(i) = \int \Pr(i \mid \beta) G(\beta, \theta) \mathrm{d}\beta$$

McFadden and Train (2000) show that, if correctly specified, mixture of logit models can approximate any other random utility choice model to arbitrary precision. This allows mixture of logit models to capture latent class membership (in which each consumer belongs with some probability to one of several classes, or segments; classes differ in the form of their utility equation and/or in the set of alternatives considered), random preference coefficients (which explicitly model the variation of preferences across the population), and some of the cognitive and behavioral deviations from classic rationality identified above. Mixture of logit models allow a richer representation of the choice process, and have found use in several sophisticated models of airline passenger choice, as described below. The Ancillary Choice Model developed in Chapter 3 is a mixture of logit model. For more detail on mixture of logit models, see McFadden and Train (2000), Walker and Ben-Akiva (2002), or Train (2009).

2.3 Previous Approaches

Aspects of the passenger choice problem with ancillary services have been studied in the airline, product bundling, and product add-on literatures, including work by economists and marketers. This section reviews relevant work in all three areas.

2.3.1 Airline Passenger Choice

Choice modeling applied to passengers in the airline industry typically utilizes random utility models, with varying degrees of complexity. Previous work has studied passenger choice of an itinerary, then of an itinerary and fare class, and finally of ancillary services.

Itinerary Choice Models

Itinerary choice models describe the selection of an airline and itinerary by a passenger, given a pair of origin and destination airports. These models typically include fares and departure and arrival times as choice determinants. Smith et al. (2007) utilize a latent-class mixture of logit model, while excluding dominated alternatives from the consideration set. Hess and Adler (2011) estimate a multinomial logit model using stated preference data, considering aircraft size and number of connections in addition to fare and departure and arrival times. A more nuanced view of connections is provided by Theis (2011), who develops an integrated choice/latent variable mixture of logit model to incorporate traveler perceptions about risk and stress.

Itinerary and Fare Class Choice Models

Several authors go beyond the selection of airline and itinerary to include the selection of fare class as an *itinerary and fare class choice model.* Algers and Beser (2001) combine stated and revealed preference datasets in a nested logit formulation to predict fare class selection within three Scandinavian markets. Carrier (2008) develops an approach for collecting the dataset necessary to estimate a revealed preference itinerary and fare class choice model. He utilizes a latent class mixture of logit formulation to account for different trip purposes.

Hopperstad (2005) proposes a model, used the Passenger Origin-Destination Simulator, for the selection of airline, itinerary, and fare class. The model is a mixture of logit with an infinite scale parameter, random coefficients, latent classes, and incorporates threshold decision rules. Consumers within a particular market assess each available itinerary and fare class option on all airlines serving that market and select the itinerary and fare class combination that maximizes their trip utility (minimizes disutility), given their travel preferences, subject to a budgetary constraint θ_n on out-of-pocket costs. Passengers who find no combinations within their budget do not fly. For each itinerary and fare class combination, utility is a function of fare f_{ik}, a disutility for β^{CNX} connecting (if applicable, $\delta_{ik}^{CNX} = 1$; 0 otherwise), a disutility β^{DTW} for departing or arriving outside of a preferred window (if applicable X_{ik}^{DTW} is the total deviation from desired time window), and disutilities β^{R_j} for coping with fare class restriction j (if applicable, $\delta_{ik}^{R_j} = 1$; 0 otherwise). Mathematically, the utility of itinerary i and fare class k for decision maker n is:

$$U_{nik} = -f_{ik} - \beta_n^{CNX}\delta_{ik}^{CNX} - \beta_n^{DTW}X_{nik}^{DTW} - \left(\sum_{\forall j}\beta_n^{R_j}\delta_{ik}^{R_j}\right)$$

$$(i_n^*, k_n^*) = \underset{(i,k): f_{ik} \leq \theta_n}{\arg\max} \; U_{nik}$$

The preferences β are randomly distributed (e.g. $\beta^{R_j} \sim \Omega_j^R$); the budgetary constraint is drawn from a shifted exponential distribution and disutilities from a normally distribution; the distribution means and variances may vary based on latent class (demand segment) membership.

Ancillary Valuation Models

In the airline context, the impact of ancillary services on passenger choices is poorly understood. Most ancillary research has used discrete choice models to study how passengers value ancillary services—not how passenger decide to purchase ancillary services in the context of the rest of their travel decisions. For example, Balcombe et al. (2009) use a stated preference survey to estimate willingness-to-pay for meals, beverages, seat legroom, seat width, and in flight entertainment; Mumbower et al. (2015) use a revealed preference dataset to assess the factors that influence purchase of extra legroom seats on JetBlue.

Garrow et al. (2007) incorporate the value of legroom into a study primarily focused on how the number and type of connections affect itinerary choice, and Espino et al. (2008) integrate meals into a study of airline choice (modeling meals a a quality differentiator between airlines). Neither model, however, incorporates separate, optional sales of the service (legroom or meals), so the results do not indicate how ancillary purchase decisions might be related to itinerary purchase decisions.

2.3.2 Bundled Goods

Economists typically assume that consumers perceive the utility of a bundle of goods as equal to the sum of the utilities of its components (e.g. Adams and Yellen (1976)). In the marketing literature, however, more complex assumptions are made. Ben-Akiva and Gershenfeld (1998) develop a probabilistic choice nested logit to model how consumers select telephone ancillary services (call waiting, caller ID, etc.) from a menu containing bundled and à la carte items. Gaeth et al. (1990) propose that consumers evaluate the utility of a bundle according to some average of the utility of each component; Yadav (1994) provides some empirical evidence to claim that consumers evaluate bundles in a multi-step process, determining a "dominant" good upon which to anchor their valuation and then adjusting the valuation up and down based on other components of the bundle.

2.3.3 Add-ons

In the economics literature, when one product (such as in-flight internet access) only provides value to a consumer if purchased in conjunction with another product (such as a plane ticket), it is referred to as an *add-on* (Lal and Matutes, 1994; Ellison, 2005; Fruchter et al., 2011; Geng and Shulman, 2015). Specific to the airline industry, Allon et al. (2011) show that a monopolist should unbundle baggage fees to reduce the airline's operating cost, and that airline baggage policies are consistent with cost reduction but not demand segmentation, based on the authors' model of airline demand segments. Brueckner et al. (2015) model consumers as perfectly informed and rational and show that, with idiosyncratic checked bag preferences, a monopolist can increase profits by simultaneously unbundling and lowering base fares; the optimal unbundled base fare plus checked bag price may be lower or higher than the original bundled fare. Cui et al. (2016) show that, with two demand segments, unbundling is profitable for a monopolist airline if it can price discriminate for fares and higher fare consumers are less likely to purchase the add-on; when the airline cannot discriminate, unbundling is profitable if higher fare consumers are more likely to purchase the add-on. As above, the conclusions of these papers are dependent on models of consumer behavior, and determining the applicability of the results to any particular situation requires understanding how airline passengers view and value ancillary services.

Two notable papers, although not airline-focused, develop models that explicitly incorporate consumer limitations in knowledge/awareness of add-on services when selecting a base goods. Gabaix and Laibson (2006) develop an analytical model examining the optimal pricing and consumer welfare implications of add-on pricing when firms can *shroud* add-on prices. The authors define a shrouded attribute as "a product attribute that is hidden by a firm, even though the attribute could be nearly costlessly revealed," drawing parallels to the add-on prices of airlines (as well as other business) that are difficult to determine (Gabaix and Laibson, 2006, pg. 512). They allow consumers to be *sophisticates*, who are classically rational, or *myopes*, who are boundedly rational. In particular, myopes do not (always) account for the price of the add-on when selecting a base good, and never account for a shrouded add-on price.

Shulman and Geng (2013) also develop an analytical model for add-on pricing for two firms. In addition, their model incorporates asymmetric quality in products between the firms and allows consumers to have a firm preference. *Base* consumers have no need for the add-on; they choose the firm that maximizes a base good utility (given by base good quality net of base good price). *Boundedly rational* consumers do value the add-on, but (erroneously) believe that it does not have an additional cost. They select the firm that maximizes a base good utility (base good quality net of base good price) *plus* the add-on quality. *Knowledgeable* consumers know add-on prices; they choose the firm that maximizes total utility (base good quality net of base good price plus add-on good quality net of add-on good price).

In both models, the presence of myopic/boundedly rational consumers results in different optimal pricing schemes by the firms than would be the case with all classically rational consumers, illustrating the importance of accounting for cognitive bounds: Gabaix and Laibson (2006) show that unlike previous models (e.g. Lal and Matutes (1994) and Ellison (2005)), if there is a sufficiently high portion of myopes, shrouding can be a price equilibrium even when advertising (which represents the cost of unshrouding) is free. In addition, in such an equilibrium, there is a "curse of debiasing:" neither firms nor sophisticates have an incentive to educate (or debias) myopes: decreasing the portion of myopes increases equilibrium base good prices (which is bad for the existing sophisticates) and decreases firm profits (which is bad for firms).

Although these approaches provide mechanisms to incorporate add-ons and bounded rationality in choice models, they have several simplifications which limit their direct applicability to the airline passenger choice problem. All consumers in the models purchase the base good from a firm (regardless of price) and the firm is limited to one add-on. In the Gabaix and Laibson (2006) model, all uninformed myopes purchase the add-on regardless of cost; all other consumers avoid the add-on at a fixed, deterministic cost. In Shulman and Geng (2013), all non-base consumers value the add-on for a given firm at a fixed, deterministic value. Quality is modeled as a single exogenous and deterministic attribute. All consumers have the same deterministic quality assessments for the goods and the add-ons, and the firm with the higher-quality base good also has the higher-quality add-on.

2.4 Key Literature Gaps

Our model, described in Chapter 3, combines concepts from previous airline itinerary choice models as well as add-on models.

Although previous work has studied factors that influence selection of itineraries and/or fare classes, and valuation of ancillary services, none of the models described above explicitly connect decisions about ancillary services to decisions about itineraries and fare classes. In addition, the airline passenger choice models all assume classically rational behavior by passengers. Finally, ancillary services differ from other dimensions in the passenger choice problem because they impose additional out-of-pocket costs, are optional, and require the purchase of a (base) itinerary/fare class. Thus, none of the models described above can be used directly for the ancillary choice problem. Our model extends itinerary/fare class choice

Table 2.1: Summary of relevant choice modeling literature

	Base good	Itinerary	Fare class	Ancillary or add-on	Bounded rationality
Passenger Choice					
Espino et al. (2008)				✓	
Balcombe et al. (2009)				✓	
Mumbower et al. (2015)				✓	
Smith et al. (2007)		✓			
Hess and Adler (2011)		✓			
Theis (2011)		✓			
Algers and Beser (2001)		✓	✓		
Hopperstad (2005)		✓	✓		
Carrier (2008)		✓	✓		
Economics					
Lal and Matutes (1994)	✓			✓	
Ellison (2005)	✓			✓	
Fruchter et al. (2011)	✓			✓	
Geng and Shulman (2015)	✓			✓	
Gabaix and Laibson (2006)	✓			✓	✓
Shulman and Geng (2013)	✓			✓	✓
This Work					
Ancillary Choice Model	✓	✓	✓	✓	✓

models to include the ancillary service dimension and to support non-rational behavior and decision making.

In the context of the add-on pricing literature, we extend base-good/add-on choice models to incorporate threshold and compensatory decision rules, to include variations in consumer preferences, and to allow itinerary/fare classes to differ in multiple dimensions. Table 2.1 summarizes the key differences between our model and previous work.

Chapter 3

Ancillary Choice Model

In this chapter we develop, calibrate, and assess a passenger choice model that describes how consumers select ancillary services in conjunction with an airline itinerary and specific fare class.[1] We propose that consumers are either classically rational, which we term *simultaneous*, or boundedly rational, which we term *sequential*. Simultaneous consumers are modeled to choose an itinerary, fare class, and set of ancillary services at the same time; sequential consumers are modeled to choose an itinerary and fare class in one phase and then choose a set of ancillary services in a second phase. We also assess the sensitivity of revenues and bookings to various ancillary fee structures, showing that ancillary bundling can result in buy-up or buy-down by simultaneous passengers, suggesting a potential benefit for integrating ancillary services into revenue management processes.

[1]Portions of this chapter were previously published as Bockelie, A. and Belobaba, P. (2017). Incorporating ancillary services in airline passenger choice models. *Journal of Revenue and Pricing Management*, 16(6):553–568.

3.1 Model Formulation and Definitions

The Ancillary Choice Model (ACM) introduced in this chapter is an integrated itinerary, fare class, and ancillary service passenger choice model that extends existing itinerary and fare class choice models. The ACM allows consumers to differ in their level of knowledge about ancillary services and allows fare classes and itineraries to differ in their ancillary offerings and prices.

We define ancillary services using the terms of Shugan et al. (2017). Namely, ancillary services are those that have a value less than the core transportation product, are occasionally bundled with the core, and, when unbundled, can be purchased only in conjunction with the core product. In addition, we restrict our attention to services that are "itinerary-related" or "trip-related," as defined in Chapter 1.

In the ACM, consumers are assumed to fall into one of two behavior types based on their knowledge of airline ancillary service policies and prices. We define *simultaneous* consumers as those who integrate information about ancillary services into their decision about itineraries and fare classes, and thus behave as classically rational consumers similar to sophisticates in Gabaix and Laibson (2006) or knowledgeable consumers in Shulman and Geng (2013). We define *sequential* consumers, on the other hand, as those who are boundedly rational in knowledge, awareness, or computing ability. Our sequential consumers correspond to Gabaix and Laibson's myopes or Shulman and Geng's boundled rational segment, and do not consider ancillary service differences until after selecting an itinerary and fare class. In the terms of Morwitz et al. (1998), simultaneous passengers perfectly calculate prices (or generalized costs) of all alternatives, while sequential passengers (initially) ignore ancillary prices completely.

Table 3.1 summarizes the notation. In a given origin-destination market, one or more airlines offer one or more itineraries (indexed by i), which may be non-stop, direct, or connecting. Each airline offers one or more fare classes for each itinerary (indexed by k), which may be subject to restrictions on their use and/or purchase. Together, the set of all possible itinerary and fare class combinations within the market (including the option to not fly) is $\mathcal{G}$. Each airline operates a revenue management system, so not all combinations are available to all consumers; when consumer n arrives, she is presented with the set $\mathcal{G}_n \subseteq \mathcal{G}$ of *available* combinations. Itinerary and fare class attributes, including fares, as well as the availability set, are determined by the airlines and assumed to be exogenous.

Each airline also offers zero or more ancillary services for each itinerary and fare class combination. Ancillary services are bundled for notational convenience; each consumer purchases exactly one bundle m, which could be the bundle containing no ancillary services.[2] Each itinerary/fare class combination (i, k) offers the set of ancillary bundles $\mathcal{M}_{ik}$, which is determined by the airlines and assumed to be exogenous.

Over the course of an Initial Booking Phase (IBP) and a Follow-up Phase (FUP) each consumer chooses a combination of itinerary, fare class, and ancillary bundle (i, k, m), with the possibility of selecting the "combination" corresponding to not flying. The IBP corresponds to *decisions* made when the ticket is initially purchased; the FUP corresponds to decisions made later. It is important to note that these phases represent the timing of purchase decisions, not necessarily the act of making the purchase itself.[3]

Consumers of both behavior types select the utility-maximizing combination from their choice set (as described below). The utility of any alternative is decomposed into an itinerary/fare class utility and an ancillary utility. The itinerary/fare class utility is provided by an existing itinerary/fare class choice model referred to as the "kernel." Although the kernel may take various forms, the ancillary utility component is an infinite-scale random coefficient mixture of logit model with threshold rules (see Chapter 2).

Each consumer has an out-of-pocket budgetary limit of $\theta_n \geq 0$, drawn from distribution Θ,

[2]This modeling notation is not meant to imply that ancillary services must be marketed to consumers or priced as bundles.

[3]Some airlines and sale channels do not allow passengers to purchase ancillary services at the same time as purchasing a ticket; this technological limitation does not prohibit simultaneous choice behavior.

as well as a disutility $\beta_{ns} \geq 0$ experienced when foregoing ancillary service s (drawn from distribution Ω_s). Consumers also have preferences about itinerary and fare class attributes (such as number of connections, airline, etc.). Multiple consumer segments (such as business or leisure) may be present, each with different budget and disutility distributions. However, for notational convenience, in this section we do not include a consumer segment index.

3.1.1 Sequential Passengers

A sequential consumer makes a decision about an itinerary and fare class during the IBP without regard to ancillary services, using the fare class/itinerary kernel, such as Hopperstad (2005). The budgetary limit constrains out of pocket expenses incurred during the IBP. Subject to the budgetary limit and RM availability, the consumer selects the utility-maximizing itinerary and fare class combination. The utility $\hat{U}_{nik}$ for sequential consumer n of itinerary i and fare class k is given by the itinerary/fare class choice model. Every passenger also has the option of not flying, which has utility $\hat{U}_{n00} = -R$ and fare $f_{00} = 0$. When $-R = -\infty$, the no-fly option is the least desirable alternative, but the zero fare ensures that it is always included in the choice set.

At the time passenger n makes a booking request, the distribution system returns the "availability set" of itinerary and fare class combinations available for purchase $\mathcal{G}_n$, including the option to not fly. The passenger, however, only considers the set of combinations $Q_n(\theta_n)$ (the "consideration set") with an out of pocket cost within their budget:

$$Q_n(\theta_n) = \{(i,k) \in \mathcal{G}_n : f_{ik} \leq \theta_n\}$$

During the IBP, passenger n chooses the utility-maximizing combination (i_n^*, k_n^*) from $Q_n(\theta_n)$:

Table 3.1: Summary of ancillary choice model notation

Counts	
N	Number of consumers in market
S	Number of ancillary services in market
Indices	
n	Consumer index
i	Itinerary index
k	Fare class index
s	Ancillary service index
m	Ancillary bundle index
i_n^*	Itinerary selected by consumer n
k_n^*	Fare class selected by consumer n
m_n^*	Ancillary bundle selected by consumer n
Consumer Preferences	
θ_n	Consumer n budgetary limit, $\theta_n \sim \Theta$
β_{ns}	Consumer n disutility for forgoing ancillary service s, $\beta_{ns} \sim \Omega_s$
Availability and Consideration Sets	
$\mathcal{G}$	Set of all possible itineraries/fare class combinations, including the option to not fly
$\mathcal{G}_n$	Availability set, set of itineraries/fare classes available to consumer n, $\mathcal{G}_n \subseteq \mathcal{G}$
$Q_n(\theta_n)$	Consideration set, subset of $\mathcal{G}_n$ with sequential OPC less than or equal to θ_n
$L_n(\theta_n)$	Consideration set, subset of $\mathcal{G}_n$ with simultaneous OPC less than or equal to θ_n
$\mathcal{M}_{ik}$	Set of ancillary bundles applicable to itinerary i and fare class k
Utility	
U_{nikm}	Utility for consumer n of itinerary i, fare class k, and ancillary combination m
$\hat{U}_{nik}$	Utility for consumer n of itinerary i and fare class k, excluding ancillary services
$-R$	Utility of no-fly option, typically set to $-R = -\infty$
Alternative Attributes	
f_{ik}	Fare of itinerary i, fare class k
p_{ikm}	Price of ancillary bundle m for itinerary i, fare class k
δ_{ims}	1 if ancillary service s is forgone for itinerary i ancillary bundle m; 0 otherwise

$$(i_n^*, k_n^*) = \underset{(i,k)\in Q_n}{\arg\max}\, \hat{U}_{nik}$$

After selecting (i^*, k^*), sequential consumers enter the FUP to evaluate ancillary services. The consumer evaluates the utility $U_{ni_n^*k_n^*m}$ of each bundle m of purchasable services for their selected itinerary and fare class from the set of all bundles $\mathcal{M}_{i_n^*k_n^*}$, which typically includes the bundle of no ancillary services:

$$U_{ni_n^*k_n^*m} = \hat{U}_{ni_n^*k_n^*} - p_{i_n^*k_n^*m} - \sum_{s=1}^{S} \beta_{ns}\delta_{i_n^*ms}$$

where p_{ikm} is the price of ancillary bundle m when purchased for itinerary i in class k; δ_{ikms} is a binary indicator variable with value 1 if ancillary service s is foregone (i.e. neither purchased nor received complimentary) by passengers selecting bundle m with itinerary i and class k, and value 0 otherwise. Note that for the bundle containing no ancillary services, the consumer will experience no additional ancillary fee but will experience a disutility for every ancillary service.

Consumer n selects the combination of ancillary services m_n^* that maximizes their utility:

$$m_n^* = \underset{m\in\mathcal{M}_{i_n^*k_n^*}}{\arg\max}\, U_{ni_n^*k_n^*m}$$

If $\mathcal{M}_{ik}$ includes all possible bundles of ancillary services and no-purchase options, and if ancillary services are priced individually (i.e., p_{ikm} is the sum of the prices of all included services, without any discounting), a sequential consumer will never purchase an ancillary service if its price is greater than their disutility for that service.

The probability that a sequential passenger selects a particular itinerary, fare class, and ancillary service combination is a function of the passenger's itinerary/fare class utility, ancillary service disutilities, budget, and the availability set. The itineraries and classes in the availability set are outputs of the various airlines' revenue management systems, and are random variables with unknown distributions. The conditional probability $\Pr((i,k) \mid \mathcal{G}_n)$ of a passenger choosing itinerary and fare class combination (i,k) given an availability set $\mathcal{G}_n$ is:

$$\Pr\left((i,k) \mid \mathcal{G}_n\right) = \int_0^{\infty} \Pr\left((i,k) \mid Q_n(\theta_n)\right) f_\theta(\theta_n) \mathrm{d}\theta_n$$

where $f_\theta(\theta_n)$ is the probability density function of a passenger's budgetary limit. Together, θ_n and consideration set $Q_n(\theta_n)$ form a latent class for each passenger. Given membership within a particular class (i.e., given that a passenger can afford a particular set of itinerary and fare class combinations), the probability of choosing itinerary and fare class combination (i,k) is:

$$\Pr\left((i,k) \mid Q_n(\theta_n)\right) = \begin{cases} \Pr\left(\hat{U}_{nik} \geq \max\limits_{(i',k') \in Q_n(\theta_n) \setminus (i,k)} \hat{U}_{ni'k'}\right) & (i,k) \in Q_n(\theta_n) \\ 0 & (i,k) \notin Q_n(\theta_n) \end{cases}$$

The probability that a sequential passenger purchases ancillary service combination m, given that they have selected itinerary and fare class combination (i,k) is:

$$\Pr\left(m \mid (i,k)\right) = \begin{cases} \Pr\left(U_{nikm} \geq \max\limits_{m' \in \mathcal{M}_{ik} \setminus m} U_{nikm'}\right) & m \in \mathcal{M}_{ik} \\ 0 & m \notin \mathcal{M}_{ik} \end{cases}$$

Thus, the probability that a sequential passenger purchases itinerary i, fare class k, and ancillary combination m, given an availability set, is:

$$\Pr((i,k,m) \mid \mathcal{G}_n) = \int_0^{\infty} \Pr(m \mid (i,k)) \Pr((i,k) \mid Q_n(\theta_n)) f_\theta(\theta_n) \mathrm{d}\theta_n$$

3.1.2 Simultaneous Passengers

A simultaneous consumer selects an itinerary, fare class, and set of ancillary services during the IBP. Because the IBP decision includes ancillary services, those prices are included in the budget-constrained out of pocket cost. The utility U_{nikm} for simultaneous consumer n of itinerary i, fare class k, and combination of ancillary services m combines the kernel and ancillary contribution:

$$U_{nikm} = \hat{U}_{nik} - p_{ikm} - \sum_{s=1}^{S} \beta_{ns}\delta_{ims} \tag{3.1}$$

Again, every passenger also has the option of not flying, which has utility $U_{n000} = -R$ and fare $f_{00} = 0$.[4] When passenger n makes a booking request the distribution system returns availability set $\mathcal{G}_n$. Each passenger n will choose the utility-maximizing combination (i_n^*, k_n^*, m_n^*) from the consideration set $L_n(\theta_n)$, which includes all available itinerary, fare class, and ancillary service combinations with an out of pocket cost within the passengers' budget:

$$L_n(\theta_n) = \{(i,k,m) : (i,k) \in \mathcal{G}_n, m \in \mathcal{M}_{ik}, f_{ik} + p_{ikm} \leq \theta_n\}$$

$$(i_n^*, k_n^*, m_n^*) = \operatorname*{arg\,max}_{(i,k,m) \in L_n(\theta_n)} U_{nikm}$$

[4]This is the same as the sequential utility for not flying because consumers are assumed to have no benefit or disutility for ancillary services when not flying.

As with sequential consumers, if $\mathcal{M}_{ik}$ includes all combinations of ancillary and no-ancillary options and if ancillary services are priced individually, a simultaneous consumer will never purchase an ancillary service if its price is greater than their disutility for that service. If $\mathcal{M}_{ik}$ includes the option to buy no ancillary services, a simultaneous passenger will fly without ancillary services before choosing not to fly at all. Because simultaneous passengers choose an itinerary, fare class, and ancillary bundle in the IBP, they have no FUP.

The probability that a simultaneous passenger selects a particular itinerary, fare class, and ancillary service combination is again a function of the a function of the passenger's itinerary/fare class utilities, ancillary service disutilities, budget, and the availability set. The conditional probability $\Pr((i,k,m) \mid \mathcal{G}_n)$ of a passenger choosing itinerary, fare class, and ancillary service combination (i,k,m) given an availability set $\mathcal{G}_n$ is:

$$\Pr\left((i,k,m) \mid \mathcal{G}_n\right) = \int_0^\infty \Pr\left((i,k,m) \mid L_n(\theta_n)\right) f_\theta(\theta_n) \mathrm{d}\theta_n$$

where the consideration set and budget again form the latent class. $\Pr\left((i,k,m) \mid L_n(\theta_n)\right)$ takes the form:

$$\Pr\left((i,k,m) \mid L_n(\theta_n)\right) = \begin{cases} \Pr\left(U_{nikm} \geq \max\limits_{(i',k',m') \in L_n(\theta_n) \setminus (i,k,m)} U_{ni'k'm'}\right) & (i,k,m) \in L_n(\theta_n) \\ 0 & (i,k,m) \notin L_n(\theta_n) \end{cases}$$

3.1.3 Key Differences in Behaviors

Simultaneous and sequential passengers differ in two critical ways. First, because different components are considered in the utility functions during the IBP, passengers with different behaviors who otherwise have identical preference may have different views on the "best" alternative, and may select different fare classes and itineraries. Second, because the budgetary constraint is applied in different ways for the two behavior types and because it is assumed to apply only to decisions made during the IBP, it constrains both fare and ancillary purchases for simultaneous passengers. For sequential passengers, however, only fare is constrained by the budget. Hence, all else equal, sequential passengers can "afford" more ancillary services and will have a higher ancillary purchase rate and will contribute more ancillary revenue.

Because simultaneous passengers incorporate ancillary services into their IBP utility, ancillary services can incentivize them to either buy-up or to buy-down—meaning that the presence of ancillary services can cause a passenger to select a more or less expensive fare class than they otherwise would have.

With a few mild assumptions, we can assess the conditions in which a simultaneous passenger would change her fare class selection from (i^*, k^*) absent ancillary services to $(i^*, k^\dagger, m^\dagger)$ when ancillary services are present. We assume that the availability set and passenger preference distributions are unaffected by the ancillary service, and we assume that ancillary prices for any given bundle are non-decreasing in fare class index and that fares are strictly decreasing in fare class index (i.e. higher number fare classes have lower fares and do not have lower ancillary prices):[5]

$$\partial p_{km}/\partial k \geq 0 \qquad \partial f_k/\partial k < 0$$

To simplify notation, we drop the consumer and itinerary indices. Note that absent the ancillary, the passenger's choice will be governed only by the itinerary/fare class kernel; therefore the class k^* has the greatest kernel utility of all affordable fare classes:

$$\hat{U}_{k^*} \geq \hat{U}_{k'} \quad \forall k' \in \mathcal{G}_n \text{ s.t. } f_{k'} \leq \theta \tag{3.2}$$

[5]The assumption of decreasing fare in fare index is common and implies no loss of generality; the assumption on ancillary prices holds for most airline ancillary fee structures in practice.

Likewise, with the ancillary, the combination $(k^\dagger, m^\dagger)$ has the highest simultaneous utility of all affordable combinations:

$$U_{k^\dagger m^\dagger} \geq U_{k',m'} \quad \forall k' \in \mathcal{G}_n, m' \in \mathcal{M}_{k'} \text{ s.t. } f_{k'} + p_{k'm'} \leq \theta$$

Thus, when $k^\dagger \neq k^*$ (the fare class choice has changed), the following equation must hold:

$$U_{k^\dagger m^\dagger} - U_{k^* m'} > 0 \quad \forall m' \in \mathcal{M}_{k^*} \text{ s.t. } f_{k^*} + p_{k^* m'} \leq \theta \tag{3.3}$$

In other words, the new combination has a higher utility than any (affordable) combination in fare class k^*. Combining Equations 3.1 and 3.3 yields:

$$\left(\hat{U}_{k^\dagger} - \hat{U}_{k^*}\right) - \left(p_{k^\dagger m^\dagger} - p_{k^* m'}\right) - \left(\sum_{s=1}^{S} \beta_s \left(\delta_{m^\dagger s} - \delta_{m' s}\right)\right) > 0 \tag{3.4}$$

$$\left(\hat{U}_{k^\dagger} - \hat{U}_{k^*}\right) - \left(v_{k^\dagger m^\dagger} - v_{k^* m'}\right) > 0 \tag{3.5}$$

where v_{km} is the *net ancillary disutility contribution* associated with selecting ancillary bundle m in fare class k:

$$v_{km} = p_{km} - \sum_{s=1}^{S} \beta_s \delta_{ms}$$

Since both $k^\dagger$ and f^* have been selected, they must both be within the budgetary constraint. Therefore, from Equation 3.2, $\hat{U}_{k^*} \geq \hat{U}_{k^\dagger}$ and first term in Equations 3.4 and 3.5 must be non-positive. For the fare class choice to change, Equation 3.4 and 3.5 must hold and the second term of Equation 3.5 must be negative: the net ancillary disutility of $(k^\dagger, m^\dagger)$ must be lower than the net ancillary disutility of any (affordable) combination in class k^*.

The intuition for buy-up ($k^\dagger < k^*$, where the ancillary entices passengers to purchase a more expensive, but lower index fare class) is clear: if some fare classes offer complimentary or discounted ancillary services, passengers who consider ancillary services during the IBP (i.e. simultaneous passengers) will find those fare classes *more appealing* when the ancillary services are present than in a case absent ancillary services.

Ancillary-driven buy-up is leveraged in airline pricing structures that entice consumers to book in higher value fare classes with complimentary services, such as the complimentary access to preferred seating or branded fare structures discussed in Section 1.1.2.

The case for buy-down ($k^\dagger > k^*$, where the ancillary entices passengers to purchase a less expensive, but higher index fare class) is more complicated: some passengers who, absent ancillary services, would purchase fares close to their budgetary constraint may find that they do not have a sufficient budget when the ancillary service is present to purchase the same fare class *as well as* the ancillary service. They will change to the lower value fare class if the improvement in their net benefit (v) of purchasing the service outweighs the decrease in their itinerary/fare class utility associated with the lower value fare class.

With sequential passengers, none of these booking shifts occur, because the IBP utility does not incorporate ancillary services. The potential for these shifts has implications for airline pricing: complimentary or discounted ancillary services in some fare classes will entice (some) simultaneous passengers to purchase that class, which may be beneficial to the airline (because of a higher profit margin, perhaps). However, with sequential passengers, the discount will not drive any changes in bookings, and will only dilute revenue from ancillary purchases. Thus, airlines have no incentive to offer complimentary ancillary services if passengers are sequential.

3.1.4 Differences from Previous Models

Our formulation is the first that explicitly links the purchase of ancillary services, or add-ons, to the selection of an itinerary/fare class, or base good, while allowing consumers to choose from many base goods that may vary in multiple quality dimensions. In the airline literature, no other works have related ancillary purchases to itinerary/fare class purchases. In addition, unlike much of the existing airline literature, our model allows consumers to vary in their level of knowledge and awareness of airline policies, capturing both classic and boundedly rational behavior.

In the economics literature, some authors (such as Gabaix and Laibson (2006) and Shulman and Geng (2013)) study the relationship between add-on purchases and base good purchases. These models, however, do not account for the multiple, varied base goods present in the itinerary/fare class choice problem and do not allow consumers to chose not to purchase the base good. Our formulation combines concepts from both the airline and economics fields to describe passenger behavior given a choice of fare classes and ancillary services.

3.2 Passenger Origin-Destination Simulator

We integrated the ACM within the Passenger Origin-Destination Simulator (PODS) to study the effects of our passenger choice model, as well as to study mechanisms for incorporating ancillary revenues into revenue management models (see Part II of this thesis).

PODS is a software package originally developed by Boeing to evaluate the interaction between passenger choice and forecasting and revenue management optimization models. It simulates the interactions between potential customers and airlines in a competitive environment and has been used in the past to assess the effects of low-fare carriers entering markets (Gorin, 2004), the influences of revenue management system users (Weatherford, 2016), and the potential revenue benefits of various dynamic pricing mechanisms (Wittman, 2018), among other topics.

Each PODS simulation consists of several independent *trials*; each trial has several hundred *samples*. A sample represents a single departure day, and all samples represent different realizations of the same departure day (meaning each sample has the same flight schedule and same expected demand, but randomized actual demand). A trial is a series of successive samples, where the bookings received for one sample are used to forecast subsequent samples. The first samples of each trial are *burned*, or excluded from all analyses, and are only used to warm up the airline demand forecasting models. The simulator reports results for a simulation as aggregate (e.g. mean, variance, etc.) metrics (e.g. ticket revenue, bookings by fare class, etc.) across all unburned samples for all trials. A typical simulation features two to five trials, a warm-up of 200 samples per trial, and an additional 400 samples per trial for analyses. Thus, the reported results are aggregates of 800 to 2,000 realizations of a particular departure day. Each simulation has the same set of random seeds, so differences between simulation results are due to changes in pricing, forecasting, or revenue management optimization models, not due to differences in underlying demand generation. Statistical testing is performed with a paired t-test, with the null hypothesis that there is no true change in revenue vs alternative hypothesis that there is a true change in revenue; see Sections 5.2.1 and 7.2 for more details. The paired test means that although the simulator has high stochasticity in terms of demand and passenger preferences, very small revenue changes can be statistically significant.

PODS features two interdependent modules, as shown in the schematic in Figure 3.1: an agent-based passenger generation and choice module, and an airline revenue management and distribution module. The passenger generation and choice module generates passengers within each origin-destination market in the simulation, and assigns randomly drawn preferences and disutilities to each passenger. Passengers book the itinerary and fare combination that is most desirable according to their preferences and choice model.

The booking window for each departure day is 63 days long and is broken into 16 *time frames*

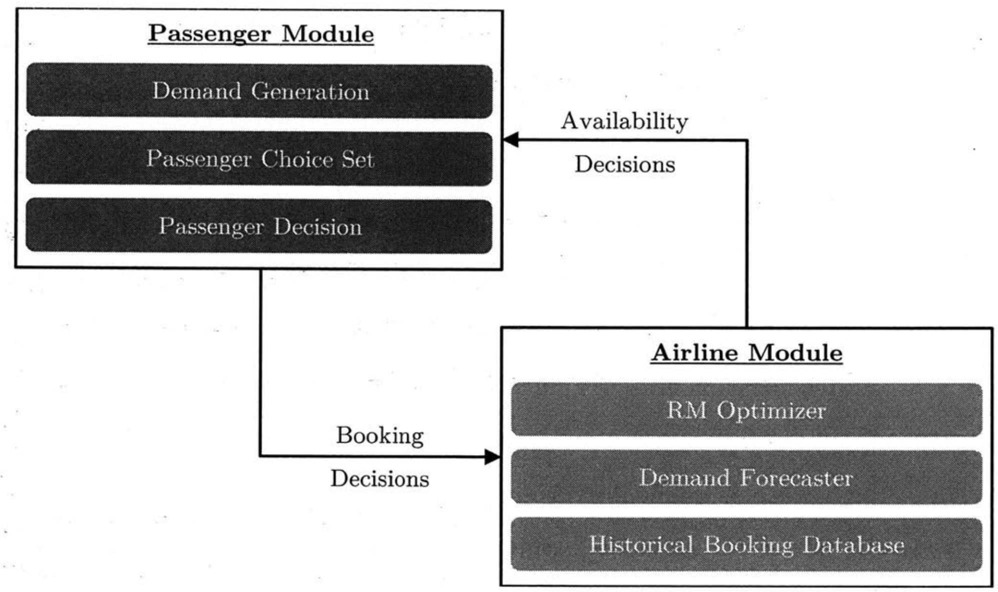

Figure 3.1: Schematic view of Passenger Origin-Destination Simulator. Modified from Wittman (2018).

or *data collection points* (DCPs) for passenger generation, booking database, forecasting, and optimization purposes. The booking window opens at the start of DCP 1, and departure occurs at the end of DCP 16. The number of days in each DCP varies, with longer DCPs early in the booking window and shorter DCPs late in the booking window. This results in more frequent and finer temporal resolution forecast and optimization updates closer to departure.

3.2.1 Passenger Module

Demand Generation

As discussed in Section 1.1.3, passenger demand is composed of different passenger segments, with each segment representing different travel purposes and/or time, cost, and quality sensitivities. PODS supports two passenger segments; passengers in the *business* segment represent travelers with low value of money and high value of time while passengers in the *leisure* segment, on the other hand, represent travelers with a high value of money and low value of time.

The overall mean demand for each market and passenger segment is specified as a simulation input. The actual mean demand for any sample is a function of several random factors:

- A single *system* factor, which increases or decreases demand for all markets and passenger segments.
- A *market* factor, drawn separately for each market, which increases or decreases demand for individual markets for all passenger segments.
- A *passenger segment* factor, drawn separately for each passenger segment, which increases or decreases demand for individual passenger segments for all markets.

The actual demand for each passenger segment and market, for a given sample, is normally distributed based on the actual mean demand described above, and variance that is a constant multiple (typically 2.0) of the actual mean demand.

As an example, consider an input demand of 100 for business passengers in market A. A system factor of 1.1 (representing high overall demand), market factor of 0.9 (meaning market A, compared to others, is less busy), and passenger segment factor of 1.0 (meaning business passengers are not more or less busy compared to other segments) would have an actual mean demand of 99. Thus business passenger demand for market A for the sample would be normally distributed with mean 99 and variance 198.

A demand arrival curve for each passenger segment controls when the actual demand arrives during the booking window. The curve specifies the mean percentage of demand arriving within a particular DCP; the actual demand arriving during the DCP is again normally distributed (with variance given by a constant multiple of the mean). The cumulative demand arrival curves for domestic business and leisure travelers in Network U10 are shown in

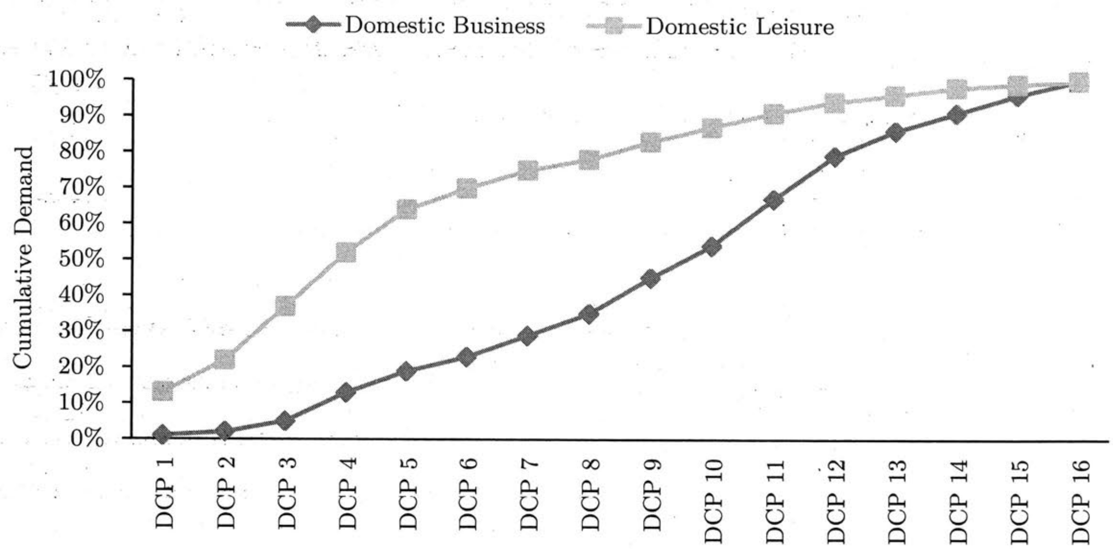

Figure 3.2: Cumulative demand arrival curves for consumers in domestic markets in Network U10

Figure 3.2. In general, business passengers arrive closer to departure and leisure passengers arrive further in advance.

Passenger Characteristics As an agent-based simulation, each passenger is modeled individually. When generated, passengers are assigned behaviors, budgets, and a set of preferences. In our simulations, passengers are randomly assigned a behavior type (simultaneous or sequential), according to an input probability. Preferences consist of disutilities for connections, for schedule deviations from a passenger's preferred departure and arrival times, for fare class purchase/use restrictions (such as a round trip or Saturday night stay requirement), and for foregoing ancillary services.

All budgetary and disutility distributions are dependent on passenger segment and origin-destination market as an input linear function of the *base fare* ($bfare$) for the segment and market. The base fare is typically the lowest fare in the market for leisure passengers and 2.5 times the lowest fare for business travelers. Budgets are distributed according to an exponential distribution (shifted by the base fare) with an input elasticity multiplier for each passenger segment. Thus all simulated passengers can afford their base fare within the market.

Disutilities are normally distributed and are also a function of base fares, such that $\Omega_{jg} = N(\mu_{jg}, \sigma_{jg}^2)$ with $\mu_{jg} = c_{jg} + d_{jg}bfare_g$ and $\sigma_{jg} = v\mu_{jg}$, where Ω_{jg} is the distribution of disutility j for passenger segment g (for connections, schedule deviation, fare class restriction, or foregone ancillary service), c_{jg} and d_{jg} are input intercept and slope parameters, and v is the input coefficient of variation.

Business passengers tend to have higher budgetary limits and higher disutilities for connections, schedule deviations, fare class restrictions, and foregoing ancillary services. Passengers from the leisure segment, on the other hand, tend to have lower budgetary limits, and lower connection, schedule deviation, fare class restriction, and foregone ancillary service disutilities. Thus, in general, business passengers are more likely than leisure passengers to spend more to purchase a less restricted but more expensive itinerary/fare class. Because preferences are random and widely distributed, though, some leisure passengers will have higher budgets and higher disutilities for fare class restrictions than some business passengers.

Passenger Choice Set and Decision

As described in Section 3.1, when a passenger makes a booking request in a particular market, each airline serving that market returns a list of available itinerary/fare class combinations (which is an output of the airlines' revenue management systems). Together, these options form the passenger's availability set. PODS passengers make a booking request when they are generated.

Ancillary services are sold and priced separately and all passengers have the option to forgo any service. Ancillary prices vary by market, but are the same for all airlines, itineraries, and fare classes within a market. Prices are a linear function of the lowest published fare within the market. For service s, price $\hat{p}_s = a_s + b_s \min_{(i,k)} f_{ik}$, where a_s and b_s are input parameters. For modeling purposes ancillary bundle prices p_{ikm} are equal to the sum of all ancillary service prices $\hat{p}_s$ included within the bundle. Within a given market, ancillary prices are the same for business and leisure passengers.

The passenger constructs their choice set as the portion of the availability set with an out-of-pocket cost less than or equal to the passenger's budgetary constraint. As described in Section 3.1, the out-of-pocket cost could include the price of ancillary services for simultaneous passengers. The ACM implementation in PODS incorporates the previous PODS passenger choice model (Hopperstad, 2005) as the itinerary/fare class choice kernel. Passengers who choose not to fly are not "regenerated" in subsequent samples.

3.2.2 Airline Module

The airline RM and distribution module stores information on passenger bookings and uses that information to generate demand forecasts. The demand forecasts are fed to an optimizer that produces booking policies; the booking policies determine the set of itineraries and fare classes made available to passengers within each market. The airline RM and distribution module only has access to the same types of information as real airlines; underlying simulation details, like demand generation parameters, are not visible to the module. Each airline in the simulation has its own booking database, forecasting system, and optimization system. This section reviews the forecasting and optimization models supported by PODS that are relevant to this thesis. For a more technical description and history of forecasting and optimization models, see Chapter 4.

Historical Database

Each booking received by an airline in the simulation is recorded in that airline's historical booking database. These databases aggregate bookings by DCP, fare class, and either itinerary or flight leg (depending on the demand forecasting model in use, see below). The simulated airlines also store data on ancillary purchases, aggregated by ancillary service, fare class, and market.[6]

Demand Forecaster

Forecasting models predict future demand based on historical booking observations. PODS supports demand forecasts generated for each itinerary/fare class combination or for each flight leg/fare class combination. Demand forecasts in PODS are generated at the start of each DCP. PODS forecasting modeling include:

- *Standard forecasting*, the simplest forecasting model, which estimates future demand for an itinerary/fare class or flight leg/fare class as an average of bookings received in the past for the same itinerary/fare class or flight leg/fare class, assuming that demand for each class and itinerary/leg are independent (Littlewood, 1972).

- *Q forecasting*, developed by Hopperstad and Belobaba (2004) and intended for use with completely unrestricted fare structures, in which the airline estimates total historical demand (at the lowest published fare, denoted "Q") for an itinerary or flight leg by scaling up bookings received in the past by estimated sell-up rates. The airline averages the historical total demands to estimate future total demand, and then uses the sell-up rates to partition that demand to each fare class.

- *Hybrid forecasting*, intended for use with mixed or partially restricted fare structures, in which the forecast is computed separately for *yieldable* demand (where passengers only purchase their preferred fare class) and *priceable* demand (where passengers purchase the least expensive fare offered, regardless of the restrictions) (Boyd and Kallesen, 2004). The yieldable demand is estimated with a standard forecast and the priceable demand with a Q forecast; the total demand is the sum of the two processes.

[6]The airlines can also record purchase data at other levels of aggregation for specialized reporting.

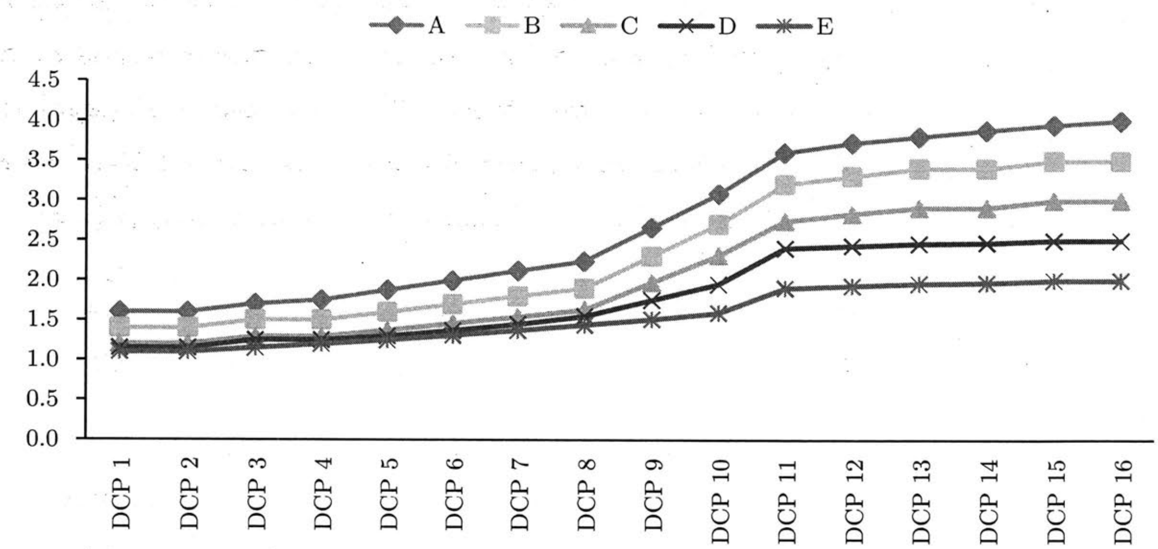

Figure 3.3: Standard FRAT5 curves in PODS

The sell-up estimates used by airlines are represented by FRAT5 (fare ratio at which 50% of passengers sell-up) curves. A FRAT5 of 2.5 implies that 50% of passengers are willing to pay 2.5 times the lowest fare in a market. Several standardized FRAT5 curves are used in PODS, referred to as curve "A" (the most aggressive, with high sell-up rates) to "E" (the least aggressive, with low sell-up rates). The curves are shown in Figure 3.3; note each curve is increasing in DCP: later DCPs have a higher portion of high-budget business travelers, so potential sell-up is higher.

Forecasting models that utilize historical booking data must include a detruncation process, which accounts for the fact that historical observations of bookings are constrained by historical booking policies (i.e. bookings are the lesser of demand and available capacity). PODS includes several detruncation models:

- *Booking curve detruncation*, which estimates booking arrival rates for unconstrained observations and then applies those rates to constrained observations (Wickham, 1995)
- *Expectation-maximization* and *projection detruncation*, which fit a statistical distribu-

tion to the unconstrained observations (Skwarek, 1996; Weatherford and Pölt, 2002)

RM Optimizer

The RM optimizer determines which itinerary/fare classes should be available at any time, given a demand forecast and given remaining capacities on each flight leg in the airline's network. The RM optimizer is run at either the start of each DCP, or at the start of each day during the booking window. Because a new forecast is only generated at the start of each DCP, daily reoptimization requires the (old) forecast to be decremented to account for demand that has already arrived.[7] The output of a revenue management model in PODS is either a booking limit for each itinerary/fare class or flight leg/fare class, or a bid price for each flight leg. A booking limit specifies the (maximum) number of bookings that may be made before the class becomes unavailable; a bid price represents an estimate of the opportunity cost of a seat on each flight. With a bid price output, a booking request for a particular itinerary/fare class is accepted if and only if its fare is greater than or equal to the sum of the bid prices for all flight legs traversed.

Leg-based models apply availability controls at the flight leg level and attempt to maximize revenue on each flight leg independently. PODS supports many models, the most relevant of which are:

- *Expected Marginal Seat Revenue* (EMSR), which protects space in higher-value (later arriving) classes against booking requests made early in the booking process; the output is a set of nested booking limits for each flight leg/fare class (Belobaba, 1987, 1989; Belobaba and Weatherford, 1996). A booking request for a particular itinerary/fare class is accepted if that fare class has a positive booking limit for each flight leg traversed.

- *Leg Dynamic Programming* (LDP), which assumes Poisson demand arrivals and pro-

[7] See Bockelie and Belobaba (2016a) for forecast decrementing methods supported by PODS.

duces a bid price for each time/capacity state for each flight leg (Lee and Hersh, 1993; Lautenbacher and Stidham, 1999).

Network-based models (or origin-destination models) apply availability controls at the itinerary/fare class level and attempt to maximize total network revenue, and require forecasts generated at the itinerary/fare class level. Network-based RM models available in PODS include:

- *Displacement-Adjusted Virtual Nesting* (DAVN), which uses a deterministic network linear program to estimate displacement costs for each flight leg, groups displacement-adjusted OD fares for each leg into several virtual classes, and then optimizes availability of virtual classes on each leg with EMSR (Smith and Penn, 1988; Williamson, 1988). A booking request for a particular itinerary/fare class is accepted if that virtual class into which the request falls has a positive booking limit for each flight leg traversed.

- *Probabilistic Bid Price* (ProBP), in which an iterative convergence algorithm prorates OD fares to legs according to the ratio of the displacement costs of the legs, and then displacement costs are computed as the marginal seat revenue determined by running EMSR on each leg. The final bid prices are the converged displacement costs (Bratu, 1998).

- *Unbucketed Dynamic Programming* (UDP), which combines the displacement cost-adjustment network LP of DAVN with the leg optimization of LDP to produce displacement-aware bid prices for each flight leg (Hopperstad, 2009).

The above models, which are intended for use with restricted fare structures and assume independence of demand between fare classes, can be paired with *fare adjustment* (and hybrid or Q forecasting) to account for demand dependencies. With fare adjustment, fares (as used for optimization and availability control, not the fare actually paid by consumers) are decreased to account for buy-down (Fiig et al., 2010; Walczak et al., 2010). The buy-down adjustment is based on an exponential sell-up model with an input FRAT5 curve and an input scaling parameter. The scaling factor (between 0 and 1) allows simulated airlines to apply less aggressive sell-up estimates to fare adjustment than in hybrid or Q forecasting. With hybrid forecasting, the buy-down adjustment is prorated to only apply to the priceable demand component.

3.2.3 Networks

A network file contains the schedules, flight capacities, and fares for each airline in a simulation and contains demand generation parameters for each market. PODS supports networks that range in size and complexity from single airline, single flight leg "toy" networks (such as A1ONE) to large, realistic networks with multiple airlines with connecting hubs, hundreds of flight legs, and hundreds of markets (such as U10). Specific networks will be described as necessary in subsequent sections and chapters.

3.3 Booking and Revenue Impacts

We can study the revenue and booking impacts of the ancillary choice model, as well as its sensitivity to behavior type and ancillary fee structure, via simulation in PODS. For this work, we use the four airline, many flight leg network U10 calibrated in Appendix A. The four competing airlines each operate a connecting hub, and some airlines offer some point-to-point flights, as shown in Figure 3.4. The airlines each offer ten economy fare classes and three ancillary services: a checked bag (BAG), an advance seat reservation (ASR), and upgraded seating (UPG). Fares vary across markets, but within a particular market all airlines offer the same set of fares. The highest fare, FC 1, is between 1.7 and 6.0 times the lowest fare, FC 10.

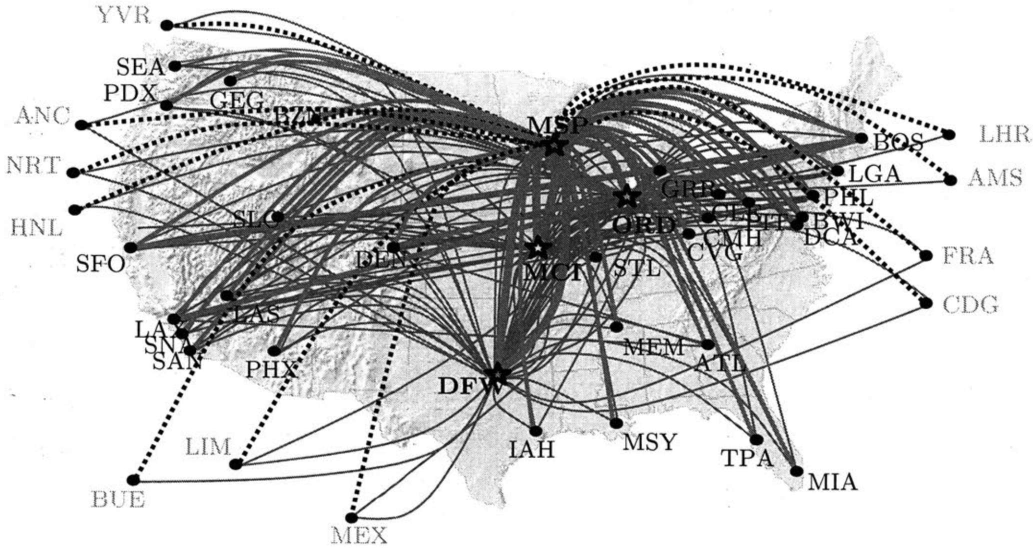

Figure 3.4: Network U10 map. Airline hubs shown with stars. Domestic destinations shown in black, international destinations shown in red. Airline 1's domestic network shown in blue, with international network in dashed black. Airline 2–4 networks shown in grey.

All airlines have a branded fare structure, meaning fare classes are grouped into three "brands" with similar purchase/use restrictions and similar ancillary product offerings. Passengers booking in FC 1 receive all services complimentary; passengers booking in FC 2-6 must pay UPG but receive BAG and ASR complimentary; passengers booking in FC 7-10 must pay for all three services. The full fare and ancillary fee structure is shown in Table 3.2, and the fare class restriction disutilities, ancillary service disutilities, and ancillary service prices are shown in Table 3.3. Note that there are two sets of calibrated disutilities and ancillary prices: one set assuming all passengers are simultaneous, and one set assuming passengers are a random mix of 50% simultaneous and 50% sequential.

We focus on Airline 1's domestic markets. We refer to the fare structure in Table 3.2 as the *partial bundling* case because some ancillary services are bundled for some of the fare brands. We will look at two additional cases: where the industry (that is, all four airlines) moves to a fully bundled offering for all ancillary services (e.g. all ancillary services are complimentary to all passengers) and a to a fully a la carte offering for all ancillary services

Table 3.2: Network U10 fare and ancillary fee structure. Advanced purchase requirement in days.

	Fare range	Average fare	Advanced purchase	Restrictions R1	R2	R3	R4	Ancillary Services BAG	ASR	UPG
Brand 1										
FC 1	\$255–3,037	\$905	None	-	-	-	-	*Free*	*Free*	*Free*
Brand 2										
FC 2	\$153–2,360	\$632	None	Yes	-	-	-	*Free*	*Free*	Paid
FC 3	\$139–2,015	\$563	3	Yes	-	-	-	*Free*	*Free*	Paid
FC 4	\$124–1,670	\$494	7	Yes	-	-	-	*Free*	*Free*	Paid
FC 5	\$115–1,463	\$450	17	Yes	-	-	-	*Free*	*Free*	Paid
FC 6	\$96–1,361	\$406	21	Yes	Yes	-	-	*Free*	*Free*	Paid
Brand 3										
FC 7	\$89–1,241	\$372	7	Yes	-	-	Yes	Paid	Paid	Paid
FC 8	\$79–1,180	\$338	14	Yes	-	-	Yes	Paid	Paid	Paid
FC 9	\$72–1,075	\$309	21	Yes	-	Yes	Yes	Paid	Paid	Paid
FC 10	\$64–979	\$280	28	Yes	Yes	-	Yes	Paid	Paid	Paid

Table 3.3: Network U10 mean disutility ranges and price ranges for each restriction and ancillary service.

	Restrictions R1	R2	R3	R4	Ancillary Services BAG	ASR	UPG
Mean disutilities							
Simultaneous calibration							
Business	\$32–169	\$16–84	\$8–42	\$24–126	\$23	\$10–23	\$28–101
Leisure	\$10–51	\$10–51	\$3–17	\$13–67	\$28	\$12–27	\$28–101
50/50 mix calibration							
Business	\$32–169	\$16–84	\$8–42	\$32–169	\$23	\$10–23	\$28–101
Leisure	\$10–51	\$10–51	\$3–17	\$16–84	\$25	\$10–23	\$25–92
Prices	*n/a*	*n/a*	*n/a*	*n/a*	\$25	\$15–33	\$40–147

Note: price and disutility ranges are shown only for domestic markets; prices are higher in simulated international markets.

(e.g. no ancillary services are complimentary to any passengers). We use cases where all airlines have the same pricing structure to minimize booking shifts between airlines, allowing us to focus on booking shifts between fare classes and brands. We each case for either all simultaneous or a 50/50 mix of simultaneous and sequential passengers.

Expected Results In both the fully bundled and fully a la carte cases, there will be no differentiation between fare classes or brands in terms of ancillary service pricing. Therefore, simultaneous passengers will have less incentive to sell-up to the higher-priced brands 1 and 2. (Recall that fare class restrictions do differ between brands, causing some passengers sell-up). Because sequential passengers do not consider ancillary services during their itinerary/fare class selection, their booking decisions will not change as ancillary fee structures change. Therefore booking changes with the 50/50 mix of simultaneous and sequential passengers should be smaller than with all simultaneous passengers.

The fully a la carte case should increase ancillary revenues, because passengers in all classes must pay for all services. The fully bundled case will have no ancillary revenue, and the net impact on total revenue will depend on how ticket revenues change.

Ticket, ancillary, and total revenues for each of the cases with each behavior type are listed in Table 3.4. Note that the partial bundling case has higher ticket revenue than either the fully bundled or fully a la carte cases for both behavior types. The partial bundling case has the highest ticket revenue because the bundling in brands 1 and 2 creates ancillary differentiation across the brands, which encourages some simultaneous passengers to sell-up to brands 1 and 2.

The fully a la carte case has the highest ancillary revenue because all passengers have to pay for all ancillary services, and the fully bundled case has no ancillary revenue because all passengers receive all services complimentary.

Although the a la carte structure increases ancillary revenue vs the partial bundling case,

Table 3.4: Network U10 Airline 1 simulation results with various ancillary bundling strategies and passenger behavior types. Average fare and average ancillary revenue expressed per passenger segment.

	Fully bundled	Partial Bundling	Fully a la carte
100% Simultaneous Passengers			
Ticket Revenue	$2,572,060	$2,776,606	$2,569,151
Ancillary Revenue	$0	$124,642	$211,280
Total Revenue	$2,572,060	$2,901,248	$2,780,431
Passenger Segments	19,200	19,250	19,205
Average Fare	$133.96	$144.24	$133.77
Average Ancillary Revenue	$0.00	$6.47	$11.00
Change from Partial Bundling			
Ticket Revenue	-7.4%		-7.5%
Ancillary Revenue	-100.0%		+69.5%
Total Revenue	-11.3%		-4.2%
Passenger Segments	-0.3%		-0.2%
Average Fare	-7.1%		-7.3%
Average Ancillary Revenue	-100.0%		+69.9%
50% Simultaneous/50% Sequential Passengers			
Ticket Revenue	$2,632,971	$2,738,155	$2,631,822
Ancillary Revenue	$0	$117,751	$177,830
Total Revenue	$2,632,971	$2,855,906	$2,809,652
Passenger Segments	19,220	19,233	19,221
Average Fare	$136.99	$142.37	$136.92
Average Ancillary Revenue	$0.00	$6.12	$9.25
Change from Partial Bundling			
Ticket Revenue	-3.8%		-3.9%
Ancillary Revenue	-100.0%		+51.0%
Total Revenue	-7.8%		-1.6%
Passenger Segments	-0.1%		-0.1%
Average Fare	-3.8%		-3.8%
Average Ancillary Revenue	-100.0%		+51.1%

Results include domestic markets only.

with these parameters the decrease in ticket revenue is greater than the increase in ancillary revenue, and therefore the a la carte case has a lower total revenue than the partial bundling case. Because the fully bundled case has lower ticket and ancillary revenues than the partial bundling case, it also a lower total revenue.

With the 50/50 mix of behavior types, the directionality of changes is consistent with the all simultaneous behavior results, but the magnitude of changes is reduced (because sequential passengers do not change their booking decisions based on changing ancillary fee structures).

Although the change in the total number of bookings (measured as enplanements) is small between the cases, the distribution of those bookings across fare brands varies widely, as shown in Figure 3.5. The fully bundled and fully a la carte cases have substantially more bookings in the lowest priced brand 3, and substantially fewer bookings in the higher priced brands 1 and 2 compared to the partial bundling case. Because there is no difference in ancillary service pricing across fare brands in either the fully bundled or fully a la carte cases, simultaneous passengers have less incentive to purchase a brand 1 or brand 2 ticket when brand 3 tickets are available. Because sequential passengers do not consider ancillary services when selecting an itinerary and fare class, their choices are not affected by the ancillary fee structure and the magnitude of booking shifts is smaller for the 50/50 mix of simultaneous and sequential passengers.

The potential for ancillary fee structures to drive buy-up amongst simultaneous passengers is clear from Figure 3.5. However, ancillary fee structures can also create buy-down amongst simultaneous passengers, in which the presence of the ancillary fee causes passengers to purchase a lower-priced fare than they otherwise would. Figure 3.6 illustrates this phenomenon, comparing the booking by fare class for the fully a la carte case against the fully bundled case. Although neither case has any differentiation in ancillary service pricing across fare classes or brands, in the fully a la carte case passengers must pay for any ancillary service in any fare class. These patterns of simultaneous passenger buy-up/down indicate that when an airline introduces, unbundles, or re-bundles an ancillary service, ticket revenues

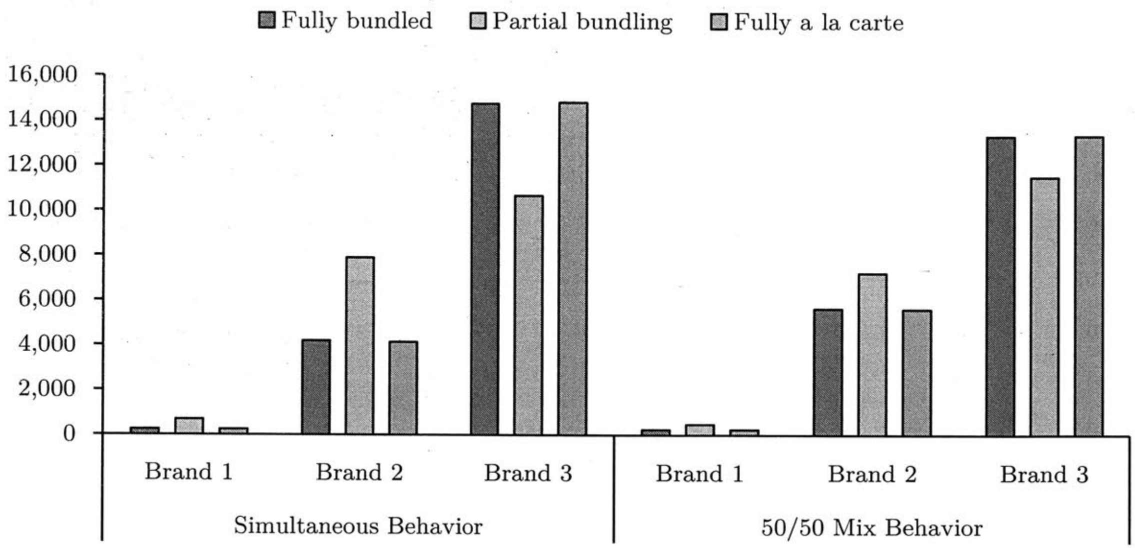

Figure 3.5: Network U10 enplanements by fare brand for fully bundled, partial bundling, and fully a la carte cases (domestic markets only).

may change (in addition to ancillary revenues) without any explicit change to the airline's revenue management processes. This also suggests that the airline may be able to improve its total revenue by accounting for ancillary services in its RM processes.

3.4 Conclusions

This chapter proposes a new Ancillary Choice Model that integrates ancillary service prices and offerings into the consumer decision process about airline itineraries and fare classes. Boundedly rational sequential passengers select an itinerary and fare class, then consider ancillary services, potentially leading to a different booking decision than classically rational simultaneous passengers, who select all three components in one utility-maximization process.

Simulation results in PODS, with the Hopperstad (2005) model as the itinerary/fare class selection kernel, illustrate that simultaneous passengers may alter their itinerary and fare

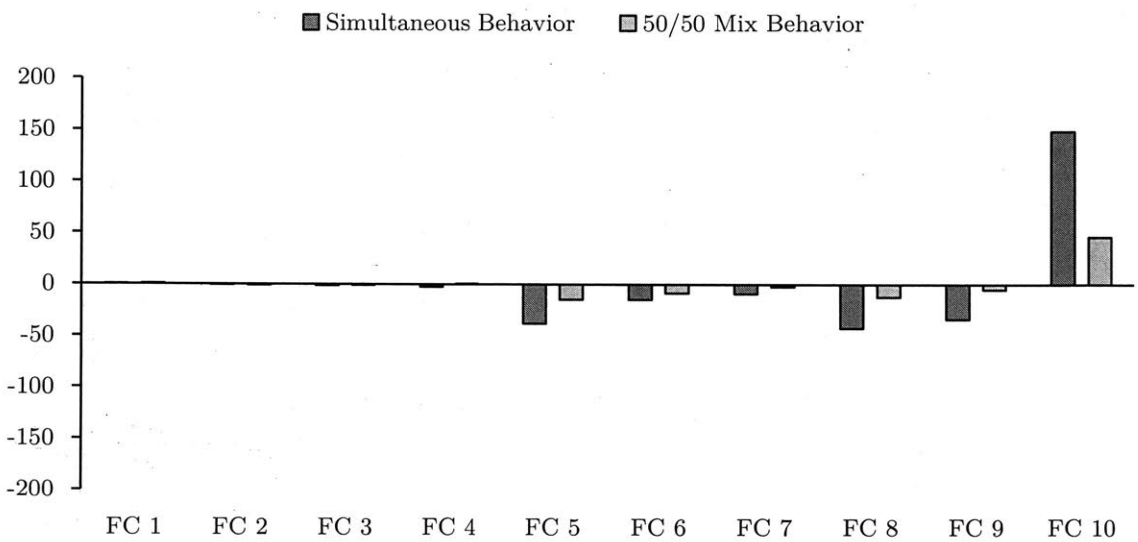

Figure 3.6: Network U10 change in enplanements by fare class for the fully a la carte case vs fully bundled case (domestic markets only).

class selection, either as buy-up or buy-down, under different ancillary fee structures. Use of other kernels could change these results. For example, a more traditional non-infinite scale kernel (containing an error term in the utility equation) would have greater choice stochasticity, which would decrease the magnitude of the impacts of incorporating ancillary services in the choice model. Other fare structures, network structures, or passenger preference distributions could also change these results.

Future work could explicitly reframe the ancillary choice model to include behavior type as a latent class, and then investigate the factors that may influence a person to be, or that indicate a person is, simultaneous or sequential. In addition, variations on the simultaneous or sequential behaviors could be considered, where budgetary constraints are implemented in different manners, or where passengers are simultaneous about some ancillary services but sequential about others.

Overall, this chapter's findings that ancillary structures and passenger behavior type influence booking decisions suggests that ancillary-aware revenue management forecasting and optimization models may need to account for ancillary-specific behavior. Part II of this thesis investigates such concepts.

Part II

Revenue Management with Ancillary Services

Chapter 4

Literature Review: Airline Revenue Management

Growth in the number of ancillary services, and the revenue they provide airlines, not only prompts the need for the new passenger choice modeling described in Part I, but also suggests the potential for new revenue optimization approaches. The economics literature contains many studies and models of optimal pricing strategies for goods, bundles, and add-ons, showing that the ability to charge separately for add-ons can result in changes in optimal prices for goods.[1] In addition, empirical studies have shown that the introduction of airline ancillary fees has reduced average ticket prices.[2] The airline revenue management problem differs significantly from the pricing models in the economics literature because prices are discrete (through filed fares), capacity is constrained, and seats are perishable assets (meaning that once a flight departs, unsold seats have no value). Little work has assessed if or how, under these conditions, airline revenue management models should be modified to attempt to optimize total revenue, not just ticket revenue.

[1] See for instance the bundling/partitioned pricing work of Oi (1971); Adams and Yellen (1976) or the add-on work of Lal and Matutes (1994); Ellison (2005); Gabaix and Laibson (2006); Fruchter et al. (2011); Shulman and Geng (2013); Geng and Shulman (2015).
[2] See Section 1.1.2, Ancarani et al. (2009), Brueckner et al. (2015), and Scotti and Dresner (2015).

In this Part of the thesis, we study the impacts of ancillary services on airline revenue management models. This chapter reviews the literature around revenue management (RM) forecasting and optimization models, with a particular focus on choice-based RM and the limited work on total (or ancillary-aware) RM. Chapter 5 expands upon a previous total revenue optimization heuristic, the optimizer increment, and develops a theoretical foundation for the heuristic, but also shows through simulation that the performance of the optimizer increment degrades when demand forecasts are uncertain and when passengers make choices. Chapter 6 provides the main contribution of this Part by developing the Ancillary Choice Dynamic Program, an ancillary-aware and choice-aware RM approach for total revenue optimization. The chapter includes two derivative heuristics, the Ancillary Marginal Demand and Ancillary Marginal Revenue transformations, that allow traditional (non-ancillary-aware and non-choice-aware) RM methods to be extended to support ancillary services and passenger choice. Challenges in operationalizing the heuristics are addressed, as are potential applications to larger networks. Chapter 7 provides simulation results showing the benefit of our heuristics compared to existing RM methods, including both single airline/single leg network settings and under competition. Finally, Chapter 8 summarizes and concludes the thesis.

4.1 Independent Demand Revenue Management

Revenue management models were developed when airlines began to introduce discounted "leisure" or "tour" fares; the central problem in revenue management is to determine the revenue-maximizing number of early booking, discounted leisure tickets to sell, while protecting seats for the later arriving full fare passengers. Initial revenue management models assumed no ancillary revenue and independent demand streams for each fare class, meaning each passenger only considers one type of fare and will not buy-up or buy-down or choose another itinerary if the desired fare is unavailable.

Leg-based RM models apply availability controls at the flight leg level and attempt to maximize (ticket) revenue on each flight leg independently. The first contribution came from Littlewood (1972), who proposed an optimal leg-based booking strategy for an airline with two nested fare classes where all low fare demand arrived before all high fare demand. Extending Littlewood's model to more than two fare classes requires either computationally complex demand function convolution or heuristic approximations. Optimal solutions, known as Optimal Booking Limits (OBL), were proposed by Wollmer (1992) and Brumelle and McGill (1993). Belobaba proposed heuristic solutions in the form of the Expected Marginal Seat Revenue models (EMSRa and EMSRb) (Belobaba, 1987, 1989; Belobaba and Weatherford, 1996). A different approach, with a Poisson demand arrival process, was taken by Lee and Hersh (1993) and Lautenbacher and Stidham (1999), who developed optimal dynamic programming (DP) solutions. Lautenbacher and Stidham showed that their approach could simplify to the static OBL solutions.

The next major development was a move toward optimizing network (not just flight leg) revenue through the use of origin-destination (OD) availability controls. Development of OD optimization began in 1983 at American Airlines with a system to group OD itineraries into virtual fare classes and then apply existing leg-based controls to the virtual classes (Smith and Penn, 1988). American revised the original algorithm to incorporate a displacement-cost adjustment into the virtual classes, known as Displacement-Adjusted Virtual Nesting (DAVN) (Smith and Penn, 1988; Williamson, 1988).

A different approach to OD control is the use of *bid prices*, which measure the opportunity cost (or reduction in expected future revenue) of selling a seat on a particular flight leg. In contrast to the complexity of the booking limit controls described above, additive bid price controls are simple: a booking request is accepted if the revenue value of the request is greater than the sum of the bid prices across all the legs traversed. Bratu (1998) developed one approach, known as Probabilistic Bid Price (ProBP), which uses an iterative convergence algorithm to prorate OD fares to legs according to the ratio of the displacement costs of the legs. The final bid prices are the converged displacement costs. Talluri and van Ryzin (1998) created a more general dynamic bid price model and identified some instances in which bid price controls might not be optimal.

Bertsimas and de Boer (2005) proposed a third approach to network revenue management: a virtual nesting solution computed through simulation. Their algorithm combines a network linear program with a simulation-based dynamic program to provide an RM model that accounts for stochastic demand, detailed network effects, and remains tractable for large networks. van Ryzin and Vulcano (2008b) extended the model with continuous approximations of demand and capacity to decrease computational complexity and to further increase tractability for large networks.

RM optimization models require a forecast of future demand; forecasting models predict future demand based on historical booking observations, either at the itinerary/fare class level (origin-destination (OD) based) or at the flight leg/fare class (leg-based) level. The simplest models, known as *standard forecasting*, estimate future demand as an average of bookings received for the same OD or leg and fare class on previous departure days, assuming that demands for each class and itinerary are independent (Littlewood, 1972). L'Heureux (1986) extended the Littlewood formulation to include data from flights that have not yet departed. Lee (1990) provided a comprehensive Poisson model of airline bookings and cancellations and illustrated several different maximum likelihood approaches for estimating the underlying Poisson rate parameters.

Forecasting models that utilize historical booking data must include a detruncation process, which accounts for the fact that historical observations of bookings are constrained by historical booking policies (i.e. bookings are the lesser of demand and available capacity). Lee (1990)'s estimation methods directly account for demand truncation. For pick-up based forecasts, multiple detruncation models have been proposed. Booking curve detruncation estimates booking arrival rates for unconstrained observations, and then applies those rates to constrained observations; see Wickham (1995) for additional details and numerical performance simulations. More sophisticated approaches include fitting a statistical distribution to the unconstrained demand estimate through expectation-maximization and projection detruncation (Skwarek, 1996; Weatherford and Pölt, 2002). Queenan et al. (2007) provided a comparison of multiple detruncation methods and developed a double exponential smoothing model. For additional review of detruncation and demand models, see Azadeh et al. (2014).

4.2 Choice-aware Revenue Management

Traditional airline fare structures contained sufficient restrictions to ensure each fare class was, more or less, purchased by a single segment of travelers. In the early 2000s, however, low cost carriers grew and implemented less restricted fare structures. The lack of restrictions permitted more travelers belonging to different demand segments to purchase the lowest offered fare. Because RM forecasts rely on historical bookings to estimate future demand, a *spiraling down* of revenue occurred as systems built on the assumption of independent demand failed to account for passenger behavior (Cooper et al., 2006). While simulations indicate that even in semi-unrestricted environments the independent demand RM models discussed above can deliver revenue gains over a first-come first-serve system (Belobaba et al., 2009), choice-aware models can deliver substantially better performance (Cléaz-Savoyen, 2005).

Work to relax the independent demand assumption has long been underway. Belobaba (1987) offered a variant of EMSRa that incorporates "passenger shift" from a lower to higher value fare class. Belobaba and Weatherford (1996) presented an extension to the EMSRb heuristic to incorporate passenger sell-up, and show simulated revenue gains over the original EMSRb.

A key development is the single-leg single-airline *choice-based dynamic program* of Talluri and van Ryzin (2004), whose dynamic programing formulation explicitly assumes a general model of passenger choice. They show that the optimal solution relies on a series of efficient sets, which are nested by fare order under certain demand models. They note, however, that their exact formulation would be impractical for networks and that some approximation method would be required. Gallego et al. (2004) suggested a Choice-based Deterministic Linear Program (CDLP) as an approximation of the true stochastic program, and approximate demand and capacity as continuous and deterministic. They show that the resulting LP can be solved with column generation. Liu and Van Ryzin (2008) further analyzed CDLP, showing that the solution LP solution relies on a network extension of the concept of Talluri and van Ryzin (2004)'s efficient sets and develop a choice-based equivalent of DAVN. van Ryzin and Vulcano (2008a) extended the simulation-based model of van Ryzin and Vulcano (2008b) to account for customer choice with a general choice model. Separately, Gallego et al. (2009) developed an extension to EMSR that incorporates customer choice via a multinomial logit model (MNL) for fare class selection. They presented solutions when demand is mixed and when demand is ordered by "first choice" preference.

Forecasting with passenger choice is more complex than with independent demand, because dependencies between classes must be considered. Hopperstad and Belobaba (2004) developed the *Q-forecasting* process for completely unrestricted fare structures, in which an airline forecasts total demand for a flight/market at the lowest fare class (denoted "Q"), and then uses estimated sell-up rates to partition that demand into higher-value fare classes. Boyd and Kallesen (2004) proposed separately forecasting *yieldable* demand (where passengers only purchase their preferred fare class) and *priceable* demand (where passengers purchase the least expensive fare offered, regardless of the restrictions), a system known as *hybrid forecasting* (HF). The yieldable demand segment is forecasted with traditional (independent demand) standard forecasting methods; the priceable segment is forecasted with a Q-forecast. Fiig et al. (2010) and Walczak et al. (2010) developed marginal revenue transformations and marginal demand transformations for the inputs to the RM system based on a model of customer choice. The transformations feed the optimizer with the expected incremental revenue or demand to be captured by opening an additional fare class, taking into account the potential for buy-down. These transformations allow independent-demand revenue management optimization models to be extended to account for passenger choice, enabling such models to operate in an unrestricted or semi-restricted fare environment. The marginal revenue transformation is often operationalized as *fare adjustment* (FA) and paired with hybrid forecasting as HF/FA. Together, these extensions and models provide a mechanism to directly incorporate choice behavior into RM systems.

Choice-aware RM and forecasting models require knowledge or estimation of passenger behavior. Guo (2008) simulated the performance of several different non-parametric methods of estimating the sell-up rates utilized by Q-forecasting, hybrid forecasting, and fare adjustment. Talluri and van Ryzin (2004) and van Ryzin and Vulcano (2017) proposed expectation-maximization methods for estimating passenger behavior for choice-based RM, and see Section 2.3 for additional approaches to estimating passenger choice models.

4.3 Ancillary-aware Revenue Management

Previous work on incorporating ancillary revenue streams into revenue management systems is limited, though the need to account for ancillary revenue streams in RM systems has long been acknowledged (Phillips, 2005). Metters et al. (2008) reports that Harrah's, a hotel casino chain, operates its hotel revenue management system on expected total nightly contributions—the amount that an individual spends on a hotel room plus their expected gambling losses. No mathematical or numerical details are provided. The only detailed theoretical work found in the literature is Zhuang and Li (2012), who develop a dynamic programming formulation for hotel casinos to dynamically price hotel rooms while accounting for expected gambling revenue. The model, however, has continuous instead of discrete price points and does not incorporate customer choice between different products.

The most relevant work is that of Hao (2014), who simulates the revenue performance of two heuristics in a two-airline competitive environment. In Hao's simulations, passengers are generated within each market and select an itinerary and fare class according to the Hopperstad (2005) model. Passengers are then assigned an ancillary "spend" based on their market and selected fare class. Airlines in the simulations always collect the full ancillary spend of every passenger, and passenger choice is not affected by the ancillary offer.

Hao's two heuristics are an optimizer increment (OI, which he terms "RM Input Fare Adjustment") and an availability increment (AI, which he terms "Availability Fare Adjustment"). Airlines employing OI augment each fare sent to the optimizer by the expected ancillary revenue per passenger. In Hao's simulations, the airlines have perfect knowledge about expected ancillary revenue. Airlines employing AI optimize based on actual fares, but incorporate ancillary revenue estimates when interpreting the optimizer outputs: the RM system calculates bid prices or virtual class boundaries based on actual fares and the distribution system then makes available any class whose fare plus expected ancillary revenue exceeds the bid price or falls within an open virtual class.[3]

Hao's simulations show revenue losses for OI when coupled with standard forecasting, but show revenue gains when coupled with hybrid forecasting and fare adjustment. AI showed a mix of gains and losses with standard forecasting (with better performance at lower demand levels) and gains with HF/FA. Hao found that both methods lead to more low-fare seat availability, which increased bookings in lower value fare classes and decreased bookings in higher value fare classes.

4.4 Key Literature Gaps

Total revenue optimization has remained an elusive goal, despite interest among airline RM practitioners. The literature shows the need to integrate information about customer choices, in particular willingness-to-pay, into revenue management, however, no known models incorporate both customer choice information and ancillary service information. We ad-

[3]Because of the reliance on bid prices or virtual classes, AI is not applicable to all RM optimization models.

dress this knowledge gap in two ways: first, with a detailed assessment of Hao's Optimizer Increment in Chapter 5, showing why the mechanism leads to revenue dilution. And second, with the development of an ancillary-aware and choice-aware dynamic program and associated heuristics in Chapter 6 that maximizes total revenue. Simulations in Chapter 7 show that our heuristics can increase revenue over existing methods in a variety of scenarios.

Chapter 5

Naïve Total Revenue Optimization: The Optimizer Increment

The simplest approach to total revenue management would be to optimize not on ticket revenues (fares), but on an estimate of total revenues, such as fares plus expected ancillary revenue—an approach we term the optimizer increment. Performance of the optimizer increment has been previously simulated by Hao (2014), and other reports suggest that the approach has been used by hotel casinos to incorporate gambling revenue into hotel room pricing systems (Metters et al., 2008).

Hao studied the optimizer increment with the Passenger Origin-Destination Simulator (PODS, described in Section 3.2). He utilized a two airline network (D6) with six fare classes. In a significant departure from the work of this thesis, in Hao's simulations *the passenger choice process does not include consideration of ancillary services.* Passengers select an itinerary and fare class according to the Hopperstad (2005) model, which incorporates price, schedule, number of connections, and fare class restrictions. *After* having selected an itinerary and fare class, passengers are *assigned* an expected "ancillary spend" based on the selected fare class (typically \$30 if booking in FC 3–6 and \$0 if booking in

FC 1 or FC 2); the airlines in the simulation always collect the expected ancillary "spend" of each passenger. Hao however does not provide a thorough theoretical assessment of the approach.

In this chapter, we contribute to the literature by developing a theoretical basis for the optimizer increment, showing that it is part of an optimal control strategy in a limited number of environments. However, we also provide simulation results illustrating that the approach leads to a decrease in total revenue in more realistic settings as it increases buy-down and leads to displacement of late-arriving, high-fare customers. These findings serve as motivation for the development of the ancillary choice dynamic program (ACDP) in Chapter 6, which accounts for both ancillary revenue and passenger choices.

The remainder of this chapter is organized as follows. Section 5.1 provides the first contribution of this chapter, with a proof of optimality for optimizer increment in a very simple two fare class environment with independent and ordered demands. The proof is then extended to cases with more than two fare classes and mixed passenger arrivals by reference to Section 6.1.3. Section 5.2 then describes simulation results in the very simple two class environment, a more general six class environment, and a larger network with multiple airlines and hundreds of flights, showing that the revenue gains of optimizer increment decrease in the more realistic conditions where airlines must forecast demand and where passengers make choices about fare classes and ancillary services. Section 5.3 provides concluding remarks and explains how this chapter motivates the central contribution of this thesis, the Ancillary Choice Dynamic Program.

5.1 Optimality of the Optimizer Increment

In this section we demonstrate that the optimizer increment is optimal under limited conditions. Consider first the simplest possible revenue management setting, known as the Littlewood Conditions: a single airline/single leg network where the airline offers two fare classes, FC 1 and FC 2, with the fare for FC 1 f_1 greater than that of FC 2. Demand for each class is stochastic but independent and drawn from a known distribution. All FC 2 demand arrives before any FC 1 demand. Under these conditions, the airline must determine at the start of the booking process the number of seats π to protect for FC 1 passengers. Littlewood (1972) and Richter (1982) show that the expected ticket-revenue maximizing solution is to choose π^* such that:

$$f_1 \Pr(X_1 \geq \pi^*) = f_2 \tag{5.1}$$

where X_1 is the stochastic demand for FC 1. There is also a stochastic demand X_2 for FC 2, but its distribution has no effect on the optimal protection level for FC 1. We now consider a slight modification: the airlines also sells an ancillary service to each passenger, with a price a_k in class k, such that total revenues remain nested $f_1 + a_1 > f_2 + a_2$.

Theorem 1. *The expected total revenue maximizing protection level π^* is chosen with the optimizer increment such that:*

$$(f_1 + a_1) \Pr(X_1 \geq \pi^*) = (f_2 + a_2) \tag{5.2}$$

Note that this is equivalent to saying that the airline optimizes its protection levels based on an optimizer increment adjusted fare.

Proof. We begin be defining the expected total revenue earned by the airline R as a sum of the expected total revenue R_k from each class. In addition, we make the assumption that demands and protection levels are continuous. Because $-R$ is convex at π^*, R is maximized with respect to π when:[1]

[1]See Appendix B.1 for the proof that $-R$ is convex at π^*.

$$\frac{\partial R}{\partial \pi} = \frac{\partial R_1}{\partial \pi} + \frac{\partial R_2}{\partial \pi} = 0 \qquad \frac{\partial R_1}{\partial \pi} = -\frac{\partial R_2}{\partial \pi} \tag{5.3}$$

We define the terms $\partial R_k / \partial \pi$ as the expected marginal total revenue earned from FC k due to protecting π seats for FC 1. Increasing π *reduces* FC 2 revenue if FC 2 demand exceeds the remaining capacity of the flight; therefore the marginal revenue for FC 2 is:

$$\frac{\partial R_2}{\partial \pi} = -(f_2 + a_2) \Pr(X_2 > c - \pi) \tag{5.4}$$

where c is the capacity of the flight. Increasing π increases FC 1 revenue if FC 1 demand exceeds π (otherwise the extra space is wasted) *and* if FC 2 demand exceeds its allocated space (otherwise the unused FC 2 space would be released to FC 1 passengers anyway); therefore the marginal revenue for FC 1 is:

$$\frac{\partial R_1}{\partial \pi} = (f_1 + a_1) \Pr(X_1 \geq \pi) \Pr(X_2 > c - \pi) \tag{5.5}$$

Combining the previous three equations yields:

$$(f_1 + a_1) \Pr(X_1 \geq \pi) \Pr(X_2 > c - \pi) = (f_2 + a_2) \Pr(X_2 > c - \pi) \tag{5.6}$$

$$(f_1 + a_1) \Pr(X_1 \geq \pi) = (f_2 + a_2) \tag{5.7}$$

□

If we introduce the assumption that customer arrivals occur due to a Poisson process, we can extend the proof to cover more than two fare classes and we can drop the requirement for ordered demands. In addition, we can allow the purchase of the ancillary service to be probabilistic, with different probabilities for each fare class (although independent fare class demands remain). In this case, incrementing by the expected ancillary revenue per passenger (by fare class) remains an optimal solution because it is equivalent to the Ancillary Marginal Revenue (AMR) transformation developed in this thesis, with independent demand. AMR is introduced in Chapter 6; see Section 6.1.3 for the proof.

5.2 Simulated Performance of Optimizer Increment

We integrated the optimizer increment in the Passenger Origin-Destination Simulator to test its performance and to understand how violations of the assumptions in the optimization model can result in revenue losses. In the PODS implementation, airlines continually estimate expected ancillary revenue per passenger based on the average historical ancillary service purchases. These estimates are computed for each market/fare class combination separately, and are averaged across all DCPs.

Consider that the airline has a historical dataset of bookings and ancillary purchases and uses observations from the n_{ob} most recent departure days (set to 26[2]) in the estimation process for $\mu_{mkt,k,s}$, the true mean ancillary revenue per passenger in market mkt and fare class k from ancillary service s. If the airline has ancillary services $1, \ldots, s, \ldots, S$ and has previously sold $x_{mkt,k,s,dep}$ units of ancillary service s to passengers in fare class k within market mkt on previous departure day dep, and received a total of $b_{mkt,k,dep}$ bookings for in class k and market mkt on previous departure day dep when the ancillary service s is priced at $r_{mkt,k,s}$ when purchased in market mkt, fare class k, then the estimated average ancillary revenue per passenger is:

[2]If each sample represents a typical Friday, for example, 26 observations corresponds to 26 weeks, or 6 months, of data.

$$\hat{a}_{mkt,k} = \begin{cases} \frac{\sum_{dep=1}^{n_{ob}} \sum_{s \in S} x_{mkt,k,s,dep} r_{mkt,k,s}}{\sum_{dep=1}^{n_{ob}} b_{mkt,k,dep}} & \sum b_{mkt,k,dep} \neq 0 \\ \tilde{a}_{mkt,k} & \text{otherwise} \end{cases}$$

Note that in the case when no bookings have been received for a given market/class during the period used for forecasting, the estimator will be undefined, and some default value $\tilde{a}_{mkt,k}$ must be used. The expected value of the estimator, conditional on some booking history, is:

$$\mathrm{E}\left[\hat{a}_{mkt,k} \mid b\right] = \frac{\sum_{dep=1}^{n_{ob}} \sum_{s \in S} \mathrm{E}\left[x_{mkt,k,s,dep} r_{mkt,k,s}\right]}{\sum_{dep=1}^{n_{ob}} b_{mkt,k,dep}} = \begin{cases} \sum_{s \in S} \mu_{mkt,k,s} & \sum b_{mkt,k,dep} \neq 0 \\ \tilde{a}_{mkt,k} & \text{otherwise} \end{cases}$$

Note that the estimator is conditionally unbiased, but has a bias of $\sum_s \mu_{mkt,k,s} - \tilde{a}_{mkt,k}$ when a market/class has no bookings (e.g. missing observations) during the historical data period. In our implementation, we set $\tilde{a}_{mkt,k} = 0$. This approach produces a biased estimator when a fare class receives no bookings in a market during the historical data period; section 5.2.4 evaluates the magnitude of this bias and the effect of an alternative estimation method.

5.2.1 Result Analysis Methodology

We measure the performance of the optimizer increment relative to a *baseline* simulation, where all airlines use standard (independent demand) forecasting and optimization models that do not account for ancillary services or revenues.

Recall form Chapter 3 that each PODS *simulation* consists of 2 to 5 independent *trials* of 600 *samples*. Each sample represents one realization of the "same" departure day (meaning the demand and passenger characteristics for all samples are drawn from the same distribution).

The first 200 samples of each trial are used to warm up the forecasting models and are "burned" and **never** included in any reported results, leaving 400 samples per trial. The seeds for all random draws (i.e. the number of passengers generated and individual passenger preferences) are the same for all simulations with the same demand parameters, so the difference in revenue or bookings between two simulations with different forecasting or optimization methods is due to the change in forecasting or optimization, not due to different demand generation.

The primary item of interest in our studies is the change in total revenue due to a change in forecasting or optimization method, which we measure as the average across all (unburned) samples of the sample-specific change in total revenue for a test simulation vs a baseline simulation. Mathematically, if X_i^j is the revenue (or other simulation output) for sample i for simulation $j \in \{\text{TEST}, \text{BASE}\}$, we are interested in the term $\bar{\Delta}$:

$$\bar{\Delta} = \frac{1}{n}\sum_{i=1}^{n} \hat{\Delta}_i \qquad\qquad \hat{\Delta}_i = X_i^{\text{TEST}} - X_i^{\text{BASE}}$$

Unless otherwise stated, all references to statistical significance and confidence intervals in this chapter are derived from a paired t-test (or a one sample t-test on the change in revenue by sample), with the null hypothesis that there is no true change in revenue ($H_0 : \Delta = 0$) vs alternative hypothesis that there is a true change in revenue ($H_a : \Delta \neq 0$). The test statistic t is:

$$t = \frac{\bar{\Delta}}{\text{se}_{\bar{\Delta}}} \sim T_{df=n-1}$$

where $\text{se}_{\bar{\Delta}}$ is the standard error of $\bar{\Delta}$.

The demand forecasting process nominally introduces a dependency between demand generated (and therefore total revenue) in one sample and booking limits (and therefore total revenue) in another sample. These dependencies potentially violate the t-test assumption of independence between samples; however, as discussed in Appendix C, our testing indicates that these correlations are minor and therefore should not affect the statistical significance of our results.

5.2.2 Littlewood Conditions (Network C2)

We start by considering a case where the optimizer increment is proven optimal, the Littlewood Conditions described above. The demand and fare parameters for the single airline/single leg network are summarized in Table 5.1: FC 1 has a fare of \$500 and provides no ancillary revenue; FC 2 has a fare of \$300 and provides \$25 to \$100 in ancillary revenue (per booking; all passengers in FC 2 purchase the ancillary service). Demand for each fare class is independent, normally distributed, and we initially assume the distribution parameters are known to the airline. FC 2, which books first, has a mean demand of 40 to 60 and standard deviation of 16 to 24 (always 40% of the mean); FC 1, which books last, has a mean demand of 20 to 80 and a standard deviation of 8 to 32 (always 40% of the mean).[3] As a baseline scenario, the airline calculates the number of seats to protect for FC 1 based on Equation 5.1; as a test scenario, the airline computes the protection limit using the optimizer increment (Equation 5.2).

The simulations are run within an Excel-based simulator (for instances where demand distribution parameters are assumed to be known) and PODS (for instances where we relax the assumption that demand distributions are known). For any specific set of demand parameters, the simulated random draws are exactly the same between all PODS and Excel-based simulations; hence, statistical comparisons between two simulations (with the same demand parameters) can still be performed on a paired basis. Reporting for each simulation covers 2,000 individual departure days, composed of 5 trials of 400 (unburned) samples each.

[3]Other demand and fares combinations were tested and showed similar trends to the results described here.

Table 5.1: Network C2 fare and ancillary fee structure.

	Fare	Ancillary price	Mean demand	k-factor	Demand arrival
FC 1	\$500	\$0	20 to 80	0.4	Last
FC 2	\$300	\$25 to \$100	40 to 60	0.4	First

In this simplest possible scenario, the optimizer increment increases total revenue, as shown in Figure 5.1. The revenue gain of the optimizer increment reaches 1.8%, with higher revenue gains at higher ancillary prices (where the benefit of ancillary-awareness increases) and at higher demand levels (where revenue management systems are generally more effective). Figure 5.2 shows 99% confidence intervals ($df = 1,999$) for the total revenue impact of the increment for \$50 and \$100 ancillary price levels, and Table 5.2 provides a more detailed breakout of results. The optimizer increment consistently decreases bookings in FC 1 while increasing bookings in FC 2. In revenue terms, the increment decreases FC 1 revenue, increases FC 2 ticket revenue, and increases ancillary revenue, with an overall decrease in ticket revenue but increase in total revenue. These effects are shown for a range of ancillary prices in Figure 5.3. The revenue impacts of the increment are statistically significant ($p < 0.01$, $df = 1,999$) at the \$100 ancillary price for a wide range of demand scenarios; at the \$50 ancillary price, the revenue changes are statistically significant ($p < 0.01$, $df = 1,999$) for higher FC 1 demand.

In reality, airlines do not know demand distribution parameters and must forecast demand based on historical bookings. When we introduce demand forecasting in these simulations, performance of the optimizer increment deteriorates dramatically. In these forecasted-demand experiments, the airline uses a standard independent-demand forecasting model (Littlewood, 1972) for *both the baseline and the test cases.* As shown in Figure 5.4, the optimizer increment results in a (small) decrease in revenue for many demand and ancillary price combinations, with total revenue changes ranging from +0.1% to -0.2%. As with the known demand cases, the increment with demand forecasting results in a decrease in FC 1 protection levels, increasing FC 2 bookings and decreasing FC 1 bookings (with a

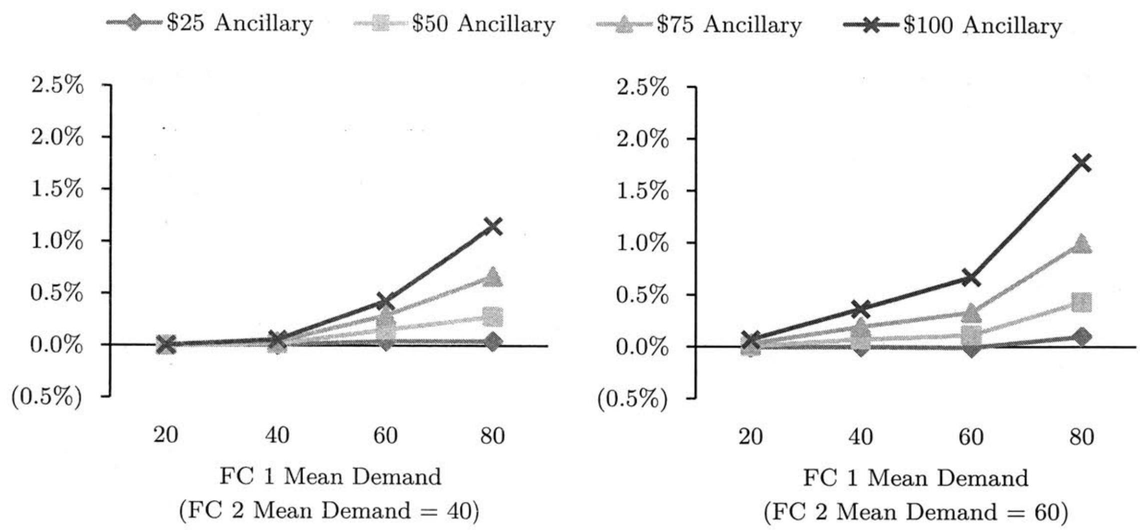

Figure 5.1: Network C2 change in total revenue due to optimizer increment vs baseline with various ancillary prices and market demand levels (known demand).

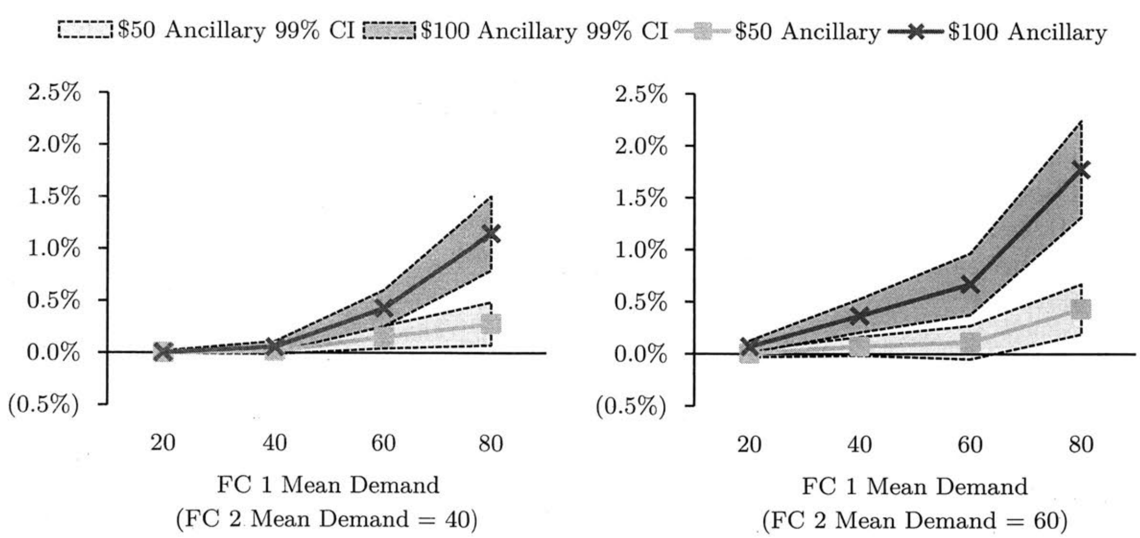

Figure 5.2: Network C2 change in total revenue due to optimizer increment and 99% confidence interval vs baseline with various ancillary prices and market demand levels (known demand). $df = 1,999$.

Table 5.2: Network C2 simulation results with various ancillary prices and market demand levels (known demand, FC 2 $\mu = 60$).

	\$50 Ancillary		\$100 Ancillary	
	FC 1 $\mu = 40$	FC 1 $\mu = 60$	FC 1 $\mu = 40$	FC 1 $\mu = 60$
Baseline				
FC 1 Ticket Revenue	\$17,223	\$24,236	\$17,223	\$24,236
FC 2 Ticket Revenue	\$15,730	\$12,842	\$15,730	\$12,842
Ticket Revenue	\$32,953	\$37,077	\$32,953	\$37,077
Ancillary Revenue	\$2,622	\$2,140	\$5,243	\$4,281
Total Revenue	\$35,574	\$39,217	\$38,196	\$41,358
Load Factor	86.9%	91.3%	86.9%	91.3%
Optimizer Increment				
FC 1 Ticket Revenue	\$16,739	\$22,857	\$16,192	\$21,106
FC 2 Ticket Revenue	\$16,167	\$14,061	\$16,608	\$15,397
Ticket Revenue	\$32,906	\$36,918	\$32,800	\$36,503
Ancillary Revenue	\$2,695	\$2,343	\$5,536	\$5,132
Total Revenue	\$35,601	\$39,261	\$38,336	\$41,636
Load Factor	87.4%	92.6%	87.7%	93.5%
Change from Baseline				
FC 1 Ticket Revenue	-2.8%	-5.7%	-6.0%	-12.9%
FC 2 Ticket Revenue	+2.8%	+9.5%	+5.6%	+19.9%
Ticket Revenue	-0.1%	-0.4%	-0.5%	-1.5%
Ancillary Revenue	+2.8%	+9.5%	+5.6%	+19.9%
Total Revenue	+0.1%	+0.1%	+0.4%	+0.7%
Load Factor	+0.5 pts	+1.3 pts	+0.9 pts	+2.3 pts
Significance of Change in Total Revenue from Baseline				
Standard Error	0.0%	0.1%	0.1%	0.1%
t-statistic	2.12	1.82	5.72	5.86
p-value	0.034	0.069	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

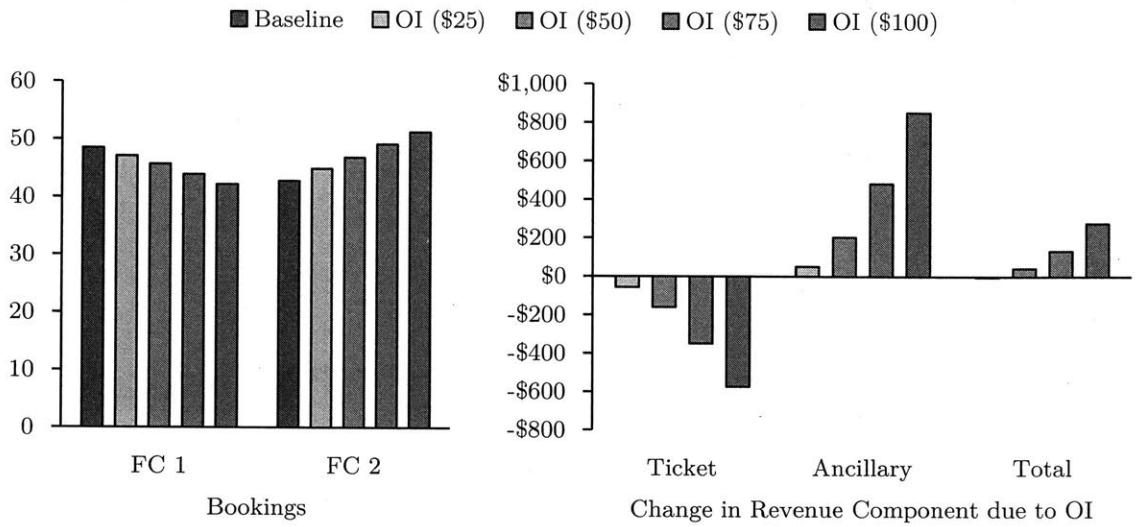

Figure 5.3: Network C2 change in bookings by fare class (left) and change in revenue component (right) due to optimizer increment vs baseline with various ancillary prices (known demand, FC 1 $\mu = 60$, FC 2 $\mu = 60$).

larger effect at higher ancillary prices), as shown in Figure 5.6. In addition, the increment increases ancillary revenue and decreases ticket revenue, however, with demand forecasting, the increase in ancillary revenue does not outweigh the decrease in ticket revenue and net revenue (typically) declines.

Figure 5.5 shows 99% confidence intervals ($df = 1,999$) for the \$50 and \$100 ancillary price scenarios; note that in general these revenue differences are not statistically significant. An exception is the highest FC 1 and FC 2 demand scenarios, in which case the revenue loss due to the optimizer increment is significant for both ancillary prices ($p < 0.01$, $df = 1,999$, for FC 1 demand of 60 or 80 and FC 2 demand of 60).

It is notable, and concerning from an implementation perspective, that the violation of only one assumption, known demand distributions, leads to such a dramatic change in the total revenue effects of the optimizer increment. Note that all other optimality conditions—independent fare class demands, ordered demand, only two fare classes—have been main-

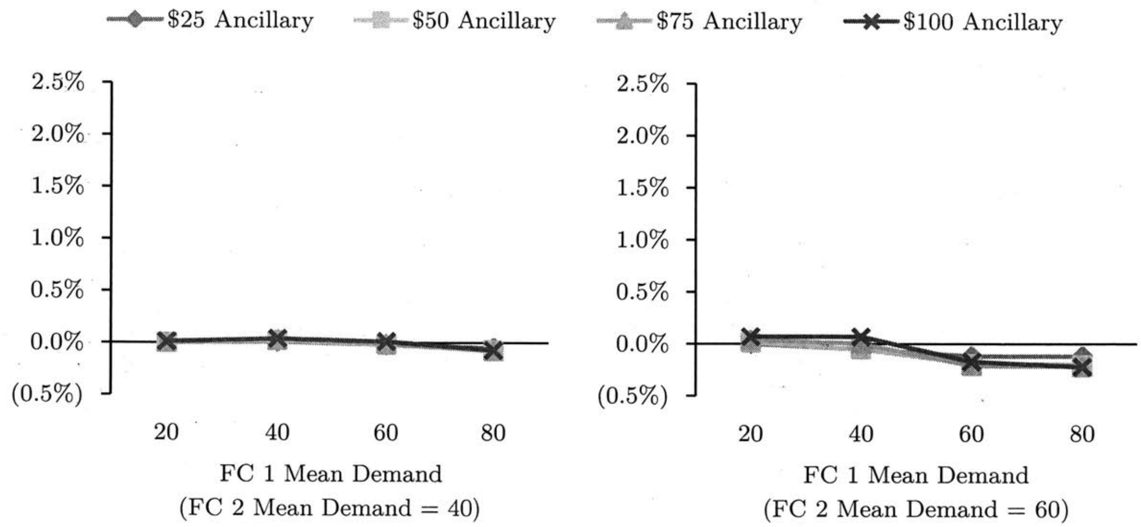

Figure 5.4: Network C2 change in total revenue due to optimizer increment vs baseline with various ancillary prices and market demand levels (forecasted demand).

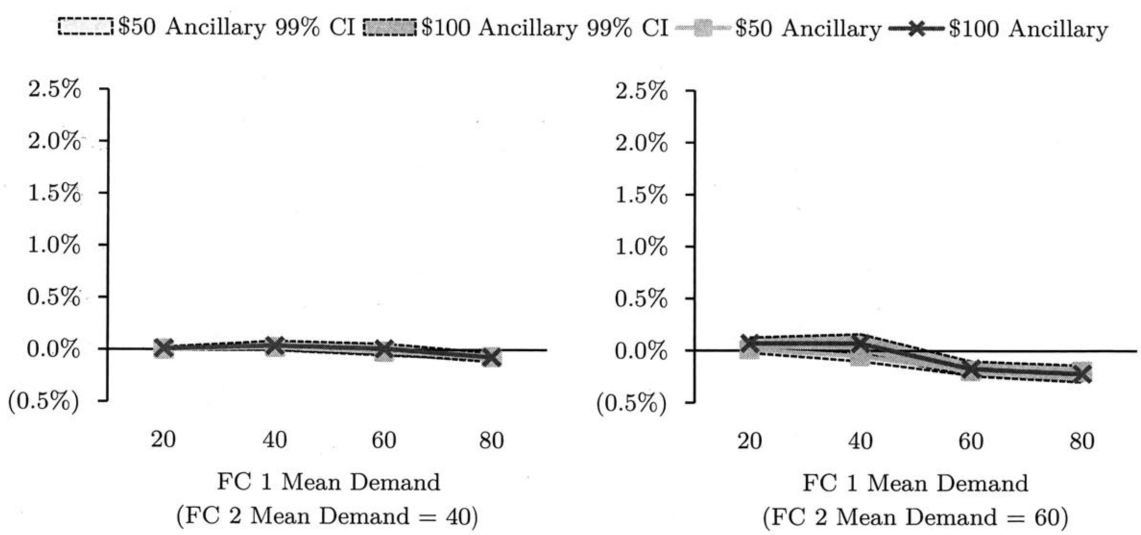

Figure 5.5: Network C2 change in total revenue due to optimizer increment and 99% confidence interval vs baseline with various ancillary prices and market demand levels (forecasted demand). $df = 1,999$.

Table 5.3: Network C2 simulation results with various ancillary prices and market demand levels (forecasted demand, FC 2 $\mu = 60$).

	$50 Ancillary		$100 Ancillary	
	FC 1 $\mu = 40$	FC 1 $\mu = 60$	FC 1 $\mu = 40$	FC 1 $\mu = 60$
Baseline				
FC 1 Ticket Revenue	$16,604	$20,581	$16,604	$20,581
FC 2 Ticket Revenue	$16,266	$15,750	$16,266	$15,750
Ticket Revenue	$32,870	$36,331	$32,870	$36,331
Ancillary Revenue	$2,711	$2,625	$5,422	$5,250
Total Revenue	$35,581	$38,956	$38,292	$41,581
Load Factor	87.4%	93.7%	87.4%	93.7%
Optimizer Increment				
FC 1 Ticket Revenue	$16,307	$20,170	$16,008	$19,723
FC 2 Ticket Revenue	$16,506	$16,038	$16,734	$16,341
Ticket Revenue	$32,813	$36,208	$32,742	$36,064
Ancillary Revenue	$2,751	$2,673	$5,578	$5,447
Total Revenue	$35,564	$38,881	$38,320	$41,511
Load Factor	87.6%	93.8%	87.8%	93.9%
Change from Baseline				
FC 1 Ticket Revenue	-1.8%	-2.0%	-3.6%	-4.2%
FC 2 Ticket Revenue	+1.5%	+1.8%	+2.9%	+3.8%
Ticket Revenue	-0.2%	-0.3%	-0.4%	-0.7%
Ancillary Revenue	+1.5%	+1.8%	+2.9%	+3.8%
Total Revenue	-0.0%	-0.2%	+0.1%	-0.2%
Load Factor	+0.2 pts	+0.1 pts	+0.4 pts	+0.3 pts
Significance of Change in Total Revenue from Baseline				
Standard Error	0.0%	0.0%	0.0%	0.0%
t-statistic	-2.29	-11.58	2.07	-6.19
p-value	0.022	< 0.001	0.038	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

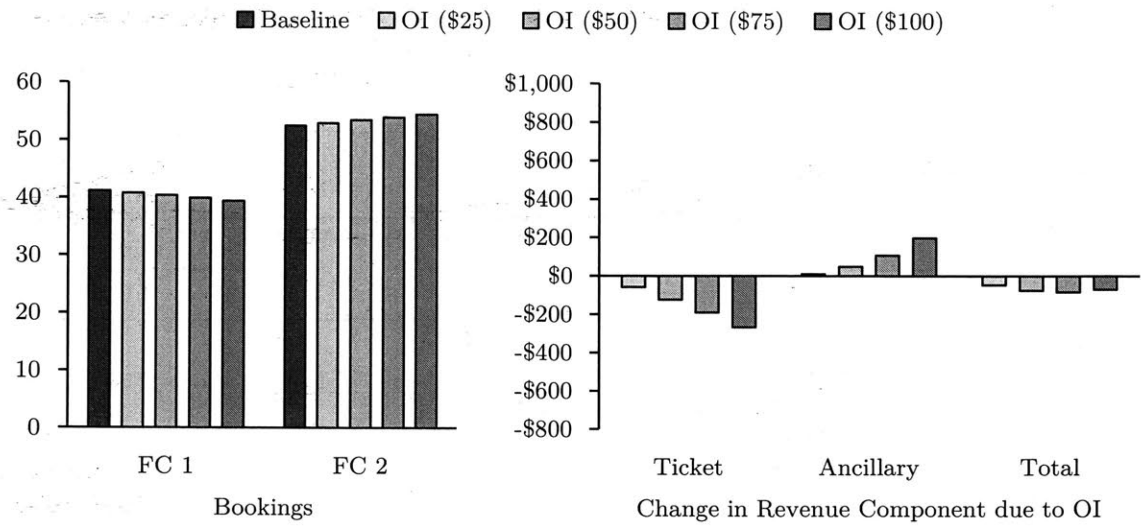

Figure 5.6: Network C2 change in bookings by fare class (left) and change in revenue component (right) due to optimizer increment vs baseline with various ancillary prices (forecasted demand, FC 1 $\mu = 60$, FC 2 $\mu = 60$).

tained throughout these experiments.

Although the incremental benefit of the optimizer increment essentially disappears (or becomes negative) when demand distributions must be forecast, as shown in Figure 5.7, the optimizer increment with forecasting can still produce a higher revenue than the known demand baseline. Recall that with a $100 ancillary price, the optimizer increment with forecasting has a lower total revenue than the baseline with forecasting (and the change is statistically significant), yet both forecasted demand cases, in this scenario, have a higher total revenue (+0.5% for baseline and +0.4% for the optimizer increment) than the known demand baseline. The known demand optimizer increment has the highest revenue. This unintuitive result is a consequence of forecasting errors that unintentionally *increase* total revenue in the forecasted demand baseline case by under-forecasting FC 1 demand and accepting more FC 2 bookings (vs the known demand baseline).

Figure 5.8 shows the baseline FC 1 demand forecast and breakout of baseline revenue by

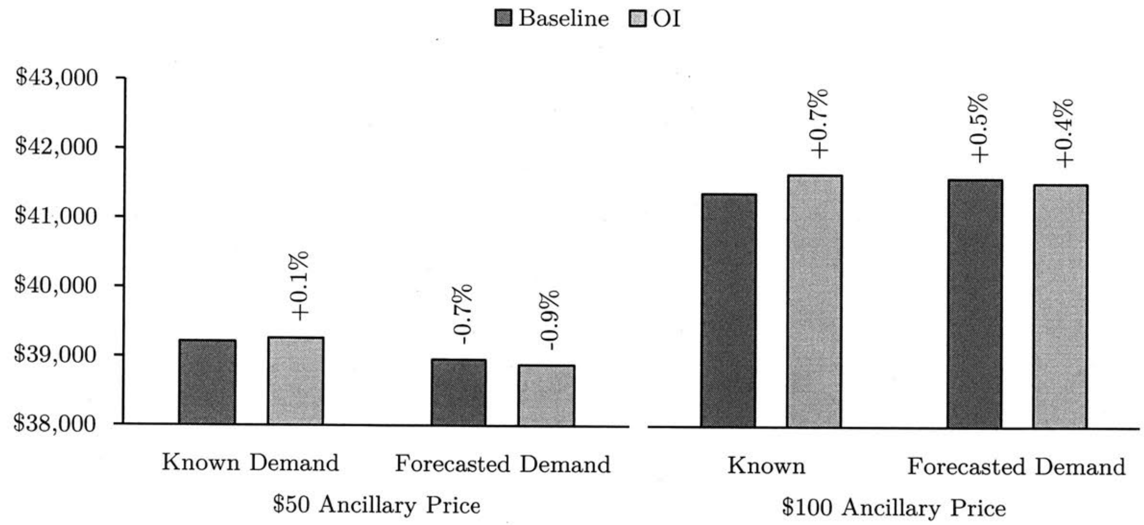

Figure 5.7: Network C2 total revenue with known and forecasted demand and with various ancillary prices (FC 1 $\mu = 60$, FC 2 $\mu = 60$). Percentages show the change in total revenue vs known demand baseline.

source when FC 1 and FC 2 mean demand are 60, the ancillary service has a price of $100, and the demand distributions are either known or forecasted. In the forecasted demand case, the airline generates a new demand forecast at the start of each departure day based on the bookings received on the most recent previous n_{ob} departure days. Each generated forecast has an estimate of the mean and an estimate of the standard deviation of underlying demand. Figure 5.8 shows the average (across all 2,000 departure days) of the FC 1 forecast mean (in the horizontal bar), the typical range (plus/minus one standard deviation) of the forecast mean (as the shaded rectangle), and the average (across all 2,000 departure days) of the forecast standard deviation (as the dashed error bars, drawn from the high/low ends of the typical range of the forecast mean). In the known demand case there is no estimation process and the "forecast" mean is always the true value of 60 and the "forecast" standard deviation is always the true value of 24.

The forecasted demand case has a much smaller FC 1 forecast mean and standard deviation than the known demand case. This leads to a lower protection level for FC 1, which

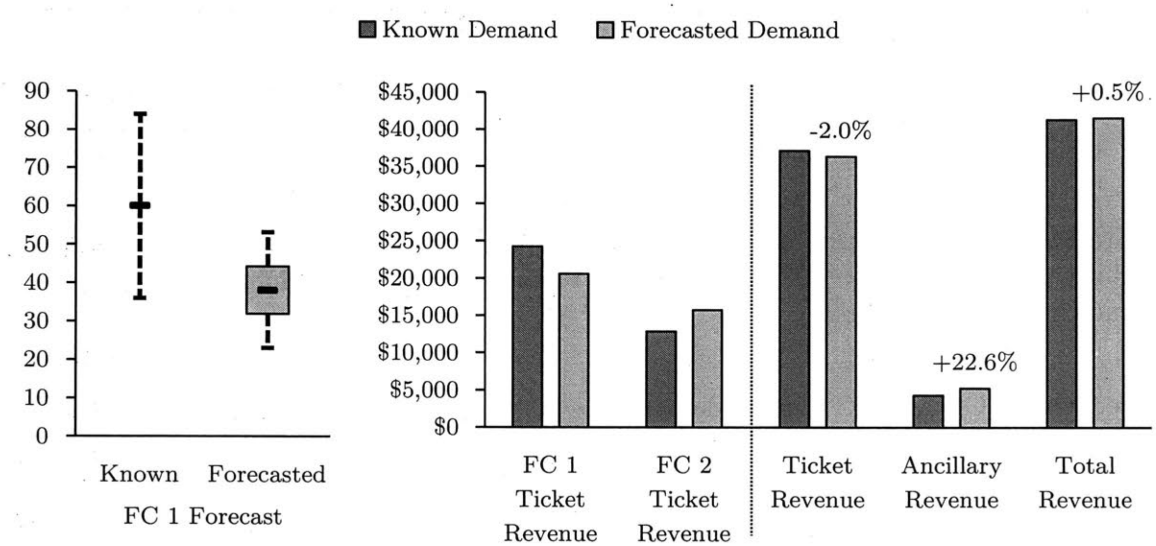

Figure 5.8: Network C2 baseline FC 1 forecast (left) and baseline revenue components (right) (FC 1 μ = 60, FC 2 μ = 60, \$100 ancillary price). The average forecast mean is shown with a horizontal line, plus and minus one standard deviation of the estimate of the forecast mean with a shaded box, and plus and minus the average estimate of the forecast standard deviation (from the high/low end of the forecast mean range) with the dashed line. Percentages show the change in revenue component vs known demand.

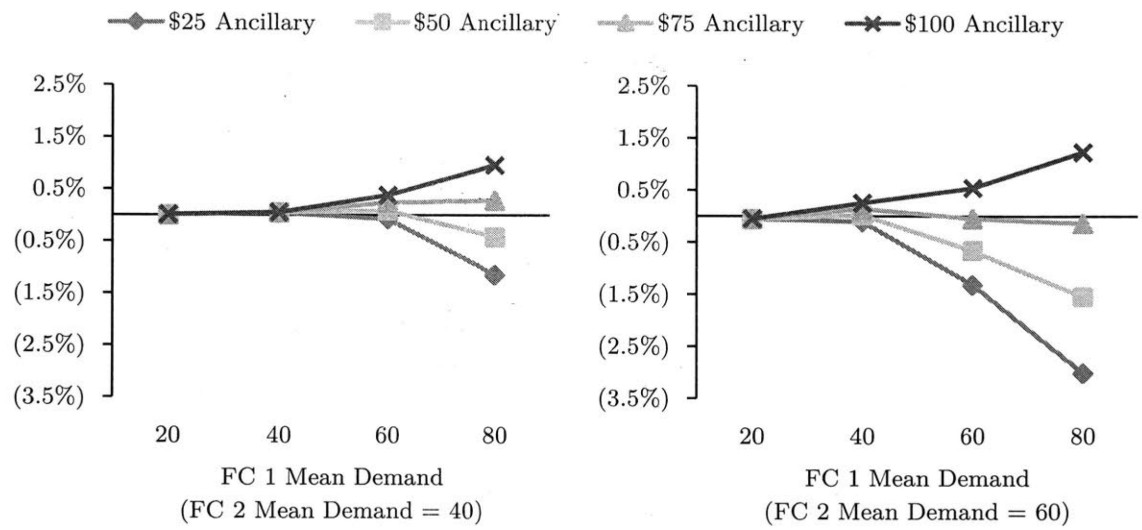

Figure 5.9: Network C2 change in baseline total revenue due to demand forecasting vs known demand with various ancillary prices and market demand levels.

allows more FC 2 bookings but reduces space available for FC 1 bookings. Recall that protection levels in the baseline case are set to maximize *ticket* revenue; the net impact of the introduction of forecasting is a reduction (by 2.0%) of baseline ticket revenue. However, with ancillary revenues present, the baseline protection levels undervalue FC 2 bookings. The forecasting errors that result in additional FC 2 bookings also result in an increase (by 23%) in ancillary revenue; in this scenario, the increase in ancillary revenue outweighs the loss in ticket revenue and results in an increase in total revenue. Note that this inadvertent total revenue increase is dependent on the value of the ancillary service: as shown in Figure 5.9, at lower ancillary prices, the baseline with demand forecasting has lower total revenue than the baseline with known demand (and the magnitude of the effect is greater at higher demand levels).

With the demand forecasting baseline accepting more FC 2 bookings (vs the known demand baseline), incorporating the optimizer increment (slightly further increasing FC 2 bookings and decreasing FC 1 bookings) results in too many FC 2 bookings, and leads to a reduction in total revenue (vs the forecasted demand baseline).

The Littlewood Conditions represent the simplest possible revenue management environment, and each of the necessary assumptions for optimality is violated in reality. The introduction of additional assumption violations (such as multiple re-optimization periods, more than two fare classes, unordered fare class demands, and passenger choice of fare class and/or ancillary services) may combine to increase, reduce, or even reverse, some of the effects outlined in this section. Nonetheless, the Littlewood Conditions provide a useful setting for examining the fundamental dynamics of an independent-demand optimization approach, such as the optimizer increment. Subsequent sections of this chapter assess the performance of OI in increasingly complex environments, with the analysis focused on illustrating how revenue impacts of OI in the more complex environments differ from those under the Littlewood Conditions.

5.2.3 One Airline, One Flight Leg Network (A1ONE)

Network A1ONE consists of one airline with one flight leg and six fare classes. Passengers make choices about fare classes and the ancillary service (according to the simultaneous or sequential models described in Chapter 3, and in contrast to the Littlewood Conditions), and therefore fare class demands are not independent and passenger arrivals are not ordered. Network A1ONE has two passengers segments, business and leisure, as described in Section 3.2; business passengers typically (although not always) have higher budgets, are more restriction-averse, and book closer to departure than leisure travelers. In addition, the airline divides the booking window into 16 data collection points (DCPs); it generates new demand forecasts and re-optimizes its protection levels/booking limits at the start of each DCP. Thus, each of the optimality assumptions of the Littlewood Conditions is violated.

As with network C2, A1ONE simulations have 2,000 total samples (5 trials of 400 unburned samples each) and the underlying stochastic demands for a given sample are the same between different simulations with the same demand generation parameters, so comparisons on performed on a pairwise basis.

Table 5.4: Network A1ONE fare and ancillary fee structure.

	Fare	Advanced purchase	Restriction Applies? R1	R2	R3
FC 1	$500	None	-	-	-
FC 2	$390	3 days	-	-	Yes
FC 3	$295	7 days	-	Yes	Yes
FC 4	$200	10 days	Yes	-	Yes
FC 5	$160	14 days	Yes	Yes	-
FC 6	$125	21 days	Yes	Yes	Yes

The fare and ancillary fee structure is shown in Table 5.4. FC 1 has the highest fare at $500 and has no restrictions; FC 6 has the lowest fare at $125 and has several advance purchase and other restrictions in place. The airline sells one ancillary service, priced at $25 to $100 depending on the scenario. We consider several different mean ancillary disutility scenarios: an "equally appealing" case where both segments have a mean ancillary disutility equal to its price, and a "business-oriented" and "leisure-oriented" case where the business (leisure resp.) segment has a mean ancillary disutility equal to 125% of its price and the leisure (business resp.) segments has a mean disutility equal to 75% of its price. These configurations are summarized in Table 5.5.

Table 5.6 lists booking and ancillary purchase data for a baseline case with 100% simultaneous passengers, medium demand, a $50 ancillary price, and equally appealing disutility scenario. In the baseline scenario, about 33% of passengers purchase the ancillary service, with a much higher purchase rate in the higher fare classes (46% for FC 1) than in the lower

Table 5.5: Ancillary disutility scenarios. Mean ancillary disutility for each passenger segment shown as a function of ancillary price.

Scenario	Business	Leisure
Leisure-oriented	75%	125%
Equally appealing	100%	100%
Business-oriented	125%	75%

Table 5.6: Network A1ONE baseline bookings and ancillary purchase data by fare class (medium demand, 100% simultaneous passengers, \$50 ancillary price, \$50 ancillary disutility).

	FC 1	FC 2	FC 3	FC 4	FC 5	FC 6
Bookings	6	14	10	8	16	56
Booking Mix	5%	12%	9%	8%	15%	52%
Average Fare	\$500	\$390	\$295	\$200	\$160	\$125
Average Ancillary per Passenger	\$23	\$22	\$21	\$22	\$19	\$12
Portion of Total Revenue from Ancillary	4%	5%	7%	10%	11%	9%
Ancillary Sales Rate	46%	44%	41%	44%	39%	25%

fare classes (25% in FC 6); likewise, the average ancillary revenue per booking is highest in FC 1 (\$23) and lowest in FC 6 (\$12). Despite the lower ancillary purchase rate and lower average ancillary revenue in lower fare classes, because fares are lower in FC 6 than FC 1, the portion of total revenue derived from ancillary sales is highest in the lower fare classes (11% in FC 5) and lowest in higher fare classes (4% in FC 1). Recall that major US airlines report about 8% of total revenue from ancillary services, according to the US Department of Transportation.[4]

The lower ancillary purchase rate in the lower fare classes is driven by a fundamental behavioral assumption in the Simultaneous choice model: passengers have an overall budgetary constraint that limits their spending on the combination of fare and ancillary services. Passengers booking in the lower fare (and highly restricted) fare classes tend to have smaller budgets, which constrains their ability to afford ancillary services. These basic ancillary purchase and revenue trends are similar in the other baseline and experimental cases.

Given these characteristics, the change in total revenue when the airline implements the optimizer increment is shown in Figure 5.10 for various ancillary prices, demands, and disutility configurations. The low, medium, and high demand cases correspond to baseline load factors of 77%, 84%, and 88%, respectively. The optimizer increment consistently decreases total revenue in these simulations; while the change is small, it is statistically

[4]US DOT Form 41, Schedule P-1.2

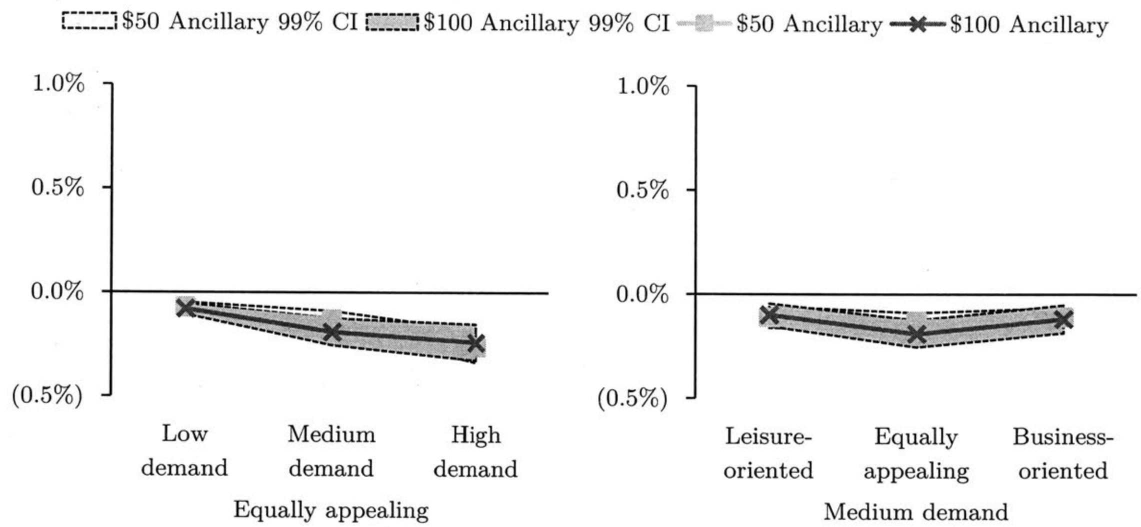

Figure 5.10: Network A1ONE change in total revenue due to optimizer increment and 99% confidence interval vs baseline with various ancillary prices, market demand levels, and disutility scenarios (100% simultaneous passengers). $df = 1,999$.

significant: the 99% confidence intervals ($df = 1,999$) lie exclusively below zero in Figure 5.10.

Additional results are shown in Table 5.7 for the range of ancillary prices with medium demand and equally appealing disutility configuration. Although the changes in total revenue due to the optimizer increment are small, they are statistically significant ($p < 0.001$, $df = 1,999$) and they are consistent with the findings in Section 5.2.2 when the airline forecasted demand: the optimizer increment decreases ticket revenue, increases load factor, decreases yield (meaning the booking mix shifts toward the lower-value fare classes), and increases ancillary revenue. In these cases, the net impact is a consistent decrease in total revenue, with a (slightly) larger decrease at higher ancillary prices. Although the results are not shown here, the same trends exist when all passengers exhibit sequential behavior or when passengers have a mix of 50% simultaneous and 50% sequential behavior types.

Table 5.7: Network A1ONE simulation results with various ancillary prices (medium demand, 100% simultaneous passengers, equally appealing disutility scenario).

	\$25 Ancillary	\$50 Ancillary	\$75 Ancillary	\$100 Ancillary
Baseline				
Ticket Revenue	\$22,123	\$22,098	\$22,067	\$22,034
Ancillary Revenue	\$1,096	\$1,808	\$2,290	\$2,618
Total Revenue	\$23,219	\$23,906	\$24,357	\$24,652
Load Factor	83.8%	83.8%	83.8%	83.8%
Total Yield	21.32	21.95	22.36	22.63
Ancillary Purchase Rate	40.3%	33.2%	28.0%	24.0%
Optimizer Increment				
Ticket Revenue	\$22,107	\$22,068	\$22,029	\$21,987
Ancillary Revenue	\$1,096	\$1,808	\$2,288	\$2,618
Total Revenue	\$23,203	\$23,876	\$24,317	\$24,605
Load Factor	83.8%	83.9%	83.9%	83.9%
Total Yield	21.29	21.90	22.30	22.57
Ancillary Purchase Rate	40.2%	33.2%	28.0%	24.0%
Change from Baseline				
Ticket Revenue	-0.1%	-0.1%	-0.2%	-0.2%
Ancillary Revenue	+0.0%	+0.0%	-0.1%	+0.0%
Total Revenue	-0.1%	-0.1%	-0.2%	-0.2%
Load Factor	+0.1 pts	+0.1 pts	+0.1 pts	+0.1 pts
Total Yield	-0.1%	-0.2%	-0.3%	-0.3%
Ancillary Purchase Rate	-0.0 pts	-0.0 pts	-0.1 pts	-0.0 pts
Significance of Change in Total Revenue from Baseline				
Standard Error	0.0%	0.0%	0.0%	0.0%
t-statistic	-4.59	-7.69	-6.73	-7.66
p-value	< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

Table 5.8: Network D6 fare and ancillary fee structure.

	Fare range	Average fare	Advanced purchase	Restriction Applies? R1	R2	R3
FC 1	$188–743	$413	None	-	-	-
FC 2	$137–515	$293	3 days	-	Yes	-
FC 3	$88–297	$179	7 days	Yes	-	Yes
FC 4	$76–248	$153	10 days	Yes	Yes	-
FC 5	$65–198	$127	14 days	Yes	Yes	Yes
FC 6	$54–153	$101	21 days	Yes	Yes	Yes

5.2.4 Two Airline, Many Flight Leg Network (D6)

Our final study of the optimizer increment utilizes the two airline, multiple flight leg network D6. In D6, shown in Figure 5.11, the airlines each operate a connecting hub, with a total of 252 flight legs serving 482 different origin-destination markets. Each airline offers three itineraries in each market and each airline has six fare classes and sells one ancillary service. Fares vary by market but both airlines have the same fares within a given market. The FC 1 fare for a market is between 3 and 5 times the FC 6 fare. Fares and demands were previously calibrated based on data provided by airline members of the MIT PODS Research Consortium

Because both airlines operate connecting hubs, in all our simulations they utilize network-based revenue management optimization and forecasting models (DAVN and standard forecasting, described in Sections 3.2 and 4.1). The fare structure is shown in Table 5.8. The ancillary service is priced at 40% of the FC 6 (the same ratio as the $50 ancillary price in network A1ONE), which corresponds to ancillary prices between $22 and $62. We use the same disutility scenarios as in the previous section (see Table 5.5). D6 simulations have two trials of 400 unburned samples, leading to a total of 800 simulated departure days.

Baseline revenues for both airlines are shown in Table 5.9 for a medium demand, equally appealing case with simultaneous passengers. Airline 2's more southern hub makes its

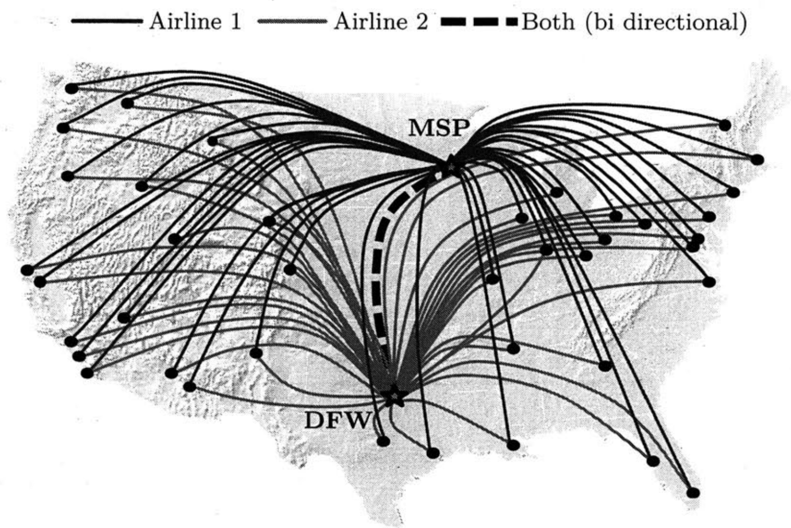

Figure 5.11: Network D6 map. Traffic flows from west to east, except for flights between the two hubs which are bi-directional. Each airline operates three connecting banks.

flights slightly longer and slightly less appealing to most connecting traffic; it therefore has a slightly smaller load factor and revenues than Airline 1. Airline 1 has a system load factor of 83.7%, and 42.6% of its passengers purchase the ancillary service.

A breakout of booking and ancillary purchase data by fare class for the same case is shown in Table 5.10 for Airline 1. Results for Airline 2 are not shown but are similar. The ancillary service has a relatively even purchase rate, between 40% and 46% in all fare classes. The higher average budgets in this network result in similar ancillary purchase rates across fare classes, even with simultaneous passengers. Between 4% (for FC 1) and 14% (for FC 6) of total revenue from each fare class comes from the ancillary service, leading to a network average of 8%, again in line with US DOT estimates of ancillary revenue for US airlines.

We conduct two sets of experiments to assess the effect of the optimizer increment in this large, competitive network. First, we evaluate a *symmetrically* competitive case where both airlines use the optimizer increment, and second an *asymmetric* case where Airline 1 only implements the optimizer increment (and Airline 2 retains its standard forecasting and

Table 5.9: Network D6 baseline statistics (medium demand, 100% simultaneous passengers, equally appealing disutility scenario).

	Airline 1	Airline 2
Ticket Revenue	\$1,326,442	\$1,307,530
Ancillary Revenue	\$115,904	\$115,431
Total Revenue	\$1,442,346	\$1,422,961
Load Factor	83.7%	83.0%
Total Yield	14.04	13.46
Ancillary Sales Rate	42.6%	42.4%
Portion of Total Revenue from Ancillary	8.0%	8.1%

Note: Total yield in cents per mile.

optimization systems from the baseline case); these cases are described in Table 5.11.

Symmetric Optimizer Increment

As shown in Figure 5.12, when both airlines implement the optimizer increment they both see decreases in total revenue for low, medium, and high demand (corresponding to 78.2%, 83.7%, and 87.6% Airline 1 baseline load factors with 100% simultaneous passengers and the equally appealing disutility scenario). The change in total revenue vs baseline cases where the airline optimizes only on the fare are small—between -0.2% and -0.4%—but are statistically significant (as shown by the 99% confidence interval, $df = 799$). As with networks C2 and A1ONE, use of the optimizer increment results in a decrease in ticket

Table 5.10: Network D6 Airline 1 baseline bookings and ancillary purchase data by fare class (medium demand, 100% simultaneous passengers, equally appealing disutility scenario).

	FC 1	FC 2	FC 3	FC 4	FC 5	FC 6
Bookings	987	1,334	992	938	545	2,506
Booking Mix	14%	18%	14%	13%	7%	34%
Average Fare	\$363	\$278	\$169	\$145	\$123	\$91
Average Ancillary per Passenger	\$17	\$17	\$17	\$15	\$14	\$15
Portion of Total Revenue from Ancillary	4%	6%	9%	9%	10%	14%
Ancillary Sales Rate	45%	44%	46%	40%	36%	42%

Table 5.11: Network D6 experimental outline.

Case	Airline 1	Airline 2
Baseline	Filed fares	Filed fares
Symmetric OI	Optimizer increment	Optimizer increment
Asymmetric OI	Optimizer increment	Filed fares

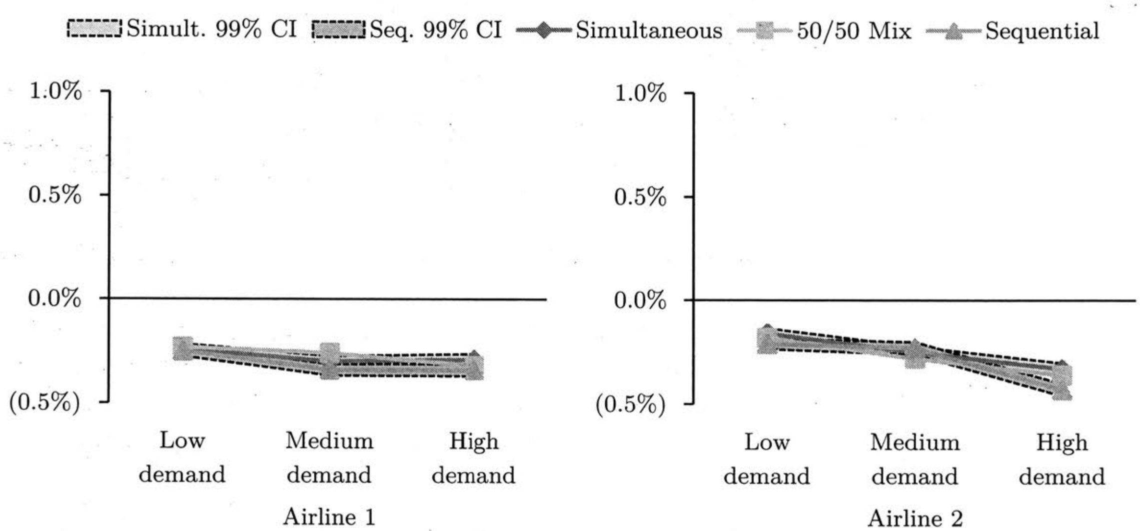

Figure 5.12: Network D6 change in Airline 1 and Airline 2 total revenue revenue due to symmetric optimizer increment and 99% confidence interval vs baseline with various market demand levels and passenger behavior types (equally appealing disutility scenario). $df = 799$.

revenue, an increase in ancillary revenue, and a decrease in total revenue, as shown for Airline 1 in Figure 5.13. The decrease in ticket revenue is driven by a shift in the booking mix toward lower-value fare classes and *not* by a reduction in the number of bookings. In fact, as the optimizer increment leads to more availability in the lowest-value fare classes, the airline accepts more bookings. The increase in total bookings, and the magnitude of the shift in the booking mix toward FC 6, increase at higher demand levels. Results for Airline 2 show similar patterns.

Figure 5.14 shows how bookings by fare class change within each DCP for Airline 1 (with

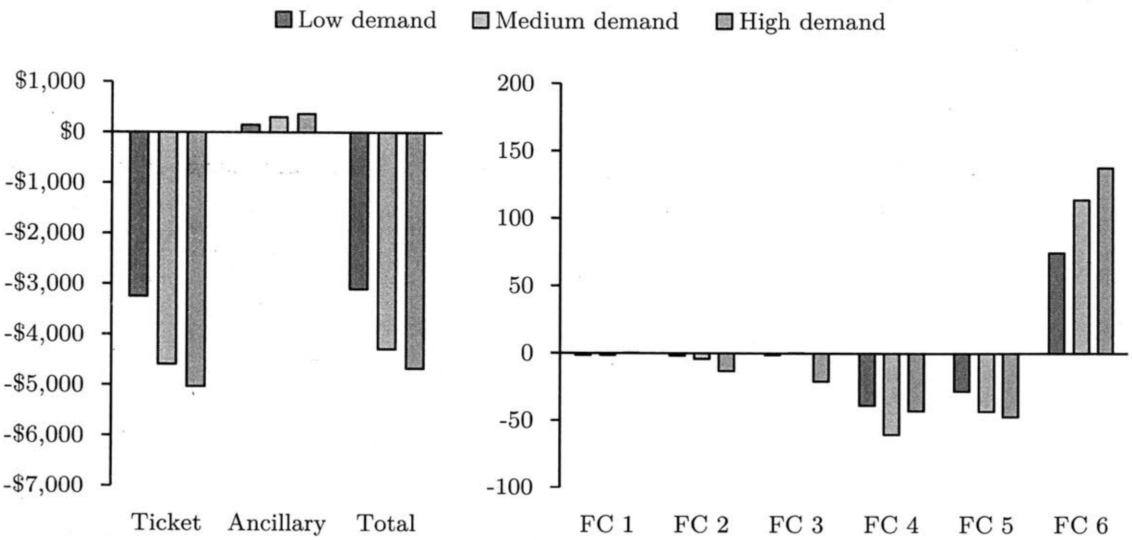

Figure 5.13: Network D6 change in Airline 1 revenue component (left) and change in bookings by fare class (right) due to symmetric optimizer increment vs baseline with various market demand levels (100% simultaneous passengers, equally appealing disutility scenario).

the total industry change plotted for comparison). The figure suggests that a primary driver of optimizer increment-induced revenue losses is *buy-down*. In the early DCPs, for each FC 6 booking gained, Airline 1 loses bookings in higher-value fare classes (in particular FC 4 and FC 5), driving ticket revenue losses (recall that FC 4 and 5 have higher fares than FC 6, so trading one FC 4 booking for one FC 6 booking reduces revenue). Airline 1 accepts more bookings during the first portion of the booking window and therefore can accept fewer bookings in the last few DCPs, resulting in *displacement* of late-arriving passengers. Although the magnitude of booking changes is small in the DCPs near departure, the changes occur entirely in the highest value fare classes, so each incremental booking has a large effect on total revenue.

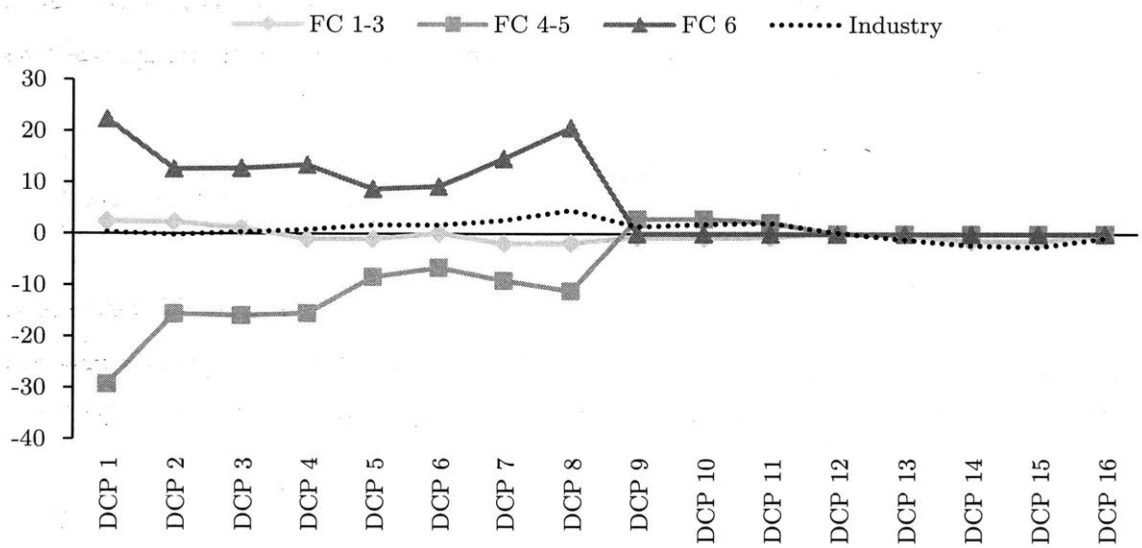

Figure 5.14: Network D6 change in Airline 1 bookings by fare class and DCP due to symmetric optimizer increment vs baseline (medium demand, 100% simultaneous passengers, equally appealing disutility scenario). Industry total includes all Airline 1 and Airline 2 fare classes.

Asymmetric Optimizer Increment

Results for *asymmetric* competition, where Airline 1 uses the optimizer increment and Airline 2 maintains its baseline approach of optimizing only on the fare, are similar for Airline 1 as the symmetric case. As shown in Figure 5.15, the optimizer increment again decreases total revenue for Airline 1 for low, medium, and high demand and with all three passenger behavior scenarios, with small but statistically significant changes vs baseline (as shown by the 99% confidence interval, $df = 799$). Again, the optimizer increment results in a decrease in ticket revenue, an increase in ancillary revenue, and a decrease in total revenue. Airline 2 (results not shown) is largely unaffected by Airline 1's use of the optimizer increment, and sees total revenue increases less than 0.1%.

Empirical cumulative distribution functions for the change in total revenue (vs baseline) are shown in Figure 5.16 for both airlines for both the symmetric and asymmetric cases.

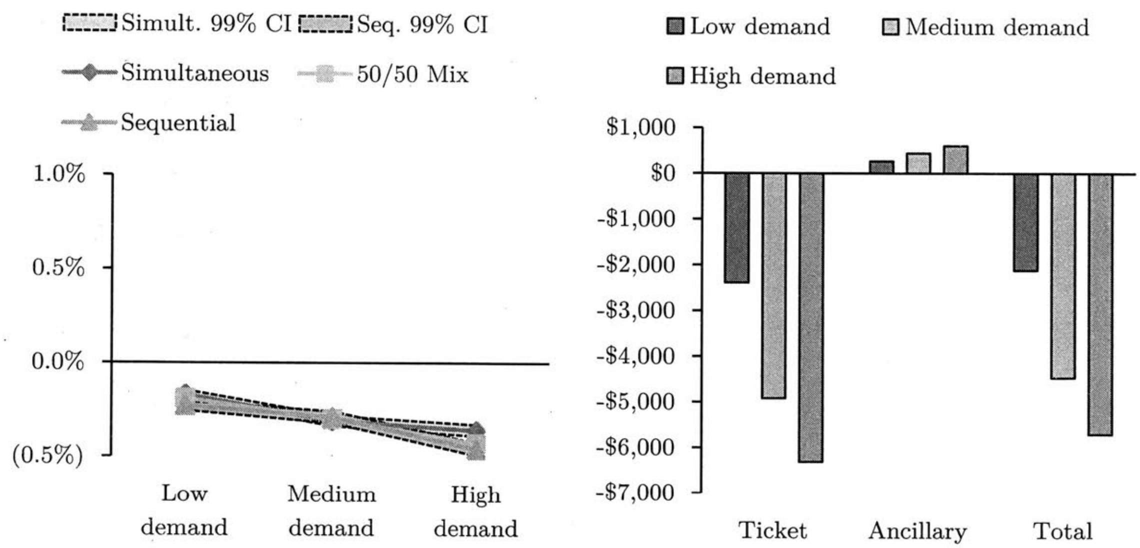

Figure 5.15: Network D6 change in Airline 1 total revenue (left) and change in Airline 1 revenue component (right) due to asymmetric optimizer increment vs baseline with various market demand levels and passenger behavior types (100% simultaneous passengers for revenue component change, equally appealing disutility scenario). $df = 799$.

For the symmetric optimizer increment, both airlines have median and mean total revenue changes below zero; for the asymmetric optimizer increment (where Airline 2 maintains its traditional optimization based only on filed fares), Airline 1 has a median and mean total revenue change less than zero while Airline 2 has a median and mean total revenue change near zero.

Estimation Quality

Some markets and fare class combinations may receive very few bookings, in which case the historical database may not contain any booking records for these combinations during some of the historical data periods used for estimating average ancillary revenue per passenger.[5] In such a case, use of the default estimate $\tilde{a}_{mkt,k}$ will lead to a biased estimate.

[5] United Airlines, for example, states that it generates nearly 8 million forecasts per day and that "nearly all have [less than] one passenger forecasted" (United Airlines Investor Day Presentation (2016), pg. 46–47).

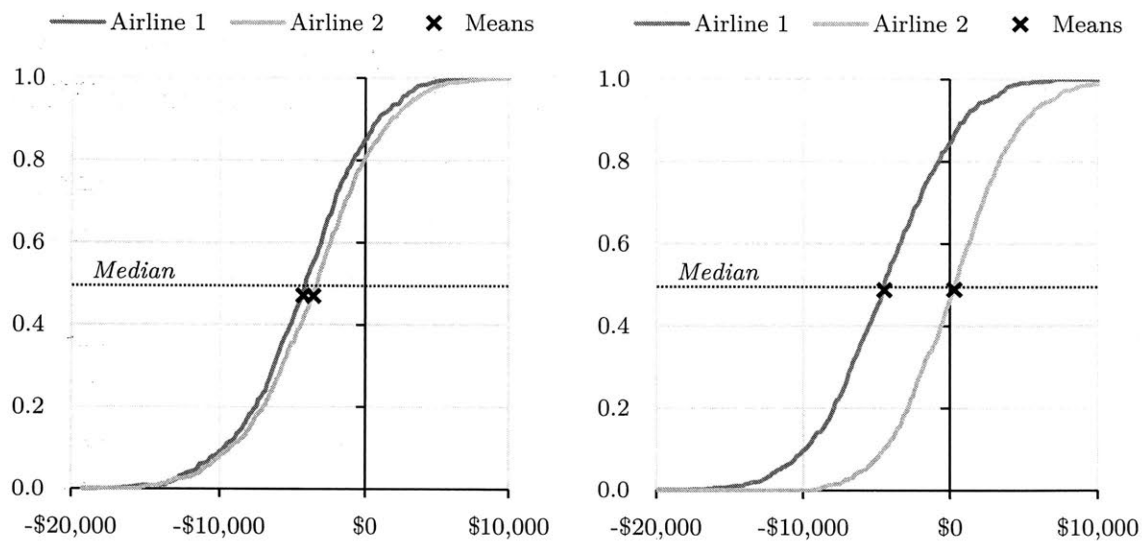

Figure 5.16: Network D6 empirical cumulative distribution function of change in total revenue by sample due to symmetric optimizer increment (left) and asymmetric optimizer increment (right) vs baseline (medium demand, 100% simultaneous passengers, equally appealing disutility scenario).

To reduce the bias in $\hat{a}_{mkt,k}$ caused by the use of the default value $\tilde{a}_{mkt,k}$ for the periods without bookings, we tested an alternative estimation method ("Method 2") in which we estimate ancillary purchase rates for each ancillary service and fare class across a group of markets, where the true purchase rate is $\rho_{mkt,k,s}$ for market mkt (not group), fare class k, and ancillary service s. The market groups have similar fare class restrictions, ancillary service offerings, and competitive characteristics (in D6 all markets are combined into one group); the group-level purchase rates are combined with market-specific ancillary prices to get market and class-specific estimates of average ancillary revenue. The estimated purchase rate $\hat{\rho}_{g,k,s}$ for market group g, fare class k, and ancillary service s is given by:

$$\hat{\rho}_{g,k,s} = \begin{cases} \frac{\sum_{dep=1}^{n_{ob}} \sum_{mkt \in MG_g} x_{mkt,k,s,dep}}{\sum_{dep=1}^{n_{ob}} \sum_{mkt \in MG_g} b_{mkt,k,dep}} & \sum b_{mkt,k,dep} \neq 0 \\ \tilde{\rho}_{g,k,s} & \text{otherwise} \end{cases}$$

where MG_g is the set of markets in group g, and $\tilde{\rho}_{g,k,s}$ us a default value (set to 0 in our implementation). The market/class specific estimate is computed as:

$$\hat{a}_{mkt,k} = \sum_{s \in S} \hat{\rho}_{g_{mkt},k,s} r_{mkt,k,s}$$

where g_{mkt} is the group of market mkt. Note that this estimation method will also be biased if different markets (within a market group) have different ancillary purchase characteristics, even assuming a market group receives some bookings during the historical data period:

$$\begin{aligned} \mathrm{E}\left[\hat{a}_{mkt,k} \mid b\right] &= \sum_{s \in S} \mathrm{E}\left[\hat{\rho}_{g_{mkt},k,s} \mid b\right] r_{mkt,k,s} \\ &= \frac{\sum_{s \in S} \rho_{mkt,k,s} r_{mkt,k,s} \sum_{dep=1}^{n_{ob}} b_{dep,mkt,k} + \sum_{q \in (MG_{g_{mkt}} \setminus mkt)} \left(\sum_{s \in S} \rho_{q,k,s} \sum_{dep=1}^{n_{ob}} b_{dep,q,k} \right)}{\sum_{dep=1}^{n_{ob}} b_{dep,mkt,k} + \sum_{q \in (MG_{g_{mkt}} \setminus mkt)} \sum_{dep=1}^{n_{ob}} b_{dep,q,k}} \end{aligned} \tag{5.8}$$

The left hand term in both the numerator and the denominator of Equation 5.8 simplify to an unbiased estimate of the mean ancillary revenue per passenger:

$$\frac{\sum_{s \in S} \rho_{mkt,k,s} r_{mkt,k,s} \sum_{dep=1}^{n_{ob}} b_{dep,mkt,k}}{\sum_{dep=1}^{n_{ob}} b_{dep,mkt,k}} = \sum_{s \in S} \rho_{mkt,k,s} r_{mkt,k,s} = \sum_{s \in S} \mu_{mkt,k,s}$$

The right hand terms in both the numerator and denominator of Equation 5.8 reflect the bias introduced in the estimator by integrating purchase rates from other markets within the same market group. The number and type of markets within each group will affect this bias: larger groups increase bias because $\hat{\rho}$ mixes effects from more markets, but also decreases bias because fewer estimate instances will have zero bookings (and therefore rely

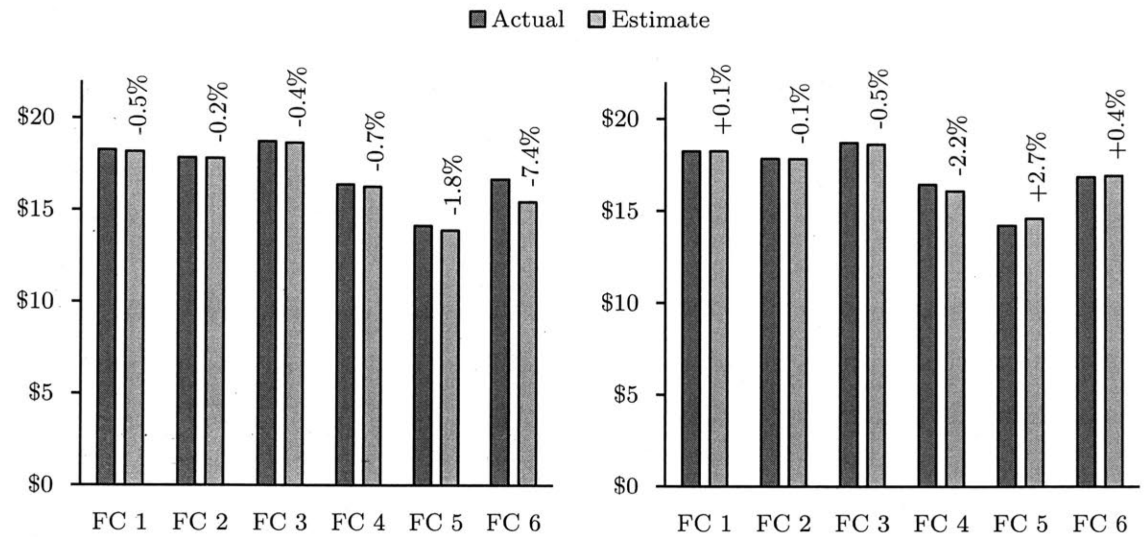

Figure 5.17: Network D6 actual and estimated Airline 1 average ancillary revenue per passenger by fare class with symmetric optimizer increment using historical average ancillary revenue per booking (left, "Method 1") and historical average ancillary purchase rate (right, "Method 2") (medium demand, 100% simultaneous passengers, equally appealing disutility scenario). Note that unlike Table 5.10, the averages reported here are not weighted by bookings within each market. Percentages show the change vs actual average per passenger by fare class.

on default purchase rate $\tilde{\rho}$).

The difference between the actual ancillary revenue per passenger by fare class and the estimated value using estimation Method 1 or Method 2 is shown in Figure 5.17 for the symmetric optimizer increment (the asymmetric case has similar results and is not shown). Using Method 1, the estimated values (and therefore the magnitude of the increment) generally match actual values, with the largest difference (of less than 8%) occurring in FC 6. FC 6 is most likely to have zero bookings (and therefore a biased estimate) because the revenue management system may close it entirely. Using estimation Method 2 produces a closer match with the actual values, with the maximum error reduced to less than 3%. While Method 1 consistently underestimates ancillary revenue, Method 2 shows a mix of over and under estimates. Although there are slight differences in the estimated values, both method produce similar revenue performance: a decrease in total revenue for Airline 1 vs baseline of 0.3% with either symmetric or asymmetric optimizer increment with medium demand, 100% simultaneous passengers, and the equally appealing disutility scenario. Thus, the alternative estimation method does not appear to have a substantial benefit in these tests.

5.3 Conclusions

The optimizer increment is an easy and intuitive approach to optimization of total revenues (ticket plus ancillary). Previous studies have simulated its performance in environments where passengers do not make choices about their ancillary services, but no prior work has provided a thorough theoretical assessment of the mechanism, nor has any prior work evaluated its performance when passengers *do* make choices about ancillary services. In this chapter, we have shown that the optimizer increment is an optimal approach to maximizing total expected revenue under limited conditions, and under such conditions its use leads to higher ancillary and total revenues, with the increase in ancillary revenue partially offset by a decrease in ticket revenue.

Two additional important findings from this analysis are that the optimizer increment is not an optimal revenue management optimization approach when passenger make choices about fare classes, and that introducing demand forecasting can substantially degrade the optimizer increment's revenue performance (even in cases where it is otherwise an optimal strategy). Our simulations show decreases in total revenue driven by ticket revenue losses that more than offset gains in ancillary revenue in variety of environments, including competitive networks with connecting flights and hundreds of origin-destination markets. These key findings provide the motivation for Chapter 6, where we develop an ancillary-aware revenue management optimization and forecasting model that explicitly accounts for passenger choices, and for Chapter 7, where we verify that our total revenue optimization proposal remains functional even when demand forecasting and other complexities of more realistic environments are introduced.

Chapter 6

Ancillary Choice Dynamic Program

For decades, airlines have invested in revenue management (RM) systems to maximize the proceeds from ticket sales.[1] Since the mid-2000s, however, airlines have been developing a secondary, ancillary, revenue stream by both unbundling their product and offering new services for sale. As the number, price, and value to airlines of ancillary services grow, so does the potential benefit of a new generation of revenue management models that attempt to maximize total revenue, not just ticket revenue. In this chapter, we develop a new approach to total revenue optimization with the Ancillary Choice Dynamic Program (ACDP), which explicitly incorporates ancillary revenues and the passenger choice impacts of ancillary services. We then show that the model leads to two heuristics, an Ancillary Marginal Demand transformation (AMD) and an Ancillary Marginal Revenue transformation (AMR), which together can transform existing RM models to be both ancillary-aware and choice-aware in their calculation of fare class booking limits. Finally, we discuss how emerging technologies

[1]Portions of this chapter were previously awarded the 2018 Anna Valicek bronze medal as Bockelie, A. and Belobaba, P. (2018). Total revenue optimization with the Ancillary Marginal Demand and Ancillary Marginal Revenue transformation heuristics. Presented at the 58th Annual AGIFORS Symposium, Tokyo, Japan.

in airline distribution will allow further extensions to our model, and describe how our work could provide a platform for an offer generation engine.

As with Chapter 3, we focus specifically on optional services sold to passengers in conjunction with a particular itinerary, such as checked baggage, seating upgrades and seat assignments, inflight meals and entertainment, priority boarding, and lounge access. We develop a single leg RM optimization and forecasting algorithm, using a flexible customer choice model. We discuss the significant practical marketing and distribution constraints that restrict the types of offers and that airlines can sell, and show how our model can be restricted to produce booking policies that can be implemented under these conditions. We develop two heuristics that can be used to convert existing RM models to total revenue management (i.e. maximizing the sum of ticket and ancillary revenue) and we show that under specific choice models our heuristics are equivalent to existing total revenue optimization methods, but that in general our models provide an additional level of specificity. We also describe additional approximations and processes that may be required to operationalize our models.

The remainder of this chapter is organized as follows: Section 6.1 presents our model formulation, discusses practical constraints, introduces our two heuristics, and investigates their equivalence to previous approaches. Section 6.2 addresses challenges that may be encountered in implementing our heuristics and proposes two solution processes: one addressing the presence of inefficient booking policies, another providing a demand forecasting module. Section 6.3 describes how the heuristics could be used in a network setting, despite the leg-based formulation of ACDP. We conclude with a summary in Section 6.5 as well as thoughts on potential future work.

6.1 Model Formulation

We consider a single airline, single flight leg network, with multiple fare classes and multiple ancillary services. The fare classes are indexed $1, \ldots, k, \ldots, n_{FC}$ and ordered by decreasing fare; the fare for class k is f_k, so $f_1 \geq f_2 \geq f_k \geq f_{n_{FC}}$. The airline has grouped its ancillary services into purchasable combinations $0, 1, \ldots, m, \ldots, n_{COMB}$, where set 0 corresponds to the set of no ancillary services. These combinations are formulated subject to the airline's marketing policies and goals (and ensure that passengers combine ancillary services in a sensible manner, such as a prohibiting a second checked bag without also buying a first checked bag); the set $M_k \subseteq \{0, 1, \ldots, n_{COMB}\}$ lists the combinations that are permitted in class k. Combination m of ancillary services purchased in conjunction with fare class k has price a_{km}. For modeling convenience, we assume that the fare class 0 corresponds to a decision by the passenger to not fly.

Time is discrete and counts down to departure, which occurs at $t = 0$. We assume that demand has a Poisson distribution, that time slices are small enough that there is at most one arrival per slice, and the probability of an arrival during slice t is λ_t. Capacity x is the number of unsold seats, which constrains the total number of sales. We assume that individual ancillary services have no capacity constraints and have a negligible marginal cost to the airline. We assume no cancellations or overbooking.

The airline's booking policy for each time and capacity state (t, x) is an *offer set* O. We define an offer $(k, m) \in O$ as a specific fare class k and combination $m \in M_k$ of ancillary services; a consumer purchases exactly one offer in its entirety from the offer set. The offer $(0, 0)$, which corresponds to the no-fly option, is always included in the offer set. Ancillary services are considered "optional" for a fare class if there are offers for that class in the offer set both with and without the ancillary service.

We assume that consumers make choices according to a flexible and general choice model, where the probability that a consumer chooses a particular offer is a function of the booking policy in effect and the time at which the consumer arrives in the booking process. We specify this choice model in terms of choice probabilities $P_{kmt}(O)$, which is the probability that a consumer arriving at time t chooses offer (k,m) when presented with offer set O, with $P_{kmt}(O) = 0$ if $(k,m) \notin O$. The exact structure of this choice function can vary by context; two potential models are the sequential or simultaneous behaviors described in Chapter 3.

Multiple consumer demand segments may be present, each with their own choice function. However, we assume that the airline cannot provide different booking policies to these segments, and so the probabilities $P_{kmt}(O)$ reflect a weighted average of the segments arriving at time t.

The probability that the airline sells offer (k,m) during time t, given that it has booking policy O in effect, is $\lambda_t P_{kmt}(O)$. The probability that the airline sells nothing at time t is $(1-\lambda_t)+\lambda_t P_{00t}(O)$, which reflects that the lack of sale may be due to no arrival, or because the consumer chose not to fly. The total probability of sale TP_t for a particular booking policy is the probability that an arriving consumer purchases anything from the policy, and is:

$$TP_t(O) = \sum_{(k,m)\in O} P_{kmt}(O) \tag{6.1}$$

Likewise, the total expected revenue TR_t from an arriving customer presented with policy O is:

$$TR_t(O) = \sum_{(k,m)\in O} P_{kmt}(O)(f_k + a_{km}) \tag{6.2}$$

The airline selects the booking policy that maximizes the total expected revenue to come V in future time periods via the Ancillary Choice Dynamic Program (ACDP):

$$V(t,x) = \max_{O} \left\{ \sum_{(k,m)\in O} \lambda_t P_{kmt}(O)\left(f_k + a_{km} + V(t-1,x-1)\right) \right.$$
$$\left. + \left(\lambda_t P_{00t}(O) + 1 - \lambda_t\right) V(t-1,x) \right\} \quad (6.3)$$

We define a bid price function $\Delta V(t,x) = V(t,x) - V(t,x-1)$ as the marginal cost of capacity and can rewrite Equation 6.3 in simpler terms:

$$V(t,x) = \max_{O} \left\{ \sum_{(k,m)\in O} \lambda_t P_{kmt}(O)(f_k + a_{km} - \Delta V(t-1,x)) \right\} + V(t-1,x) \quad (6.4)$$

The airline chooses the one offer set, or booking policy, O in each time and capacity state that maximizes the total expected revenue earned from that time and capacity state, minus the cost of consumed capacity, plus the maximum expected revenue to come in future time periods. For modeling convenience, it is possible that the airline would only "offer" the no-fly option. This function is solved recursively, with the boundary conditions $V(0,x) = 0$ and $V(t,0) = 0$: no future revenue can be earned when there is no time and/or capacity remaining.

We can redefine an offer (k,m) as a *product* j sold by the airline, with price $r_j = f_{k(j)} + a_{k(j),m(j)}$ where $k(j)$ and $m(j)$ map the product index to the fare class and ancillary combination indices. Product j has purchase probability $P_{jt}(O) = P_{k(j),m(j),t}(O)$ when the airline is selling the set of products O. We can now rewrite Equation 6.4 as:

$$V(t,x) = \max_{O} \left\{ \sum_{j\in O} \lambda_t P_{jt}(O)(r_j - \Delta V(t-1,x)) \right\} + V(t-1,x) \quad (6.5)$$

which recovers the time-varying extension of the choice-based dynamic program proposed by Talluri and van Ryzin (2004), although our use case contains the addition ancillary dimension. Talluri and van Ryzin prove several important points about the optimal solutions to Equation 6.5 (in the non-time-varying case):

1. When input parameters λ_t, $P_{jt}(O)$, and r_j are known and accurate, and when demand is Poisson, the solutions to Equation 6.5 are optimal booking policies.

2. The optimal policies are always one of the *efficient sets.*

An efficient set, or efficient policy, is one which maximizes total expected revenue for any given total sale probability, or is part of a linear combination of policies that maximizes total expected revenue for any given total sale probability; the linear combination is, in practice, equivalent to alternating between booking policies. Talluri and van Ryzin define a policy O as *inefficient* in time slice t if there exists a set of policy weights $\alpha_t(S)$, with $\sum_{\forall S} \alpha_t(S) = 1$, such that:

$$TP_t(O) \geq \sum_{\forall S} \alpha_t(S) TP_t(S) \quad \text{and} \quad TR_t(O) < \sum_{\forall S} \alpha_t(S) TR_t(S)$$

otherwise O is efficient in time t (Talluri and van Ryzin, 2004, Section 3.1). In other words, a set is inefficient if there is a weighted combination of other sets with a smaller (or equal) sale probability and a greater total expected revenue. The no-fly-only policy (offering $(0,0)$ in our notation) is always efficient.

We can index and order the efficient sets in time t by increasing sale probability, such that $O_{1,t}, \ldots, O_{n_{ES_t},t}$ have $TP_t(O_{i,t}) \leq TP_t(O_{i+1,t})$; the sets are therefore also ordered in terms of increasing expected revenue (Talluri and van Ryzin, 2004, Proposition 3). Fiig et al. (2010), Walczak et al. (2010), and Gallego (2013) develop marginal revenue and marginal demand transformations to convert Equation 6.5 into an equivalent independent demand formulation. We extend those ideas to the ancillary dimension as the Ancillary

Marginal Demand transformation (AMD, Equation 6.6) and Ancillary Marginal Revenue transformation (AMR, Equation 6.7):

$$d_{i,t} = \lambda_t \left(TP_t(O_{i,t}) - TP_t(O_{i-1,t}) \right) \tag{6.6}$$

$$f'_{i,t} = \frac{TR_t(O_{i,t}) - TR_t(O_{i-1,t})}{TP_t(O_{i,t}) - TP_t(O_{i-1,t})} \tag{6.7}$$

where d_{it} is the marginal, or additional, demand that can be accommodated and f'_{it} is the marginal total revenue per unit of capacity that can be earned by moving from policy $O_{i-1,t}$ to O_{it} in time t, and we define $TR_t(O_{0,t}) = TP_t(O_{0,t}) = 0$. Fiig et al. (2010) show that these marginal unit revenues are decreasing in i, and that the optimal policy in time t is to offer the set $O_{i_t^*,t}$ with the smallest expected revenue greater than or equal to the bid price:

$$i_t^* = \max\{i \mid f'_{it} \geq \Delta V(t,x)\}$$

6.1.1 Practical Constraints and Limitations

Airline marketing policies and distribution technology impose significant practical constraints on the types of offers and offer sets that airlines can sell. Traditionally, airlines have sold tickets to consumers directly (through the airline website, call center, and ticket offices) and indirectly (through travel agents and online travel retailers). Airline sales through indirect channels are subject to the technical limitations of distribution technology, and sales through all channels are potentially subject to commercial agreements with various retailers. In practice, for many airlines, the same booking policy must be in place for consumers shopping in all channels (direct and indirect).

Indirect sales are often made through a Global Distribution System (GDS), which serves as a content aggregator. Approximately 50% of bookings worldwide are made through a GDS, and therefore the structure of GDSs has a significant impact on how airlines sell travel (Taubmann, 2016). A consumer makes a shopping request to a travel retailer (such as a human travel agent, or an online travel agent like Expedia), which then requests search results from a GDS. The GDS draws upon schedule data from a third party (Official Airline Guide, OAG), fare and ancillary pricing data from a third party (Airline Tariff Publishing Company, ATPCO), and availability data from airlines. The airline availability data lists the number of seats available in each booking class. The GDS then combines all of this information to assemble the travel options returned to the retailer (and in turn to the consumer). Figure 1.4 shows a schematic of this process.

Airlines are responsible for providing their schedules to OAG and fares to ATPCO, and can update the data in batch processes as necessary. However, the only "real time" control that the airline has in this process is the availability response. Thus, the airline can only control products at the fare class level. In the example in Figure 1.4, the airline has responded that fare classes Y, B, and M are available, and that the airline sells two ancillary products, BAG1 and BAG2. The airline cannot dictate in real time how those booking classes and ancillary services can be combined: if the airline wants to sell BAG1 and BAG2 to passengers booking in all three fare classes in general, the booking policy must always permit an optional BAG1 and BAG2 for classes Y, B, and M.[2] We term this constraint *fare class completeness*, and say that, to comply with traditional distribution system architecture, the offer sets considered by ACDP must be fare class complete.

Definition 1 (Fare Class Completeness). An offer set O is fare class complete (FCC) if and only if, for every fare class included in the set, all possible offers based on that fare class are also included in the set: if $(k, m) \in O$ for some $m \in M_k$, then O is FCC if and only if $(k, m') \in O$ for all $m' \in M_k$.

[2]It is possible that the airline marketing policies would offer some ancillary services complimentary to some fare classes, or would prohibit their purchase in other fare classes. These restrictions can be implemented through fare filing (and would affect the composition of M_k), but cannot be generated in real time.

In addition, we impose the practical constraint that offer sets must be nested by fare order (a common assumption in many of today's revenue management systems):

Definition 2 (Nesting by Fare Order)**.** An offer set O is nested by fare order (NFO) if and only if, for every fare class included in the set, all fare classes with a higher priced fare are also included in the set: if $(k, m) \in O$ for some $m \in M_k$, then O is NFO if and only if $(j, m') \in O$ for all $j \leq k$ and for some $m' \in M_j$.

Offer sets that are both fare class complete and nested by fare order are marketable by airlines within traditional distribution frameworks.

6.1.2 Ancillary Marginal Demand Transformation and Ancillary Marginal Revenue Transformation Heuristics

With the constraints described in the previous section, we can simplify our notation: the airline uses Equation 6.3 to choose a fare class complete, nested by fare order booking policy, which is equivalent to choosing a *lowest available fare class* with no RM-imposed limitations on ancillary purchase options.[3] We denote this policy simply as k, which in our earlier notation is equivalent to the set $O = \{(i, m) \mid \forall i \leq k, \forall m \in M_i\}$. Following Fiig et al. (2010) and Walczak et al. (2010), we propose using AMD and AMR, which are optimal as inputs to an independent demand dynamic program, as heuristics to transform the demand and revenue inputs for existing (independent demand) *static* RM optimizers, such as EMSR. These systems are typically designed for fare-class level control, which is why we focus on the case where the airline must choose a fare class complete, nested by fare order policy to implement.

Using AMD and AMR as heuristic input modifiers provides an easy method to obtain the benefits of ancillary-awareness and choice-awareness without the need to significantly modify

[3]To reiterate, ancillary purchase options for a given fare class could be limited by pre-specified marketing policies.

the core processes of the RM optimizer. The AMD demand corresponds to the incremental demand accommodated by opening class k, and the AMR fare to the incremental total revenue per unit of capacity earned by opening class k. We initially assume that each of these FCC/NFO policies is efficient, which allows the AMD and AMR outputs to give an adjusted demand and fare for each fare class, minimizing the need to modify the optimizer structure. We discuss the case of inefficient policies in Section 6.2.1. We express the total sale probability and total expected revenue for these booking policies as $TP_t(k)$ and $TR_t(k)$, and the heuristic versions of AMD and AMR are:

Ancillary Marginal Demand Transformation. The heuristic marginal demand $d_{k,t}$ associated with moving from fare class complete and nested by fare order booking policy $k-1$ to k in time slice t is:

$$d_{k,t} = \lambda_t \left(TP_t(k) - TP_t(k-1)\right) \tag{6.8}$$

Ancillary Marginal Revenue Transformation. The heuristic marginal total revenue $f'_{k,t}$ associated with moving from fare class complete and nested by fare order booking policy $k-1$ to k in time slice t is:

$$f'_{k,t} = \frac{TR_t(k) - TR_t(k-1)}{TP_t(k) - TP_t(k-1)} \tag{6.9}$$

In practice the airlines collect booking data, generate forecasts, and estimate parameters at various Data Collection Points (DCPs), which aggregate many time slices. In our simulations (Chapter 7), we break the booking window into 16 DCPs. We assume that choice probabilities and demand arrival rates are equal for each time slice within a DCP, and we express these parameters in terms of DCPs:

$$
\begin{gathered}
\lambda_t = \lambda_{dcp_t}, \quad P_{imt}(k) = P_{i,m,dcp_t}(k), \quad TP_t(k) = TP_{dcp_t}(k) \\
TR_t(k) = TR_{dcp_t}(k), \quad f'_{kt} = f'_{k,dcp_t}
\end{gathered} \tag{6.10}
$$

where dcp_t is the DCP that contains time slice t. We will typically refer to these quantities by DCP only, without reference to any particular time slice.

Gallego (2013) cautions that this heuristic approach is no longer an optimal solution, as the assumptions of the existing RM model are likely violated. For example, EMSR assumes that all low-fare demand arrives before high-fare demand, but under the general choice model we have incorporated there is no requirement that this assumption will hold. In addition, EMSR (typically) assumes normal demand distributions instead of Poisson demand distributions.

Despite the misalignment of these assumptions and loss of optimality, we believe (and our simulation results indicate) that the AMD and AMR heuristics can still provide a significant revenue benefit over traditional RM models when passengers make choices among fare classes and ancillary services.

6.1.3 Equivalence to Other Models

In this section we show that, under certain choice model conditions, the AMD and AMR heuristics are equivalent to the optimizer increment (OI) and OI combined with the (non-ancillary) marginal demand and revenue transformations.

The optimizer increment is a total revenue optimization heuristic for existing RM models based on supplying the optimizer with an adjusted fare the includes the expected ancillary revenue per passenger:

$$
f'^{OI}_{k,dcp} = f_k + \bar{a}_{k,dcp} \tag{6.11}
$$

where $f'^{OI}_{k,dcp}$ is the OI-adjusted fare for class k in DCP dcp and $\bar{a}_{k,dcp}$ is the expected ancillary revenue per booking in class k in DCP dcp. In practice, $\bar{a}$ would likely be estimated based on historical ancillary purchase data and could be aggregated across DCPs. No change is made to demand estimates.

The (non-ancillary) marginal demand and marginal revenue transformations (MD) and (MR) have the same structure as AMD and AMR Equations 6.8 and 6.9, except they rely on $TP^{MR}_{dcp}(k)$ and $TR^{MR}_{dcp}(k)$, which exclude the ancillary dimension:

$$TP^{MR}_{dcp}(k) = \sum_{i=1}^{k} \sum_{m \in M_k} P_{i,m,dcp}(k) \tag{6.12}$$

$$TR^{MR}_{dcp}(k) = \sum_{i=1}^{k} \sum_{m \in M_k} P_{i,m,dcp}(k) f_k \tag{6.13}$$

When these approaches are combined, the optimizer increment occurs before the expected revenue calculation:

$$TR^{OI+MR}_{dcp}(k) = \sum_{i=1}^{k} \sum_{m \in M_k} P_{i,m,dcp}(k) f'^{OI}_{k,dcp} = \sum_{i=1}^{k} \sum_{m \in M_k} P_{i,m,dcp}(k)(f_k + \bar{a}_{k,dcp}) \tag{6.14}$$

There is no additional change to total sale probability, so:

$$TP^{OI+MR}_{dcp}(k) = TP^{MR}_{dcp}(k) \tag{6.15}$$

Independent Demand Model

We first consider the independent demand choice model, in which each arriving customer has one preferred fare class k^* and combination of ancillary services m^*. If (k^*, m^*) is included in the airline's booking policy, the customer purchases it. Otherwise, they choose the no-fly option. We denote by $q_{k,m,dcp}$ the probability that an arriving consumer in DCP dcp has preference (k, m). Then, the independent demand model is specified by the choice probabilities:

$$P_{i,m,dcp}(k) = \begin{cases} q_{i,m,dcp} & i \leq k \\ 0 & \text{otherwise} \end{cases} \tag{6.16}$$

Theorem 2. *With an independent demand model, the AMR heuristic has the same expected value as the optimizer increment: $f'^{AMR}_{k,dcp} = f'^{OI}_{k,dcp}$, where f'^{AMR} is the AMR adjusted fare and f'^{OI} is the optimizer increment adjusted fare.*

Proof. We first note that with the independent demand model total sale probability is given by $TP_{dcp}(k) = \sum_{i=1}^{k} \sum_{m \in M_i} q_{i,m,dcp}$ and total expected revenue by $TR_{dcp}(k) = \sum_{i=1}^{k} \sum_{m \in M_i} q_{i,m,dcp}(f_i + a_{i,m})$. The expected ancillary revenue $\bar{a}_{k,dcp}$ from a booking in k class during DCP dcp is:

$$
\begin{aligned}
\bar{a}_{k,dcp} &= \sum_{m \in M_k} a_{km} \operatorname{Pr}_{dcp}(\text{buy } m \mid \text{book } k) \\
&= \sum_{m \in M_k} a_{km} \frac{\operatorname{Pr}_{dcp}(\text{buy } m \cap \text{book } k)}{\operatorname{Pr}_{dcp}(\text{book } k)} \\
&= \sum_{m \in M_k} a_{km} \frac{q_{k,m,dcp}}{\sum_{m' \in M_k} q_{k,m',dcp}} \\
&= \frac{\sum_{m \in M_k} a_{k,m} q_{k,m,dcp}}{\sum_{m \in M_k} q_{k,m,dcp}}
\end{aligned}
$$

We can now show that the AMR fare is equal to the filed fare plus the (DCP-specific) expected ancillary revenue per passenger:

$$
\begin{aligned}
f'^{AMR}_{k,dcp} &= \frac{TR_{dcp}(k) - TR_{dcp}(k-1)}{TP_{dcp}(k) - TP_{dcp}(k-1)} \\
&= \frac{\sum_{i=1}^{k} \sum_{m \in M_i} q_{i,m,dcp}(f_i + a_{i,m}) - \sum_{i=1}^{k-1} \sum_{m \in M_i} q_{i,m,dcp}(f_i + a_{i,m})}{\sum_{i=1}^{k} \sum_{m \in M_i} q_{i,m,dcp} - \sum_{i=1}^{k-1} \sum_{m \in M_i} q_{i,m,dcp}} \\
&= \frac{f_k \sum_{m \in M_k} q_{k,m,dcp} + \sum_{m \in M_k} q_{k,m,dcp} a_{k,m}}{\sum_{m \in M_k} q_{k,m,dcp}} \\
&= f_k + \frac{\sum_{m \in M_k} a_{k,m} q_{k,m,dcp}}{\sum_{m \in M_k} q_{k,m,dcp}} \\
&= f_k + \bar{a}_{k,dcp}
\end{aligned}
\tag{6.17}
$$

Thus, $f'^{AMR}_{k,dcp} = f'^{OI}_{k,dcp}$ and the two formulations lead to equivalent adjusted fares. □

Remark. With the independent demand model, the ancillary marginal demand transformation has no effect on forecasting: the marginal demand associated with opening a class is the entire demand for the class, since passengers either purchase their (one) preferred class or do not fly.

Conditionally Independent Demand

We now consider the case when passengers exhibit *conditionally independent* behavior, that is, where all passengers booking in class i in DCP dcp have a constant conditional probability $w_{m|i,dcp}$ of purchasing ancillary combination m, regardless of which booking policy $k \geq i$ is in effect. Therefore, the choice probabilities satisfy the following equation:

$$\begin{aligned} w_{m|i,dcp} &= \Pr_{dcp}(\text{buy } m \mid \text{book } i \text{ under policy } k) \\ &= \frac{P_{i,m,dcp}(k)}{\sum_{m' \in M_i} P_{i,m',dcp}(k)} \quad \forall i \leq k \end{aligned} \tag{6.18}$$

Theorem 3. *With a conditionally independent demand model that satisfies Equation 6.18, the AMR heuristic has the same expected value as the optimizer increment combined with the (non-ancillary) Marginal Revenue transformation defined by Fiig et al. (2010) and Walczak et al. (2010):* $f'^{AMR}_{k,dcp} = f'^{OI+MR}_{k,dcp}$*, where* f'^{AMR} *is the AMR adjusted fare and* f'^{OI+MR} *is the optimizer increment with marginal revenue transformation adjusted fare.*

Proof. With conditionally independent demand, the expected ancillary revenue in DCP dcp for a booking in class i is $\bar{a}_{i,dcp} = \sum_{m \in M_i} a_{i,m} w_{m|i,dcp}$; therefore the optimizer increment adjusted fare is $f'^{OI}_{k,dcp} = f_k + \sum_{m \in M_i} a_{i,m} w_{m|i,dcp}$. For the remainder of this section, we will refer to Equations 6.1 and 6.2 (and the DCP variants in Equation 6.10) as $TP^{AMR}_{dcp}(k)$ and $TR^{AMR}_{dcp}(k)$; the AMR adjusted fare is therefore:

$$f'^{AMR}_{k,dcp} = \frac{TR^{AMR}_{dcp}(k) - TR^{AMR}_{dcp}(k-1)}{TP^{AMR}_{dcp}(k) - TP^{AMR}_{dcp}(k-1)}$$

The optimizer increment/marginal revenue transformation adjusted fare has the same structure:

$$f'^{OI+MR}_{k,dcp} = \frac{TR^{OI+MR}_{dcp}(k) - TR^{OI+MR}_{dcp}(k-1)}{TP^{OI+MR}_{dcp}(k) - TP^{OI+MR}_{dcp}(k-1)}$$

To prove $f'^{AMR}_{k,dcp} = f'^{OI+MR}_{k,dcp}$, we show that the formulations have equivalent total sale probabilities and total expected revenues: $TP^{AMR}_{dcp}(j) = TP^{OI+MR}_{dcp}(j)$ and $TR^{AMR}_{dcp}(j) = TR^{OI+MR}_{dcp}(j)$. Total sale probability for the marginal revenue transformation is given by Equation 6.15:

$$TP^{OI+MR}_{dcp}(j) = \sum_{i=1}^{j} P_{i,dcp}(j) = \sum_{i=1}^{j} \sum_{m \in M_i} P_{i,m,dcp}(j)$$

which is equivalent to $TP^{AMR}_{dcp}(j)$ as defined in Equations 6.1 and 6.10. Next we show that the total expected revenues are equivalent:

$$\begin{aligned}
TR^{OI+MR}_{dcp}(k) &= \sum_{i=1}^{k} P_{i,dcp}(k) f'^{OI}_{i,dcp} \\
&= \sum_{i=1}^{k} \left(\left(\sum_{m \in M_i} P_{i,m,dcp}(k) \right) \left(f_i + \sum_{m \in M_i} a_{i,m} w_{m|i,dcp} \right) \right) \\
&= \sum_{i=1}^{k} \left(f_i \left(\sum_{m \in M_i} P_{i,m,dcp}(k) \right) + \frac{\sum_{m \in M_i} a_{i,m} P_{i,m,dcp}(k)}{\sum_{m' \in M_i} P_{i,m',dcp}(k)} \sum_{m \in M_i} P_{i,m,dcp}(k) \right) \\
&= \sum_{i=1}^{k} \sum_{m \in M_i} P_{i,m,dcp}(k)(f_i + a_{im}) \\
&= TR^{AMR}_{dcp}(k)
\end{aligned}$$

Since $TP^{AMR} = TP^{OI+MR}$ and $TR^{AMR} = TR^{OI+MR}$ for the conditionally independent choice model, the two approaches have the same expected adjusted fares (and same adjusted demands). $\square$

Under a general choice model, Equation 6.18 will not hold: the probability that a passenger purchases a particular ancillary combination, even given that they book in class k, may vary based on the other classes offered. In this case, AMR and the optimizer increment plus marginal revenue transformation will lead to different adjusted fares and, potentially, different booking policies. Our AMD and AMR formulations, by explicitly including the ancillary dimension of passenger choice, can more precisely measure the marginal total revenue (or demand) associated with opening a fare class than the previous approaches.

6.2 Operationalization

In this section we discuss additional processes and assumptions that are necessary or helpful to implement the AMD and AMR heuristics. First, as noted above, the AMD and AMR transformations described in Equations 6.8 and 6.9 are only valid if the booking policies k and $k-1$ are optimal solutions to the original Equation 6.3. As described by Talluri and van Ryzin (2004), the solutions to Equation 6.3 must always be *efficient sets*, and it is possible that the booking policies k and/or $k-1$ are not efficient. Borrowing from previous Fare Family research, we propose several convex hull approximation methods, which we will refer to as "gap-filling" methods, to cope with inefficient booking policies (see Hopperstad (2008) and Fiig et al. (2012)). Second, the ACDP formulation assumes that demand arrival rates λ_{dcp} are known; in reality, forecasting demand is a significant challenge for airlines. We develop an AMD Forecasting Model below to help cope with this challenge, particularly when the heuristics are coupled with an optimizer that requires a forecast of demand-to-come by fare class.

6.2.1 Gap-Filling

In cases where a fare class complete, nested by fare order policy k is not efficient, AMR adjusted fares will be inverted (i.e. the adjusted fare for class k may be *less than* the adjusted fare for class $k+1$).

The optimal approach to dealing with these inefficient policies is, of course, to not offer them. However, our proposed use case for the AMD and AMR heuristics is to feed adjusted demands and adjusted fares to an existing RM optimizer (such as EMSR) to provide ancillary and choice-awareness without significantly revising the core of the optimization procedure. As existing RM optimizers are based on fare class demands and revenues, it is important that the output of AMD/AMR can be expressed in terms of fare classes as well: each class needs an AMD demand and an AMR fare.

We propose three different mechanisms, which we term *gap-filling* methods, to deal with inefficient booking policies while maintaining compatibility with the structure of existing RM optimizers. The first two, vertical and horizontal gap-filling, involve approximating the airline's computed choice probabilities to move the inefficient policies onto the efficient frontier. The third mechanism, exclusion gap-filling, involves strategically modifying the AMD/AMR outputs to prevent the optimizer from producing an inefficient booking policy.

The outputs of gap-filling are the adjusted total sale probabilities $TP'_{dcp}(k)$ and adjusted total expected revenues $TR'_{dcp}(k)$ for each DCP and for each booking policy. These adjusted TP' and TR' values are used in subsequent AMD/AMR processes.

To illustrate each of these mechanisms, we consider a running example with six fare classes and one ancillary service (corresponding to DCP 2 of the simulation results presented in Section 7.3). The computed total sale probabilities and total expected revenues are listed in the upper-left portion of Table 6.1, along with the associated EMSR booking limits. Note that without any gap-filling, the AMR adjusted fares for FC 2 and FC 5 are inverted.

Figure 6.1 shows a plot of the efficient frontier and the fare class complete/nested by fare order booking policies $\{1\}$, $\{1,2\}$, and $\{1,2,3\}$; the policy of offering classes 1 and 2 is inefficient. Graphically, the AMR adjusted fare is the slope of the line segment connecting two adjacent policies. Vertical and horizontal gap-filling operate by shifting the policy FC

1–2 until it falls on the segment connecting FC 1 and FC 1–3. Vertical gap-filling shifts the policy up; horizontal gap-filling shifts the policy left. Other (diagonal) gap-filling policies would be possible, but are not investigated here.

Vertical Gap-Filling

With the vertical gap-filling approximation, the airline's computed total expected revenue $TR_{dcp}(k)$ is increased for inefficient policies until the policy falls on the efficient frontier. AMD forecasting is not directly affected by this process, as the total sale probabilities remain unchanged. It is important to note that this approximation changes only the airline's perception of expected revenue; the airline does not change fares, fare class restrictions, or ancillary prices, so actual passenger choice will not be affected.

The vertical gap-filling approximations are:

$$TP'_{dcp}(k) = TP_{dcp}(k)$$

$$TR'_{dcp}(k) = \begin{cases} TR_{dcp}(k) & \text{if } k \text{ is efficient} \\ TR_{dcp}(k-1) + \frac{TR_{dcp}(k+1) - TR_{dcp}(k-1)}{TP_{dcp}(k+1) - TP_{dcp}(k-1)} & \\ \qquad \times (TP_{dcp}(k) - TP_{dcp}(k-1)) & \text{otherwise} \end{cases}$$

The result of vertical gap-filling is shown in upper-right section of Table 6.1. With the example parameters, vertical gap-filling increases TR for FC 2 by 2.3% and for FC 5 by 0.4%. These small changes, however, translate into a 27% increase in the adjusted fare of FC 2, as well as smaller changes for the adjusted fares of FC 3, 5, and 6. In addition, the EMSR booking limits change; one fewer seat is permitted to FC 3 and one additional seat is permitted to FC 5.

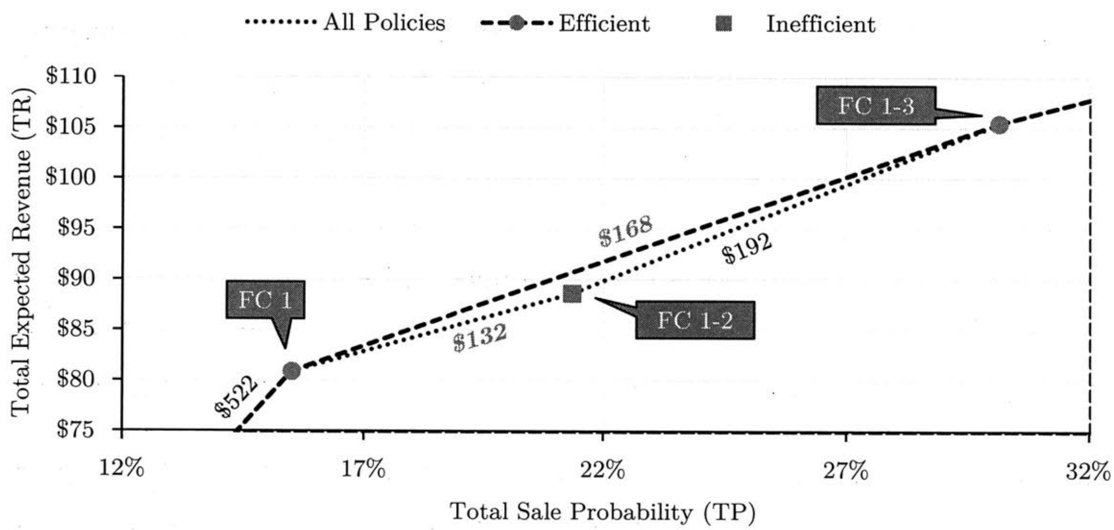

Figure 6.1: Portion of an example convex hull, showing efficient policies in blue circles and inefficient policies in red squares. The slope (AMR adjusted fare) of various segments is indicated, with emphasis in red for inverted fares and in blue for gap-filled fares.

Table 6.1: Example AMD mean demand $\tilde{\mu}_k$, AMR fares f'_k, and associated EMSR booking limits (BL) for DCP 2 with various gap-filling methods. Inverted fares are emphasized in red; changes from no gap-filling emphasized in blue with the size of the change indicated in parentheses. (Forecast volume mean $\tilde{\mu} = 85$, forecast volume standard variance $\tilde{\sigma}^2 = 900$, and capacity remaining $x = 110$).

k	$TP(k)$	$TR(k)$	$\tilde{\mu}_k$	f'_k	Booking Limit
No Gap-Filling					
1	15.5%	\$81	13	\$522	110
2	21.4%	\$89	5	*\$132*	102
3	30.1%	\$105	7	\$192	101
4	52.8%	\$134	19	\$126	96
5	71.7%	\$148	16	*\$77*	85
6	100.0%	\$172	24	\$82	78
Vertical Gap-Filling					
1	15.5%	\$81	13	\$522	110
2	21.4%	***\$91 (+2.3%)***	5	***\$168 (+27%)***	102
3	30.1%	\$105	7	***\$168 (-12%)***	***100 (-1)***
4	52.8%	\$134	19	\$126	96
5	71.7%	***\$149 (+0.4%)***	16	***\$80 (+4%)***	***86 (+1)***
6	100.0%	\$172	24	***\$80 (-2%)***	78
Horizontal Gap-Filling					
1	15.5%	\$81	13	\$522	110
2	***20.1% (-1.2pts)***	\$89	***4 (-21%)***	***\$168 (+27%)***	102
3	30.1%	\$105	***9 (+14%)***	***\$168 (-12%)***	101
4	52.8%	\$134	19	\$126	96
5	***71.0% (-0.7pts)***	\$148	***15 (-4%)***	***\$80 (+4%)***	***86 (+1)***
6	100.0%	\$172	***25 (+2%)***	***\$80 (-2%)***	78
Exclusion Gap-Filling					
1	15.5%	\$81	13	\$522	110
2	*n/a*	*n/a*	***0 (-100%)***	***\$168 (+27%)***	102
3	30.1%	\$105	***12 (+67%)***	***\$168 (-12%)***	***102 (+1)***
4	52.8%	\$134	19	\$126	96
5	*n/a*	*n/a*	***0 (-100%)***	***\$80 (+4%)***	85
6	100.0%	\$172	***40 (+67%)***	***\$80 (-2%)***	***85 (+7)***

Horizontal Gap-Filling

With the horizontal gap-filling approximation, the total sale probability $TP_{dcp}(k)$ for inefficient policies is decreased until the policy lies on the efficient frontier. AMD forecasting is directly affected by this process. While the airline does not adjust $TR_{dcp}(k)$, because of the change to TP_{dcp}, AMR adjusted fares also change. The AMR fares from horizontal gap-filling will be equal to those of vertical gap-filling because adjusted fares are dictated by the properties of the efficient booking policies in both cases.

The horizontal gap-filling approximations are:

$$TP'_{dcp}(k) = \begin{cases} TP_{dcp}(k) & \text{if } k \text{ is efficient} \\ TP_{dcp}(k-1) + \frac{TP_{dcp}(k+1) - TP_{dcp}(k-1)}{TR_{dcp}(k+1) - TR_{dcp}(k-1)} & \\ \qquad \times \left(TR_{dcp}(k) - TR_{dcp}(k-1)\right) & \text{otherwise} \end{cases}$$

$$TR'_{dcp}(k) = TR_{dcp}(k)$$

In our running example, shown in the bottom-left of Table 6.1, the change to total sale probabilities due to horizontal gap-filling is small—a decrease of 1.2 percentage points for FC 2, and of 0.7 points for FC 5. However, these small changes again have large impacts on AMD and AMR: adjusted demand decreases 21% for FC 2 and increases 24% for FC 3; there are also smaller changes for FC 5 and FC 6. The net result of horizontal gap-filling, in this example, is an increase by 1 in the FC 5 booking limit compared to the no gap-filling case. Recall that vertical gap-filling increased the FC 5 booking limit by 1, but also reduced the FC 2 booking limit.

Exclusion Gap-Filling

Exclusion gap-filling is the most mathematically-correct approach to dealing with inefficient policies when AMD and AMR are used as heuristic input modifiers for existing RM optimizers. Exclusion gap-filling is a multistage process that attempts to produce AMD demands and AMR fares that prevent the RM optimizer from selecting inefficient booking policies.

In the first stage, inefficient policies are discarded, and AMD demands and AMR fares are computed using only the efficient policies. Next, demands and adjusted fares are filled in for the inefficient policies: a demand of zero and an adjusted fare equal to the adjusted fare of the next efficient policy. This process ensures that, at the time of optimization with EMSR, the inefficient policies (in the case of this example, offering FC 1–2 or FC 1–5) will never be selected.

It is important to note that "FC 1–2" as an inefficient policy does not mean that consumers should be prohibited from booking in FC 2; when FC 1–3 are available (which is an efficient policy, as shown in Figure 6.1), consumers are free to choose to buy-up to FC 2 if they wish.

The bottom-right of Table 6.1 shows the effect of exclusion gap-filling on our running example. Note that the AMR fares are equal to those produced by vertical and horizontal gap-filling. However, the AMD demands differ from both methods; the exclusion gap-filling forecast has less high-fare demand and more lower-fare demand (note that all demand from FC 2 in the no gap-filling case gets moved to FC 3 with exclusion gap-filling; the same applies for FC 5 and 6). This leads exclusion gap-filling to have less aggressive booking limits than the no, vertical, and horizontal gap-filling: compared to no gap-filling, exclusion gap-filling increases the FC 6 booking limit by 7, and increases the FC 3 booking limit by 1.

The adjusted total sale probability and adjusted total expected revenue from any of the gap-filling methods will be used to compute AMR fares and AMD demands.

6.2.2 Demand Forecasting

Our AMD Forecasting Model provides an approach for estimating parameters for the distribution of demand volume, based on historical bookings. This model is an extension of Q-Forecasting (Hopperstad and Belobaba, 2004) to support generic fare structures and generic passenger choice models. We propose a four step process, in which we convert historical booking observations into estimates of historical demand volume, and then forecast future demand based on the historical volume estimates. Note that this process occurs after any gap-filling.

1. Convert observed (historical) bookings to equivalent "Q"-bookings, which represents the number of bookings that would have been received in the past if all fare classes and all ancillary services had been available. We assume that the airline has recorded which offer set was presented to the consumer for each booking. The equivalent Q-bookings for DCP *dcp* on previous departure day *dep* is given by:

$$qb_{dcp,dep} = \sum_{k=1}^{n_{FC}} \frac{b_{dcp,dep|k}}{TP'_{dcp}(k)} \tag{6.19}$$

   where n_{FC} is the number of fare classes, $b_{dcp,dep|k}$ is the number of bookings received in DCP *dcp* on previous departure date *dep* when FCC and NFO booking policy k was offered to consumers, and $qb_{dcp,dep}$ is the equivalent Q-bookings for DCP *dcp* on previous departure date *dep*, and serves as the estimate for total demand volume for that DCP and day.

2. Detruncate historical observations for any instances where all classes were closed:

$$\hat{qb}_{dcp,dep} = d(qb_{dcp,dep}, Z) \tag{6.20}$$

   where $\hat{qb}_{dcp,dep}$ is the detruncated (unconstrained) equivalent Q-bookings for DCP *dcp* on previous departure date *dep*, $d()$ is a detruncation function, and Z is a vector of other data or parameters for the detruncation process, including whether or not all classes were closed. For details on detruncation methods, see Lee (1990), Wickham (1995), Skwarek (1996), Weatherford and Pölt (2002), and Queenan et al. (2007).

3. Forecast future demand volume distribution parameters based on detruncated historical equivalent Q-demand for n_{dep} previous departure days:

$$\mu_{dcp} = \frac{1}{n_{dep}} \sum_{dep=1}^{n_{dep}} \hat{qb}_{dcp,dep}$$

$$\sigma^2_{dcp} = \frac{1}{n_{dep}-1} \sum_{dep=1}^{n_{dep}} (\hat{qb}_{dcp,dep} - \mu_{dcp})^2$$

where μ_{dcp} is the forecast future demand volume mean for DCP dcp and σ^2_{dcp} is the forecast future demand volume variance for DCP dcp .

4. Partition demand within each DCP to each booking policy:

$$\mu_{k,dcp} = \mu_{dcp}(TP'_{dcp}(k) - TP'_{dcp}(k-1))$$

$$\sigma^2_{k,dcp} = \sigma^2_{dcp}(TP'_{dcp}(k) - TP'_{dcp}(k-1))$$

For RM optimizers such as EMSR that utilize demand-to-come forecasts, instead of DCP-specific forecasts, the demands must be aggregated across DCPs. The forecast of all future demand to come, generated at the start of DCP dcp, for policy k, is:

$$\tilde{\mu}_{k,dcp} = \sum_{i=dcp}^{n_{DCP}} \mu_{k,i} \qquad \tilde{\sigma}^2_{k,dcp} = \sum_{i=dcp}^{n_{DCP}} \sigma^2_{k,i}$$

and the total future demand to come from the start of DCP dcp (aggregating across all booking policies) is:

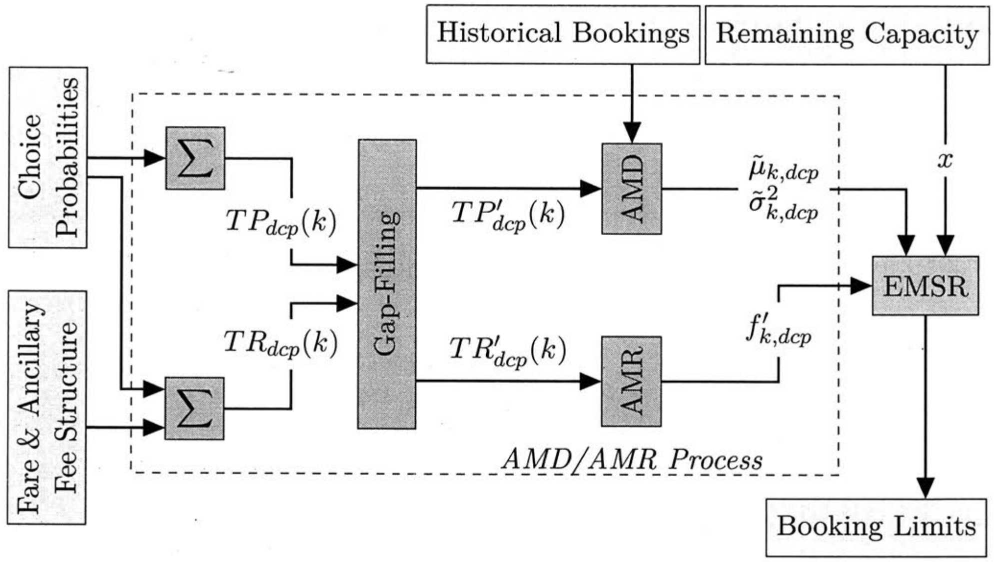

Figure 6.2: AMD/AMR process diagram when used as a heuristic in conjunction with EMSR. Inputs are shown in grey, AMD/AMR processes and operations in blue, the RM optimizer in red, and the RM output in green.

$$\tilde{\mu}_{dcp} = \sum_{k=1}^{n_{FC}} \tilde{\mu}_{k,dcp} \qquad \tilde{\sigma}^2_{dcp} = \sum_{k=1}^{n_{FC}} \tilde{\sigma}^2_{k,dcp}$$

We use these aggregated-across-DCP demand parameters ($\tilde{\mu}_{k,dcp}$ and $\tilde{\sigma}^2_{k,dcp}$) in our implementation of AMD and AMR with EMSR.

6.2.3 Process Summary

An overview of the complete AMD and AMR methodology, when used as a heuristic with EMSR, is illustrated in Figure 6.2. The process requires input choice probabilities ($P_{i,m,dcp}(k)$) for each of the fare class complete/nested by fare order booking policies, fares and ancillary prices, and historical bookings. The process returns AMR fares and AMD demand estimates (mean and variance) for each class, generated at the start of each DCP. These adjusted fares and demands are fed to the RM optimizer (along with remaining flight capacity), which then returns a booking policy in the form of booking limits.

6.3 Use in Larger Networks

The ACDP model and AMD/AMR heuristics described in the preceding sections are designed for a single-airline, single-leg environment. With multiple airlines and/or multiple flight legs, the ACDP formulation becomes much more complicated:

- A network extension of ACDP, as with other dynamic programs, would suffer from "exploding dimensionality," meaning the size of the state space increases with the number of possible itinerary/fare class combinations, which grows much faster than just the number of flight legs.
- The choice probabilities $P_{kmt}(O)$ in Equation 6.3 would need to be redefined. As discussed in Chapter 3 of this thesis, the probability that a consumer chooses a particular itinerary/fare class/ancillary service is conditional on the set of options available to the consumer. A true network extension of ACDP would account for these choice dependencies between different itineraries and between the airline and its competitors—the set O would include information not just on the fare class/ancillary service combinations to offer for a particular itinerary, but also on similar itineraries offered by the airline and its competitors.

Fully extending ACDP to address these issues is beyond the scope of this thesis and is left to future work. However, we will consider a more limited application of the AMD/AMR heuristics in larger networks. Our simple approach relies on combining leg and network optimization, as with other origin-destination RM heuristics. We propose forecasting demand with AMD and adjusting fares with AMR for each itinerary/fare class in isolation (i.e. without including multiple itineraries in O); we then feed the AMD demands and AMR fares to a network optimization heuristic (such as DAVN). While this approach does not account for choice dependencies between itineraries or between airlines, it is easy and computationally tractable. In general, revenue management forecasting and optimization models do not consider demand dependencies between different itineraries and/or flights.

6.4 Estimating Choice Probabilities

Our model description has, thus far, assumed that the airline has accurate knowledge of a customer choice model and can compute the choice probabilities $P_{i,k,dcp}(k)$. In reality, however, airlines would need to estimate these probabilities. The airline could estimate the parameters of a customer choice model, as in Appendix A, and then apply the model to the airline's fare structure to compute the probabilities. This approach has the advantage of providing a behavioral explanation for customer decision making, and would allow the airline to understand how changes in the fare or ancillary fee structure would affect customer booking choices.

An alternative approach, which is potentially simpler to implement, would be to directly estimate choice probabilities from historical booking data, without attempting to devise a functional form to explain why customer make the choices that have been observed. We propose two forms of this simple approach, incorporating different levels of aggregation and external information. In both cases we assume that the airline will use a long-term offline calibration process (i.e. incorporating data from month or years of departure days, with the estimation process performed infrequently) due to substantial complications of missing observations/sparse data.

6.4.1 Raw Estimates

Our first proposal, which we refer to as "raw estimates," directly estimates the probabilities $P_{i,m,dcp}(k)$ from historical booking data with minimal processing or adjustment:

$$\hat{P}_{i,m,dcp}(k) = \begin{cases} \frac{\sum_{\forall dep} b_{i,m,dcp,dep|k}}{\sum_{\forall dep} h_{dcp,dep|k}} & \sum h \neq 0 \\ 0 & \text{otherwise} \end{cases} \tag{6.21}$$

where $\hat{P}$ is the estimate of P, $b_{i,m,dcp,dep|k}$ is the number of bookings received in class i with ancillary option m during DCP dcp for previous departure day dep when class k was the lowest available fare class, and $h_{dcp,dep|k}$ is the number of *booking requests* the airline received during DCP dcp for previous departure day dep when class k was the lowest available fare class. Note that the estimate aggregates across "all" previous departure days, denoting that the historical data utilized is much greater than that used by a demand forecasting model.

6.4.2 Heuristic Estimates

The raw estimate described above utilizes a simple, straightforward procedure, but in practice could have several significant flaws. First, even utilizing data from many previous departure days, there may be very few (or zero) booking requests for some DCP/lowest available class combinations. For example, FC 1 is almost never the lowest available class at the start of the booking window (DCP 1), so $\hat{P}_{i,m,1}(1)$ will have a small sample size (and therefore high variance), or will have no data and therefore be 0. Second, the formulation is not market specific, so in a multi-leg network variations in choice probabilities caused by variations in fare ratios will be averaged out, and choice probabilities will be biased for any particular market.

Third, and potentially most problematic, the approach requires observing *booking requests*, which are created not just for each booking, but also for each consumer who searches and decides not to book. Although airlines could track the number of booking requests processed by their own systems, consumers shopping via third parties who do not end up booking may not generate a booking request to the airline (the request could be answered by the third party's cache of previous shopping results). Alternatively, one consumer shopping on a "meta-search" engine (such as Kayak or Google Flights) could trigger multiple booking requests to the airline (because each party contacted by the meta search engine could pass the request to the airline). Airlines in 2015 experienced more than 1,000 booking requests for each actual booking received, and expected that to rise to 10,000 or more requests per booking as more consumer use meta search engines (tnooz, 2015). Finally, some airlines already have estimates of price elasticity (used for hybrid forecasting and fare adjustment), which are not incorporated in the raw estimate described above.

To address these challenges, we propose a second estimation approach, which we refer to as the "heuristic estimate." The airline separates the estimate $\hat{P}_{i,m,dcp}(k)$ into two components, a *sell-up estimate* (a), or the probability that a consumer will book anything, and a *conditional purchase estimate* (b), or the probability that a consumer will book a particular offer, given that they book something from the offer set:

$$\hat{P}_{i,m,dcp}(k) = \underbrace{\hat{P}_{dcp}(\text{Book} \mid k)}_{\text{(a)}} \cdot \underbrace{\hat{P}_{i,m,dcp|\text{Book}}(k)}_{\text{(b)}} \tag{6.22}$$

This separation allows the airline to use an external sell-up estimate for the first component, decreasing the challenges of estimation with small numbers and missing observations, and eliminating the need to estimate the "no-go" alternative. The second component is estimated based on historical purchases and requires no assumptions about the no-go alternative. The external sell-up estimate could take many forms, but a common approach (used by Q-forecasting and fare adjustment) is an exponential curve:

$$\hat{P}_{dcp}(\text{Book} \mid k) = \exp\left(\frac{-\ln(2)}{\varepsilon_{dcp} - 1}\left(\frac{f_k}{f_Q} - 1\right)\right) \tag{6.23}$$

where ε_{dcp} is the (externally estimated) price elasticity multiplier, and f_Q denotes the lowest published fare in the market (thus $f_k \geq f_Q \; \forall k$). The conditional purchase estimate is estimated from the historical booking database:

$$\hat{P}_{i,m,dcp|\text{Book}}(k) = \begin{cases} \dfrac{\sum_{\forall dep} b_{i,m,dcp,dep|k}}{\sum_{\forall dep} b_{dcp,dep|k}} & \sum b \neq 0 \\ \tilde{P}_{i,m,dcp|\text{Book}}(k) & \text{otherwise} \end{cases}$$

where $b_{dcp,dep|k}$ is the total number of bookings received during DCP dcp for departure day dep when class k was the lowest available fare class.

We apply an additional processing step of "neighbor matching" for instances where no bookings have been received ($\sum_{\forall dep} b_{dcp,dep|k} = 0$). In those cases, we use the conditional purchase estimate for the same booking policy from the nearest DCP dcp^* that has received bookings in the past:

$$\tilde{P}_{i,m,dcp|\text{Book}}(k) = \hat{P}_{i,m,dcp^*|\text{Book}}(k)$$

$$dcp^* = \arg\min_{dcp'} \left\{ |dcp' - dcp| \text{ s.t. } \sum_{\forall dep} b_{dcp',dep|k} > 0 \right\}$$

The heuristic estimation method allows the airline to aggregate multiple markets in estimating conditional purchase probabilities, but also allows the airline to integrate market-specific fares in the sell-up estimate to produce choice probabilities that are less susceptible to missing observations than the raw estimates while still scaling to market-specific conditions. The performance of these estimation approaches will be tested via simulation in Sections 7.4 and 7.5.

6.5 Conclusions

In this chapter, we have developed a new dynamic programming model for total revenue optimization that incorporates fares, ancillary revenues, and passenger choices. The model produces an optimal set of offers, which we define as a fare class and a combination of ancillary services, to be presented at any given time. Following previous work on choice-based RM, we use the Ancillary Marginal Demand transformation and Ancillary Marginal Revenue transformation to convert our ancillary and choice-aware DP into an equivalent independent demand formulation. After addressing practical distribution constraints, we devised a series of processes to utilize AMD and AMR as total revenue optimization heuristics in conjunction with existing RM optimizers, and we developed the Ancillary Marginal Demand forecasting model to provide demand volume estimates.

Numerous extensions and enhancements to our work are possible. For example, while exact network formulations for dynamic programs suffer from exploding dimensionality, our heuristics could be more formally extended to support a network setting. A key challenge will be determining the level of detail and specificity necessary in the input choice probabilities to maintain reasonable revenue performance; the extent to which these probabilities could be aggregated and/or scaled across different markets and fare structures is unknown. In addition, competitive effects in a network setting raise questions about the degree to which competitor offerings should be explicitly incorporated in the model: failing to account for competitor offerings could lead to availability decisions that are too aggressive, while explicitly modeling competitor actions increases data and computation requirements. Further work is required to develop more efficient choice probability estimation methods, and to understand how inaccuracies in the input choice probabilities affect revenue performance.

Finally, the rise of New Distribution Capability could allow airlines to have significantly more control over the offers they produce. We believe that this work could be extended to generate offers based on filed fares and prices. A separate, potentially larger challenge, would be devising *dynamic* offer generation engines that also incorporate a dynamic pricing aspect.

Chapter 7

Simulated Performance of AMD and AMR

In the previous chapter we developed the Ancillary Choice Dynamic Program (ACDP) for total revenue management in a single-airline, single-leg environment.[1] We then described two heuristics, the Ancillary Marginal Demand and Ancillary Marginal Revenue transformations (AMD and AMR), for utilizing the dynamic program. In this chapter, we test the performance of the heuristics via simulation in the Passenger Origin-Destination Simulator (PODS) described in Section 3.2. We examine how AMD and AMR affect booking and ancillary purchase patterns across a range of demand and passenger preferences.

In the first section, we consider a single-airline, single-leg network (A1ONE) and show that AMD and AMR increase revenue over a variety of existing RM approaches. Then, we examine the performance of the heuristics when competition is introduced, with a two-airline, two-flight-leg network (A2TWO). A2TWO provides a platform for assessing the impacts of

[1] Portions of this chapter were previously awarded the 2018 Anna Valicek bronze medal as Bockelie, A. and Belobaba, P. (2018). Total revenue optimization with the Ancillary Marginal Demand and Ancillary Marginal Revenue transformation heuristics. Presented at the 58th Annual AGIFORS Symposium, Tokyo, Japan.

different competitive strategies on AMD and AMR. In all cases we compare the performance of AMD and AMR against existing methods for total revenue optimization. Although both A1ONE and A2TWO are small networks that significantly simplify the airline RM problem, they provide useful insights into the interactions between RM optimization and forecast, and between competitors, when AMD and AMR is used.

In a final set of tests we utilize the much larger network D6, which features two airlines with hundreds of flight legs and markets. D6 illustrates the performance of the heuristics in a complex environment with connecting flights, and where the ability to scale choice probabilities across markets becomes important.

7.1 AMD and AMR Implementation within PODS

A version of the AMD and AMR heuristics were programmed into PODS, supporting one ancillary service with booking policies that are both fare class complete and nested by fare order (see Section 6.1.1). The airline's marketing policy may provide the ancillary complimentary or a la carte in any fare class, but no class is prohibited from purchasing the ancillary. One set of choice probabilities $P_{i,m,dcp}(k)$ is entered per airline; the airline uses those probabilities for all itineraries in all markets (unless the choice probability heuristic is activated as described below). Because choice probabilities can vary by DCP, the airline performs its gap-filling computations for each DCP. All three gap-filling mechanisms described in Section 6.2.1 are supported.

PODS does not include an inline choice probability estimation process, meaning that in the course of a simulation, airlines cannot estimate choice probabilities and use those estimates in the same simulation. Instead, PODS supports a simple batch estimation process: at the end of a simulation, each airline reports the portion of all generated passengers who booked a particular fare class with or without the ancillary service, conditional on offered booking policy and DCP. These estimates are aggregated over all markets and provide estimates of

$P_{i,m,dcp}(k)$ using Equation 6.21, where "$\forall dcp$" means all trials and all un-burned samples. The estimated probabilities are then used as inputs for subsequent simulations. The same simulation reporting is used to generate the estimates $\hat{P}_{i,m,dcp|\text{Book}}(k)$ if the airline is to use the heuristic estimation process (Section 6.4.2). For the heuristics, the price elasticity multipliers ε_{dcp} input from a standard PODS FRAT5 curve, matching the price elasticity values used for hybrid forecasting and fare adjustment. PODS computes the $\hat{P}_{dcp}(\text{Book} \mid k)$ terms (Equation 6.23) separately for each market, based on the fares of the market; the values are then combined with the input conditional purchase probabilities to produce market-specific estimates of choice probabilities.

7.2 Result Analysis Methodology

As in Chapter 5, we measure the performance of RM forecasting and optimization methods relative to a *baseline* simulation, where all airlines use standard (independent demand) forecasting and optimization models that do not account for ancillary services or revenues.

Each PODS *simulation* consists of 2 to 5 independent *trials* of 400 *samples.* Each sample represents one realization of the "same" departure day (meaning the demand and passenger characteristics for all samples are drawn from the same distribution). The first 200 samples of each trial are used to warm up the forecasting models and are "burned" and **never** included in any reported results, leaving 400 samples per trial. The seeds for all random draws (i.e. the number of passengers generated and individual passenger preferences) are the same for all simulations with the same demand parameters, so the difference in revenue or bookings between two simulations with different forecasting or optimization methods is due to the change in forecasting or optimization, not due to different demand generation.

Our analytical approach in this chapter matches that of Chapter 5. The primary item of interest in our studies is the change in total revenue due to a change in forecasting or optimization method, which we measure as the average across all (unburned) samples of the sample-specific change in total revenue for a test simulation vs a baseline simulation. Mathematically, if X_i^j is the revenue (or other simulation output) for sample i for simulation $j \in \{\text{TEST}, \text{BASE}\}$, we are interested in the term $\bar{\Delta}$:

$$\bar{\Delta} = \frac{1}{n}\sum_{i=1}^{n}\hat{\Delta}_i \qquad\qquad \hat{\Delta}_i = X_i^{\text{TEST}} - X_i^{\text{BASE}}$$

Unless otherwise stated, all references to statistical significance and confidence intervals in this chapter are derived from a paired t-test (or a one sample t-test on the change in revenue by sample), with the null hypothesis that there is no true change in revenue $(H_0 : \Delta = 0)$ vs alternative hypothesis that there is a true change in revenue $(H_a : \Delta \neq 0)$. The test statistic t is:

$$t = \frac{\bar{\Delta}}{\text{se}_{\bar{\Delta}}} \sim T_{df=n-1}$$

where $\text{se}_{\bar{\Delta}}$ is the standard error of $\bar{\Delta}$.

The demand forecasting process nominally introduces a dependency between demand generated (and therefore total revenue) in one sample and booking limits (and therefore total revenue) in another sample. These dependencies potentially violate the t-test assumption of independence between samples; however, as discussed in Appendix C, our testing indicates that these correlations are minor and therefore should not affect the statistical significance of our results.

7.3 One Airline, One Flight Leg Network (A1ONE)

We utilize the single airline, single flight leg network A1ONE within PODS for our initial studies. A1ONE is a simple network, and because it lacks competition it closely matches

Table 7.1: Network A1ONE and A2TWO fare and ancillary fee structure.

		Advanced	Restriction Applies?		
	Fare	purchase	R1	R2	R3
FC 1	\$500	None	-	-	-
FC 2	\$390	3 days	-	-	Yes
FC 3	\$295	7 days	-	Yes	Yes
FC 4	\$200	10 days	Yes	-	Yes
FC 5	\$160	14 days	Yes	Yes	-
FC 6	\$125	21 days	Yes	Yes	Yes

the ACDP assumptions. The only significant assumption violations are that demand is unknown (and must be forecast) and that the underlying demand generation process is not Poisson, but is instead based on a series of normal distributions (as described in Section 3.2). The flight has a capacity of 130 seats. The airline offers one optional ancillary service and six economy fare classes, denoted FC 1 (the most expensive and least restricted, \$500) through FC 6 (the least expensive, most restricted, and subject to advance purchase requirements, \$125), as shown in Table 7.1. The airline divides its booking window into 16 DCPs, with the end of DCP 16 corresponding to departure. Forecasts are generated and booking limits are re-optimized at the start of each DCP.

This network features two consumer demand segments, business and leisure. Business passengers tend, although do not always, to have higher budgetary constraints, to book closer to departure, and to be more averse to fare class restrictions (such as a Saturday-night stay, or non-refundability). The average business passenger and leisure passenger booking curves, as well as the business/leisure mix, are shown in Figure 7.1. Note that the early DCPs have a low proportion of business travelers shopping, while shoppers in the later DCPs are predominantly from the business segment.

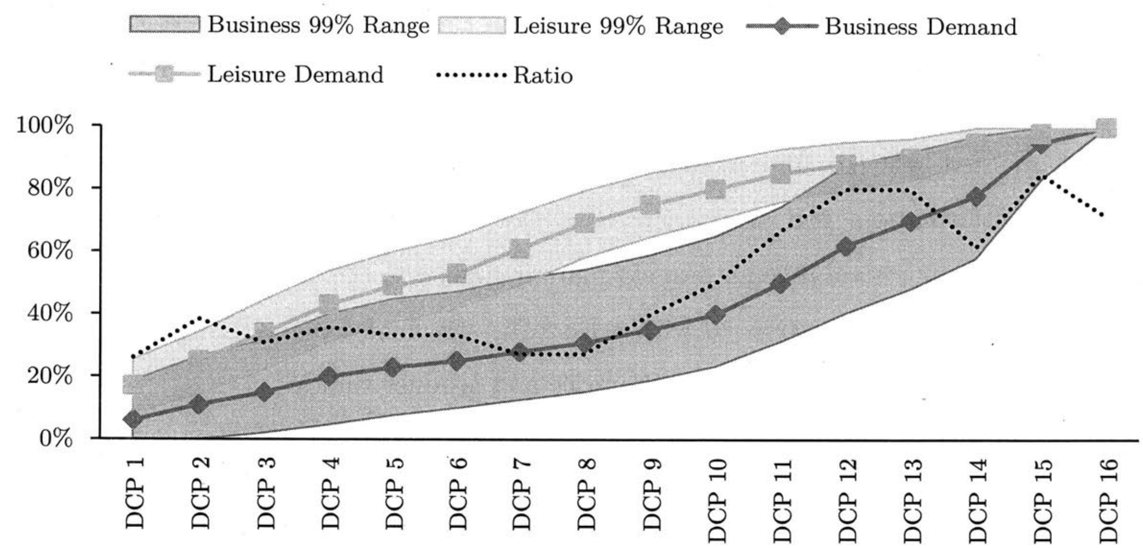

Figure 7.1: Cumulative average business and leisure consumer arrival curves, and the ratio of average business to leisure arrivals within each DCP. Ratio of average business to leisure arrivals within each DCP indicated in dotted line. Shaded region indicates 99% range of all realized arrival curves.

7.3.1 Experimental Outline

In all of our A1ONE simulations, the airline uses EMSR as its RM optimizer. As a baseline, we consider the case where the airline optimizes based on filed fares and uses an independent-demand forecasting model (see Belobaba and Weatherford (1996) and Littlewood (1972)). For AMD and AMR, we assume that the airline has accurately estimated a passenger choice model, and can therefore compute the probabilities $P_{i,m,dcp}(k)$ as described in Equation 6.10. Unless otherwise noted, the airline will employ Exclusion Gap-Filling when using AMD and AMR.

We will compare the performance of AMD and AMR against the baseline as well as three existing approaches for accounting for ancillary revenue and/or passenger choice: the optimizer increment (OI), hybrid forecasting and fare adjustment (HF/FA), and a combination of the two approaches (OI + HF/FA). In our implementation of OI, the airline estimates

$\bar{a}_{k,dcp}$ based on historical purchases aggregated across all DCPs. Recall that hybrid forecasting and fare adjustment are operationalized versions of the (non-ancillary) marginal revenue and marginal demand transformations, where demand is divided into two groups: product-oriented, which is forecasted with an independent demand model and no marginal revenue transformation, and price-oriented, which is forecasted with a marginal demand model and has a marginal revenue transformation applied to the fares (based on a negative exponential sell-up curve expressed with a FRAT5); the final demand and fare values sent to the optimizer are a combination of the price and product values.

We focus our detailed assessment of AMD and AMR on a representative case in which the ancillary service has a price of \$50, and both consumer segments have a mean disutility of forgoing the service of \$50. We assume initially that passengers in the simulation are aware of ancillary prices and incorporate ancillary prices and preferences into their fare class decision process ("simultaneous" behavior, as defined in Chapter 3).

The total probabilities $TP_{dcp}(k)$ computed by the airline are shown in Figure 7.2. Early in the booking window, the airline calculates that the probability of sell-up is low (as indicated by the low total sale probability for FC 1, and reflecting the low proportion of business travelers). However, later in the booking window, when the portion of business passengers is higher, the total sale probability for higher-value classes increases, suggesting that more sell-up is possible. Note that the total sale probability for FC 6 is always 100%, reflecting that all generated passengers in the simulation can afford to purchase the lowest published fare.

7.3.2 Initial Results

Table 7.2 lists booking and ancillary purchase data for the baseline simulation. In the baseline scenario, about 33% of passengers purchase the ancillary service, with a much higher purchase rate in the higher value fare classes (46% for FC 1) than in the lower value

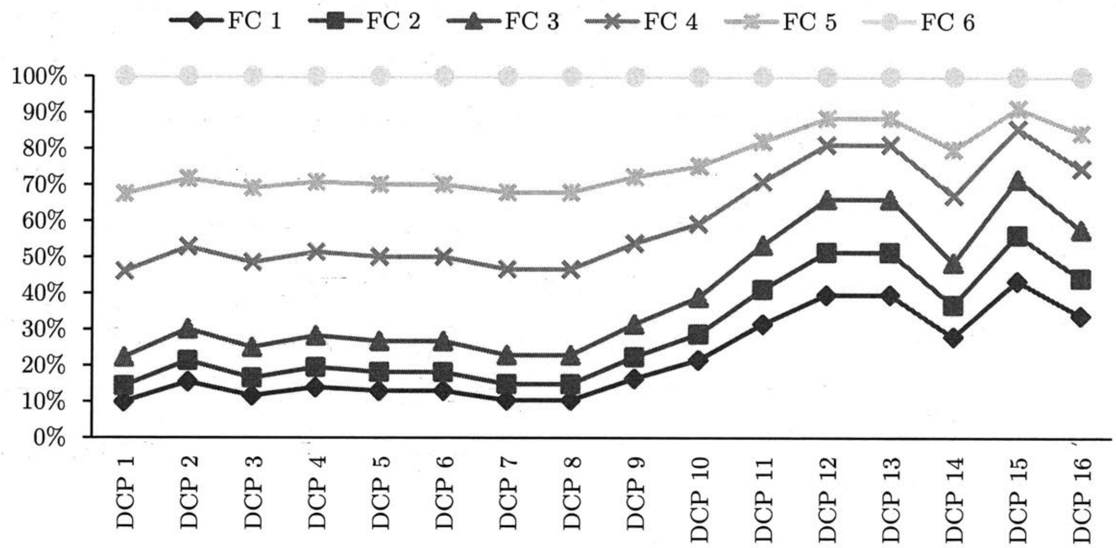

Figure 7.2: Total sale probability $TP_{dcp}(k)$ for each class as computed by the airline prior to any gap-filling (100% simultaneous passengers, \$50 mean ancillary disutility, \$50 ancillary price).

fare classes (25% in FC 6); likewise, the average ancillary revenue per booking is highest in FC 1 (\$23) and lowest in FC 6 (\$12). Despite the lower ancillary purchase rate and lower average ancillary revenue, because fares are lower in FC 6 than FC 1, the portion of total revenue derived from ancillary sales is highest in the lower value classes (11% in FC 5) and lowest in higher value classes (4% in FC 1). Recall that major US airlines report about 8% of total revenue from ancillary services, according to the US Department of Transportation.[2]

The lower ancillary purchase rate in the lower value classes is driven by a fundamental behavioral assumption in the Simultaneous choice model: passengers have an overall budgetary constraint that limits their spending on the combination of fare and ancillary services. Passengers booking in the lower value (and highly restricted) fare classes tend to have lower budgets, which constrains their ability to afford ancillary services. These basic ancillary purchase and revenue trends in the baseline case are similar in the other experimental cases.

[2] US DOT Form 41, Schedule P-1.2

Table 7.2: Network A1ONE baseline bookings and ancillary purchase data by fare class (medium demand, 100% simultaneous passengers, $50 ancillary price, $50 ancillary disutility).

	FC 1	FC 2	FC 3	FC 4	FC 5	FC 6
Bookings	6	14	10	8	16	56
Booking Mix	5%	12%	9%	8%	15%	52%
Average Ancillary per Passenger	$23	$22	$21	$22	$19	$12
Portion of Total Revenue from Ancillary	4%	5%	7%	10%	11%	9%
Ancillary Sales Rate	46%	44%	41%	44%	39%	25%

The total revenue, load factor, and yield (expressed as revenue per passenger mile) for each of the experimental cases (with AMD and AMR using exclusion gap-filling) are shown in Table 7.3. AMD and AMR produces a revenue increase of 1.8% over baseline; HF/FA produces a gain of 1.2%, 0.6 pts less than AMD and AMR. Both AMD and AMR and HF/FA decrease load factor by 1.3 pts, and both increase total yield, although the increase is larger for AMD and AMR (3.4%) than for HF/FA (2.7%).

The optimizer increment has a small negative effect on revenue: an 0.1 pt decrease compared to baseline, and a reduction of the benefit of HF/FA by 0.1 pt. OI increases load factor by 0.1 pt when used alone, and decreases the load factor loss due to HF/FA by 0.1 pt. Although the revenue and load factor changes due to OI are small, they are directionally consistent with the results seen in numerous studies within the MIT PODS Research Consortium (e.g. Chapter 5).

These revenue and load factor changes are driven by shifts in the booking mix, as shown in Figure 7.3. Both HF/FA and AMD/AMR reduce bookings in the lowest value class (FC 6), while increasing bookings in the highest value class (FC 1). The methods differ in the magnitudes of changes: AMD/AMR reduces bookings in FC 6 by about 2, while HF/FA reduces by 5. Changes in bookings in higher value fare classes have a disproportionate effect on total revenue: recall that FC 1 has a fare of $500, and FC 6 has a fare of $125, so each FC 1 booking is worth four times the ticket revenue of an FC 6 booking. The FC 1

Table 7.3: Network A1ONE simulation results for baseline and experimental cases (medium demand, 100% simultaneous passengers, \$50 ancillary price, \$50 ancillary disutilities).

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
Ticket Revenue	\$22,098	\$22,068	\$22,383	\$22,372	\$22,548
Ancillary Revenue	\$1,808	\$1,808	\$1,798	\$1,799	\$1,788
Total Revenue	\$23,906	\$23,876	\$24,181	\$24,171	\$24,336
Load Factor	83.8%	83.9%	82.5%	82.6%	82.5%
Total Yield	21.95	21.90	22.54	22.50	22.70
Ancillary Sales Rate	33.2%	33.2%	33.5%	33.5%	33.4%
Change from Baseline					
Ticket Revenue		-0.1%	+1.3%	+1.2%	+2.0%
Ancillary Revenue		+0.0%	-0.6%	-0.5%	-1.1%
Total Revenue		-0.1%	+1.2%	+1.1%	+1.8%
Load Factor		+0.1 pts	-1.3 pts	-1.2 pts	-1.3 pts
Total Yield		-0.2%	+2.7%	+2.5%	+3.4%
Ancillary Sales Rate		-0.0 pts	+0.3 pts	+0.3 pts	+0.2 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.0%	0.1%	0.1%	0.1%
t-statistic		-7.69	11.43	11.61	13.10
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

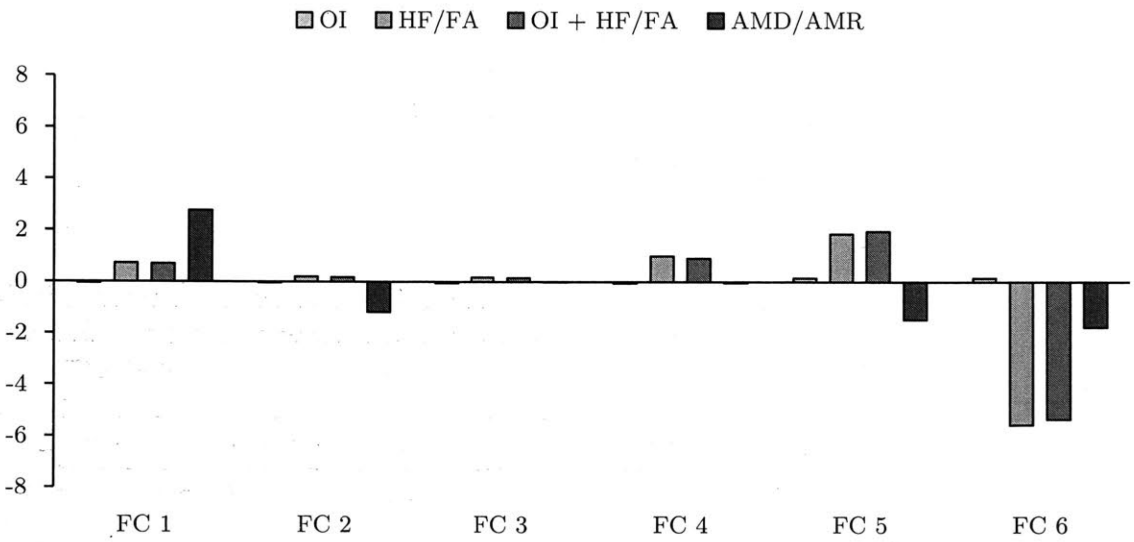

Figure 7.3: Change in bookings by fare class vs baseline (medium demand, 100% simultaneous passengers, \$50 mean ancillary disutility, \$50 ancillary price).

booking increase with AMD/AMR is worth six times the revenue loss associated with the FC 6 booking decrease.

While both AMD/AMR and HF/FA reduce FC 6 bookings (which saves space for later arriving, higher value FC 1 bookings), only AMD/AMR also reduces FC 5 bookings; HF/FA leads to an increase in FC 5 (as well as FC 4). The booking changes by fare class due to OI are minimal, and the changes with OI + HF/FA are approximately equal to the sum of the changes in the OI case and in the HF/FA case.

The initial demand and ticket revenue forecast (generated at the start of DCP 1) is shown in Figure 7.4. Compared to the baseline, AMD provides a lower demand forecast mean (by 15%), but a higher demand forecast standard deviation (by 33%). Overall, the probability that demand exceeds the aircraft capacity of 130 (a rough indicator of whether the capacity constraint should restrict availability) is 24% for AMD, a reduction from the 39% of the baseline or 34% of HF/FA. Despite the lower volume of demand with AMD, however, the value of demand with AMD is greater because the composition of the AMD forecast is shifted toward FC 1. AMD produces an initial forecast of ticket revenue (computed as a sum of demand for each class multiplied by the fare for each class) 13% higher than the baseline, with a standard deviation 65% higher. A higher value forecast, for the same demand mean and standard deviation and same optimizer fares, will lead to more aggressive availability decisions and fewer low-value bookings.

In addition to the higher value forecast, the AMR adjusted fares are lower than filed fares, especially for the low value fare classes. The higher value forecast of AMD combined with the reduced optimizer fares of AMR reduces availability of lower-value fare classes, as shown in Figure 7.5. Reducing lower class availability forces consumers to buy-up to higher value classes, reducing load factor and increasing yields (and in this case increasing total revenue). By explicitly accounting for ancillary revenue and passenger choices, AMD/AMR can more precisely close classes. Note that AMD/AMR has a smaller reduction in FC 6 availability than HF/FA, but a larger reduction in FC 5 availability (especially in DCP 7 and 8), leading to the booking shifts seen in Figure 7.3.

7.3.3 Effect of Gap-Filling

Table 7.4 lists performance data for each of the gap-filling mechanisms. Exclusion gap-filling produces a revenue increase of 1.8% over baseline; AMD/AMR with the other gap-filling methods has a smaller revenue gain, with the worst revenue performance when gap-filling is not used. Load factor changes with AMD/AMR are inverse to total revenue changes: the largest revenue gain (exclusion gap-filling) has the smallest load factor loss (-1.3 pts), while the lowest revenue gain (1.4%, no gap-filling) has the largest load factor loss (-3.1 pts). All of the gap-filling methods have higher revenue than HF/FA.

As shown in Figure 7.6, exclusion gap-filling, which has the highest revenue and is most mathematically correct, has the smallest reduction on FC 6 bookings; the other gap-filling methods all have FC 6 booking reductions of similar magnitude to HF/FA. In addition,

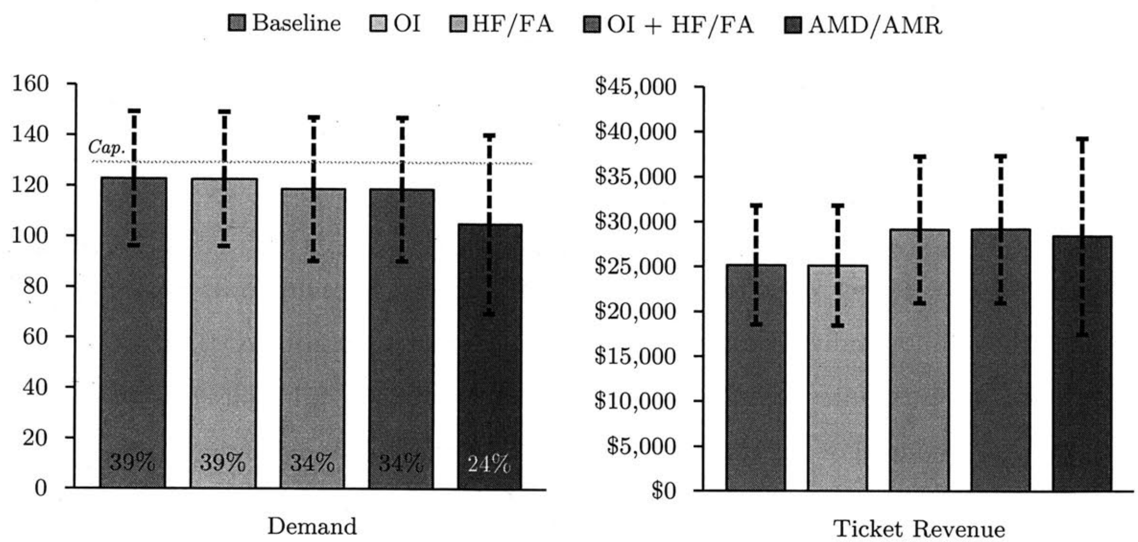

Figure 7.4: Initial (DCP 1) demand and ticket revenue forecasts (medium demand, 100% simultaneous passengers, $50 mean ancillary disutility, $50 ancillary price). Solid bars show forecast mean; dashed error lines show forecast standard deviation. The forecast probability that demand exceeds capacity (of 130) is indicated at the base of the solid bars.

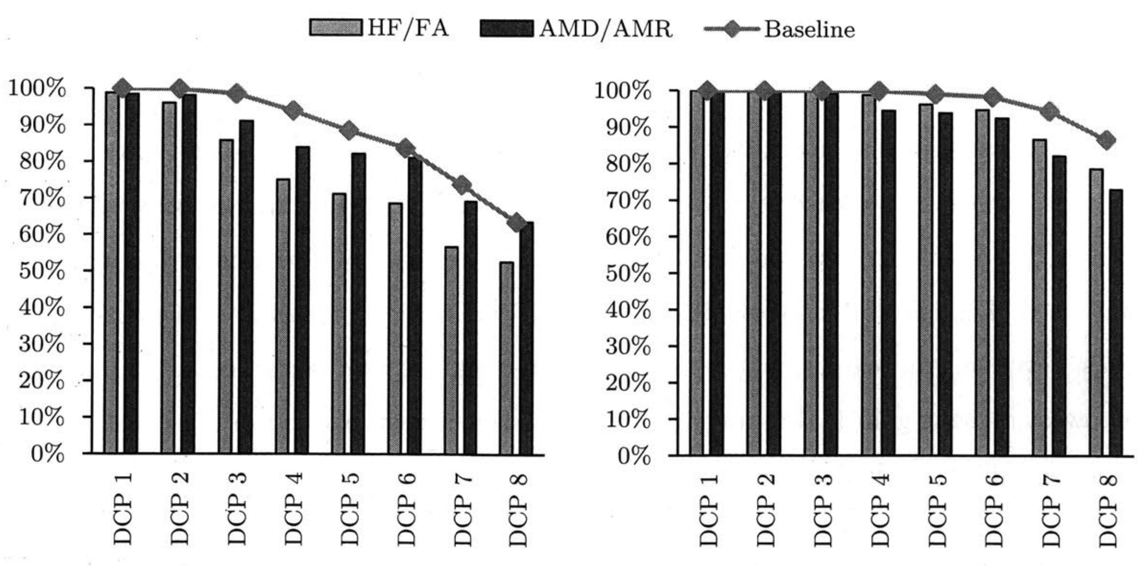

Figure 7.5: Availability for first 8 DCPs (medium demand, 100% simultaneous passengers, $50 mean ancillary disutility, $50 ancillary price). Measured as the portion of time a class is available for sale for FC 6 (left) and FC 5 (right).

Table 7.4: Network A1ONE simulation results for baseline and AMD and AMR cases with various gap-filling mechanisms (medium demand, 100% simultaneous passengers, $50 ancillary price, $50 ancillary disutilities).

	Baseline	None	Vertical	Horizontal	Exclusion
Ticket Revenue	$22,098	$22,488	$22,523	$22,540	$22,548
Ancillary Revenue	$1,808	$1,760	$1,776	$1,781	$1,788
Total Revenue	$23,906	$24,248	$24,299	$24,321	$24,336
Load Factor	83.8%	80.7%	81.5%	81.9%	82.5%
Total Yield	21.95	23.12	22.92	22.85	22.70
Ancillary Sales Rate	33.2%	33.6%	33.5%	33.5%	33.4%
Change from Baseline					
Ticket Revenue		+1.8%	+1.9%	+2.0%	+2.0%
Ancillary Revenue		-2.7%	-1.8%	-1.5%	-1.1%
Total Revenue		+1.4%	+1.6%	+1.7%	+1.8%
Load Factor		-3.1 pts	-2.2 pts	-1.9 pts	-1.3 pts
Total Yield		+5.3%	+4.4%	+4.1%	+3.4%
Ancillary Sales Rate		+0.4 pts	+0.3 pts	+0.3 pts	+0.2 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.2%	0.2%	0.1%	0.1%
t-statistic		8.61	10.80	11.75	13.10
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

exclusion gap-filling is the only method that reduces FC 5 bookings; the other methods trade large FC 6 losses for smaller FC 5 gains. This is an expected result; as shown in Table 6.1, exclusion gap-filling protects no seats for FC 5, and therefore FC 5 and FC 6 have the same booking limit. In that example, exclusion gap-filling has an FC 6 booking limit 7 seats greater than any of the other gap-filling methods, include no gap-filling. Despite the greater booking limit for FC 6, though, exclusion gap-filling has the second greatest increase in FC 1 bookings vs baseline (the additional space to accommodate FC 1 customers is provided by accepting fewer bookings in FC 5 and 4 compared to the other gap-filling mechanisms).

The differences in fare class booking changes amongst the gap-filling methods illustrate that no gap-filling is the most aggressive form of AMD/AMR, followed by vertical gap-filling, then horizontal gap-filling, and finally exclusion gap-filling. This variation in aggressiveness is a function of the AMD forecast generated by each approach: recall that vertical, horizontal, and exclusion gap-filling all have the same AMR fares. Horizontal gap-filling partially reduces demand forecasts for inefficient policies (and increases the forecast for the next efficient policy); exclusion gap-filling completely eliminates inefficient policy forecasts, and shifts all demand to the next efficient policy. Thus, exclusion gap-filling will always have less aggressive availability, and will accept more FC 6 bookings than the other gap-filling methods. The increase in FC 6 availability (relative to other gap-filling) means less space is protected for FC 5 bookings, which produces the decrease in FC 5 (relative to both the baseline and other gap-filling methods) seen in Figure 7.6.

7.3.4 Sensitivity to Ancillary Prices and Disutilities

In this section we assess the sensitivity of our results to the input ancillary prices and ancillary disutility parameters. We vary ancillary prices from \$25 to \$100, with mean passenger ancillary disutilities equal to 75%, 100%, or 125% of the ancillary price. We consider three combinations of disutilities, listed in Table 5.5: a leisure-oriented service, where leisure passenger mean disutility is 125% of the ancillary price and business passenger

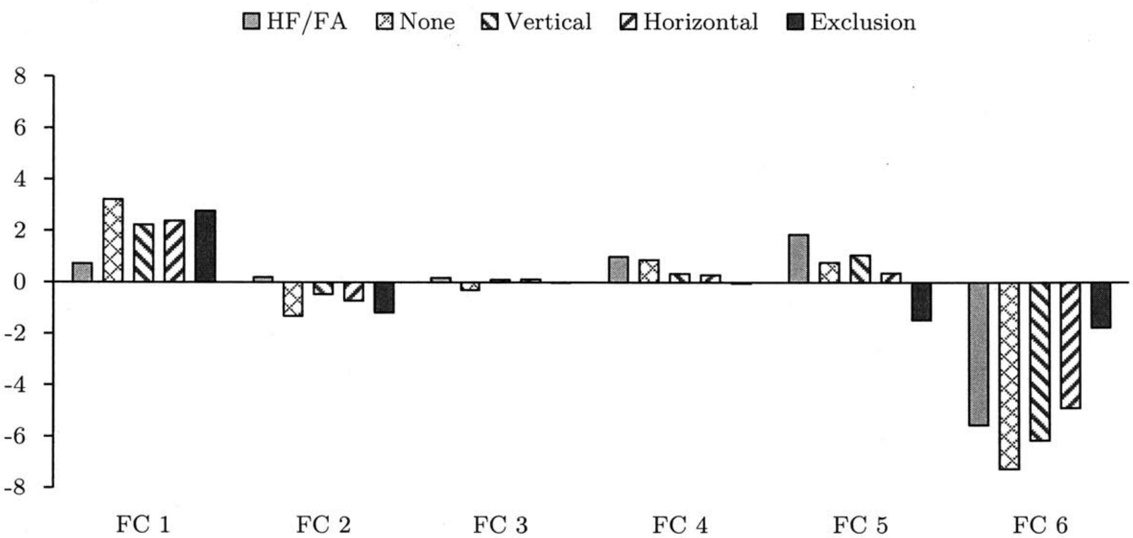

Figure 7.6: Change in bookings by fare class vs baseline with various gap-filling settings (medium demand, 100% simultaneous passengers, $50 mean ancillary disutility, $50 ancillary price).

disutility is 75% of the ancillary price, a business-oriented service where the percentages are reversed, and an equally appealing service, with a mean disutility of 100% of the ancillary price for both segments.

The results of the previous section are consistent across the range of parameters tested. The change in total revenue (vs baseline) is shown in Figure 7.7. The optimizer increment leads to revenue losses on the order of 0.1% in all cases. Hybrid forecasting and fare adjustment provides revenue increases of 1.0–1.2% over baseline in each of the cases. AMD and AMR provide an additional 0.6–0.7 pts of revenue benefit over HF/FA, for a total gain of about 1.8% over baseline when the ancillary service is optional in all classes. The benefit of AMD and AMR over HF/FA or baseline is relatively stable over the range of ancillary prices and ancillary disutilities tested. Figure 7.8 lists the airline's load factor for each simulation; in general, the optimizer increment slightly increases load factor, while HF/FA and AMD/AMR decrease load factor compared to the baseline. HF/FA has a slightly greater load factor decrease (around -1.2 pts) than AMD/AMR (around -0.9 pts)

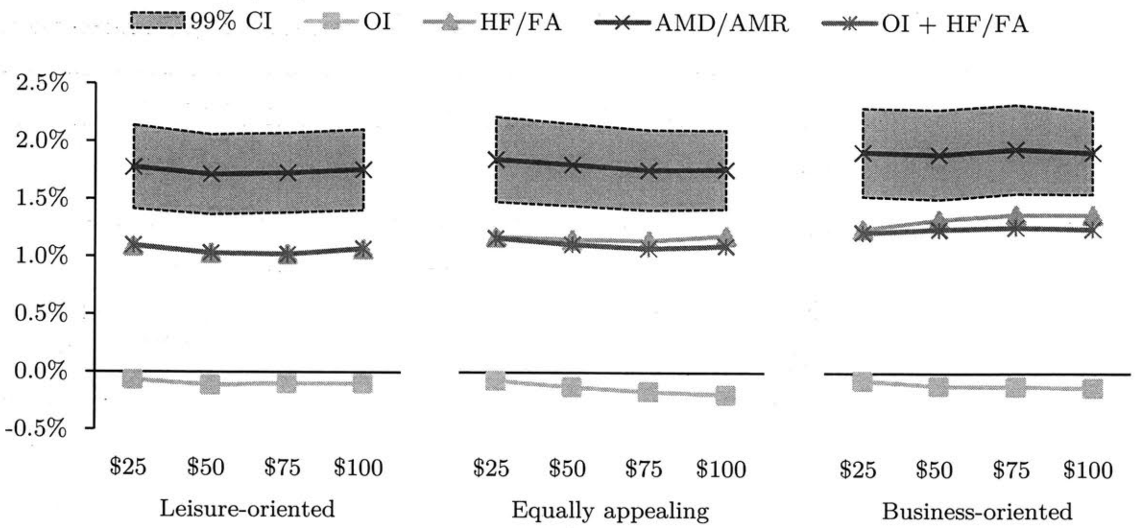

Figure 7.7: Network A1ONE change in total revenue vs baseline with various ancillary prices and disutility scenarios (medium demand, 100% simultaneous passengers). 99% confidence interval for change in total revenue due to AMD and AMR. $df = 1,999$.

vs baseline.

Ancillary purchase rates, average ancillary revenue by passenger, and the portion of total revenue generated by ancillary services (by class and overall) vary widely across these parameter ranges. However, the effect of AMD and AMR on booking mix is similar across price ranges, as shown in Figure 7.9, but has more variability as relative disutilities change (from the leisure-oriented to equally appealing to business-oriented cases). The leisure-oriented disutilities result in a decrease in bookings in all but the highest fare class. The business-oriented disutilities decrease bookings in FC 2, 5 and 6, but increase bookings in FC 3 and 4; the magnitude of the FC 6 decrease (and the total decrease across all fare classes) is greater with the business-oriented disutilities. The business-oriented disutilities have a larger decrease in FC 6 for two reasons. First, with the higher business disutility, more passengers booking in later DCPs and more passengers booking in higher value fare classes will purchase the ancillary service, making it more important to save space for late arriving, high value customers. Second, with the lower leisure disutility, fewer passengers booking in

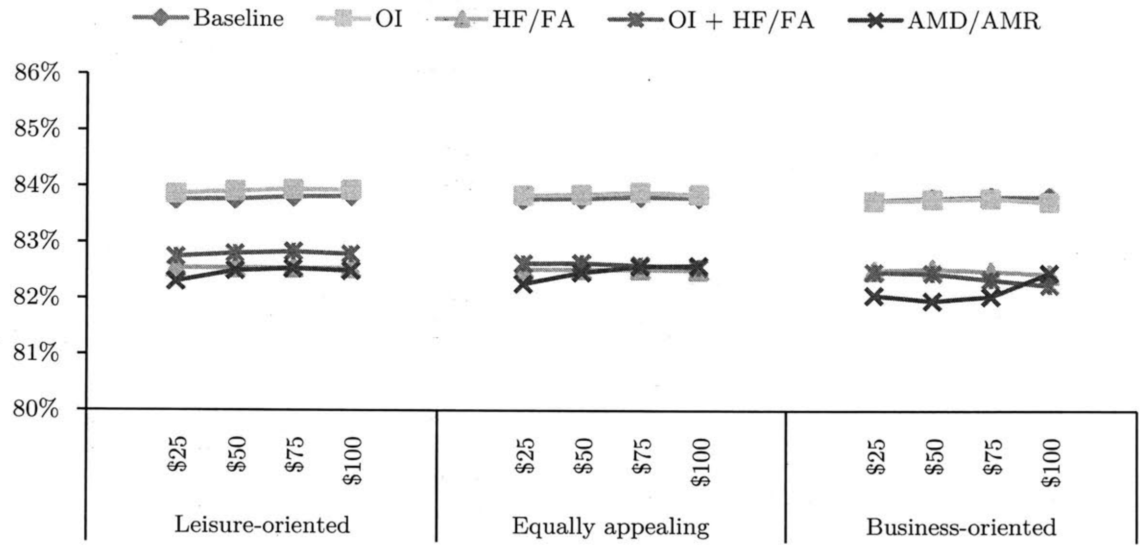

Figure 7.8: Load factor with various ancillary disutility scenarios and prices (medium demand, 100% simultaneous passengers). Ancillary disutility scenario (see Table 5.5) and ancillary price ($25–$100) listed on horizontal axis.

early TFs and fewer passengers booking in lower value classes will purchase the ancillary, making the value of early FC 6 bookings low, and compounding the booking limit effects of the higher business disutility.

7.3.5 Sensitivity to Demand Level

We now consider the sensitivity of AMD/AMR to varying overall levels of demand. We will consider the previous simulations, in which the base case had an 84% load factor, as *medium demand*, and will assess both higher and lower demand levels. We increase the mean demand (number of passengers generated) for both business and leisure passengers by 15% for a *high demand* set of scenarios, and decrease mean demand for both segments by 10% for a *low demand* scenario (matching the studies in Section 5.2.3). Baseline results for these three demand scenarios are shown in Table7.5. The low demand baseline has a load factor of 77% and the high demand baseline has a load factor of 88%. As demand grows, so does

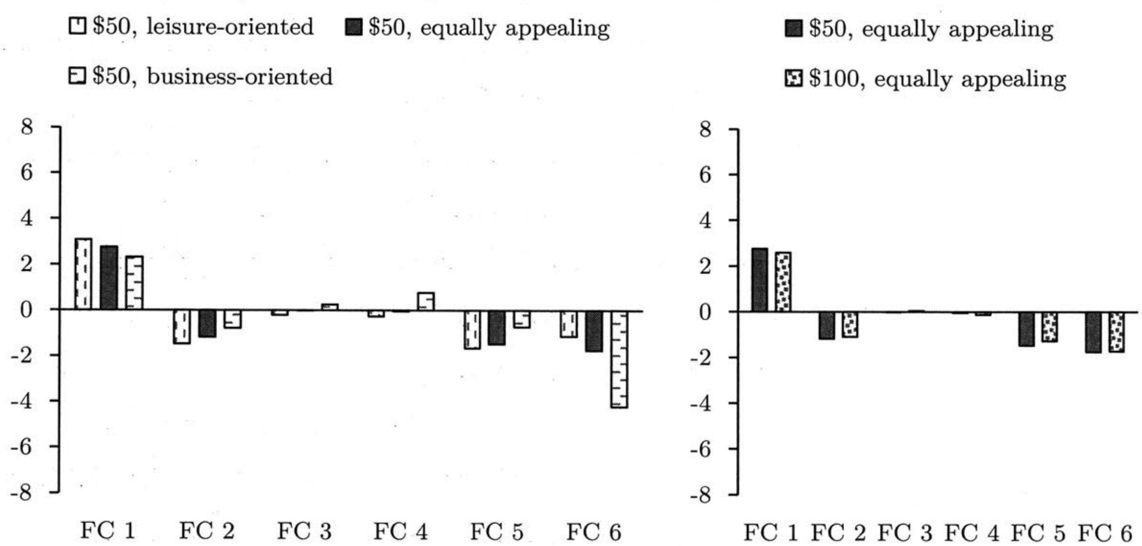

Figure 7.9: Change in bookings by fare class vs relevant baseline due to AMD and AMR with various ancillary disutility scenarios and prices (medium demand, 100% simultaneous passengers). Left: various disutility disutility scenarios (see Table 5.5) and a $50 ancillary price. Right: various ancillary prices when both consumer segments have a mean ancillary disutility equal to price. The $50 equally appealing case corresponds to the results in the previous portion of this section.

Table 7.5: Network A1ONE simulation results for baseline case with various market demand levels (100% simultaneous passengers, $50 ancillary price, $50 ancillary disutilities).

	Low demand	Medium demand	High demand
Ticket Revenue	$20,279	$22,098	$24,325
Ancillary Revenue	$1,678	$1,808	$1,927
Total Revenue	$21,957	$23,906	$26,252
Load Factor	77.1%	83.8%	87.7%
Total Yield	21.90	21.95	23.02
Ancillary Sales Rate	33.5%	33.2%	33.8%
Change from medium demand			
Ticket Revenue	-8.2%		+10.1%
Ancillary Revenue	-7.2%		+6.6%
Total Revenue	-8.2%		+9.8%
Load Factor	-6.7 pts		+4.0 pts
Total Yield	-0.2%		+4.9%
Ancillary Sales Rate	+0.3 pts		+0.6 pts

Note: Total yield in cents per mile.

load factor, although at a diminishing rate of return (a 10% decrease in demand results in a 6.7pt decrease in load factor, but a 15% increase in demand only results in a 4pt increase in load factor). Revenues—ticket, ancillary, and total—also increase with demand. Ancillary revenue grows as a function of load factor, while ticket revenue grows due to higher load factor (more bookings) and higher yield (more revenue per booking). Even though demand increases for both segments, the RM system protects seats for and preferentially accepts bookings from the higher-yielding fare classes (with primarily business passengers), so an increase in overall demand increases yield.

The revenue benefit of AMD/AMR increases with load factor, as shown in Figure 7.10. At lower demand, for a range of ancillary prices (with mean disutility equal to price for both demand segments), AMD/AMR increases revenue by about 0.3% over HF/FA, and by about 0.6% over baseline. Recall that at medium demand the benefit of AMD/AMR is about 0.6% over HF/FA and 1.8% over baseline. At higher demand, AMD/AMR increases revenue by about 0.9% over HF/FA and by 2.6% over baseline. Both HF/FA and AMD/AMR are

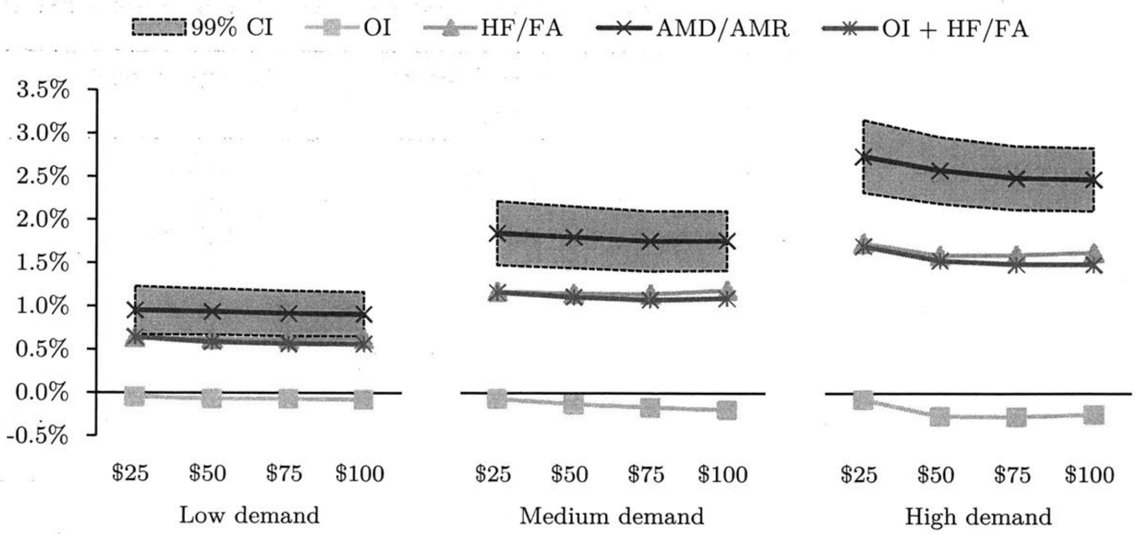

Figure 7.10: Network A1ONE change in total revenue vs baseline with various ancillary prices and market demand levels (100% simultaneous passengers, $50 ancillary disutilities). 99% confidence interval for change in total revenue due to AMD and AMR. $df = 1,999$. Note the change in vertical axis scale from other plots in this section.

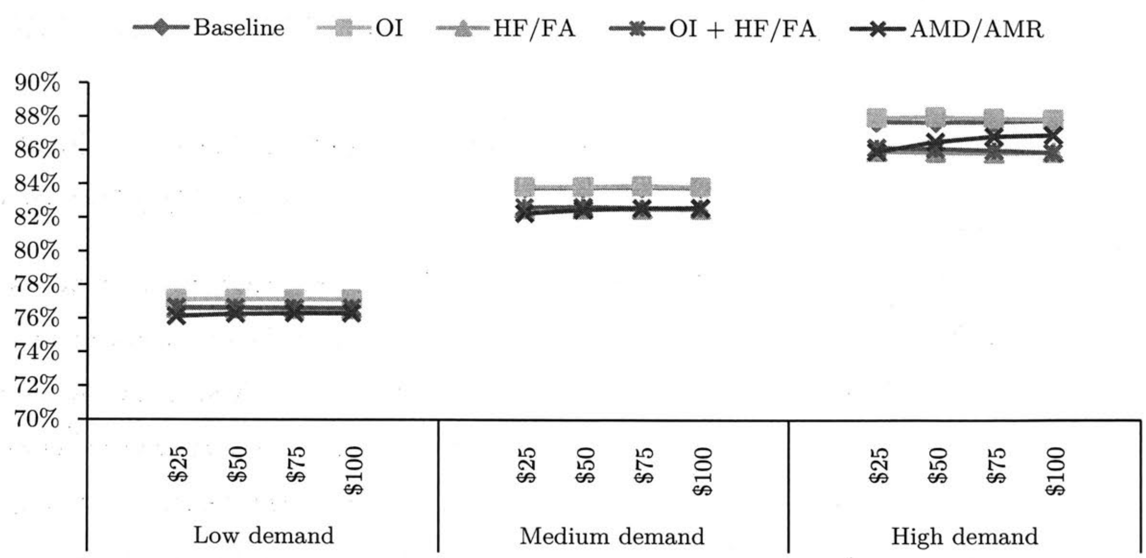

Figure 7.11: Load factor with various ancillary prices and market demand levels (100% simultaneous passengers, mean ancillary disutility equal to ancillary price).

more effective at higher demand levels because both mechanisms function by increasing protection levels for late-arriving, high-value passengers; at higher demand, there are more such customers. Figure 7.11 shows the load factor for each method; note that there is a larger load factor *decrease* associated with HF/FA and AMD/AMR at higher demand levels.

The change in bookings due to AMD/AMR for each demand scenario (compared to the relevant baseline) is shown in Figure 7.12. At higher demand levels, the increase in FC 1 bookings due to AMD/AMR grows, from +1.3 at low demand to +4.7 at high demand. As the highest-fare class, changes in FC 1 bookings have a partcularly large effect on revenue; Figure 7.13 shows the change in *ticket* revenue by fare class due to AMD/AMR for each of the demand levels, and illustrates that the primary ticket revenue benefit comes from an increase in FC 1 sales. In contrast to FC 1, at higher demand levels AMD/AMR has a larger *decrease* in bookings (and therefore ticket revenue) in FC 2 and FC 5. Recall from Section 6.2.1 that FC 2 and FC 5 in network A1ONE are inefficient, and are never offered as the lowest available class. A change in demand level is equivalent to changing λ in Equation 6.10, and therefore has no effect on the efficiency or inefficiency of any class. AMD/AMR shows larger booking reductions in FC 2 and FC 5 at higher demand levels because the relevant baselines have *more* bookings in these classes—especially FC 5—at higher demand levels. Finally, FC 6 changes show an inconsistent trend as demand level changes, with AMD/AMR decreasing FC 6 bookings at low and medium demand levels but increasing FC 6 bookings at high demand. This is again a symptom of variations in booking mix in the baseline cases: at high demand, the RM optimizer in the baseline case closes FC 6 and shifts bookings into FC 5; AMD/AMR, being aware of the inefficiency of FC 5 however, trades those FC 5 bookings for additional FC 4 and FC 6 bookings.

7.3.6 Sensitivity to Passenger Behavior Type

We now consider the performance of AMD/AMR when behavior is not 100% simultaneous. We will look at two alternative behavior types: 100% sequential, and a 50%/50% mix of

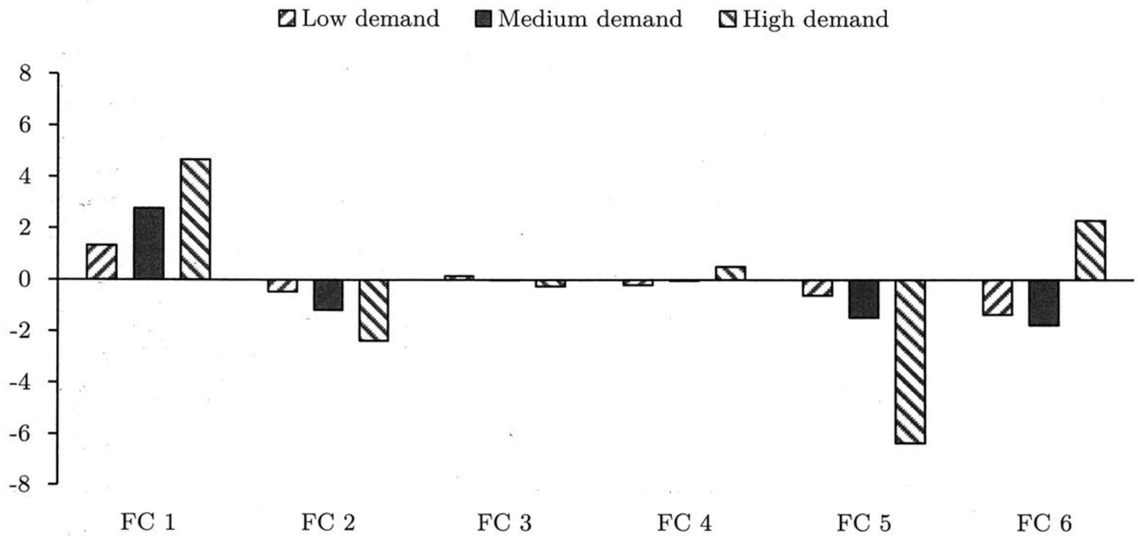

Figure 7.12: Change in bookings by fare class vs relevant baseline due to AMD and AMR with various market demand levels (100% simultaneous passengers, $50 mean ancillary disutility, $50 ancillary price).

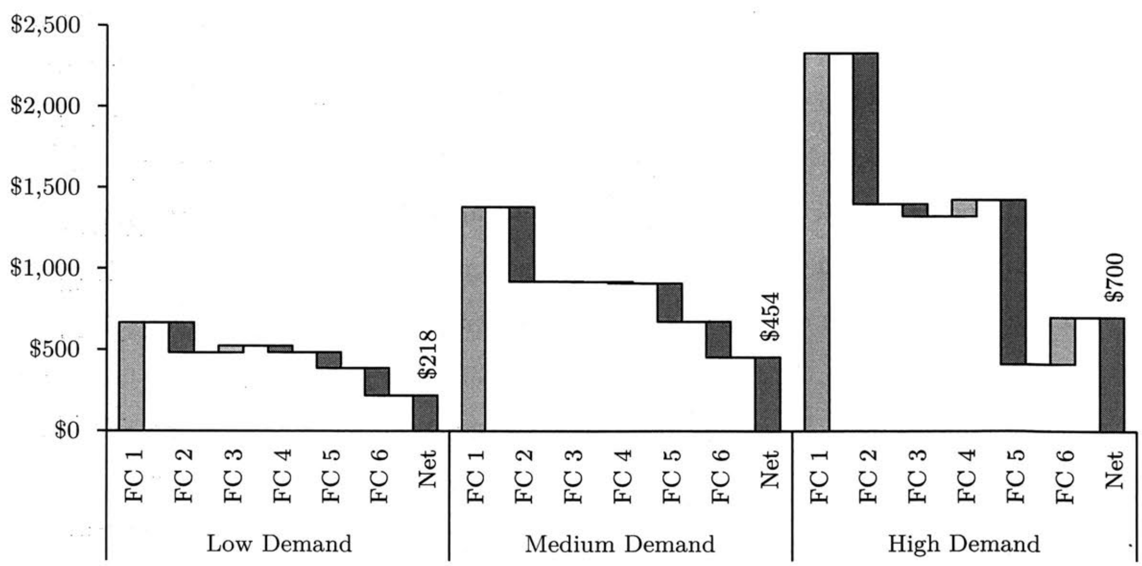

Figure 7.13: Change in ticket revenue by fare class vs relevant baseline due to AMD and AMR with various market demand levels (100% simultaneous passengers, $50 mean ancillary disutility, $50 ancillary price).

Table 7.6: Network A1ONE simulation results for baseline cases with various passenger behavior types (medium demand, \$50 ancillary price, \$50 ancillary disutilities).

	100% Simultaneous	50% Simultaneous & 50% Sequential	100% Sequential
Ticket Revenue	\$22,098	\$22,114	\$22,132
Ancillary Revenue	\$1,808	\$2,278	\$2,729
Total Revenue	\$23,906	\$24,392	\$24,861
Load Factor	83.8%	83.8%	83.8%
Total Yield	21.95	22.40	22.83
Ancillary Sales Rate	33.2%	41.8%	50.1%
Change from 100% Simultaneous			
Ticket Revenue		+0.1%	+0.2%
Ancillary Revenue		+26.0%	+50.9%
Total Revenue		+2.0%	+4.0%
Load Factor		+0.0 pts	-0.0 pts
Total Yield		+2.1%	+4.0%
Ancillary Sales Rate		+8.6 pts	+16.9 pts
Significance of Change in Total Revenue from 100% Simultaneous			
Standard Error		0.0%	0.0%
t-statistic		73.60	108.03
p-value		< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

simultaneous and sequential behavior. Recall that in the PODS simulator, passengers are randomly assigned a behavior in the 50/50 mix case.

Summary metrics for the baseline case with all three passenger behavior mixes are shown in Table 7.6. As the portion of sequential passengers increases (from 0% to 100%), there is a small increase in ticket revenue. As described in Section 3.1.3, simultaneous behavior can lead to buy-down as (some) passengers choose lower-value tickets in order to afford ancillary services. As the portion of simultaneous passengers decreases, this effect diminishes, which increases ticket revenue. Note that this effect has essentially no impact on the load factor.

The largest effect of increasing the portion of sequential passengers is that ancillary revenue

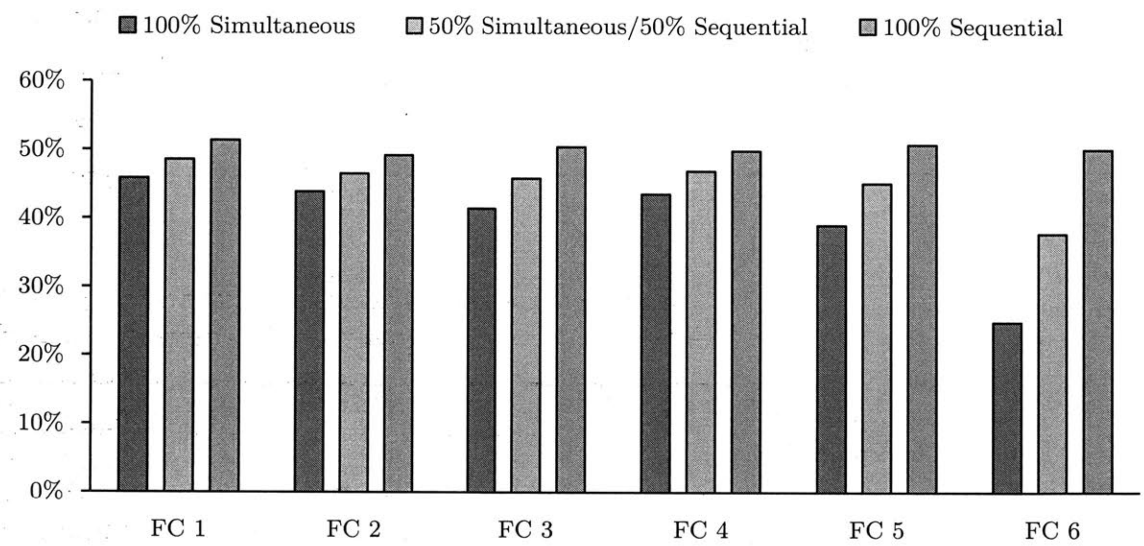

Figure 7.14: Ancillary purchase rate by fare class for baseline with various passenger behavior types (medium demand, \$50 mean ancillary disutility, \$50 ancillary price).

grows dramatically. Because sequential passengers do not have their ancillary purchases limited by the overall budgetary constraint, they can afford more ancillary services and the ancillary purchase rate increases to about 50%. When the ancillary price is equal to the ancillary mean disutilities (as in these cases), half of passengers will desire the ancillary service. With sequential behavior, all passengers who desire the ancillary service will purchase it, and purchase rates should be equal to 50%. As a result of the higher purchase rate, revenue from ancillary services increase by up to 51%. Driven primarily by the increase in ancillary revenue, the total revenue and the total yield increase as the portion of sequential passengers increases (by 4% when all passengers are sequential).

The increase in ancillary purchase rates is greatest in the lowest value classes, as shown in Figure 7.14. Because passengers who book in FC 6 tend to have the lowest budgetary constraints, when exhibiting simultaneous behavior many cannot also afford the ancillary service. Hence, when all passengers are simultaneous, ancillary purchase rates in FC 6 are only about 25%. However, as the portion of sequential passengers increases, and therefore as the portion of passengers whose ancillary purchases are not constrained by an overall

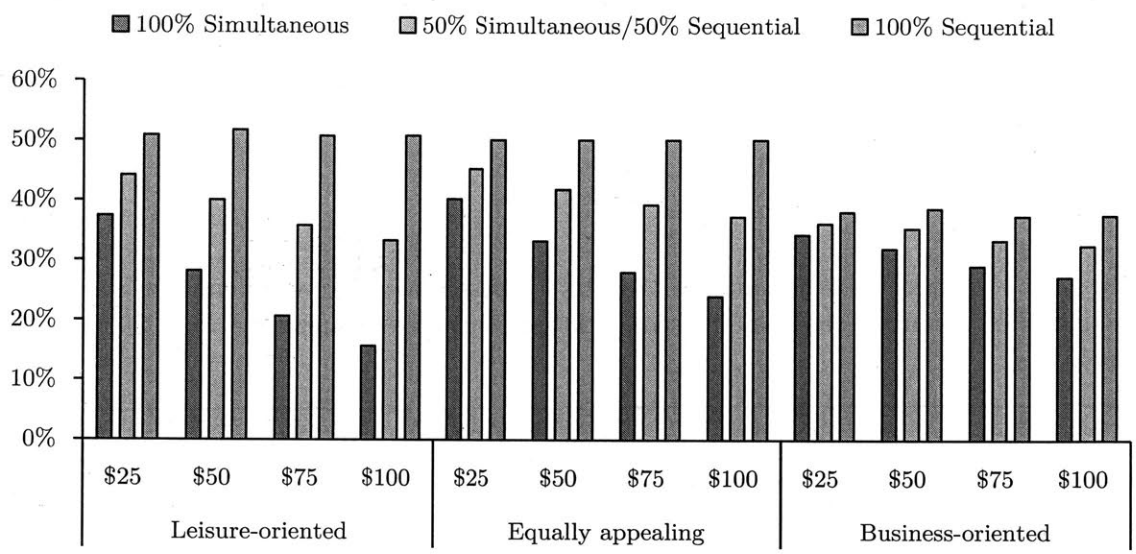

Figure 7.15: Ancillary purchase rate for baseline with various passenger behavior types and with various ancillary disutilities and prices (medium demand). Ancillary disutility configuration (see Table 5.5) and ancillary price ($25–$100) listed on horizontal axis.

budget increases, ancillary purchase rates increase. With 100% sequential behavior, about 50% of passengers in any fare class will purchase the service.

Note that these trends are influenced heavily by the relationship between ancillary prices, ancillary mean disutilities, and budgetary constraints. For a given budgetary distribution, as ancillary prices increase, fewer simultaneous passengers will be able to afford the service regardless of the number of passengers who desire it, and purchase rates for simultaneous passengers will decrease, as shown in Figure 7.15. With sequential passengers, an increase in ancillary price together with a corresponding increase in ancillary disutility leaves ancillary purchase rates unchanged. As ancillary mean disutilities increase, more passengers will desire the service, and purchase rates will increase.

With a mix of behavior types, the revenue impacts of the optimizer increment, hybrid forecasting/fare adjustment, a combination of the two, and AMD/AMR are similar to the case when all passengers are simultaneous, as shown in Tables 7.7 and 7.8: the optimizer

Table 7.7: Network A1ONE simulation results for baseline and experimental cases (medium demand, 50% simultaneous and 50% sequential passengers, \$50 ancillary price, \$50 ancillary disutilities).

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
50% Simultaneous/50% Sequential Passengers					
Ticket Revenue	\$22,114	\$22,075	\$22,399	\$22,385	\$22,554
Ancillary Revenue	\$2,278	\$2,281	\$2,250	\$2,255	\$2,249
Total Revenue	\$24,392	\$24,356	\$24,649	\$24,640	\$24,803
Load Factor	83.8%	83.9%	82.5%	82.7%	82.7%
Total Yield	22.40	22.33	22.98	22.91	23.08
Ancillary Sales Rate	41.8%	41.8%	41.9%	41.9%	41.9%
Change from Baseline					
Ticket Revenue		-0.2%	+1.3%	+1.2%	+2.0%
Ancillary Revenue		+0.1%	-1.2%	-1.0%	-1.3%
Total Revenue		-0.1%	+1.1%	+1.0%	+1.7%
Load Factor		+0.1 pts	-1.3 pts	-1.0 pts	-1.1 pts
Total Yield		-0.3%	+2.6%	+2.3%	+3.0%
Ancillary Sales Rate		-0.0 pts	+0.1 pts	+0.1 pts	+0.0 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.0%	0.1%	0.1%	0.1%
t-statistic		-7.66	10.74	11.26	12.96
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

Table 7.8: Network A1ONE simulation results for baseline and experimental cases (medium demand, 100% sequential passengers, \$50 ancillary price, \$50 ancillary disutilities).

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
100% Sequential Passengers					
Ticket Revenue	\$22,132	\$22,088	\$22,414	\$22,397	\$22,555
Ancillary Revenue	\$2,729	\$2,736	\$2,689	\$2,700	\$2,698
Total Revenue	\$24,861	\$24,824	\$25,103	\$25,097	\$25,253
Load Factor	83.8%	83.9%	82.5%	82.8%	82.8%
Total Yield	22.83	22.75	23.40	23.31	23.45
Ancillary Sales Rate	50.1%	50.2%	50.1%	50.1%	50.1%
Change from Baseline					
Ticket Revenue		-0.2%	+1.3%	+1.2%	+1.9%
Ancillary Revenue		+0.3%	-1.5%	-1.1%	-1.1%
Total Revenue		-0.1%	+1.0%	+0.9%	+1.6%
Load Factor		+0.2 pts	-1.2 pts	-0.9 pts	-0.9 pts
Total Yield		-0.4%	+2.5%	+2.1%	+2.7%
Ancillary Sales Rate		+0.0 pts	+0.0 pts	+0.0 pts	-0.0 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.0%	0.1%	0.1%	0.1%
t-statistic		-7.35	10.18	11.04	12.98
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

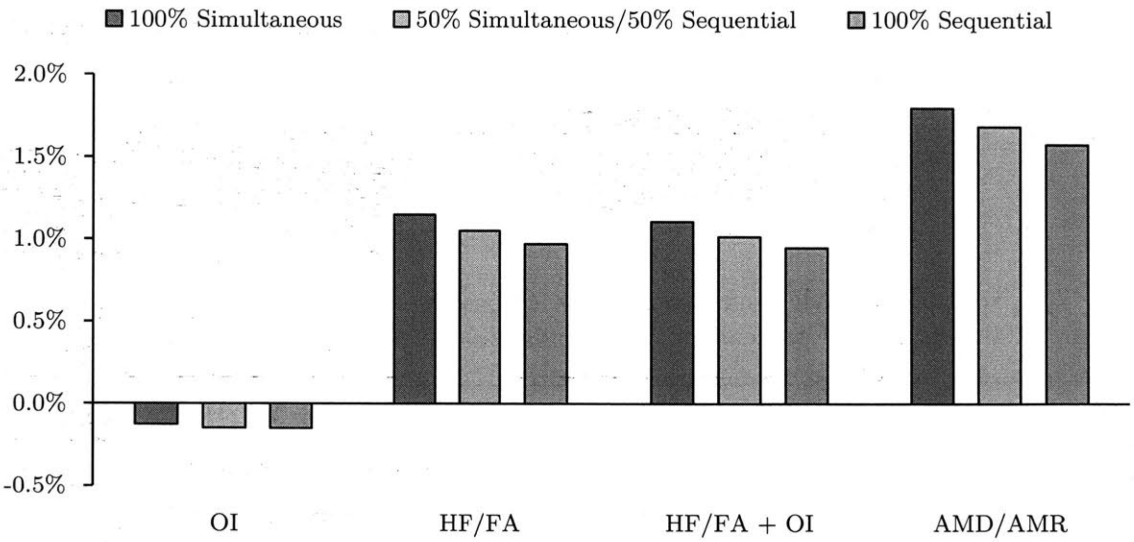

Figure 7.16: Change in total revenue vs baseline with various passenger behavior types (medium demand, \$50 mean ancillary disutility, \$50 ancillary price).

increment decreases total revenue and increases load factor (whether applied alone or in conjunction with HF/FA). Hybrid forecasting/fare adjustment and AMD/AMR both increase total revenue and decrease load factor; AMD/AMR has a greater increase in total revenue and a smaller decrease in load factor. With higher revenue and lower load factor, total yield increases with both HF/FA and AMD/AMR. The increase in yield due to HF/FA or AMD/AMR is smaller with 100% sequential passengers than with 100% simultaneous passengers, and the increase in total revenue due to HF/FA or AMD/AMR is smaller with 100% sequential passengers than with 100% simultaneous passengers. Figure 7.16 summarizes the total revenue change of each method for each passenger behavior.

Figure 7.17 shows the change in bookings by fare class due to HF/FA and due to AMD/AMR for both 100% simultaneous and 100% sequential passenger behaviors, relative to the corresponding standard forecasting baselines. Passenger behavior has relatively little impact on the performance of HF/FA: for both behavior types, HF/FA reduces FC 6 bookings by about 5.5 (flight leg capacity is 130 seats), and increases bookings in higher value classes. However, for AMD/AMR, which explicitly accounts for passenger ancillary choice behavior,

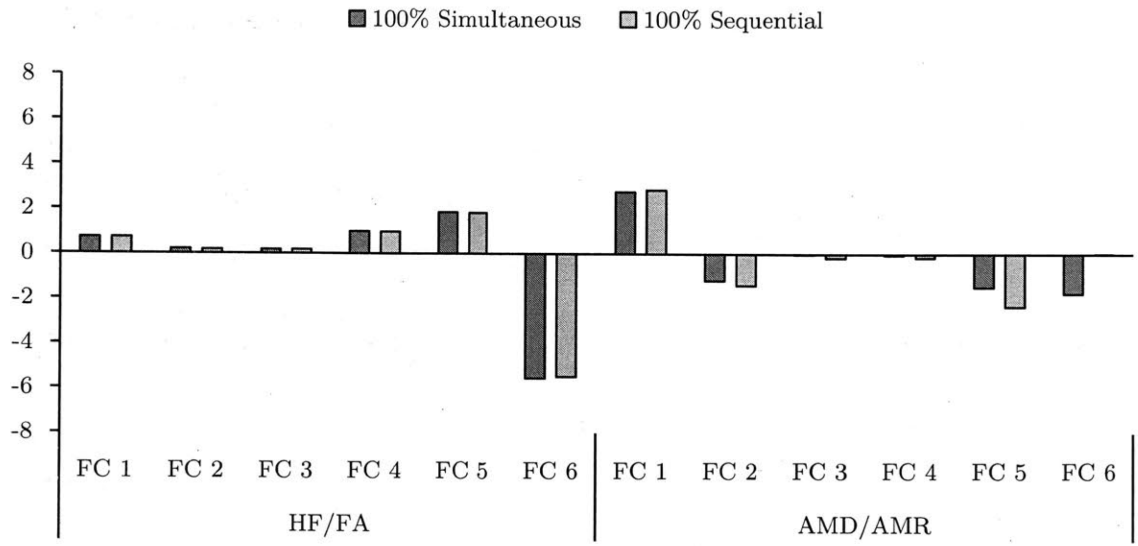

Figure 7.17: Change in bookings by fare class vs baseline with various passenger behavior types (medium demand, \$50 mean ancillary disutility, \$50 ancillary price).

there are greater differences. AMD/AMR with 100% simultaneous passengers reduces FC 6 bookings by 1.8, but with 100% sequential passengers there is essentially no change in FC 6 bookings. FC 5 bookings show a greater decrease due to AMD/AMR with 100% sequential passengers (2.3) than with 100% simultaneous passengers (1.5). The mean absolute difference of differences is 0.02 for HF/FA, but 0.54 for AMD/AMR.[3]

As indicated in Figure 7.18, the total revenue benefit of AMD/AMR decreases with higher ancillary prices, for all three tested ancillary disutility scenarios (see Table 5.5). AMD/AMR increases revenue over the baseline in all tested price/disutility combinations, providing a total revenue benefit of 1.2% to 1.9%. AMD/AMR increases revenue over HF/FA as well, by an incremental 0.4 pts to 0.9 pts. The optimizer increment used alone decreases revenue in all cases, and decreases revenue when paired with HF/FA in the equally-appealing and business-oriented disutility configurations. In the leisure-oriented disutility configuration,

[3]The mean absolute difference of differences compares variation in the impact of RM methods against a baseline due to other parameter changes; in this case it is defined as: $\frac{1}{nfcls}\sum_{k=1}^{nfcls}|(b_{i,1,1}-b_{i,1,2})-(b_{i,2,1}-b_{i,2,2})|$, where $b_{i,j,k}$ is the number of bookings received in fare class i for parameter setting j (either 100% simultaneous or 100% sequential) and for RM method k (1: HF/FA or AMD/AMR, 2: Baseline).

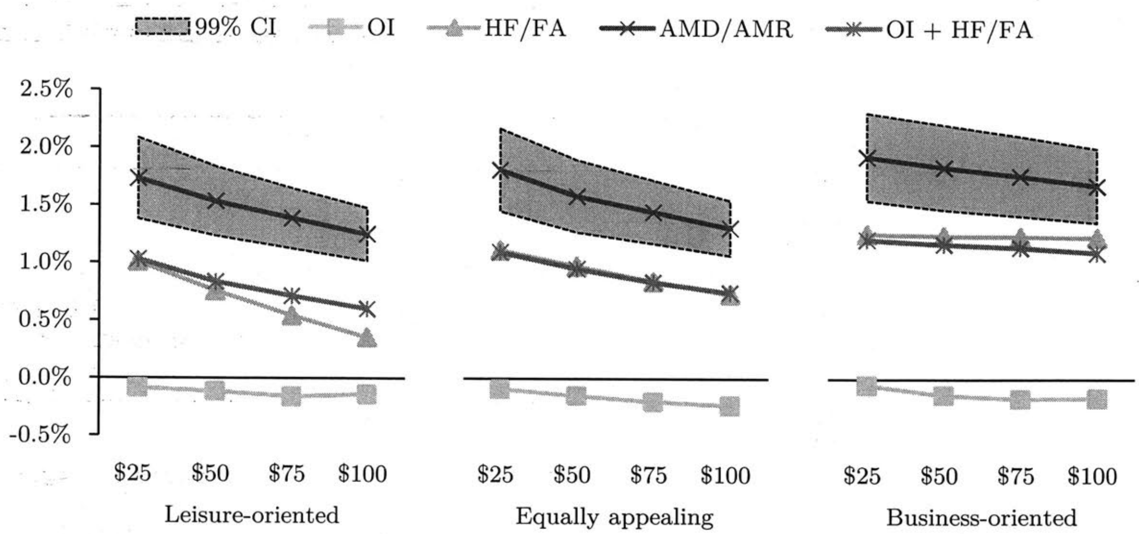

Figure 7.18: Network A1ONE change in total revenue vs baseline with various ancillary prices and disutility scenarios (medium demand, 100% sequential passengers). 99% confidence interval for change in total revenue due to AMD and AMR. $df = 1,999$.

OI + HF/FA increases revenue slightly compared to HF/FA alone, and has a greater increase at higher ancillary prices.

A summary of the total revenue impact of AMD/AMR across a range of passenger behaviors, ancillary prices, and disutility scenarios is shown in Figure 7.19. AMD and AMR have statistically significant increases in total revenue in all tested cases ($p < 0.01$, $df = 1,999$). On average, AMD and AMR increase total revenue more with simultaneous passengers than with sequential passengers; Figure 7.20 shows the difference in differences in total revenue change (measured as a percentage from AMD/AMR case relative to baseline) between 100% sequential passengers and 100% simultaneous passengers. In nearly all cases the mean is negative, meaning that the revenue gain with AMD/AMR is less when all passengers are sequential than when all passengers are simultaneous. The difference in performance is statistically significant for higher ancillary prices (as shown by the 99% confidence interval, $df = 1,999$, for price $\geq$ \$50 (leisure-oriented and equally appealing disutility scenario) or price $\geq$ \$75 (business-oriented disutility scenario)) and the difference is larger in the

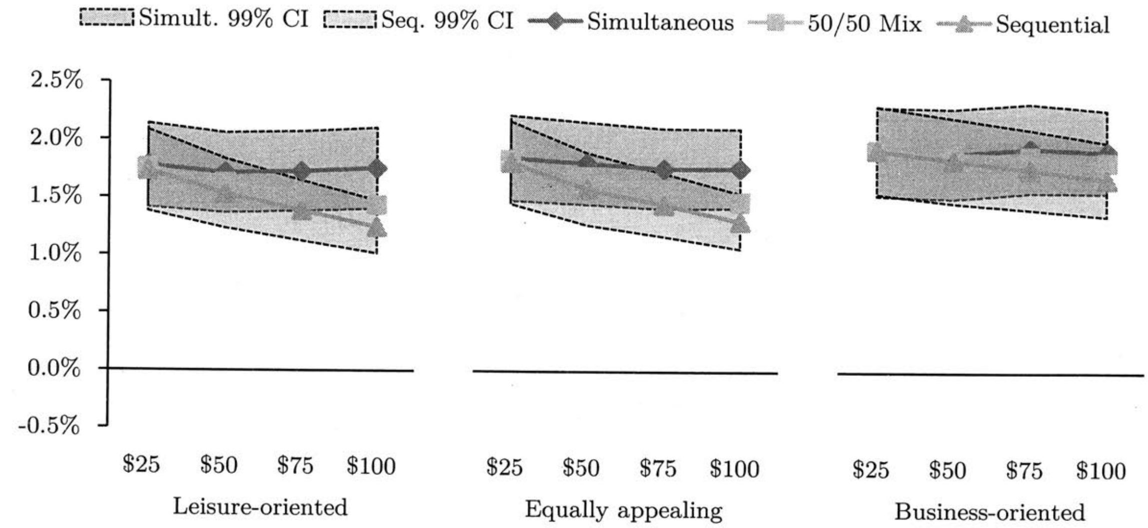

Figure 7.19: Network A1ONE change in total revenue due to AMD and AMR vs baseline with various passenger behaviors, ancillary prices, and disutility scenarios (medium demand). 99% confidence interval for change in total revenue with all simultaneous or all sequential passengers. $df = 1,999$.

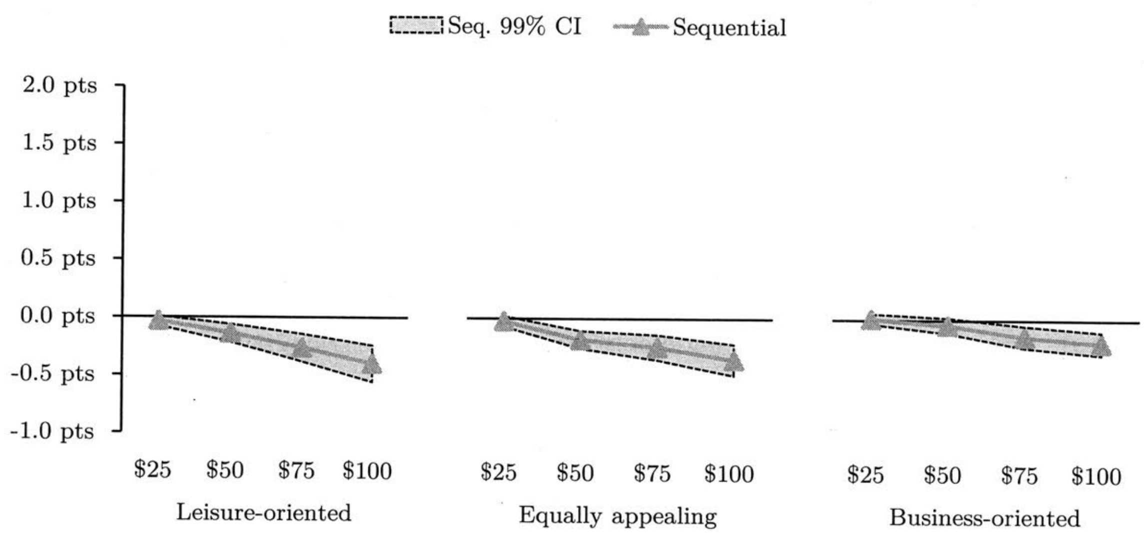

Figure 7.20: Network A1ONE difference in differences for percentage change in total revenue due to AMD and AMR vs baseline for 100% sequential passengers vs 100% simultaneous passengers with various ancillary prices and disutility scenarios (medium demand). Vertical axis has same scale as Figure 7.19 but is shifted. 99% confidence interval. $df = 1,999$.

Table 7.9: Network A2TWO baseline statistics (medium demand, 100% simultaneous passengers, $50 ancillary price, equally appealing ancillary disutility scenario).

	Airline 1	Airline 2
Ticket Revenue	$22,386	$22,388
Ancillary Revenue	$1,825	$1,822
Total Revenue	$24,211	$24,210
Load Factor	84.3%	84.2%
Total Yield	22.09	22.13
Ancillary Sales Rate	33.3%	33.3%
Portion of Total Revenue from Ancillary	7.5%	7.5%

Note: Total yield in cents per mile.

leisure-oriented disutility scenario.

7.4 Two Airline, Two Flight Leg Network (A2TWO)

We now introduce a competitor airline, offering the same six economy fare classes with the same fare class restrictions and advanced purchase requirements. There is no quality or schedule differentiation between the two airlines—both offer one non-stop flight in the market, departing and arriving at the same time. Price differentiation only arises as a result of revenue management action. As with A1ONE, simulations in network A2TWO feature five trials of 400 unburned samples each, for a total of 2,000 observations in the reported outputs (each trial has additional samples used only to warm up the forecasting models). Again, comparisons are made between two simulations on a pairwise basis because the underlying stochastic demand is the same for each simulation (for the same demand, passenger behavior, ancillary price, and ancillary disutility parameters). The total number of passengers generated in A2TWO is double that of A1ONE, and we use the same combinations of market demand levels, ancillary prices, and ancillary disutility scenarios as in A1ONE.

Baseline results for A2TWO with medium demand, 100% simultaneous passengers, and a

Table 7.10: Network A2TWO Airline 1 baseline bookings and ancillary purchase data by fare class (medium demand, 100% simultaneous passengers, \$50 ancillary price, \$50 ancillary disutility).

	FC 1	FC 2	FC 3	FC 4	FC 5	FC 6
Bookings	6	14	10	8	15	57
Booking Mix	6%	12%	9%	7%	14%	52%
Average Fare	\$500	\$390	\$295	\$200	\$160	\$125
Average Ancillary per Passenger	\$22	\$22	\$20	\$23	\$20	\$12
Portion of Total Revenue from Ancillary	4%	5%	6%	10%	11%	9%
Ancillary Sales Rate	45%	44%	41%	45%	39%	25%

\$50 ancillary service with equally appealing disutilities are shown in Table 7.9. The results are similar to the equivalent A1ONE baseline, and both airlines in A2TWO have nearly identical performance. Approximately 8% of total revenue comes from sales of the ancillary service, in line with US DOT estimates of ancillary revenue for US airlines. A breakout of bookings and ancillary purchases by fare class for Airline 1 are shown in Table 7.10. As with A1ONE, most bookings occur in the lowest value fare class. Purchase rates for the ancillary service range broadly by fare class, with lower purchase rates in the lower-value fare classes (where many passengers have tight budgetary constraints) and higher purchase rates in the higher-value fare classes. The portion of total revenue coming from ancillary services, however, is higher in lower-value classes (11% in FC 5) and lower in higher-value classes (4% in FC 1). Results for Airline 2 are similar.

7.4.1 Symmetric Competition

We first consider cases of symmetric competition, i.e. both airlines use the same revenue management forecasting and optimization configurations. We use the same revenue management configurations as with network A1ONE: the optimizer increment (OI), hybrid forecasting with fare adjustment (HF/FA), a combination of OI and HF/FA, and AMD/AMR. We assume that the airline has accurate knowledge of single-airline passenger choice probabilities but does not attempt to account for competition. In other words, for given passenger

Table 7.11: Network A2TWO Airline 1 simulation results for baseline and symmetric experimental cases (medium demand, 100% simultaneous passengers, $50 ancillary price, $50 ancillary disutilities).

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
Airline 1					
Ticket Revenue	$22,386	$22,376	$22,585	$22,585	$22,798
Ancillary Revenue	$1,825	$1,827	$1,809	$1,812	$1,807
Total Revenue	$24,211	$24,203	$24,394	$24,397	$24,605
Load Factor	84.3%	84.4%	82.6%	82.9%	82.9%
Total Yield	22.09	22.07	22.71	22.65	22.84
Ancillary Sales Rate	33.3%	33.3%	33.7%	33.6%	33.5%
Change from baseline					
Ticket Revenue		-0.0%	+0.9%	+0.9%	+1.8%
Ancillary Revenue		+0.1%	-0.9%	-0.7%	-1.0%
Total Revenue		-0.0%	+0.8%	+0.8%	+1.6%
Load Factor		+0.0 pts	-1.7 pts	-1.4 pts	-1.4 pts
Total Yield		-0.1%	+2.8%	+2.5%	+3.4%
Ancillary Sales Rate		+0.0 pts	+0.4 pts	+0.3 pts	+0.2 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.0%	0.1%	0.1%	0.2%
t-statistic		-1.20	5.98	6.51	10.71
p-value		0.229	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

behavior, ancillary price, and ancillary disutility settings, the choice probabilities used in A2TWO are the same as those in A1ONE.

Results for these symmetric competition cases are shown in Tables 7.11 and 7.12. AMD and AMR result in an increase of total revenue of 1.6% for both airlines, driven by an increase in ticket revenue (+1.8%) and a decrease in ancillary revenue (-1.0%). The decrease in ancillary revenue is a result of a decrease in load factor (-1.4 pts for Airline 1; -1.3 pts for Airline 2), which is slightly offset by a higher ancillary purchase rate (+0.2 pts for Airline 1; +0.1 pts for Airline 2). The difference in total revenue due to AMD/AMR vs baseline is statistically significant for both airlines (t = 10.71 (Airline 1) and 10.55 (Airline 2),

Table 7.12: Network A2TWO Airline 2 simulation results for baseline and symmetric experimental cases (medium demand, 100% simultaneous passengers, $50 ancillary price, $50 ancillary disutilities).

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
Airline 2					
Ticket Revenue	$22,388	$22,341	$22,684	$22,670	$22,786
Ancillary Revenue	$1,822	$1,821	$1,815	$1,816	$1,803
Total Revenue	$24,210	$24,162	$24,499	$24,486	$24,589
Load Factor	84.2%	84.2%	82.9%	83.0%	82.9%
Total Yield	22.13	22.07	22.73	22.69	22.81
Ancillary Sales Rate	33.3%	33.3%	33.7%	33.7%	33.5%
Change from baseline					
Ticket Revenue		-0.2%	+1.3%	+1.3%	+1.8%
Ancillary Revenue		-0.1%	-0.4%	-0.3%	-1.0%
Total Revenue		-0.2%	+1.2%	+1.1%	+1.6%
Load Factor		+0.1 pts	-1.3 pts	-1.2 pts	-1.3 pts
Total Yield		-0.3%	+2.7%	+2.5%	+3.1%
Ancillary Sales Rate		-0.0 pts	+0.4 pts	+0.4 pts	+0.1 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.0%	0.1%	0.1%	0.1%
t-statistic		-7.76	9.59	9.75	10.55
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

$p < 0.001$, $df = 1,999$). Trends for other RM configurations are similar in direction to those in network A1ONE: OI decreases total revenue when applied to the baseline or to HF/FA, and HF/FA increases total revenue compared to the baseline, though not to the same degree as AMD/AMR. Symmetric use of AMD/AMR increases total revenue by 0.8 pts (Airline 1) to 0.4 pts (Airline 2) over HF/FA, a statistically significant increase for both airlines ($t = 9.36$, $p < 0.001$ for Airline 1; $t = 4.09$, $p < 0.001$ for Airline 2; $df = 1,999$; using HF/FA as the baseline).

The difference in total revenue between Airline 1 and 2 when both use the same RM configuration is not statistically significant ($t = 1.14$, $p = 0.256$, $df = 1,999$ when both use HF/FA, which has the largest differences between airlines).

As with network A1ONE, AMD/AMR increases ticket and total revenue by reducing bookings in the lowest-value FC 6, and by reducing bookings in the inefficient FC 2 and FC 5, as shown in Figure 7.21. The booking reductions in these classes allow the airline to accept more bookings in the highest-value FC 1. The booking changes observed in A2TWO are very similar to those in the equivalent A1ONE case.

The sensitivity of these results to variations in passenger behavior type, ancillary price, and disutility scenario are shown in Figure 7.22. AMD/AMR consistently produce revenue gains, with an increase in total revenue vs baseline of 1.2% to 1.8%. These revenue differences are all statistically significant ($p < 0.01$, $df = 1,999$), as shown by the 99% confidence interval. In general, the revenue gains of AMD/AMR are slightly greater with simultaneous passengers than with sequential passengers, and the variation in performance by passenger type is exaggerated at higher ancillary prices and when the ancillary service is more favorable to leisure passengers (leisure-oriented). Airline 2 results show similar trends.

In summary, in this simple competitive network, when both airlines use AMD/AMR results are generally similar to the single airline network (A1ONE), even though the airlines have not accounted for the competitor in the AMD/AMR choice probabilities.

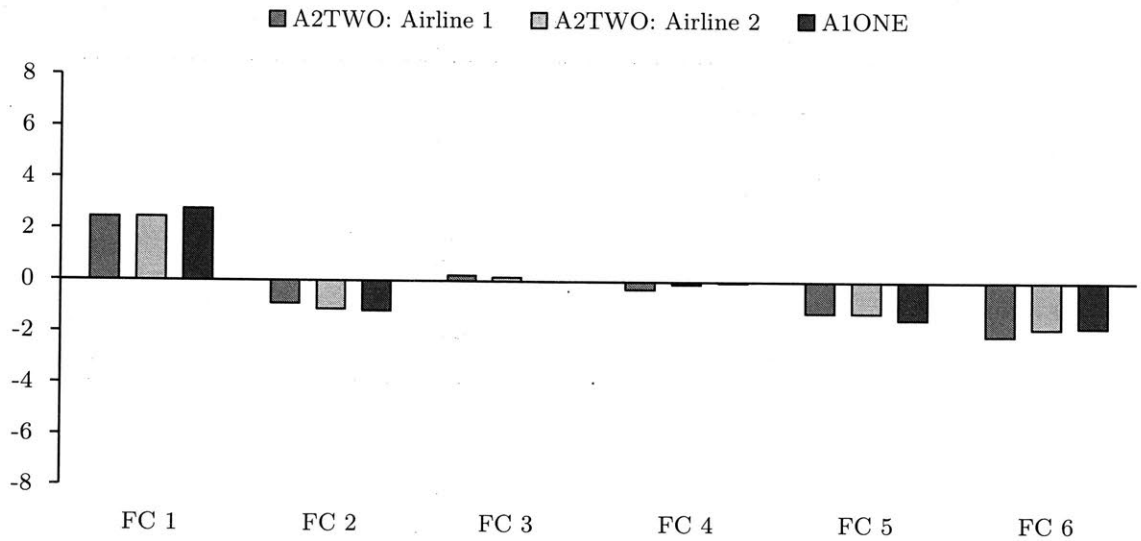

Figure 7.21: Network A2TWO change in bookings by fare class due to symmetric AMD and AMR vs baseline (medium demand, 100% simultaneous passengers, \$50 ancillary price, \$50 ancillary disutility). Equivalent network A1ONE result shown in grey for comparison.

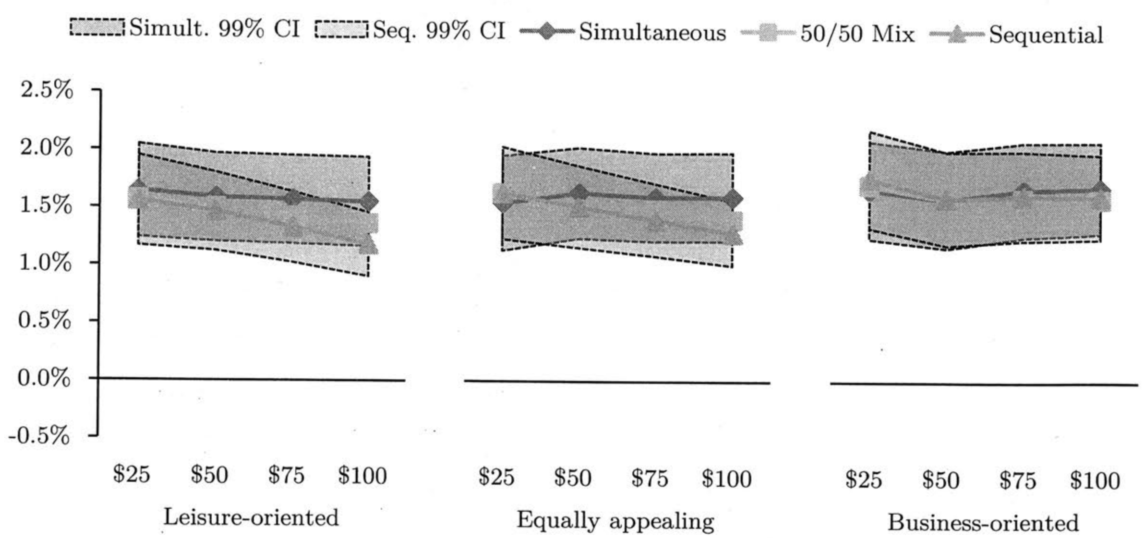

Figure 7.22: Network A2TWO change in Airline 1 total revenue due to symmetric AMD and AMR vs baseline with 99% confidence interval with various passenger behaviors, ancillary prices, and disutility scenarios (medium demand). $df = 1,999$.

Table 7.13: Network A2TWO Airline 1 simulation results for baseline and asymmetric experimental cases (medium demand, 100% simultaneous passengers, $50 ancillary price, $50 ancillary disutilities). Airline 2 always uses standard forecasting and optimizes on filed fares; Airline 1 uses specified forecasting and optimization model.

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
Airline 1					
Ticket Revenue	$22,386	$22,361	$22,678	$22,676	$22,788
Ancillary Revenue	$1,825	$1,826	$1,803	$1,807	$1,811
Total Revenue	$24,211	$24,187	$24,481	$24,483	$24,599
Load Factor	84.3%	84.4%	81.8%	82.0%	83.4%
Total Yield	22.09	22.04	23.03	22.96	22.70
Ancillary Sales Rate	33.3%	33.3%	33.9%	33.9%	33.4%
Change from baseline					
Ticket Revenue		-0.1%	+1.3%	+1.3%	+1.8%
Ancillary Revenue		+0.1%	-1.2%	-1.0%	-0.8%
Total Revenue		-0.1%	+1.1%	+1.1%	+1.6%
Load Factor		+0.1 pts	-2.5 pts	-2.3 pts	-1.0 pts
Total Yield		-0.2%	+4.3%	+3.9%	+2.8%
Ancillary Sales Rate		-0.0 pts	+0.6 pts	+0.6 pts	+0.1 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.0%	0.1%	0.1%	0.2%
t-statistic		-3.69	7.79	8.40	10.35
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1{,}999$.

7.4.2 Asymmetric Competition

In asymmetric competition, the airlines use different optimization and/or forecasting models. We assume initially that Airline 2 uses a standard (independent demand) forecasting and optimizes on filed fares (the same as the baseline). Airline 1 only changes its revenue management and/or forecasting configuration to the optimizer increment, hybrid forecasting and fare adjustment, a combination of the two approaches, or AMD/AMR. We continue to assume that choice probabilities do not account for competition.

Table 7.14: Network A2TWO Airline 2 simulation results for baseline and asymmetric experimental cases (medium demand, 100% simultaneous passengers, $50 ancillary price, $50 ancillary disutilities). Airline 2 always uses standard forecasting and optimizes on filed fares; Airline 1 uses specified forecasting and optimization model.

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
Airline 2					
Ticket Revenue	$22,388	$22,381	$22,449	$22,429	$22,587
Ancillary Revenue	$1,822	$1,821	$1,834	$1,834	$1,833
Total Revenue	$24,210	$24,202	$24,283	$24,263	$24,420
Load Factor	84.2%	84.1%	85.3%	85.2%	84.2%
Total Yield	22.13	22.13	21.91	21.90	22.31
Ancillary Sales Rate	33.3%	33.3%	33.1%	33.1%	33.5%
Change from baseline					
Ticket Revenue		-0.0%	+0.3%	+0.2%	+0.9%
Ancillary Revenue		-0.1%	+0.7%	+0.7%	+0.6%
Total Revenue		-0.0%	+0.3%	+0.2%	+0.9%
Load Factor		-0.1 pts	+1.1 pts	+1.1 pts	+0.0 pts
Total Yield		+0.0%	-1.0%	-1.0%	+0.8%
Ancillary Sales Rate		+0.0 pts	-0.2 pts	-0.2 pts	+0.2 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.0%	0.1%	0.1%	0.1%
t-statistic		-1.23	4.01	3.02	10.43
p-value		0.218	< 0.001	0.003	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

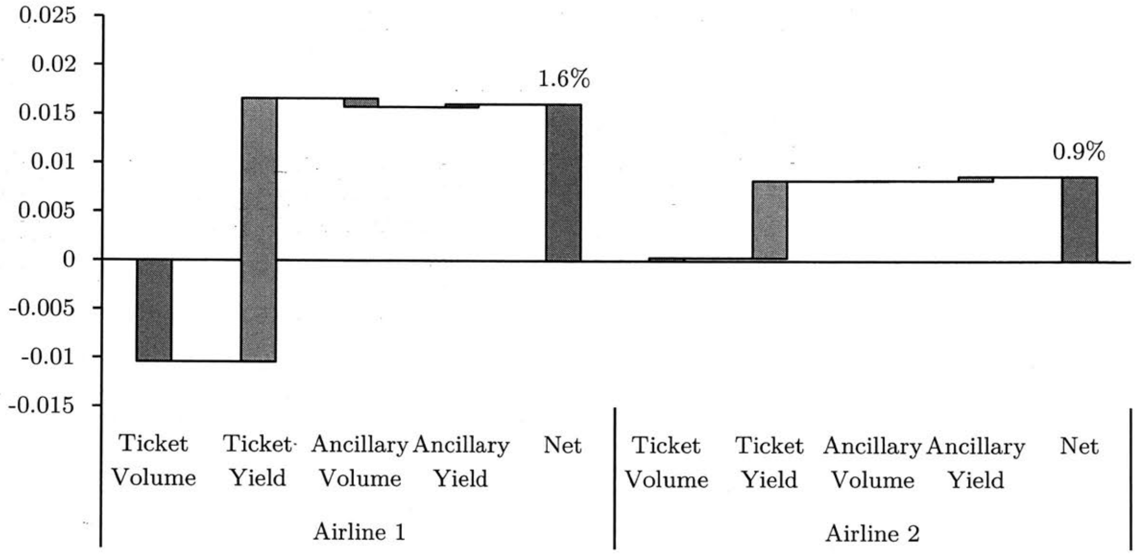

Figure 7.23: Network A2TWO change in Airline 1 and Airline 2 total revenue due to change in ticket and ancillary sales volumes and yields due to asymmetric AMD and AMR vs baseline (medium demand, 100% simultaneous passengers, \$50 ancillary price, \$50 ancillary disutility). "Volume" columns are portion of ticket or ancillary revenue change attributable to change in load factor; "yield" columns are portion of ticket or ancillary revenue change attributable to change in average fare or average ancillary revenue earned per passenger. All values expressed as percentage change from baseline total revenue.

Tables 7.13 and 7.14 shows results for the asymmetric cases. The results for Airline 1 are directionally similar to those with symmetric competition: AMD/AMR result in an increase in ticket revenue (+1.8%), a decrease in ancillary revenue (-0.8%), and an increase in total revenue (+1.6%). The reduction in ancillary revenue is driven by a reduction in load factor (-1.0 pts). These changes in total revenue are statistically significant ($t = 10.35$ for Airline 1 and 10.43 for Airline 2, $p < 0.001$, $df = 1,999$), with the 99% confidence interval for the total revenue change equal to (1.2%, 2.0%). Note that a pairwise comparison of AMD/AMR vs HF/FA shows a statistically significant increase in total revenue due to AMD/AMR ($t = 4.90$, $p < 0.001$, $df = 1,999$, using HF/FA as the baseline). Although Airline 2 makes no changes to its forecasting or optimization processes in these asymmetric experiments, it sees a substantial increase in total revenue (+0.9%) vs baseline when Airline 1 implements AMD/AMR.

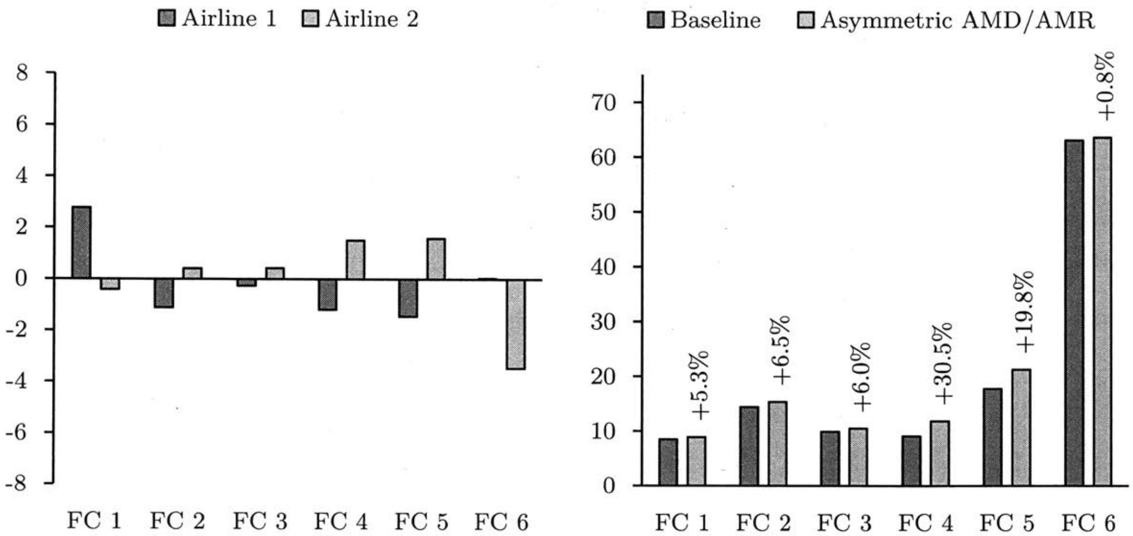

Figure 7.24: Network A2TWO change in bookings by fare class (left) and change in Airline 2 initial demand forecast mean (right) due to asymmetric AMD and AMR vs baseline (medium demand, 100% simultaneous passengers, \$50 ancillary price, \$50 ancillary disutility). Percentages show change in demand forecast vs baseline.

Figure 7.23 shows how changes in load factor, average fare, and average ancillary revenue per passenger contribute to the change in total revenue for Airline 1 and Airline 2 when Airline 1 only implements AMD/AMR. AMD/AMR produces more aggressive booking limits for Airline 1, which reduces its load factor and increases its ticket yields; the reduction in ticket revenue due to the lower load factor is more than offset by the increase in ticket revenue due to higher average fares. Likewise, the reduction in load factor has a (small) negative effect on ancillary revenue, which is partially offset by (slightly) higher ancillary purchase rates (and therefore higher "ancillary yield" or average ancillary revenue per passenger). Airline 2 sees a (small) increase in load factor, and a larger increase in ticket yield, resulting in higher ticket revenues and higher total revenues.

In contrast to the symmetric competition case, where both Airline 1 and 2 showed similar booking changes due to AMD/AMR as did the airline in network A1ONE, in the asymmetric case trends differ from A1ONE. Figure 7.24 shows the change in bookings by fare class for

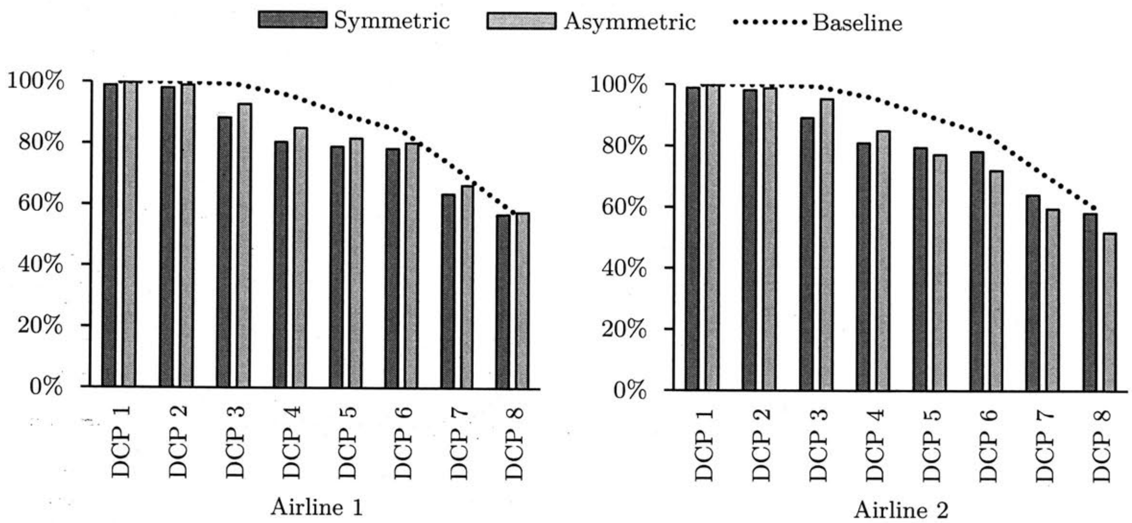

Figure 7.25: Network A2TWO FC 6 availability with symmetric and asymmetric AMD and AMR (medium demand, 100% simultaneous passengers, $50 ancillary price, $50 ancillary mean disutility). Availability measured as the portion of samples for which FC 6 is for sale.

both airlines when Airline 1 only implements AMD/AMR. Under asymmetric competition, Airline 2, which continues to use the same forecasting and optimization methods as in the baseline case, sees a decrease in FC 6 bookings and an increase in FC 2 through FC 5 bookings. Airline 1, on the other hand, experiences no decrease in FC 6 bookings (in contrast to the effect of AMD/AMR under symmetric competition). This difference in results between the symmetric and asymmetric cases is a result of competitive booking and forecast feedbacks between the two airlines, as shown in Figure 7.24. In the asymmetric case, AMD/AMR reduces availability for Airline 1; consumers therefore are more likely to book on Airline 2. The increase in Airline 2 bookings results in an increase Airline 2's forecast, particularly for FC 4 and FC 5.

The increased forecast results in less Airline 2 availability in FC 6, which in turn drives additional up-sell of passengers from FC 6 to higher value classes, reinforcing the effects of Airline 1's use of AMD/AMR. As shown in Figure 7.25, the result of these forecasting feedbacks is that under asymmetric competition, Airline 1 has a *greater* FC 6 availability

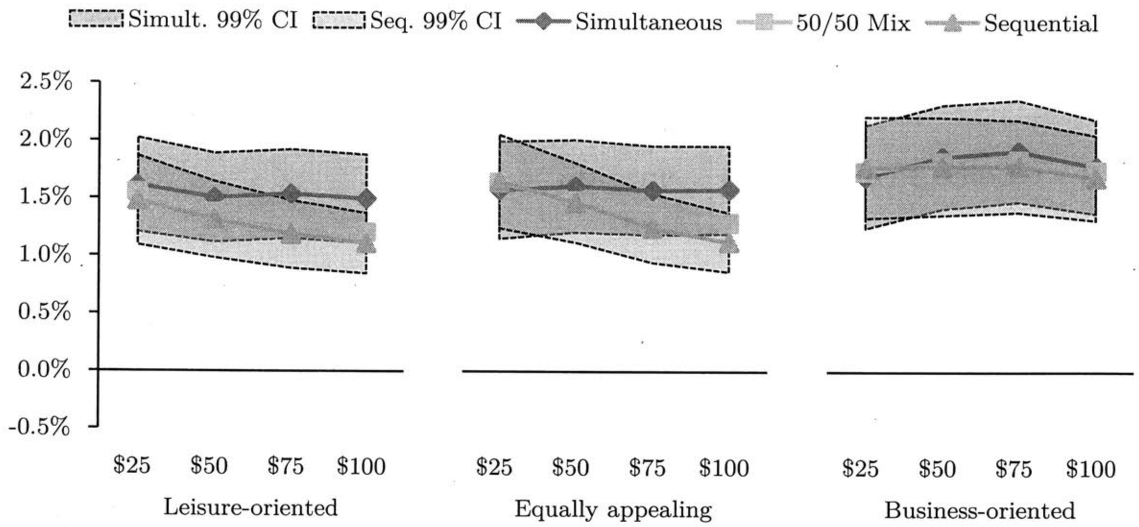

Figure 7.26: Network A2TWO change in Airline 1 total revenue due to asymmetric AMD and AMR vs baseline with 99% confidence interval with various passenger behaviors, ancillary prices, and disutility scenarios (medium demand). $df = 1,999$.

than Airline 2 during the last DCPs in which FC 6 can be sold (recall that FC 6 has an advance purchase requirement, so it is never sold after DCP 8).

The sensitivity to variations in passenger behavior type, ancillary price, and ancillary disutility scenario of the change in Airline 1 total revenue due to asymmetric AMD/AMR is shown in Figure 7.26. In all cases AMD/AMR increases total revenue over the baseline and the change is statistically significant ($p < 0.01$, $df = 1,999$), as indicated by the 99% confidence intervals. Total revenue changes for Airline 1 range from +1.1% to +1.9%. Total revenue changes for Airline 2 are not shown but are between +0.4% and +1.0%, although Airline 1 sees a revenue increase at least 0.5 pts greater than Airline 2 in all cases. As with the symmetric case, the benefit of AMD/AMR is smaller with sequential passengers than with simultaneous passengers, and the difference in performance is greater at higher ancillary prices and when the service is more favorable to leisure passengers. The improvement in Airline 1 total revenue vs asymmetric use of HF/FA (i.e., Airline 1 using HF/FA and Airline 2 using standard forecasting and optimizing on filed fares) is not shown, but ranges between +0.2% and +1.3%, with the change statistically significant ($p < 0.01$, $df = 1,999$) in all cases.

7.4.3 Choice Probability Estimation

Simulation result to this point have assumed that the airline has knowledge of the passenger choice probabilities $P_{i,m,dcp}(k)$, allowing our analysis to focus solely on the effects of our proposed forecasting and optimization processes, without the confounding effects of estimation errors. In reality, however, airlines do not know passenger choice probabilities and must estimate them. In this section, we test the performance of AMD/AMR when using choice probabilities that have been estimated from historical booking data with either the "raw" or "heuristic" approach, as described in Section 6.4. For these simulations, all choice probabilities are estimated from the *baseline* case, and each airline performs its own estimation process.[4] For the heuristic approach, the airline uses one of the standard PODS FRAT5 curves as its estimate of passenger price elasticity. These curves are lettered "A" through "E," with A leading to the most aggressive sell-up estimates, and E to the least aggressive, as shown in Figure 3.3.

Consider initially the scenario with medium demand, 100% simultaneous passengers, a \$50 ancillary price, and \$50 mean ancillary disutilities for both passenger segments. The total sale probabilities $TP_{dcp}(k)$ for Airline 1 from both the raw and the heuristic estimation process (using FRAT5 D) are shown in Figure 7.27. Although the batch estimation process has a total of 2,000 samples, as shown by the raw total sale probability estimates, there are a number of lowest available class/DCP combinations with no booking in any of the samples, which leads to missing data points in the figure (for AMD/AMR computation purposes the missing values are treated as zero). In addition, the values that can be computed show substantial variability between DCPs—the total sale probability of FC 4, for example,

[4]We also explored estimating choice probabilities from the hybrid forecasting and fare adjustment cases, but did not find meaningful differences in trends or conclusions.

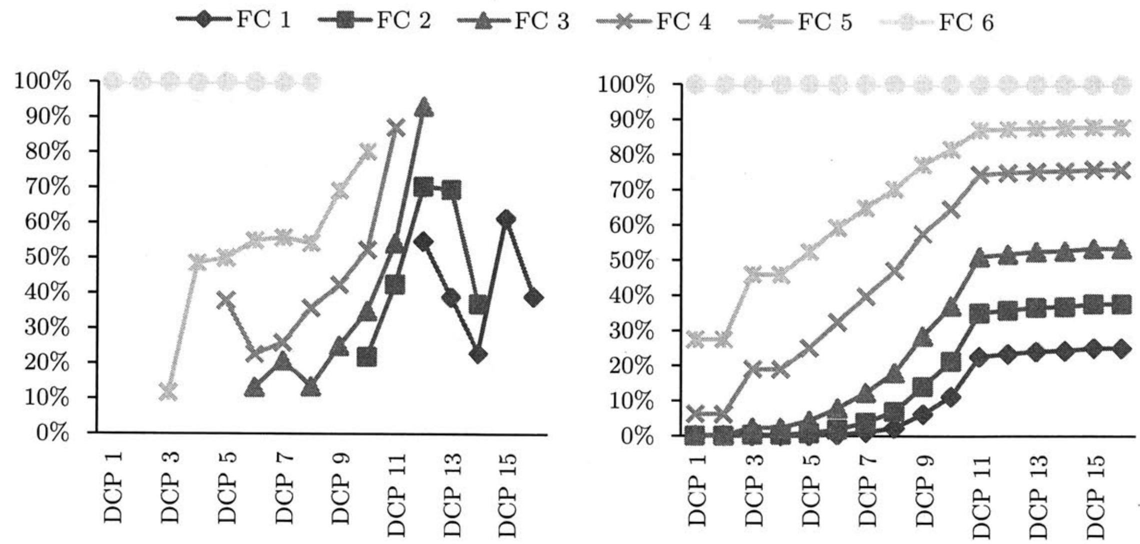

Figure 7.27: Network A2TWO estimated total sale $TP_{dcp}(k)$ for Airline 1 using the raw estimator (left) or heuristic estimator with FRAT5 D (right) (medium demand, 100% simultaneous passengers, \$50 ancillary price, \$50 mean ancillary disutilities). Choice probabilities estimated from baseline case. Total sale probabilities shown prior to any gap-filling.

drops from 38% in DCP 5 to 23% in DCP 6, before climbing to 87% in DCP 11. The total sale probabilities from the heuristic estimator are derived entirely from the sell-up estimate (Equation 6.23), which relies on the input FRAT5 curve. Therefore, the heuristic total sale probability curves are much smoother and have no missing data points. Comparing the estimated total sale probabilities to the true values that would be obtained if the airline knew the customer choice model (Figure 7.2) shows several important differences. The heuristic estimates are in general much cleaner and more reflective of the trends in the true values, but the heuristic estimates have lower total sale probabilities for all fare classes in early DCPs, and have lower total sale probabilities for the highest value fare classes in the DCPs close to departure. This will, all else equal, lead to less aggressive availability decisions by AMD/AMR.[5]

Figure 7.28 shows the change in Airline 1's AMR fare vs using true choice probabilities

[5]The heuristic estimate total sale probabilities are a function of the selected FRAT5 curve. We display FRAT5 D because it produces the highest AMD/AMR revenue for Airline 1.

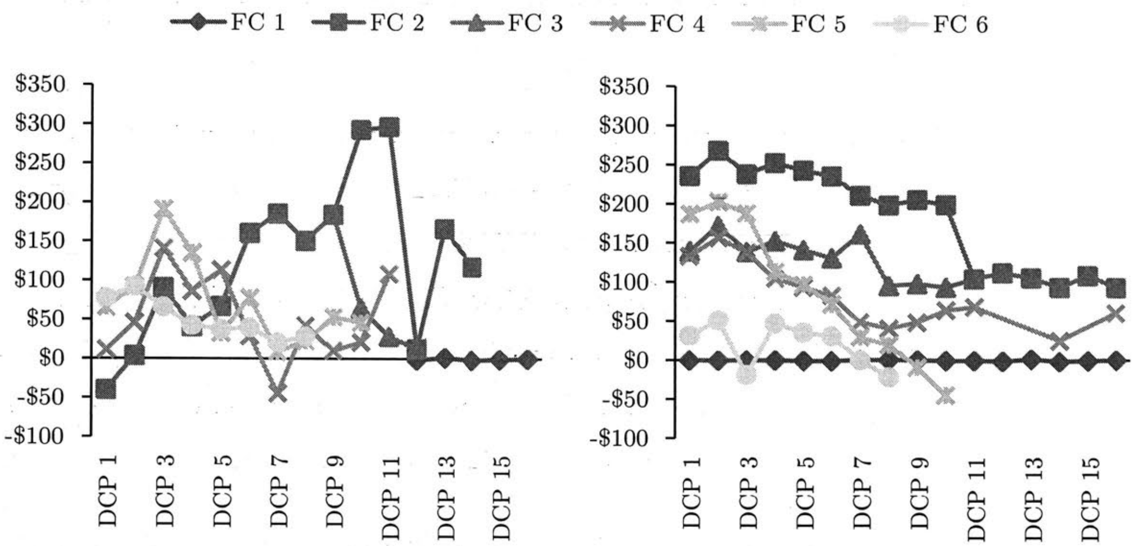

Figure 7.28: Network A2TWO change in Airline 1 AMR fare due to estimated choice probabilities vs true choice probabilities using the raw estimator (left) or heuristic estimator with FRAT5 D (right) (medium demand, 100% simultaneous passengers, \$50 ancillary price, \$50 mean ancillary disutilities). Choice probabilities estimated from baseline case.

when the airline uses either the raw estimator or the heuristic estimator with FRAT5 D. With either estimator, adjusted fares tend to be higher. When the adjusted fare for class i increases, its availability also increases (all else equal), and the availability of lower value classes may decrease (all else equal). Thus it is difficult to determine directly the impact of choice probability estimation on availability of lower value classes, but Figure 7.28 makes clear that the estimation processes will substantially affect adjusted fares.

Symmetric AMD/AMR simulation results when both airlines use true choice probabilities or estimated choice probabilities are shown in Tables 7.15 and 7.16. AMD/AMR with true choice probabilities has the highest total revenue, with a gain of +1.6% over baseline. Use of raw estimated choice probabilities leads to a substantial decrease in total revenue (-3.1% for Airline 1, -1.2% for Airline 2) vs the baseline, with decreases in both ticket revenue and ancillary revenue for both airlines. Load factor for both airlines increases while total yield decreases, suggesting that the raw estimated probabilities lead to less aggressive booking

Table 7.15: Network A2TWO Airline 1 simulation results for baseline and symmetric AMD and AMR cases with various choice probability estimation methods (medium demand, 100% simultaneous passengers, $50 ancillary price, $50 ancillary disutilities). "True" refers to choice probabilities computed with accurate knowledge of the customer choice model; "raw" and "heuristic" are estimated choice probabilities.

	Baseline	True	Raw	Heuristic FRAT5 C	Heuristic FRAT5 D
Airline 1					
Ticket Revenue	$22,386	$22,798	$21,657	$22,682	$22,710
Ancillary Revenue	$1,825	$1,807	$1,811	$1,796	$1,824
Total Revenue	$24,211	$24,605	$23,468	$24,478	$24,534
Load Factor	84.3%	82.9%	84.7%	82.1%	83.5%
Total Yield	22.09	22.84	21.31	22.93	22.59
Ancillary Purchase Rate	33.3%	33.5%	32.9%	33.6%	33.6%
Change from baseline					
Ticket Revenue		+1.8%	-3.3%	+1.3%	+1.4%
Ancillary Revenue		-1.0%	-0.8%	-1.6%	-0.1%
Total Revenue		+1.6%	-3.1%	+1.1%	+1.3%
Load Factor		-1.4 pts	+0.4 pts	-2.2 pts	-0.8 pts
Total Yield		+3.4%	-3.5%	+3.8%	+2.3%
Ancillary Purchase Rate		+0.2 pts	-0.4 pts	+0.3 pts	+0.3 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.2%	0.2%	0.1%	0.1%
t-statistic		10.71	-15.22	7.67	12.17
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

Table 7.16: Network A2TWO Airline 2 simulation results for baseline and symmetric AMD and AMR cases with various choice probability estimation methods (medium demand, 100% simultaneous passengers, \$50 ancillary price, \$50 ancillary disutilities). "True" refers to choice probabilities computed with accurate knowledge of the customer choice model; "raw" and "heuristic" are estimated choice probabilities.

	Baseline	True	Raw	Heuristic FRAT5 C	Heuristic FRAT5 D
Airline 2					
Ticket Revenue	\$22,388	\$22,786	\$22,109	\$22,736	\$22,621
Ancillary Revenue	\$1,822	\$1,803	\$1,818	\$1,800	\$1,812
Total Revenue	\$24,210	\$24,589	\$23,927	\$24,536	\$24,433
Load Factor	84.2%	82.9%	84.5%	82.3%	83.3%
Total Yield	22.13	22.81	21.77	22.92	22.56
Ancillary Purchase Rate	33.3%	33.5%	33.1%	33.6%	33.5%
Change from baseline					
Ticket Revenue		+1.8%	-1.2%	+1.6%	+1.0%
Ancillary Revenue		-1.0%	-0.2%	-1.2%	-0.5%
Total Revenue		+1.6%	-1.2%	+1.3%	+0.9%
Load Factor		-1.3 pts	+0.4 pts	-1.8 pts	-0.8 pts
Total Yield		+3.1%	-1.6%	+3.6%	+1.9%
Ancillary Purchase Rate		+0.1 pts	-0.2 pts	+0.3 pts	+0.2 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.1%	0.2%	0.1%	0.1%
t-statistic		10.55	-6.29	9.38	8.59
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 1,999$.

limits than even the baseline case, and certainly less aggressive than AMD/AMR with the true choice probabilities.

The heuristic estimates produce better revenue performance than the raw estimates, increasing total revenue for both airlines when they both use FRAT5 C (+1.1% for Airline 1 and +1.3% for Airline 2) or FRAT5 D (+1.3% for Airline 1 and +0.9% for Airline 2). FRAT5 C performs better for Airline 2 than for Airline 1, and FRAT5 D performs better for Airline 1 than for Airline 2. With either FRAT5, AMD/AMR increases ticket revenue and decreases ancillary revenue vs baseline. The decrease in ancillary revenue is driven by a decrease in load factor that is only slightly offset by an increase in the ancillary service purchase rate.

Although this is a symmetric competition case, in the sense that both airlines use the same revenue management forecasting and optimization methods, and use the same approach to calculating or estimating choice probabilities, it is important to note that the two airlines perform the choice probability estimation separately, on their own booking databases. The difference in total revenue *between the two airlines* for AMD/AMR with heuristic choice probabilities is not statistically significant ($t = 0.63$, $p = 0.526$, $df = 1,999$ for FRAT5 C; $t = -1.14$, $p = 0.255$, $df = 1,999$ for FRAT5 D; using Airline 2 as the "test" case and Airline 1 as the "baseline," pairing by sample as usual). However, clearly, there are some differences in the results of the estimation processes between the airlines, and the raw choice probabilities show a significant difference between the two airlines ($t = 4.67$, $p < 0.001$, $df = 1,999$; using Airline 2 as the "test" case and Airline 1 as the "baseline," pairing by sample as usual). The small random differences in booking histories for the two airlines lead to slightly different choice probability estimates, with greater differences for the lowest available class/DCP combinations that receive few bookings. In an extreme case, with the raw estimation heuristic, one airline could receive one booking for a particular lowest available class/DCP combination while the other receives none; the airline that receives the booking would be able to estimate choice probabilities for that combination and could compute non-zero AMD demands and AMR fares, leading to different availability decisions.

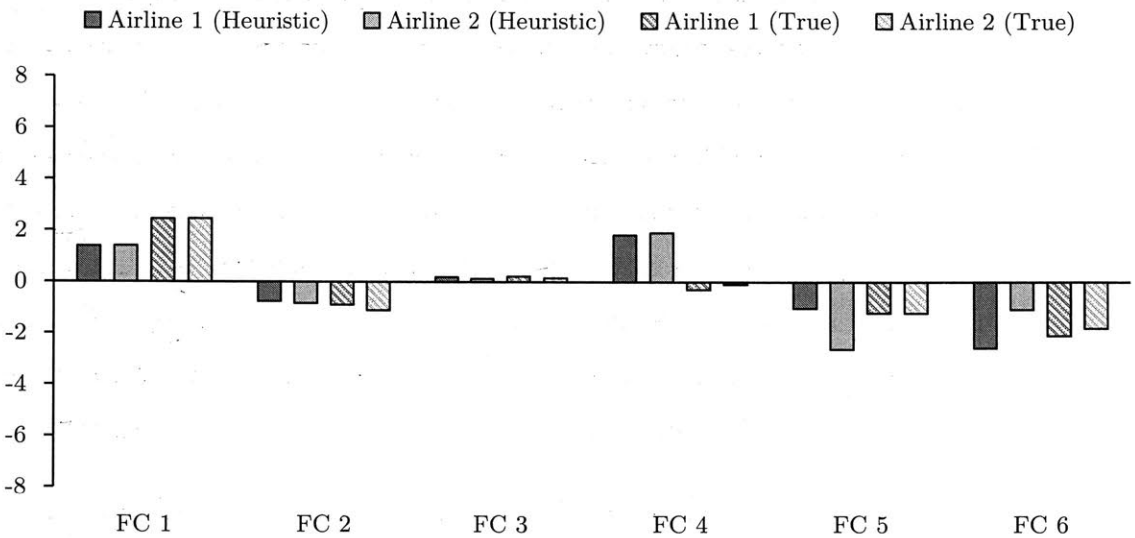

Figure 7.29: Network A2TWO change in bookings by fare class due to symmetric AMD and AMR with heuristic estimated (FRAT5 D) or true choice probabilities vs baseline (medium demand, 100% simultaneous passengers, \$50 ancillary price, \$50 ancillary disutility).

The raw estimation method is more susceptible to these influences because it relies only on the historical database; the heuristic estimation method is more stable because it includes both the "neighbor matching" process and sell-up estimates based on the input FRAT5 value. The remainder of this section focuses on the heuristic estimation methods because of their stability.

Figure 7.29 shows the change in bookings vs baseline when the airlines implement AMD and AMR with either the heuristic estimated choice probabilities (with FRAT5 D) or with the known true choice probabilities. In all cases, the result is a reduction in FC 5 and FC 6 bookings, and an increase in FC 1 bookings. However, the heuristic estimated probabilities have a smaller increase in FC 1 bookings, and instead have an increase in FC 4 bookings.

The sensitivity of AMD/AMR revenue gains with heuristic estimated choice probabilities is shown in Figure 7.30. AMD/AMR increases total revenue vs baseline for Airline 1 in all the tested scenarios, including with 100% simultaneous and 100% sequential passenger

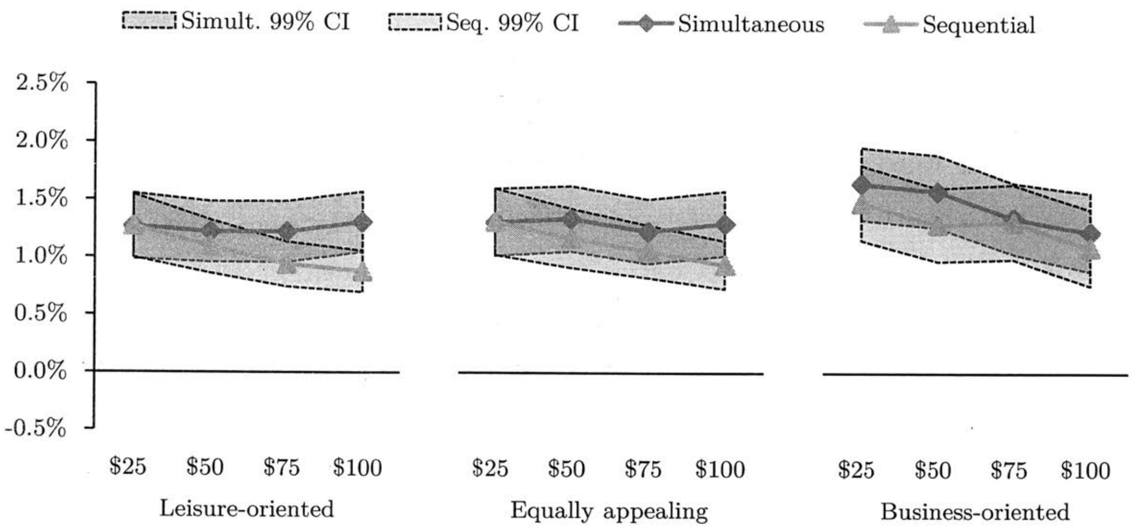

Figure 7.30: Network A2TWO change in Airline 1 total revenue due to symmetric AMD and AMR with heuristic estimated choice probabilities (FRAT5 D) vs baseline with 99% confidence interval with various passenger behaviors, ancillary prices, and disutility scenarios (medium demand). $df = 1,999$.

behavior, and the change is statistically significant ($p < 0.01$, $df = 1,999$). As with previous AMD/AMR results, the revenue benefits are higher with simultaneous passengers, and the difference in performance between the two passenger types increase with higher ancillary prices and when leisure passengers value the service more than business passengers.

These results illustrate that even with estimated choice probabilities, the AMD/AMR formulation can increase total revenue compared to the baseline case. However, the revenue benefit of AMD/AMR is reduced when choice probabilities are estimated, and the estimation process can be sensitive to small variations in historical bookings. Choice probabilities estimated based on a combination of input price elasticity (FRAT5) and historical bookings produce booking limits that are not as effective at increasing ticket revenue as the true choice probabilities, and choice probabilities estimated only from historical data can lead to revenue decreases.

Overall, our simulations in network A2TWO illustrate that AMD/AMR is capable of producing revenue gains in a simple competitive environment, even when the airlines do not adjust their choice probabilities to explicitly account for competition. Results under symmetric competition are similar to the results from network A1ONE: AMD and AMR increase ticket revenue, decrease ancillary revenue, and increase total revenue, with the decrease in ancillary revenue driven by a decrease in load factor that is only slightly offset by an increase in ancillary purchase rate. Although the theoretical formulation does not include any form of competition, our tests in this limited environment illustrate that increases in total revenue remain possible when only one of the airlines implements AMD/AMR, while the other maintains its simplistic independent demand (baseline) forecasting and optimization models.

7.5 Two Airline, Many Flight Leg Network (D6)

In our final set of experiments, we apply AMD and AMR in a large network with competition and connecting flights. The focus of this section is on comparing the performance of AMD and AMR in this more realistic setting to the findings in the toy networks A1ONE and A2TWO. Network D6, used for studies of the optimizer increment in Section 5.2.4, has two hypothetical airlines operating connecting hubs. Each airline has 252 flight legs which serve a total of 482 different origin-destination markets. Both airline serve every market, and both airlines offer six economy fare classes and one ancillary service. Figure 5.11 shows a map of the network, and Table 5.8 lists fare information for the network. Fares and demands in D6 were previously calibrated based on data provided by airline members of the MIT PODS Research Consortium. In general, passengers in network D6 have higher budgets than in A1ONE or A2TWO, so the benefit of revenue management methods that correct for buy-down (i.e. hybrid forecasting/fare adjustment and AMD/AMR) is greater.

In all our simulations, both airlines use network-based revenue management optimization models (DAVN, see Sections 3.2 and 4.1). As a baseline, the airline uses a standard (independent demand) forecasting model and optimizes on filed fares. We primarily focus on a medium demand case (about 83% baseline load factor) where the ancillary service is priced at 40% of the FC 6 fare, which is the same ratio as the $50 ancillary price in networks A1ONE and A2TWO. We use the same disutility scenarios as in networks A1ONE and A2TWO (see Table 5.5). All D6 simulations have two trials of 400 unburned samples each (with an additional 200 samples per trial "burned" to warm up the forecasting models), producing a total of 800 samples per simulation. As with all other PODS studies, comparisons are made on a pairwise basis (by sample) between different simulations that have the same demand generation parameters (i.e. overall demand level and ancillary disutilities).

Baseline statistics are shown in Table 5.9. As with network A1ONE, we consider four other revenue management configurations: optimizer increment (OI), hybrid forecasting and fare adjustment (HF/FA), a combination of the two approaches (OI + HF/FA), and AMD/AMR. We conduct tests with symmetric competition, where both airlines use the same RM methods, and asymmetric competition, where Airline 1 varies its RM method and Airline 2 always uses its baseline independent demand forecasting and optimization.

When using the optimizer increment the airline estimates average ancillary revenue per passenger according to Method 1, as described in Chapter 5. When using hybrid forecasting and fare adjustment the airline uses FRAT5 curve A. Note that curve A is more aggressive than curve C used in network A1ONE and A2TWO; A was selected because of the higher budgets of passengers in network D6.

When using AMD/AMR in network D6, the airlines always estimate choice probabilities using the heuristic estimator described in Section 6.4.2. The airlines estimate conditional purchase probabilities separately, using their own historical booking databases, based on the booking and ancillary purchase data from the baseline case. The airlines always use the same FRAT5 curves for sell-up estimates.

We begin by "tuning" the FRAT5 curve used by the airlines for sell-up estimation with

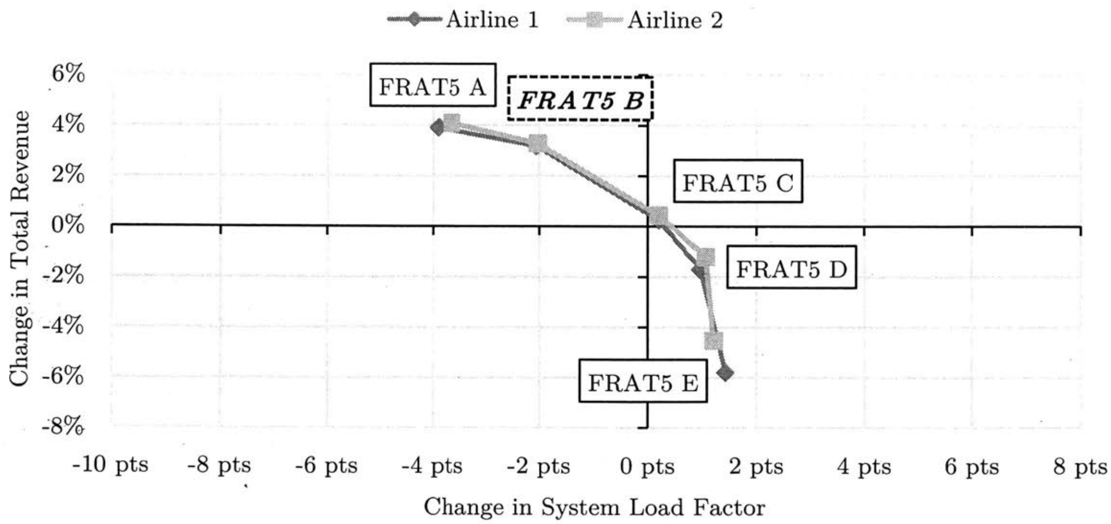

Figure 7.31: Network D6 change in total revenue and change in system load factor due to symmetric AMD and AMR with heuristic estimated choice probabilities using various FRAT5 curves (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario).

AMD/AMR. Figure 7.31 shows the change in total revenue for Airline 1 and Airline 2 plotted against the change in system load factor for the symmetric AMD/AMR cases with various FRAT5 curves (compared to the baseline case). As expected, the least aggressive FRAT5 curves result in load factor increases and total revenue decreases for the airlines. The most aggressive FRAT5 curve, A, produces the highest total revenue, but also the largest load factor decrease. Because FRAT5 B produces a total revenue increase nearly as large as FRAT5 A, but has only half the load factor decrease, we select to use FRAT5 B *for all subsequent AMD/AMR simulations in this section.*

The AMR adjusted fares for Airline 1 (averaged across all markets) when the airline uses FRAT5 B as its sell-up estimate are shown in Figure 7.32. Note that AMR fares are both higher and lower than the filed fares—the increase comes from accounting for ancillary revenue, and the decrease comes from accounting for the risk of buy-down. In contrast to networks A1ONE and A2TWO, where the booking policies of offering FC 2 or FC 5 as the

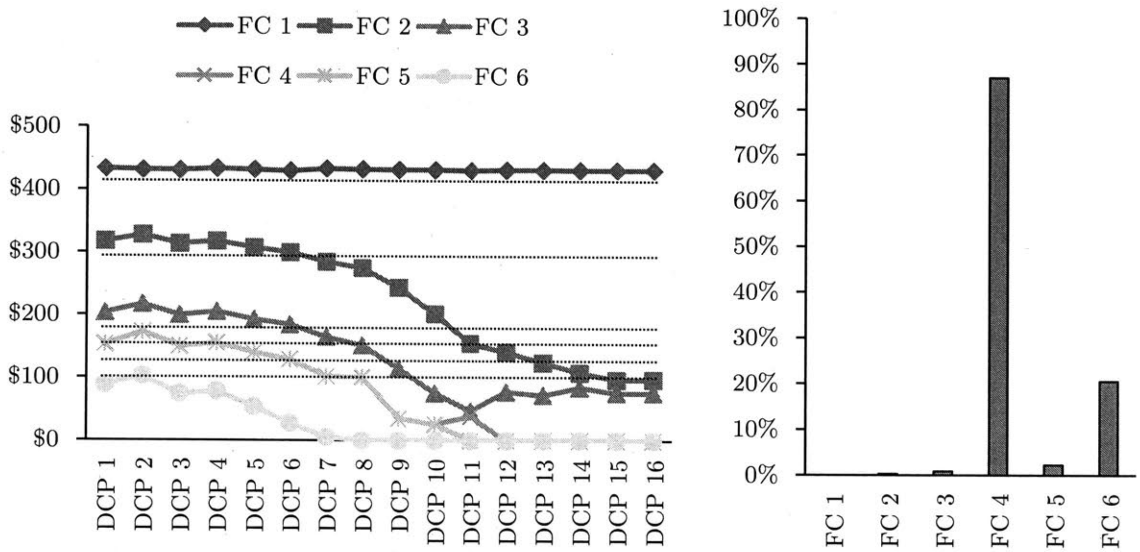

Figure 7.32: Network D6 Airline 1 network-averaged AMR fares (left) and fare class inefficiency (right) (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario). Network-averaged fares are unweighted. Dotted black lines indicate filed fares (with FC 6 lowest and FC 1 highest). Fare class inefficiency measured as percentage of itineraries and DCPs in which the fare class is closed by gap-filling.

lowest available fare are inefficient, in network D6 the policy of offering FC 4 as the lowest available fare is often inefficient. Because fares vary by market, classes may be efficient in one market but inefficient in other markets. On a network average basis, FC 4 is only efficient during DCP 11, indicated by the FC 4 AMR fare being equal to the FC 5 AMR fare for earlier DCPs in Figure 7.32.

7.5.1 Symmetric Competition

Tables 7.17 and 7.18 show simulation results for each of the four symmetric test cases as well as the baseline case. As with other networks tested, the optimizer increment leads to a decrease in total revenue whether applied alone (-0.3% for both airlines) or applied in conjunction with HF/FA (reduces HF/FA gain 0.5 pts for Airline 1 and by 0.4 pts for Airline 2). The optimizer increment produces an increase in ancillary revenue which is consistently outweighed by a decrease in ticket revenue. Hybrid forecasting and fare adjustment decrease load factor and increase total yield, which increases ticket revenue, decreases ancillary revenue, and increases total revenue.

AMD and AMR also have similar effects as in the smaller networks, but in D6 the magnitude of effects is exaggerated. AMD and AMR increase ticket revenue (3.8% for Airline 1; 3.9% for Airline 2), decrease ancillary revenue (-3.6% for Airline 1; -3.2% for Airline 2), and increase total revenue (3.2% for Airline 1; 3.3% for Airline 2). In networks A1ONE and A2TWO, AMD and AMR produced revenue gains of only 1.6—1.8%. However, the incremental improvement in total revenue over HF/FA is similar between D6 (where AMD/AMR has a total revenue gain 0.6 pts higher than HF/FA for Airline 1 and 0.7 pts higher for Airline 2) and networks A1ONE and A2TWO (where the AMD/AMR total revenue gain is 0.4 pts to 0.8 pts higher than the total revenue gain of HF/FA).

The primary effects of AMD/AMR—increased ticket revenue, decreased bookings, and decreased ancillary revenue—are seen in many, but not all, markets across the network. As

Table 7.17: Network D6 Airline 1 simulation results for baseline and symmetric experimental cases (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario).

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
Airline 1					
Ticket Revenue	$1,326,442	$1,321,852	$1,366,610	$1,358,186	$1,376,276
Ancillary Revenue	$115,904	$116,207	$113,255	$114,195	$111,721
Total Revenue	$1,442,346	$1,438,059	$1,479,865	$1,472,381	$1,487,997
Load Factor	83.7%	83.8%	82.3%	82.7%	81.7%
Total Yield	14.04	13.99	14.67	14.52	14.85
Ancillary Sales Rate	42.6%	42.6%	42.3%	42.5%	42.2%
Change from baseline					
Ticket Revenue		-0.3%	+3.0%	+2.4%	+3.8%
Ancillary Revenue		+0.3%	-2.3%	-1.5%	-3.6%
Total Revenue		-0.3%	+2.6%	+2.1%	+3.2%
Load Factor		+0.1 pts	-1.5 pts	-1.1 pts	-2.1 pts
Total Yield		-0.4%	+4.5%	+3.4%	+5.8%
Ancillary Sales Rate		+0.1 pts	-0.3 pts	-0.1 pts	-0.4 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.0%	0.1%	0.1%	0.1%
t-statistic		-29.31	37.83	31.98	44.01
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 799$.

Table 7.18: Network D6 Airline 2 simulation results for baseline and symmetric experimental cases (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario).

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
Airline 2					
Ticket Revenue	$1,307,530	$1,303,664	$1,347,674	$1,339,838	$1,357,881
Ancillary Revenue	$115,431	$115,710	$112,868	$113,830	$111,756
Total Revenue	$1,422,961	$1,419,374	$1,460,542	$1,453,668	$1,469,637
Load Factor	83.0%	83.0%	81.6%	82.1%	81.0%
Total Yield	13.46	13.42	14.05	13.90	14.25
Ancillary Sales Rate	42.5%	42.5%	42.2%	42.4%	42.2%
Change from baseline					
Ticket Revenue		-0.3%	+3.1%	+2.5%	+3.9%
Ancillary Revenue		+0.2%	-2.2%	-1.4%	-3.2%
Total Revenue		-0.3%	+2.6%	+2.2%	+3.3%
Load Factor		+0.0 pts	-1.4 pts	-0.9 pts	-2.0 pts
Total Yield		-0.3%	+4.4%	+3.3%	+5.9%
Ancillary Sales Rate		+0.1 pts	-0.2 pts	-0.1 pts	-0.2 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.0%	0.1%	0.1%	0.1%
t-statistic		-23.37	38.15	33.20	44.33
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 799$.

Table 7.19: Network D6 distribution across network of symmetric AMD and AMR increases in ticket revenue, decreases in bookings, and decreases in ancillary revenue vs baseline (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario).

Portion of markets with...	Airline 1	Airline 2
Increased ticket revenue	83.6%	81.3%
Decreased bookings	68.7%	62.9%
Decreased ancillary revenue	73.9%	67.0%

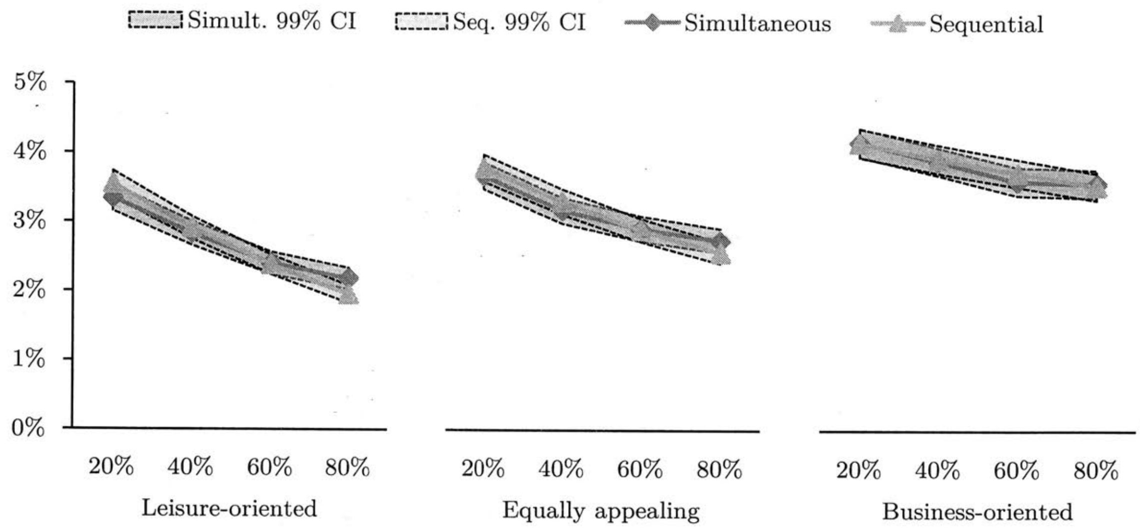

Figure 7.33: Network D6 change in Airline 1 total revenue vs baseline due to symmetric AMD and AMR with 99% confidence interval with various passenger behaviors, ancillary prices, and disutility scenarios. $df = 799$. Ancillary price expressed as percentage of FC 6 fare (20% to 80%).

shown in Table 7.19, both airlines see increased ticket revenue due to AMD/AMR in more than 80% of their markets. More than 60% of markets have fewer bookings, and more 70% of Airline 1's markets produce less ancillary revenue. It is important to note that the goal of AMD/AMR is *not* to increase ancillary revenue, but to increase total revenue.

Although this analysis has focused on the case where the ancillary service is priced at 40% of the FC 6, both passenger segments value the service equally, and all passengers are simultaneous, symmetric AMD/AMR in network D6 produces revenue increases in a variety of settings. Figure 7.33 shows the change in Airline 1 total revenue for lower ancillary prices (equal to 20% of the FC 6 fare) and higher ancillary prices (equal to 60% and 80% of the FC 6 fare), for the leisure-oriented and business-oriented disutility scenarios, as well as for sequential behavior. In all cases the change in total revenue vs baseline is statistically significant ($p < 0.01$, $df = 799$).

7.5.2 Asymmetric Competition

In asymmetric competition, only Airline 1 uses AMD/AMR; Airline 2 maintains its baseline independent demand forecasting and optimization models. Asymmetric competition results in substantial competitive feedback between the two airlines. Figure 7.34 shows the change in total revenue and change in system load factor vs baseline for the two airlines when Airline 1 uses AMD and AMR with various FRAT5 curves for sell-up estimates. The symmetric case is also shown on the figure for comparison. With asymmetric competition, Airline 1 has lower total revenues with AMD/AMR and has lower load factors—asymmetric FRAT5 B produces only two thirds the revenue gain but more than twice the load factor loss as symmetric FRAT5 B, for example. Airline 2, however, experiences an increase in load factor and in total revenues for all cases where Airline 1 implements AMD/AMR. As in the symmetric case, we restrict the remainder of our AMD/AMR analysis to cases where Airline 1 uses FRAT5 B.

The change in system bookings by fare classes vs baseline for both airlines for symmetric and asymmetric AMD/AMR are shown in Figure 7.35. In the symmetric case, both airlines experience reductions in bookings in FC 6 (the lowest value class) and FC 4 (which is frequently inefficient). In the asymmetric case, Airline 1 booking changes show similar patterns, but with typically smaller magnitudes: the reduction in FC 6 and increase in FC 5 bookings *decreases* by 11% and 41%, respectively. Airline 1 does accept more FC 1 bookings with asymmetric competition (by 77%). Airline 2, which does not change its forecasting or optimization processes, has an increase in bookings vs baseline in FC 2 through FC 6, and a reduction in bookings in FC 1.

The increase in Airline 1 total revenue is driven by an increase in ticket yield, which is offset by the decrease in system load factor, as shown in Figure 7.36. Airline 2, on the other hand, has opposite trends: ticket yields decrease (reflecting the increase in bookings in FC 2 to FC 6 and decrease in FC 1 bookings) but system load factors increase. Revenue changes for both airlines are dominated by ticket revenue changes; the contribution of changes in

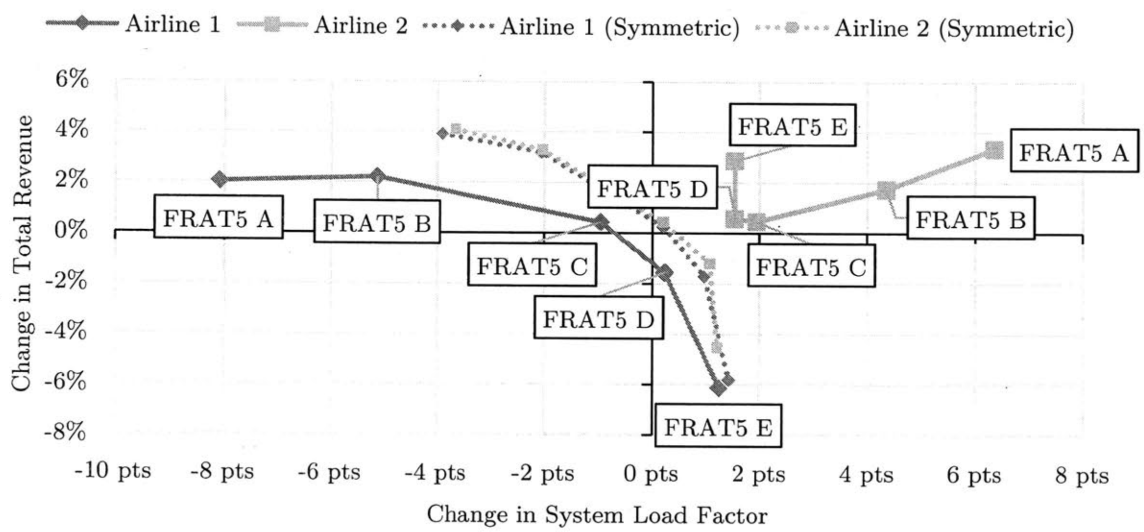

Figure 7.34: Network D6 change in total revenue and change in system load factor due to asymmetric AMD and AMR with heuristic estimated choice probabilities using various FRAT5 curves (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario). Airline 2 always uses independent demand (baseline) forecasting and optimization; FRAT5 labels refer to Airline 1 configuration.

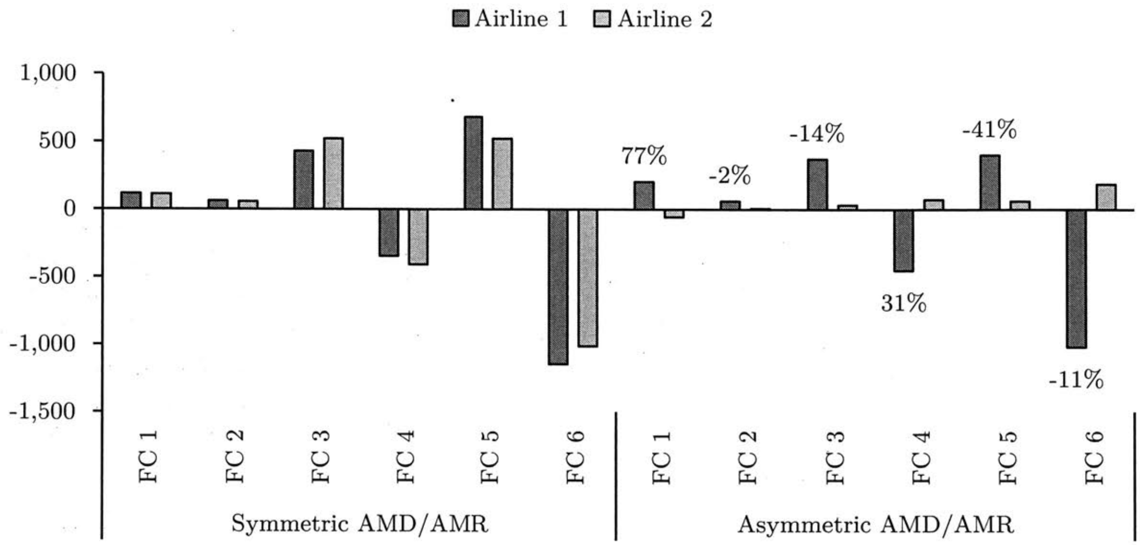

Figure 7.35: Network D6 change in system bookings by fare class for symmetric AMD and AMR and asymmetric AMD and AMR (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario). Percentages indicate difference in booking change vs symmetric AMD and AMR.

ancillary revenue to the change in total revenue is small.

Tables 7.20 and 7.21 shows complete simulation results for the asymmetric case. In contrast to previous studies, Airline 1 has the highest total revenue when it uses hybrid forecasting and fare adjustment, not AMD/AMR. Although AMD/AMR is more theoretically appealing because it accounts for ancillary revenue and accounts for the effects of ancillary services on passenger choice of fare class, in this competitive environment it does not provide the highest revenue. Note that Airline 2 experiences an increase in its total revenue when Airline 1 uses AMD/AMR or HF/FA, but the increase is greater when Airline 1 uses AMD/AMR.

Figure 7.37 illustrates the difference in performance between HF/FA and AMD/AMR for Airline 1. Both methods increase ticket yield and decrease load factor, but the much larger load factor loss (indicated as ticket volume) with AMD/AMR results in the lower overall revenue increase.

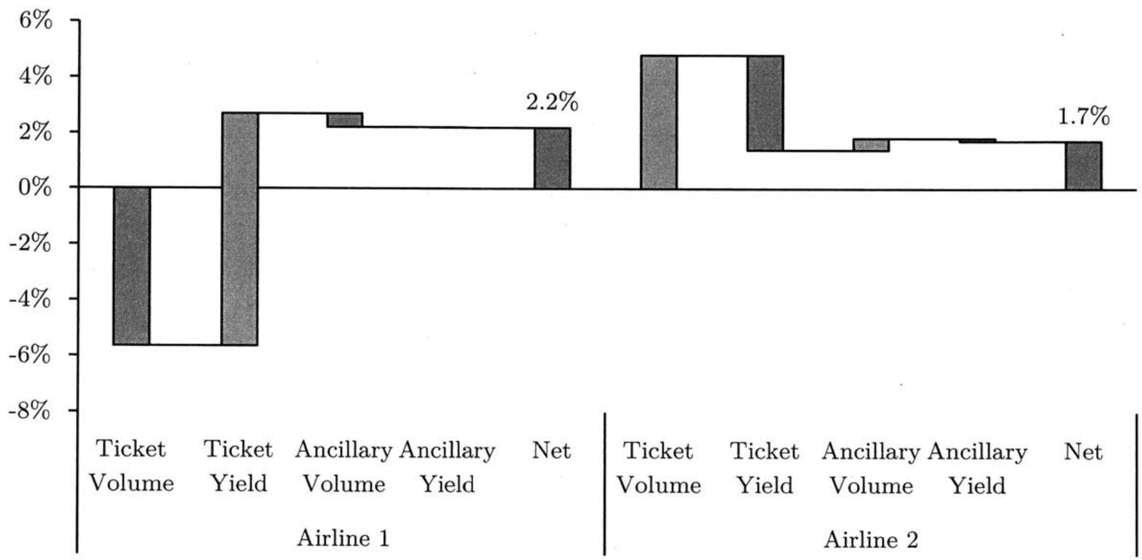

Figure 7.36: Network D6 change in Airline 1 and Airline 2 total revenue by change in ticket and ancillary sales volumes and yields due to asymmetric AMD and AMR vs baseline (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario). "Volume" columns are portion of ticket or ancillary revenue change attributable to change in load factor; "yield" columns are portion of ticket or ancillary revenue change attributable to change in average fare or average ancillary revenue earned per passenger. All values expressed as percentage change from baseline total revenue.

Table 7.20: Network D6 Airline 1 simulation results for baseline and asymmetric experimental cases (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario). Airline 2 always uses independent demand (baseline) forecasting and optimization; labels refer to Airline 1 configuration.

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
Airline 1					
Ticket Revenue	\$1,326,442	\$1,321,524	\$1,366,790	\$1,362,133	\$1,365,474
Ancillary Revenue	\$115,904	\$116,345	\$111,747	\$113,311	\$108,661
Total Revenue	\$1,442,346	\$1,437,869	\$1,478,537	\$1,475,444	\$1,474,135
Load Factor	83.7%	84.0%	80.5%	81.6%	78.6%
Total Yield	14.04	13.96	14.96	14.75	15.29
Ancillary Sales Rate	42.6%	42.6%	42.6%	42.6%	42.5%
Change from baseline					
Ticket Revenue		-0.4%	+3.0%	+2.7%	+2.9%
Ancillary Revenue		+0.4%	-3.6%	-2.2%	-6.2%
Total Revenue		-0.3%	+2.5%	+2.3%	+2.2%
Load Factor		+0.2 pts	-3.2 pts	-2.1 pts	-5.1 pts
Total Yield		-0.6%	+6.6%	+5.1%	+8.9%
Ancillary Sales Rate		+0.0 pts	-0.0 pts	+0.1 pts	-0.1 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.0%	0.1%	0.1%	0.1%
t-statistic		-29.72	27.86	29.55	20.22
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 799$.

Table 7.21: Network D6 Airline 2 simulation results for baseline and asymmetric experimental cases (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario). Airline 2 always uses independent demand (baseline) forecasting and optimization; labels refer to Airline 1 configuration.

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
Airline 2					
Ticket Revenue	\$1,307,530	\$1,307,938	\$1,314,417	\$1,310,701	\$1,327,608
Ancillary Revenue	\$115,431	\$115,328	\$117,465	\$116,594	\$119,993
Total Revenue	\$1,422,961	\$1,423,266	\$1,431,882	\$1,427,295	\$1,447,601
Load Factor	83.0%	82.9%	85.3%	84.5%	87.3%
Total Yield	13.46	13.48	13.18	13.26	13.01
Ancillary Sales Rate	42.5%	42.5%	42.2%	42.3%	42.2%
Change from baseline					
Ticket Revenue		+0.0%	+0.5%	+0.2%	+1.5%
Ancillary Revenue		-0.1%	+1.8%	+1.0%	+4.0%
Total Revenue		+0.0%	+0.6%	+0.3%	+1.7%
Load Factor		-0.1 pts	+2.3 pts	+1.5 pts	+4.4 pts
Total Yield		+0.1%	-2.1%	-1.5%	-3.3%
Ancillary Sales Rate		+0.0 pts	-0.2 pts	-0.2 pts	-0.2 pts
Significance of Change in Total Revenue from Baseline					
Standard Error		0.0%	0.0%	0.0%	0.0%
t-statistic		2.27	24.18	13.82	47.52
p-value		0.023	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = 799$.

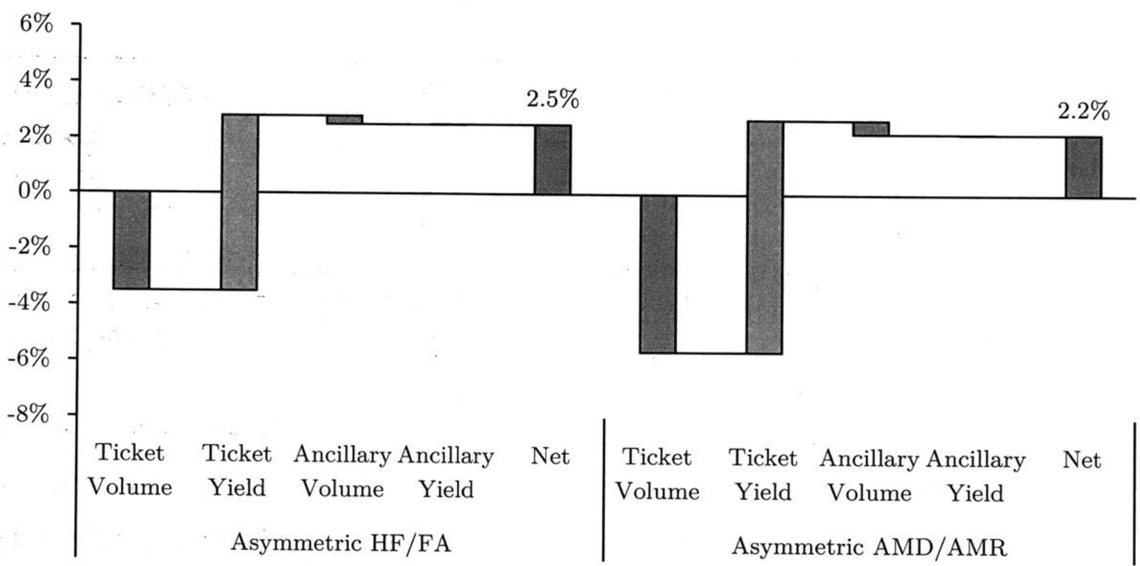

Figure 7.37: Network D6 change in Airline 1 total revenue by change in ticket and ancillary sales volumes and yields due to asymmetric HF/FA or AMD and AMR vs baseline (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario). "Volume" columns are portion of ticket or ancillary revenue change attributable to change in load factor; "yield" columns are portion of ticket or ancillary revenue change attributable to change in average fare or average ancillary revenue earned per passenger. All values expressed as percentage change from baseline total revenue.

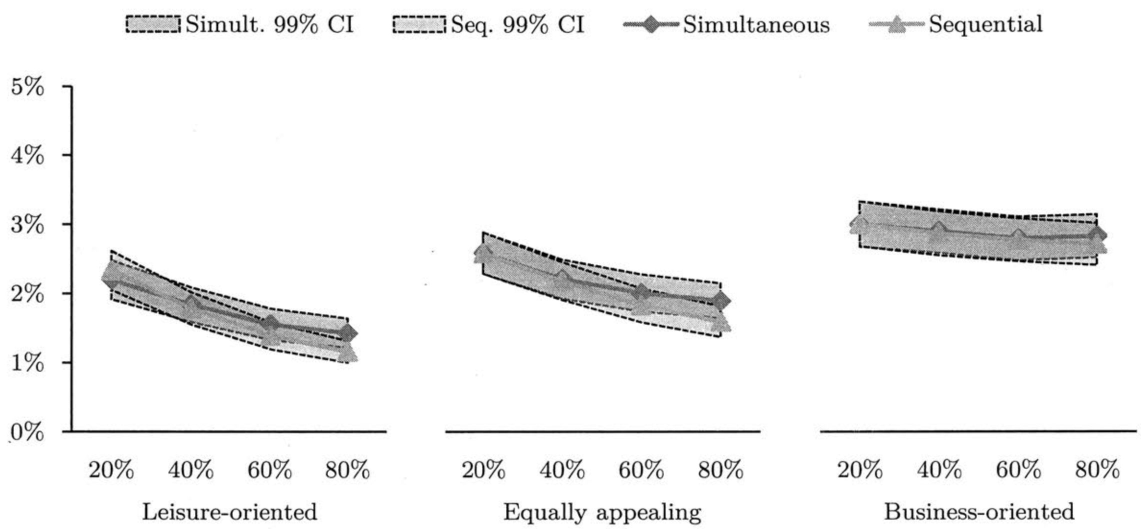

Figure 7.38: Network D6 change in Airline 1 total revenue vs baseline due to asymmetric AMD and AMR with 99% confidence interval with various passenger behaviors, ancillary prices, and disutility scenarios. $df = 799$. Ancillary price expressed as percentage of FC 6 fare (20% to 80%). Airline 2 always uses independent demand (baseline) forecasting and optimization.

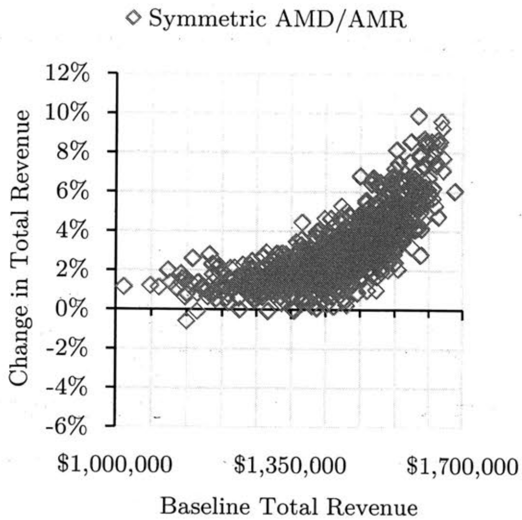

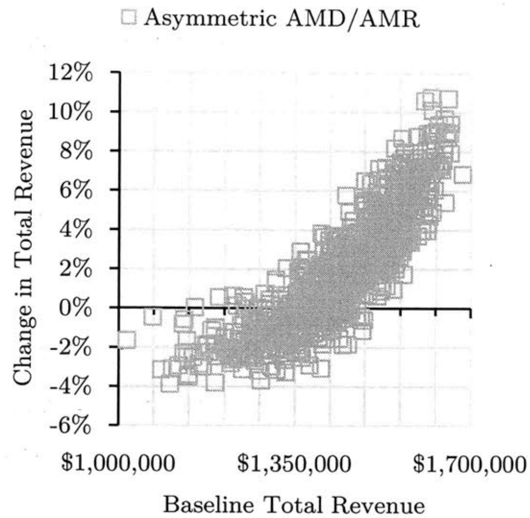

Figure 7.39: Network D6 change in Airline 1 total revenue by sample vs baseline due to symmetric (left) or asymmetric (right) AMD and AMR and baseline total revenue (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario). Each point on the figure represents the outcome of one sample (or departure day).

The change in total revenue provided by AMD/AMR is statistically significant ($t = 20.2$, $p < 0.001$, $df = 799$ for 100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario). These revenue impacts are present for asymmetric competition across a range of passenger behavior types, ancillary prices, and disutility scenarios, as shown in Figure 7.38, although the revenue increase of AMD/AMR is greater with symmetric competition.

7.5.3 Sample-by-Sample Analysis

Thus far, our analysis has primarily assessed the aggregate performance of AMD/AMR, averaging the effects across all trials and (unburned) samples in each simulation. In this section, we look at sample-by-sample changes in total revenue to gain additional insight into the performance of the algorithm.

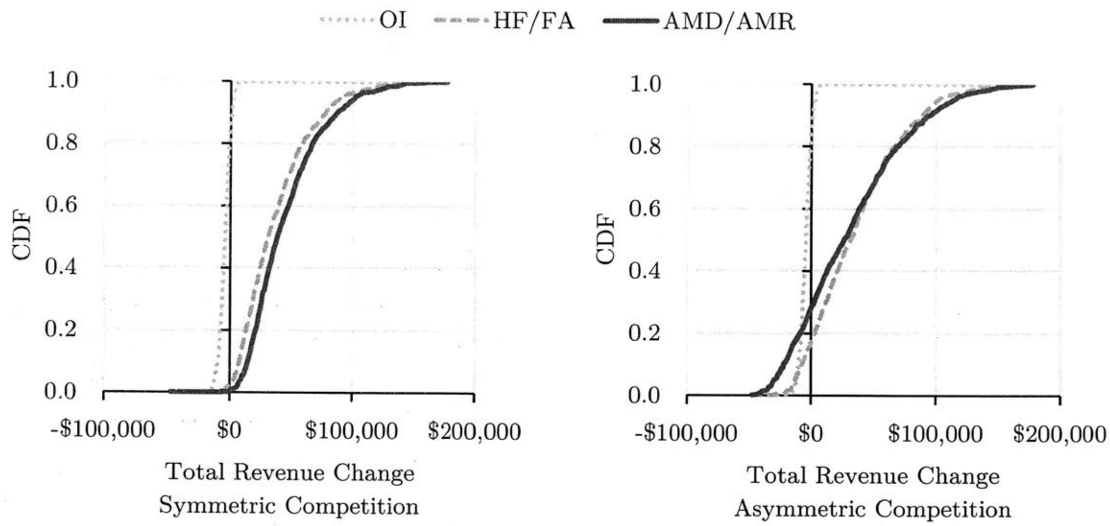

Figure 7.40: Network D6 empirical CDF of Airline 1 total revenue change vs baseline for symmetric (left) and asymmetric (right) cases (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario).

Figure 7.39 shows the change in Airline 1 total revenue due to AMD/AMR and the baseline total revenue for each (unburned) sample in the simulation. In the symmetric case, AMD/AMR produces a higher total revenue for nearly every sample (99.5%). In the asymmetric case, where the aggregate performance of AMD/AMR is reduced, the method increases total revenue in many samples, but a substantial number of samples (28%) see revenue reductions. The increase in total revenue due to AMD/AMR is clearly higher for samples with higher baseline revenue (i.e. samples with higher total demand and/or higher budget passengers), and the effect appears to be non-linear. The largest revenue increase for a given sample occurs in the asymmetric case (10.8%), but the largest revenue loss also occurs in the asymmetric case (-3.9%).

The empirical cumulative distribution function for the change in total revenue by sample is shown in Figure 7.40. In the symmetric case, AMD/AMR clearly outperforms HF/FA, as the AMD/AMR CDF is consistently below/to the right of the HF/FA CDF. Results with asymmetric competition are less clear: the AMD/AMR CDF is above/left of the HF/FA

CDF for lower values and for revenue losses, but is below/right of the HF/FA CDF for the largest changes in total revenue (implying that AMD/AMR has a wider distribution of revenue changes than HF/FA). As shown in Table 7.20, though, asymmetric HF/FA has a greater mean than asymmetric AMD/AMR.

7.6 Experiment Limitations

The results presented in this chapter are dependent on the simulation environment, including network and RM configuration, passenger behavior, and competition. We have used simple but reasonable forecasting and optimization models as the baselines for these tests, but other baseline configurations could have higher or lower revenues and load factors, which would then affect the *relative* performance of AMD/AMR to the modified baseline.

As with the simulations in Chapter 5, we have tested a variety of networks and demand profiles to show how our algorithms perform under different conditions. However, the Passenger Origin-Destination Simulator simplifies the real world processes it represents, and thus our studies have several limitations.

First, real customers may "learn," based on past purchases, about when in the booking window an airline typically offers low fares, and may shift their behavior to shop for future flights when they expect low fares to be available. Our simulations do not include this customer learning feedback between revenue management and the demand arrival process, and it is difficult to estimate the impact of customer learning on our proposed algorithms.

In addition, while the demand for each departure day and the characteristics of each generated customer are stochastic, the distributions for these parameters are constant throughout each simulation; there are no long-term trends like seasonality or macroeconomic shifts. Finally, forecasting or optimization models that result in strong competitive feedbacks (such as in Section 7.5.2) could prompt a competitor to modify its forecasting or optimization approach, or to adjust its flight network to mitigate such effects. Because our simulations are designed to test steady-state conditions, such strategic shifts are not included in our simulations.

7.7 Conclusions

In this chapter we have simulated the performance of our proposed AMD and AMR heuristics in a variety of environments to understand their effect on airline bookings and revenues, in particular in settings in which the fundamental assumptions of the algorithm are violated.

In the single airline, single flight leg network, our formulation produced an increase in total revenue across a wide range of ancillary price and passenger behaviors. AMD and AMR reduced load factor to increase yields, ticket revenues, and total revenues, with small decreases in ancillary revenue driven by the reduced load factor.

We introduced competition with the two airline, two flight leg network A2TWO, and tested AMD/AMR in both a symmetric format, where both airlines use the algorithm, and an asymmetric format, where only one airline uses the algorithm. Although the AMD and AMR formulations do not include any provision for competition, our results indicated that increase in total revenue are possible in both the symmetric and asymmetric environments. We also tested the use of *estimated* choice probabilities, in lieu of assuming that the airline has accurate knowledge of the passenger choice model. Although directly extracting choice probability estimates from the historical booking database ("raw" estimates) resulted in poor revenue performance, total revenue gains with AMD/AMR were possible by incorporating additional processing of the choice probability estimates and integrating external price elasticity estimates (the "heuristic" estimates).

In the much larger network D6, with hundreds of flight legs and hundreds of origin-destination markets, AMD and AMR using the heuristic estimated choice probabilities showed revenue increases under symmetric and asymmetric competition, although competitive feedbacks were substantial in the asymmetric case and other RM methods (hybrid forecasting and fare adjustment) provided higher total revenues. A sample-by-sample analysis showed that the benefit of AMD/AMR is greater for samples that produce higher revenues in the baseline case, and that the relationship between baseline revenue and AMD/AMR revenue increase is non-linear.

A consistent theme in these results is that, *given the constraints imposed by traditional distribution systems*, increasing total revenue does *not* necessarily require increasing ancillary revenue, and in many of our results, ancillary revenue decreases as total revenue increases. In this sense, an important benefit of a revenue management formulation like AMD/AMR that integrates both passenger choice and ancillary revenue information is the ability to successfully balance the competing objectives of more aggressive booking limits (decreasing ticket revenue dilution) and greater load factor (increasing the number of passengers, which increases the number of ancillary services that can be sold). Revenue management methods that attempt to increase ancillary revenue without accounting for passenger choice (such as the optimizer increment) risk incorrectly favoring booking policies that stimulate ancillary revenue by increasing load factor, but have the unintended consequence of diluting ticket revenue.

Chapter 8

Conclusions

As ancillary revenues have grown, airlines have become more interested in shifting the objective of their revenue management optimization from ticket revenue maximization to total revenue maximization. At the same time, the proliferation of ancillary services means that passengers can face substantially different ancillary service offerings and prices based on their selection of itinerary and fare class. This thesis has explored both of these issues with the goal of providing a better understanding of how ancillary services affect the airline industry.

8.1 Research Findings and Contributions

Part I of this thesis focuses on modeling airline passenger choices. Previous studies of airline passenger choice have modeled the selection of airline, itinerary, and/or fare class, and other studies have attempted to measure customer valuations of various ancillary services. However, no research has integrated these dimensions to describe how the presence or cost of ancillary services affects itinerary or fare class choice. In Chapter 3, we develop the Ancillary Choice Model, which combines these dimensions into a single choice framework. We specify two different behavioral paradigms, which we call simultaneous and sequential choice. Under the simultaneous choice model, passengers are classically rational and integrate all ancillary, fare class, and itinerary attributes and selections into a single unified decision process. Under the sequential choice model, passengers are boundedly rational and initially unaware of airline ancillary policies/prices/offerings. In an initial phase, passengers select an itinerary and fare class without considering any ancillary attributes. After selecting the itinerary and fare class, sequential passengers enter a follow-up phase where they choose ancillary services.

We describe how the unified choice process of simultaneous passengers means that ancillary prices and options can affect their fare class choice, and describe the conditions in which the presence of an ancillary service will cause a simultaneous passenger to either buy-up or buy-down to a different fare class. We demonstrate that even when all classes have the same ancillary service pricing, the mere presence of an ancillary service can change the fare class choice of a simultaneous passenger. We integrate both of our models into the Passenger Origin-Destination Simulator (PODS), which we use to assess the booking and revenue impacts on airlines of different ancillary service bundling strategies for different mixes of passenger behavior.

Part II of this thesis focuses on the integration of ancillary services into airline revenue management processes. Specifically, it studies total revenue optimization methods, which are intended to shift the focus of revenue management away from ticket revenue maximization and toward total (ticket plus ancillary) revenue maximization. Prior research in this area is quite limited, and no detailed theoretical formulations exist that integrate ancillary revenue into an optimization algorithm that selects fare classes to offer to passengers.

In Chapter 5, we study one previous approach to total revenue optimization, which we refer to as the Optimizer Increment (OI). While other authors have simulated the performance of the optimizer increment (Hao, 2014), its optimality has not been proven. We start by demonstrating that the approach is an optimal total revenue maximizing solution in the very limited "Littlewood Conditions," with only two fare classes that have independent, ordered demands drawn from distributions with known parameters. Subsequent development of the Ancillary Choice Dynamic Program in Chapter 6 shows that the Optimizer Increment is also optimal when multiple fare classes are present and fare class demands are unordered (but still independent and drawn from a known Poisson distribution).

Despite the theoretical optimality of the approach, however, we demonstrate that even in extremely simplistic simulation environments, incorporating something as fundamental as demand forecasting (which all airlines must do) essentially eliminates any optimizer increment revenue benefit. Simulations in larger, more complex environments show that the optimizer increment tends to accept too many low fare bookings, diluting ticket and decreasing total revenue (although increasing ancillary revenue).

Our findings illustrate two phenomena: first, that "optimal" solutions in theory do not necessarily translate to revenue increases in more realistic simulated environments when there is a mismatch between the assumptions of the optimizer and the environment being simulated, and second, that *buy-down* is a threat to revenue performance when forecasting and optimization models do not correctly account for passenger choice. This effect has been shown in many previous RM studies in the context of ticket revenue, but has not been previously demonstrated in studies with ancillary services, passenger choices, and total revenue optimization.

In Chapter 6, we develop the principal contribution of this thesis: the Ancillary Choice Dynamic Program (ACDP). ACDP builds on the growing body of choice-based revenue management, and, for the first time, explicitly incorporates both ancillary service revenue and ancillary service choice impacts in the revenue management forecasting and optimization models.

We describe how the solutions to ACDP must be restricted to booking policies that are *fare class complete* and *nested by fare order* to comply with current distribution system limitations faced by most airlines. We develop a forecasting process to produce demand volume estimates from a historical booking database, and design a set of adjustment mechanisms to correct for inefficiencies (in terms of expected revenue vs expected sale probability) in the airline's fare structure. We build on the existing literature on marginal revenue and marginal demand transformations to develop the Ancillary Marginal Demand (AMD) and Ancillary Marginal Revenue (AMR) transformations, which convert ACDP into an equivalent independent demand model and which we propose using as heuristics to feed the demand and fare inputs to a traditional revenue management optimization model. Finally, we propose two estimation methods for obtaining from a historical booking database the customer choice probabilities that are required for ACDP, AMD, and AMR.

We integrate AMD and AMR into PODS and, in Chapter 7, simulate the performance of the heuristics in a variety of networks and environments, including a range of ancillary prices and disutilities, demand levels, and passenger behavior types. We demonstrate that even in large networks with competing airlines and hundreds of flights serving both local and connecting passengers, AMD and AMR can increase total revenue over existing RM forecasting and optimization models that do not explicitly account for ancillary services or revenues, including both independent demand models and hybrid forecasting/fare adjustments models that account for passenger willingness-to-pay. However, our simulations, particularly in the large network D6, also show that under asymmetric competition, AMD and AMR produce booking and forecast feedbacks with competing airlines that can adversely affect total revenue performance. If an airline's use of AMD and AMR leads it to reject bookings in mid or higher fare classes, those bookings may *spill* to the competitor airline. As the competitor accepts more of the (spilled) mid or higher fare class bookings, it will forecast more mid/higher fare demand, leading it to set more aggressive booking limits, reject more low fare class demand, and accept even more high fare class demand—a phenomenon known as *spiral-up* that can increase the (competitor) airline's total revenue.

One key conclusion from our simulations is that the forecasting and optimization model that maximizes total revenue is often *not* the model that maximizes ancillary revenue. For example, in network D6 (Table 7.17), AMD/AMR generates the highest total revenue (3.5% higher than the optimizer increment), but the optimizer increment has the highest ancillary revenue (a full 4.0% higher than AMD/AMR). This potentially surprising result is a function of conflicting drivers of the two revenue streams: more aggressive booking limits typically decreasing ticket revenue dilution and increase ticket revenue, but reduce the number of passengers and therefore the total ancillary revenue. On the other hand, less aggressive booking limits increase load factor, which increases the number of ancillary services that can be sold, but often lead to ticket revenue losses. The implications are clear for revenue management practitioners: (1) the benefit of a total revenue optimization model should be measured based on its impact on total revenue, not on ancillary revenue, and (2) a successful total revenue optimization model must balance the potentially competing dimensions of ticket and ancillary revenue.

Our simulation results, and the revenue benefit of AMD/AMR relative to existing RM forecasting and optimization models, are dependent on the simulation environment. Different network configurations, different baseline revenue management forecasting and optimization models, and incorporating effects like customer learning and changes in airline network or strategy could affect the relative revenue gains.

8.2 Future Research Directions

The Ancillary Choice Model developed in this thesis provides a framework for understanding airline passenger choices involving ancillary services, fare classes, and itineraries. Additional work could further develop and test the models in several ways.

- The categorization as simultaneous or sequential behavior described in this thesis was assumed to apply to all ancillary services, and passengers were assumed to be fully aware of all other airline schedule and/or quality differentiators. A more nuanced application of bounded rationality could describe additional behavior types in which passengers are simultaneous about some ancillary services (or quality dimensions, such as fare class restrictions) and sequential about other ancillary services or quality dimensions.

- We assumed that an exogenous budgetary constraint applies to all out of pocket costs incurred during the Initial Booking Phase. Future work could explore the extent to which budgetary constraints could vary as a function of the set of alternatives presented to the passenger, or could vary across alternatives within a set based on alternatives attributes.
- We performed a heuristic calibration of our proposed choice models in Appendix A, but a more extensive calibration process could attempt to assess the degree to which one of our proposed behaviors better explains observed booking and purchase patterns.

Total revenue optimization, although long recognized as an area of promise, has very little previous work. Our approach represents an initial proposal for integrating the revenue value and passenger choice impacts of ancillary services into revenue management forecasting and optimization models. Future research could improve and extend these models to increase their practicality for airlines.

- Our process for estimating choice probabilities could likely be improved to increase both the stability and accuracy of the resulting AMD demands and AMR fares. One potential approach would be to forgo estimation of the individual choice probabilities $P_{i,m,dcp}(k)$ and to instead directly estimate total expected revenue and total sale probabilities ($TR_{dcp}(k)$ and $TP_{dcp}(k)$) based on historical bookings. In additional, future studies could model the impact of different choice probability estimation errors on AMD demands and AMR fares, and on total airline revenue, to understand which probabilities are most important to estimate correctly.

- The models presented in this thesis assume that flight capacity is constrained, but that ancillary capacity is unconstrained—a reasonable approximation for some services, but not for others. Extensions to the Ancillary Choice Dynamic Program could introduce additional state dimensions and variables for ancillary capacity; the resulting value function would then jointly compute bid prices for the flight and for the ancillary service(s). The joint computation of bid prices and the reliance on the combination of both bid prices when determining availability would substantially increase the complexity of the model.

Industry efforts are underway to reduce, or eliminate, the fare class completeness constraint described in Section 6.1.1. The International Air Transport Association (IATA) is leading the development of New Distribution Capability (NDC), which is a suite of new distribution technologies and standards. One relevant aspect of NDC is that it would enable airlines to display and sell more content through more channels, particularly indirect channels. When NDC is implemented by an airline and its distribution partners, GDSs would no longer be required to aggregate schedules, availability, and fares to assemble sets of booking options, as described in Figure 1.4. Instead, GDSs (or other content aggregators) could request offers directly from airlines.[1] Most importantly for this work, these offers will no longer be limited to fare class availability. The airline could respond to each request with a specifically designed offer or set of offers—each offer would consist of an itinerary, a set of zero or more ancillary services, various purchase/use restrictions, and a price.

The shift from a traditional distribution environment to NDC has significant implications for both passenger choice modeling and total revenue optimization.

From a passenger choice perspective, NDC increases the ability of airlines to display and sell ancillary services in indirect channels, and therefore increases the importance of understanding how ancillary services affect the way passengers make decisions about itineraries and fare classes. In addition, NDC allows airlines to increase the branding and marketing

[1] Airlines and GDSs could still choose to have the GDS assemble offers.

around their product, even on indirect distribution channels. This potentially provides new dimensions for airlines to attempt to differentiate their products, and allows airlines to produce new product combinations. Airlines expect these distribution improvements to increase ancillary revenues.

NDC could enable a new frontier of total revenue optimization. In addition to the expected increase in ancillary revenue generation, because NDC moves away from fare class-centered availability control, airlines could develop *offer generation systems.* NDC allows more detailed booking requests (such as indicating round trip travel, or frequent flyer status), and does not require the airline to utilize filed fares or ancillary prices, which could enable offer generation systems to create individualized sets of offers. At the limit, each offer could be personally and dynamically constructed and priced for each consumer.

Recent research has explored approaches to utilize NDC for dynamically pricing flights, either as dynamic a price adjustment to a filed fare (Wittman, 2018; Wittman and Belobaba, 2018), a dynamic adjustment to fare class availability (Wittman and Belobaba, 2017), or even without any fare classes. One exciting prospect for additional work is to combine the ideas of dynamic pricing with this study on total revenue management to produce a *dynamic* offer generation engine, responsible for determining simultaneously and in real time whether or which ancillary services to bundle with a flight, and how the flight or bundle as well as any a la carte ancillary services should be priced. Initial research in this area shows promising results for a single ancillary service (Bockelie and Wittman, 2017; Bockelie, 2018), but many practical and theoretical questions remain.

In the near term, however, we imagine a more limited view of NDC that allows airlines to control the availability of specific fare class and ancillary service combinations in real time, but still relies on filed fares and filed ancillary prices and does not personalize or individualize offers. In such an environment, the Ancillary Choice Dynamic Program we have developed in this paper could serve as the basis for an offer generation system. In such an implementation, the solutions to Equation 6.3 need not be fare class complete or potentially even nested by fare order. The solution space for the problem would grow significantly, and the task of determining which offer sets are efficient would be more complex. However, significant gains in total revenue could also be possible. As one example, consider a case where demand is very high relative to capacity. With a traditional distribution system, the airline's highest expected unit revenue booking policy is to offer only FC 1. However, with NDC and offer generation, it would be possible to create an even higher unit revenue offer: FC 1 with a requirement that the passenger purchase any typically optional ancillary services. While requiring passengers to purchase ancillary services (likely marketed as a bundle, with a total price equal to the fare plus the ancillary fees) would reduce the probability of a consumer buying the offer, in very high demand to capacity scenarios (i.e. when the bid price is very high), it is a revenue-maximizing strategy. Additional development and experimentation is required to assess the revenue potential of this approach.

Airline business models have changed dramatically over the last ten to fifteen years as airlines have broadened the product and revenue portfolios by unbundling fares and by introducing new services to sell to passengers. As a result, consumers today have more types of fares and more types of travel experiences to choose from than ever before, and airlines have a greater potential to increase total revenue by accounting for the total spend of each passenger. With new distribution technologies on the horizon, ancillary services will become even easier for airlines to market and sell, further increasing ancillary revenues, and further increasing the competitive advantage of airlines that thoroughly understand how these changes impact both passengers and revenue management systems.

Appendix A

Calibrating the Ancillary Choice Model

In this appendix, we demonstrate that our proposed Ancillary Choice Model can produce realistic booking and ancillary service purchase patterns by calibrating it within a large, complex, multi-airline network in PODS. Specifically, our objective is to estimate parameter values for ancillary and fare class restriction disutilities for an assumed mix of simultaneous and sequential passengers, using booking and purchase data provided by a major airline.

One approach to discrete choice model calibration is parameter estimation though a maximum likelihood estimator on a disaggregate revealed or stated preference dataset. However, this approach is currently infeasible for the ancillary choice model because airlines do not typically collect or store individual-level purchase data that includes both booking information and individual ancillary purchases—these purchase figures are typically only recorded and stored at more aggregated levels. Thus, we must use an aggregate calibration framework. In addition, because the choice probability output of the ACM is a function of preferences (to be estimated) as well as itinerary and fare class availability (which is a complicated function of forecasted demand, which is based on decision making by previous customers), there is no closed form expression that relates the estimated parameters to booking or ancillary purchase metrics. Thus, we use the PODS simulator with integrated ACM to perform the calibration.

A.1 Airline Data

A major airline provided a dataset covering bookings and ancillary purchases on their short and medium-haul network for a one year period. The data includes all short and medium-haul bookings in all fare classes and brands, and all ancillary purchases for the specified ancillary services. The data was provided aggregated by month, booking class, fare brand, and geographic region. Each record in the dataset includes the number of passengers (measured as enplanements[1]), total ticket revenue, and total number of purchases and associated revenue for four ancillary services: first checked bag, advanced seat reservation, extra legroom seating upgrade, and cabin upgrade (economy to premium cabin). The dataset contained 9,066 records for more than 16 million enplanements.

During the period covered by the data, the airline offered 14 economy fare classes grouped into three fare *brands*. All fare classes in a brand share the same purchase/use restrictions and ancillary service pricing, and the brands are marketed by name to consumers on the airline's website. Brand 3, the least expensive five fare classes, is the most restricted, and provides no free ancillary services. Brand 2, which has moderate prices, moderate restrictions, and some complimentary ancillary services consists of the middle seven fare classes. Brand 1, which has the highest prices, fewest restrictions, and the most complimentary ancillary services, consists of the most expensive two fare classes.[2] The airline's website allowed customers to choose amongst the three brands (subject to availability controls). See Vinod and Moore (2009) for a more detailed description of branded fares.

[1]One enplanement is one take off and one landing for one passenger, so one passenger on an itinerary with one connection would produce two enplanements.

[2]Brand names and fare class identifiers have been masked or modified to preserve the anonymity of the airline that provided the data.

A.2 Calibration Network

We start with an existing PODS network for the calibration. The existing network, U10, had fares and demands previously calibrated based on data provided by PODS member airlines and features four hub-and-spoke airlines offering 10 economy fare classes. Not all airlines operate in all origin-destination markets, but most markets have competition from multiple airlines; there are 44 cities with 442 flight legs per day and 572 origin-destination markets (with more than 4,000 different itineraries). All airlines in a given market offer the same fares, but different markets have different fares. The range of fares for each class is shown in Table 3.2. Because the data provided by the airline for calibration only covers short and medium-haul markets, we focus our calibration process on the short and medium haul ("domestic") markets in the PODS network. We model the airline that provided the data as Airline 1, and use Airlines 2–4 to represent (generic) competitors. Figure 3.4 shows the networks of the four simulated airlines.

Airline 1 uses UDP (a network revenue management system based on a network linear program and a leg dynamic program), similar to the system in use at the airline that supplied the data. Airlines 2 and 4 represent sophisticated major network airlines and thus employ a Displacement Adjusted Virtual Nesting revenue management model. Airlines 1, 2, and 4 incorporate estimates of customer willingness-to-pay in their revenue management forecasting and optimization models through the use of hybrid forecasting and fare adjustment. Airline 3 represents a simplified low cost carrier, and therefore uses a leg-based revenue management model and a standard (independent demand) forecasting model.[3]

The existing PODS network did not have branded fares, so we modified the airline product offerings to group fare classes into three distinct brands; each brand has a set of ancillary service policies and purchase/use restrictions. All four airlines in the simulation have the same product offerings. We focus on the four ancillary services included within the dataset.

[3]See Part II of this thesis for details on revenue management forecasting and optimization models.

Table A.1: Ancillary fee structure in airline dataset and in PODS

		Dataset			PODS		
		Brand 1	Brand 2	Brand 3	Brand 1	Brand 2	Brand 3
BAG	Checked bag	*Free*	*Free*	Paid	*Free*	*Free*	Paid
ASR	Advance seat reservation	*Free*	*Free*	Paid	*Free*	*Free*	Paid
UPG	Extra legroom	*Free*	Paid	Paid	*Free*	Paid	Paid
	Cabin upgrade	Paid	Paid	Paid			

Due to limitations in the software, we combine the extra legroom seating and cabin upgrade products as an "upgraded seating product." The ancillary fee structure, by brand, for the airline the provided the data and for the simulation are shown in Table A.1. Passengers purchasing brand 3 must pay for all three ancillary services. Passengers purchasing brand 2 receive a free first checked bag (BAG) and a free advance seat reservation (ASR) but must pay for upgraded seating (UPG); passengers purchasing brand 3 in the simulation receive all three ancillary products for free.

We also modify the default purchase/use restrictions for each fare class to match the branded fare structure offered by the airline providing the data. We use four generic restrictions (denoted R1–4) to model the differences in purchase/use conditions between the three brands, as shown in Table A.2. At the airline that provided the data, passengers booking brand 3 must pay a fee to change their flight or refund their ticket, only accrue 25% of possible frequent flyer miles, do not have priority services, and are not eligible for loyalty-based upgrades. Passengers booking brand 2 pay a smaller change/refund fee, earn full frequent flyer miles, and are eligible for loyalty upgrades, but do not receive priority services. Passengers booking brand 1 have no change/refund fee, earn 125% frequent flyer miles, are eligible for loyalty upgrades, and receive priority services.

In the simulation, we model brands 2 and 3 as having generic restriction R1, which models the primary differences between brand 1 and brands 2 and 3: full refundability and priority services. We model brand 3 as also having generic restriction R4, which models the primary

Table A.2: Fare class restriction structure in airline dataset and in PODS (advance purchase in days).

	Dataset			PODS		
	Brand 1	Brand 2	Brand 3	Brand 1	Brand 2	Brand 3
Change fee	*Free*	\$75	\$100	*None*	R1–3	R1–3
Frequent flyer miles	125%	100%	25%	*None*	*None*	R4
Priority services	*Free*	n/a	n/a	*None*	R1	R1–3
Advanced purchase	*None*	*None* to 7	7 to 21	*None*	*None* to 21	7 to 28
Fare classes	A & B	C to I	J to N	1	2 to 6	7 to 10

Note: more restrictive APs in PODS are used to model the (manual) RM analyst interventions that the airline that reduce low-fare availability

differences between brands 1 and 2 and brand 3: full frequent flyer miles and eligibility for loyalty upgrades. The value of full refundability is greater when booked further in advance (because there is more time for plans to change), however PODS does not support changing restriction disutilities over time. As a workaround, we apply additional generic restrictions (R2 and R3) to the fare classes and brands (FC 6, 9, and 10) that have change/refund fees and that are typically purchased early in the booking window. The complete simulated fare structure is shown in Table 3.2.

A.3 Calibration Process

The objective of the calibration is to determine the ancillary service disutility and fare class restriction disutility distribution parameters (Ω_s in Section 3.1 and Ω_j^R in Section 2.3, respectively) that result in the best fit between the PODS outputs and the dataset provided by the airline. We use the following process:

0. Initialization: Start with the default U10 fare class restriction and ancillary service disutilities.

1. Simulation in PODS: each simulation ran for about 30 minutes, with an additional 10

minutes to process the simulation outputs

2. Comparison of PODS outputs to provided data: we evaluated the squared differences between PODS outputs and the real airline dataset for common airline industry performance indicators. We weighted and combined the metrics into a single fit error, with weights chosen to account for both the different dimensions/scales of each metric and the relative importance of each metric to business decisions. We included these metrics at multiple levels of aggregation to balance accuracy at the fare class, fare brand, and ancillary service levels.

    - Booking mix: the portion of bookings within each fare class or within each fare brand
    - Average ancillary revenue per enplanement, aggregated or broken out by fare class and/or fare brand and/or ancillary service
    - Ancillary service purchase rate, aggregated or broken out by fare class and/or fare brand and/or ancillary service
    - Portion of total revenue from ancillary services, aggregated or broken out by fare class and/or fare brand and/or ancillary service

3. Parameter adjustment calculation: we stopped the calibration when the total fit error was smaller than 0.05. Otherwise, we determined parameters for the next calibration iteration based on our knowledge of the primary effects of each parameter (described below) and simulated PODS again (process step 1). We used minimum step sizes of \$2.50 for the end points of the ancillary service mean disutilities (which were then converted to an intercept and slope) and 0.05 for the fare class restriction mean disutility slopes (fare class restriction disutilities in PODS do not typically use the intercept parameter).

    The primary effect of each parameter is as follows:

    - R1: increasing R1 makes brands 2 and 3 less appealing to consumers, and therefore decreases the portion of bookings in brands 2 and 3 and increases the portion of bookings in brand 1

- R4: increasing R4 makes brand 3 less appealing to consumers, and therefore decreases the portion of bookings in brand 3 and increases the portion of bookings in brands 1 and 2
- R2 and R3: increasing R2 and R3 makes brands 2 and 3 less appealing to customers shopping early in the booking process, and therefore decreases the portion of bookings in the early booking fare classes of brands 2 and 3 (FC 6, 9, and 10) and increases the portion of early bookings in brand 1
- BAG, PRS, and UPG price: increasing ancillary prices makes fewer consumers interested in purchasing the service and increases the perceived value of receiving the service complimentary, and therefore decreases the ancillary purchase rate, decreases the portion of bookings in brand 3 (for BAG, PRS, and UPG) and brand 2 (for UPG), and increases the portion of bookings in brand 1 (for BAG, PRS, and UPG) and brand 2 (for BAG or PRS).
- BAG, PRS, and UPG disutility: increasing ancillary disutilities makes more consumers interested in purchasing the service and increases the perceived value of receiving the service (complimentary or paid), and therefore increases the ancillary purchase rate, increases the portion of revenue from ancillary services, decreases the portion of bookings in brand 3 (for BAG, PRS, and UPG) and brand 2 (for UPG), and increases the portion of bookings in brand 1 (for BAG, PRS, and UPG) and brand 2 (for BAG or PRS).

We performed the calibration twice: once assuming 100% simultaneous passengers, and once assuming a 50/50 mix of simultaneous and sequential behavior (where each generated consumer is randomly assigned a passenger type). We did not calibrate with a 100% sequential mix because the airline's website clearly explains the differences between brands during the booking process, making 100% sequential behavior unlikely for passengers who purchase via the airline's website. While PODS supports different disutility distributions

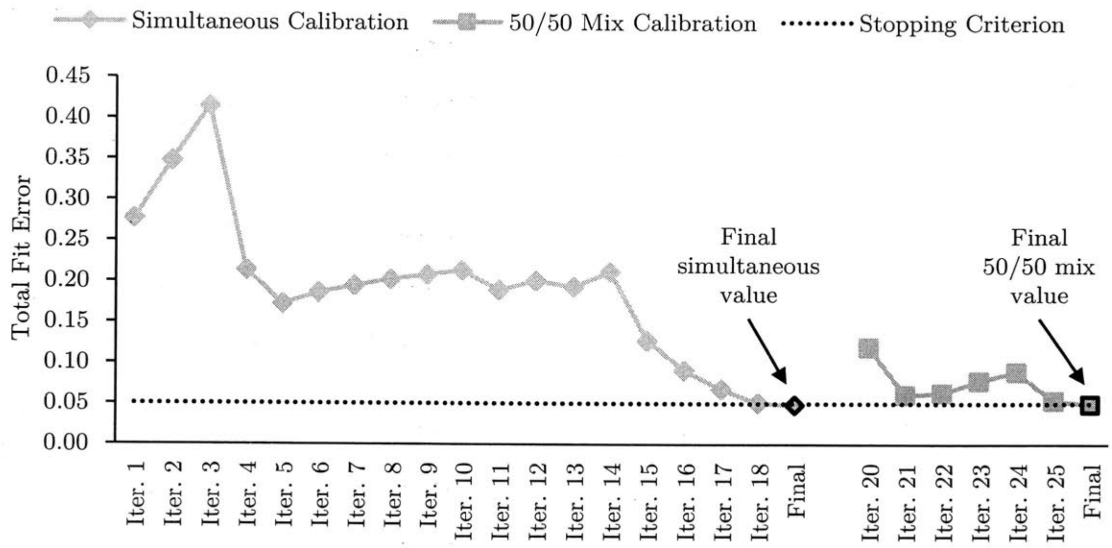

Figure A.1: Evolution of total fit error by iteration.

for different passenger segments (business and leisure), it does not support different distributions for different behavior types (simultaneous and sequential). Once the initial network modifications described in Section A.2 were complete, an additional 26 calibration iterations were required.

A.4 Results

Figure A.1 shows the total fit error by iteration for each behavior mix. The final calibrations have fit errors less than 0.05. Figure A.2 shows the portion of revenue from ancillary services, for several calibration steps for 100% simultaneous passengers.

The calibrated parameter values are shown in Table 3.3. Overall, the calibration with all simultaneous passengers has a lower R4 disutility and higher ancillary disutilities than the calibration with a 50/50 mix of simultaneous and sequential passengers. Because sequential passengers do not consider ancillary services when choosing an itinerary and fare class, they require higher fare class restriction disutilities to produce the same booking mix as

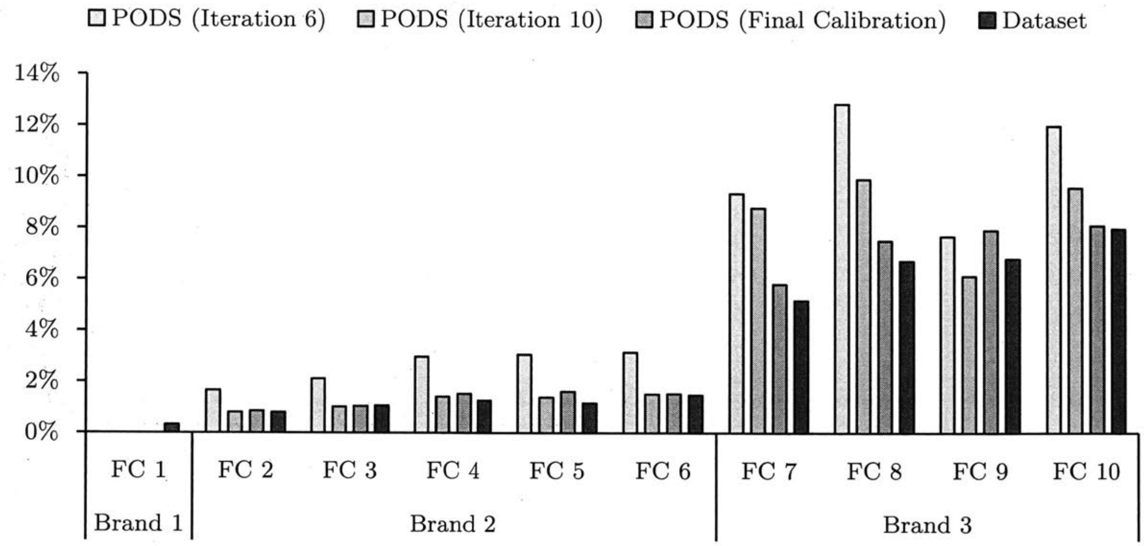

Figure A.2: Portion of revenue from ancillary services for various iterations for Simultaneous calibration (domestic markets only)

simultaneous passengers. In addition, because sequential passengers do not have an overall budgetary constraint on out-of-pocket spending, they require lower ancillary utilities to produce the same ancillary purchase rates as simultaneous passengers.

As R1–4 are generic restrictions that represent combinations of features not often sold separately, it is hard to assess the reasonableness of the calibrated values. In general, however, we see that for the simultaneous calibration, business passengers have higher mean disutilities than leisure passengers, and that the primary (R1 and R4) disutilities are larger than the secondary disutilities (R2 and R3). In addition, we see that business passengers have a higher mean R1 than R4, indicating that refundability and priority services are (on average) even more important than full frequent flyer miles and upgrade eligibility. Leisure passengers, however, have a lower mean R1 than R4, indicating that full frequent flyer miles and upgrade eligibility are more important than refundability and priority services.

The calibration results show that business travelers have a lower mean utility for a checked bag than do leisure travelers, but that mean utilities for both segments are close to the price of \$25, so purchase rates will be high. As business travelers typically have shorter stays and higher value of time (and therefore less desire to wait at baggage claim), this result makes sense. Both segments have a mean ASR utility less than the price, indicating that purchase rates will be low.

The calibrated mean ASR disutility for leisure passengers is higher than business passengers: leisure passengers are more likely to travel in groups, where sitting near friends and family (and therefore needing a seat assignment) is necessary.

Both passengers have a calibrated mean UPG utility lower than UPG price, so UPG purchase rates will also be low. For the simultaneous calibration, both segments have the same mean utility; for the 50/50 mix, the leisure segment has a lower mean utility than the business segment.

Figure A.3 shows the portion of revenue from ancillary services, as well as the booking mix, by brand for both calibrations. Most bookings occur in the lowest priced brand 3, with the fewest bookings in the highest priced brand 1. Passengers who purchase brand 3 contribute the most revenue from ancillary services (about 8% of total). Both calibrations are reasonably close to provided data, with a slightly better fit for the 50/50 mix calibration.

Ancillary purchase rates for each service, broken out by brand, are shown in Figure A.4, which indicates that both calibrations are close to the provided data. Note that passengers who purchase brand 1 and brand 2 receive BAG and ASR complimentary, and passengers who purchase brand 1 receive UPG complimentary, so there are no purchases for those services in those brands. Finally, Figure A.5 shows the average ancillary revenue per enplanement, again broken out by service and brand. Although both calibrations are reasonably close to the provided data, the all simultaneous calibration has a slightly better fit.

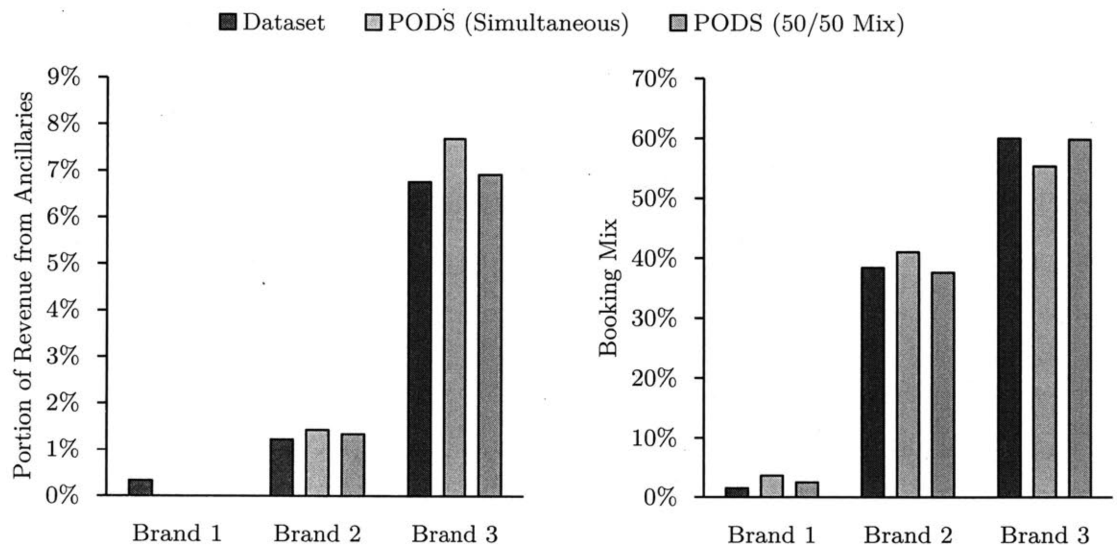

Figure A.3: Portion of revenue from ancillary services (left) and booking mix (right) broken out by fare brand (domestic markets only)

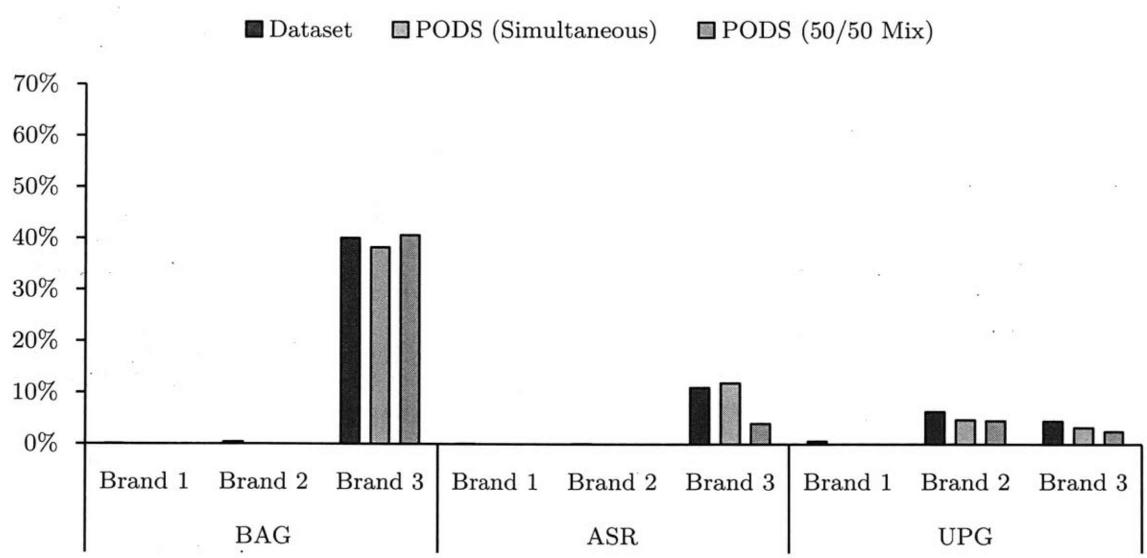

Figure A.4: Ancillary purchase rate per enplanement (domestic markets only)

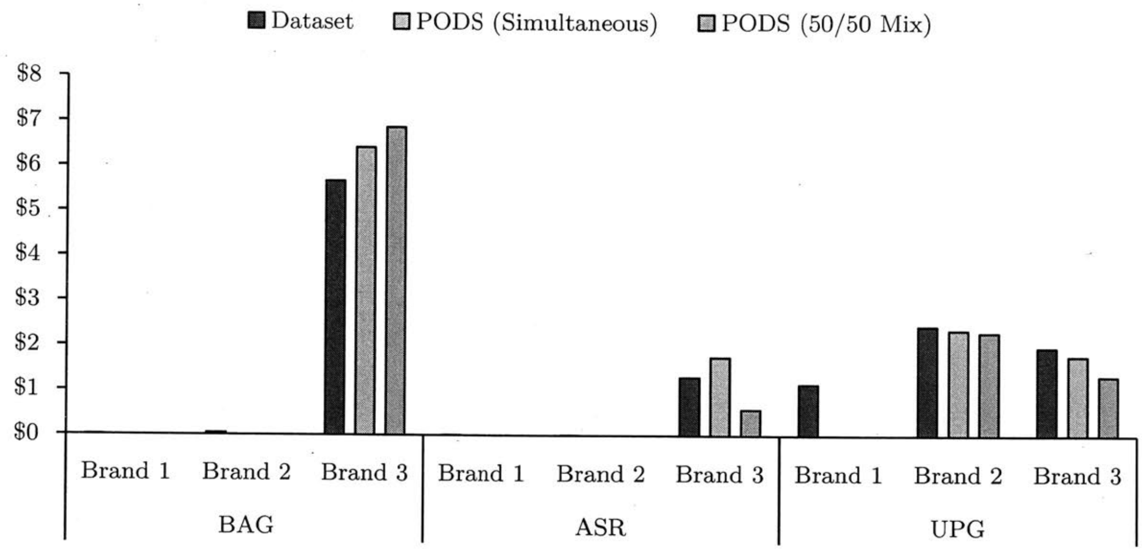

Figure A.5: Average ancillary revenue per enplanement (domestic markets only)

A.5 Conclusions

By integrating the ACM within PODS, we performed an aggregate calibration using booking and ancillary purchase data provided by a major airline for its short/medium-haul network. We calibrated two environments, one with all simultaneous passengers and one with a 50/50 mix of simultaneous and sequential passengers, to show that the ACM, with appropriate input parameters, can replicate the booking and ancillary purchase patterns observed by real airlines. The results of this calibration process are used in the Ancillary Choice Model sensitivity study in Section 3.3. Future work could include a more robust calibration, in which the portion of simultaneous vs sequential passengers is a variable. This more robust calibration could potentially estimate the behavior mix for the airline supplying the calibration data.

Acknowledgments

We would like to thank the airline that supplied the booking and ancillary purchase data for their assistance with this project.

Appendix B

Proofs

B.1 Proof of Optimizer Increment Convexity

Theorem 4. *The expected marginal total revenue R is concave ($-R$ is convex) at π^*.*

Proof. We show R is concave at π^* by showing that $\frac{\partial^2 R}{\partial \pi^2}|_{\pi^*} < 0$. First consider $\partial^2 R_2/\partial \pi^2$:

$$\begin{aligned}
\frac{\partial^2 R_2}{\partial \pi^2} &= \frac{\partial}{\partial \pi}\left(-(f_2 + a_2)\Pr(X_2 > c - \pi)\right) \\
&= -(f_2 + a_2)\frac{\partial}{\partial \pi}\Pr(X_2 > c - \pi) \\
&= -(f_2 + a_2)F'_{X_2}(c - \pi)
\end{aligned}$$

where F'_{X_k} denotes the probability density function of random variable X_k. Now consider $\partial^2 R_1/\partial \pi^2$:

$$\begin{aligned}
\frac{\partial^2 R_1}{\partial \pi^2} &= \frac{\partial}{\partial \pi}\Big((f_1 + a_1)\Pr(X_1 \geq \pi)\Pr(X_2 > c - \pi)\Big) \\
&= (f_1 + a_1)\left(\Pr(X_2 > c - \pi)\frac{\partial}{\partial \pi}\Pr(X_1 \geq \pi) + \Pr(X_1 \geq \pi)\frac{\partial}{\partial \pi}\Pr(X_2 > c - \pi)\right) \\
&= (f_1 + a_1)\Big(- \Pr(X_2 > c - \pi)F'_{X_1}(\pi) + \Pr(X_1 \geq \pi)F'_{X_2}(c - \pi)\Big) \\
&= -(f_1 + a_1)\Pr(X_2 > c - \pi)F'_{X_1}(\pi) + (f_1 + a_1)\Pr(X_1 \geq \pi)F'_{X_2}(c - \pi)
\end{aligned}$$

Combining the previous two equations yields:

$$\begin{aligned}
\frac{\partial^2 R}{\partial \pi^2} &= \frac{\partial^2 R_1}{\partial \pi^2} + \frac{\partial^2 R_1}{\partial \pi^2} \\
&= -(f_1 + a_1)\Pr(X_2 > c - \pi)F'_{X_1}(\pi) + (f_1 + a_1)\Pr(X_1 \geq \pi)F'_{X_2}(c - \pi) \\
&\qquad - (f_2 + a_2)F'_{X_2}(c - \pi) \\
&= -(f_1 + a_1)\Pr(X_2 > c - \pi)F'_{X_1}(\pi) \\
&\qquad + F'_{X_2}(c - \pi)\Big((f_1 + a_1)\Pr(X_1 \geq \pi) - (f_2 + a_2)\Big)
\end{aligned}$$

Note that at π^* the last term of the above line is zero due to the first order conditions in Equation 5.6. Therefore, $\frac{\partial^2 R}{\partial \pi^2}$ evaluated at π^* is:

$$\frac{\partial^2 R}{\partial \pi^2}|_{\pi^*} = -(f_1 + a_1)\Pr(X_2 > c - \pi)F'_{X_1}(\pi^*)$$

Because Pr() is a probability statement, and because $F'_{X_1}()$ is a probability density function, both are non-negative. Therefore, with the trivial assumption that $f_1 + a_1 > 0$, $\frac{\partial^2 R}{\partial \pi^2}|_{\pi^*} \leq 0$ and R is concave at π^*. $\square$

Appendix C

Impact of Correlations Between Samples

As discussed in Sections 5.2.1 and 7.2, the t-tests used to measure the statistical significance of revenue changes assume that each sample is independent. In PODS simulations, demand forecasts generated for one sample depend on the bookings received in the previous n_{ob} (typically 26) samples, which potentially introduces a correlation between different samples and violates the assumptions of the t-test. In this appendix we investigate the degree to which the total revenue changes due to the Optimizer Increment (OI) or due to the Ancillary Marginal Demand and Ancillary Marginal Revenue transformations (AMD/AMR) are correlated between different PODS samples (within an individual trial), and we evaluate the potential impact of an alternative experimental and analysis approach that uses simulations with more trials, fewer samples, and performs statistical tests by *trial*, not by sample.

C.1 Assessment of Correlations Between Samples

Recall the notation from Sections 5.2.1 and 7.2: X_i^j is the revenue (or other simulation output) for sample i for simulation $j \in \{\text{TEST}, \text{BASE}\}$, and $\hat{\Delta}_i$ is the change in revenue (or other simulation output) in sample i (due to a change in forecasting or optimization method), with $\hat{\Delta}_i = X_i^{\text{TEST}} - X_i^{\text{BASE}}$.

The statistical tests in the main body of the thesis use a t-test on these differences $\hat{\Delta}_i$, with the null hypothesis that there is no true change in revenue ($H_0 : \Delta = 0$) vs alternative hypothesis that there is a true change in revenue ($H_a : \Delta \neq 0$).

The forecasting process could potentially introduce a correlation between the revenue change in sample i and the revenue change in sample $i - l$ for $l \in \{1, \ldots, n_{ob}\}$. However, Figure C.1 shows the correlations between sample i and $i - l$ for the symmetric optimizer increment and AMD/AMR in network A1ONE and in network D6, using the same experimental structures as in Sections 7.3 and 7.5. Note that all of the correlations are small, and with the exception of the optimizer increment in network D6, all correlations have an absolute value less than 0.1 and show no discernible pattern. Although there is a link between samples (via the forecasting process), the variability of demand generated within each sample (which is independent) dominates the simulation outputs. Thus, the revenue changes between samples are essentially uncorrelated and the t-test is a valid statistical approach.

To further investigate the case of the symmetric optimizer increment in network D6, we visually example scatter plots of $\hat{\Delta}_i$ vs $\hat{\Delta}_{i-l}$ for $l \in \{1, \ldots, n_{ob}\}$, as shown in Figures C.2 and reffig:statistical-significance.d6-sym-oi-lag-scatter-2. In general, the plots show no obvious linear or non-linear relationships between $\hat{\Delta}_i$ and $\hat{\Delta}_{i-l}$, again supporting our use of t-test for statistical assessments.

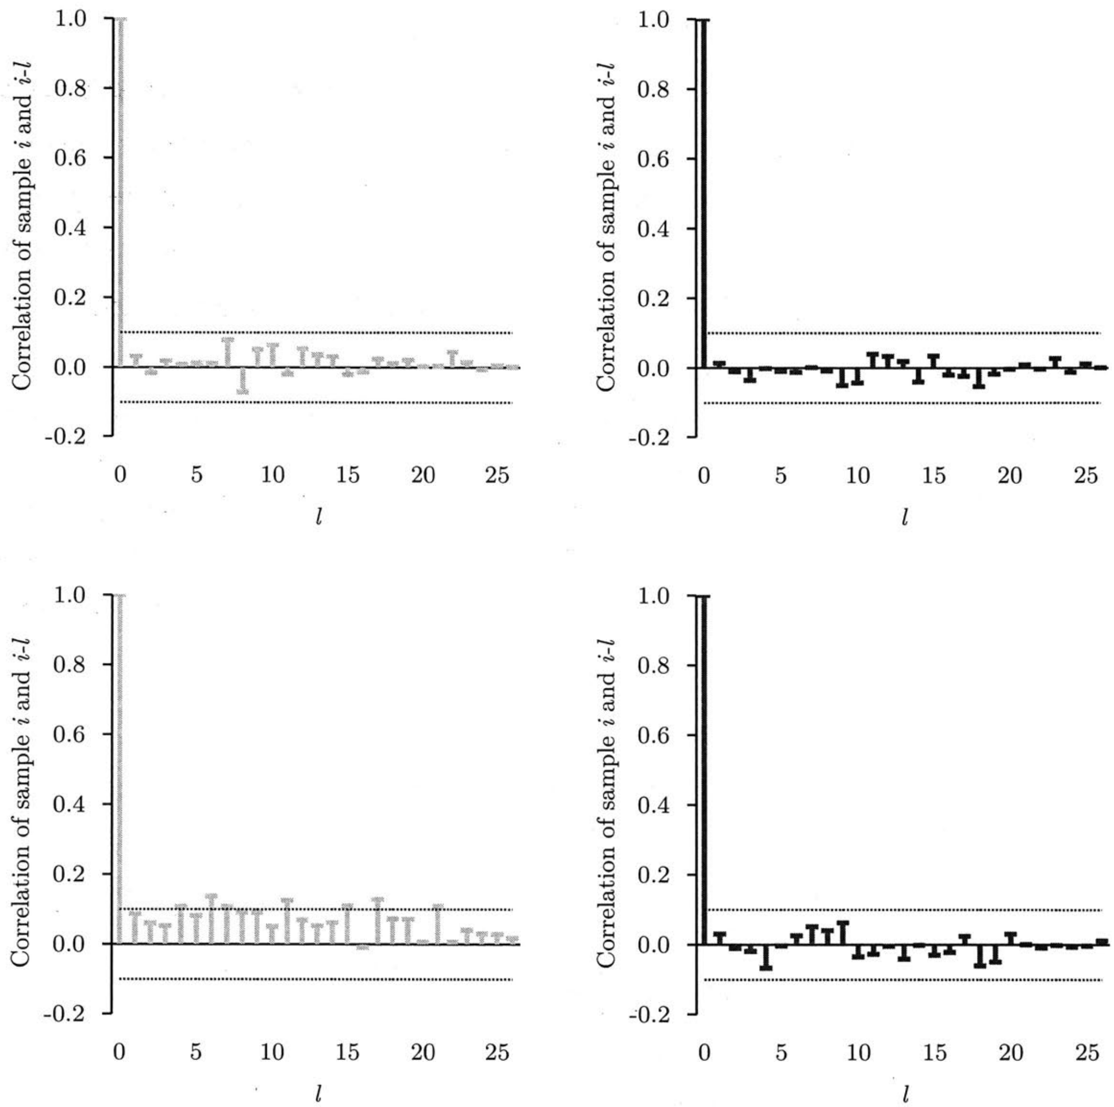

Figure C.1: Network A1ONE (top) and Network D6 (bottom) correlation between sample i and $i - l$ of total revenue changes due to OI (left) and AMD/AMR (right) vs baseline for various value of l (medium demand, 100% simultaneous passengers, \$50 ancillary price for A1ONE and ancillary price equal to 40% of FC 6 fare for D6, equally appealing disutilities). Dotted lines show ± 0.1 correlation.

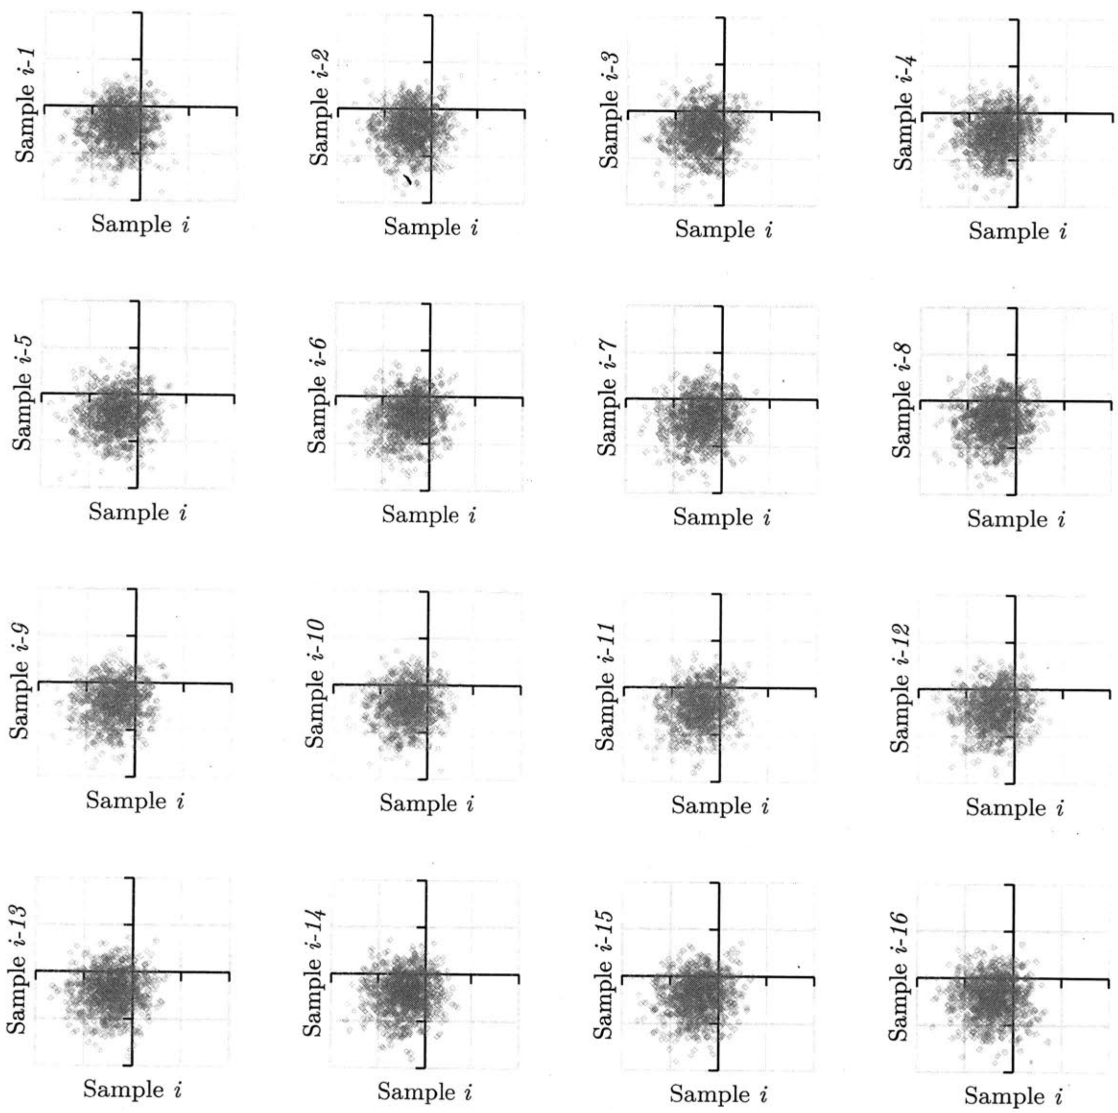

Figure C.2: Network D6 Airline 1 total revenue change due to symmetric OI vs baseline in sample $i - l$ against sample i for $l \in \{1, \ldots, 16\}$ (100% simultaneous passengers, ancillary price = 40% of FC 6 fare, equally appealing disutilities). Both the horizontal and vertical axes range from -\$20,000 to +\$20,000.

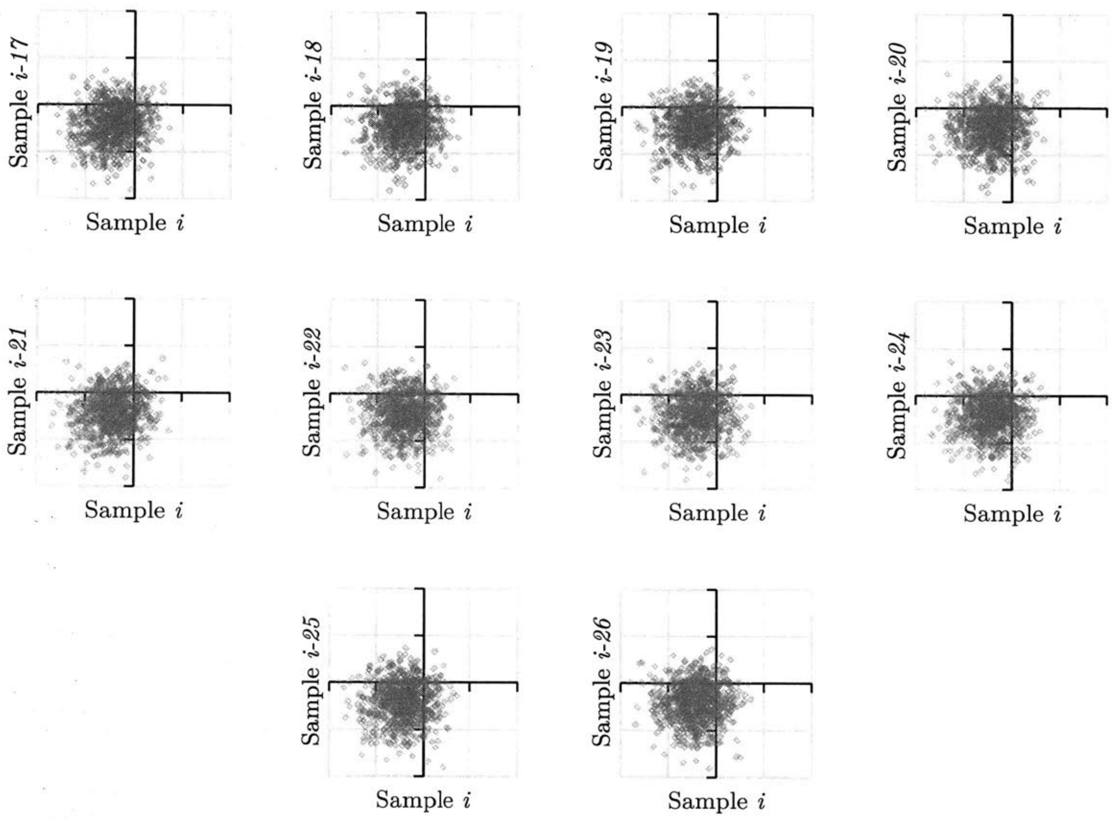

Figure C.3: Network D6 Airline 1 total revenue change due to symmetric OI vs baseline in sample $i - l$ against sample i for $l \in \{17, \ldots, 26\}$ (100% simultaneous passengers, ancillary price = 40% of FC 6 fare, equally appealing disutilities). Both the horizontal and vertical axes range from -\$20,000 to +\$20,000.

C.2 Alternative Experimental Structure: More Trials, Fewer Samples

One alternative approach to our experimental structure, which would eliminate any concerns about potential correlation between samples, would be simulating our test and baseline cases with many trials, but few samples per trial, and then using the average revenue (or other simulation output) *within each trial* as the basis of our analysis. We did not select this approach for the main body of the thesis because even with few *unburned* samples per trial, each trial requires 200 *burned* samples to warm up the forecasting models. Thus, increasing the number of trials drastically increases the computation time for each simulation.

Under this alternative approach, we use a t-test on the difference in average revenue (or other simulation output) across all (unburned) samples within an individual trial. Mathematically, if $X^j_{trl,i}$ is the revenue (or other simulation output) for sample i in trial trl for simulation $j \in \{\text{TEST}, \text{BASE}\}$, we are interested in the term $\bar{\Delta}^*$:

$$\bar{\Delta}^* = \frac{1}{n_{trl}} \sum_{trl=1}^{n_{trl}} \hat{\Delta}^*_{trl} \qquad\qquad \hat{\Delta}^*_{trl} = \frac{1}{n_{smp}} \sum_{i=1}^{n_{smp}} X^{\text{TEST}}_{trl,i} - \frac{1}{n_{smp}} \sum_{i=1}^{n_{smp}} X^{\text{BASE}}_{trl,i}$$

The null hypothesis of our t-test that there is no true change in revenue ($H_0 : \Delta = 0$) vs alternative hypothesis that there is a true change in revenue ($H_a : \Delta \neq 0$). The test statistic t is now:

$$t = \frac{\bar{\Delta}^*}{\text{se}_{\bar{\Delta}^*}} \sim T_{df=n_{trl}-1}$$

where $\text{se}_{\bar{\Delta}^*}$ is the standard error of $\bar{\Delta}^*$.

The purpose of this section is to compare the statistical conclusions obtained from these two analysis approaches. To do so, we re-run a few representative simulations with 50 trials and 50 unburned samples per trial (with an additional 200 burned samples per trial, which as usual are not included in any of the values reported here). The two approaches are mathematically equivalent in terms of the average change in revenue, bookings, and ancillary sales; they only differ in the way they measure the associated variances.

Table C.1 shows the results for network A1ONE with medium demand, 100% simultaneous passengers, $50 ancillary price, and the equally-appealing disutility scenario (see Table 5.5). The original analysis method (comparing the change in total revenue by sample, and used in the main body of the thesis) shows that all total revenue changes are statistically significant ($p < 0.001$, $df = 2,499$). The alternative analysis method (comparing the change in total revenue by trial, as described in this section) also shows that all total revenue changes are significant ($p < 0.001$, $df = 49$). It is worth noting that there is very little difference in t distribution for 49 degrees of freedom and 2,499 degrees of freedom, as shown in Figure C.4 (with such large degrees of freedom, the distribution has essentially already converged to a normal distribution). Comparing the two analysis methods, the alternative method has a larger (in absolute terms) t-statistic for the change in total revenue due to the optimizer increment, but smaller (in absolute terms) t-statistics for the change due to HF/FA, OI + HF/FA, and AMD/AMR.

In network D6, both analysis methods also show statistically changes in total revenue ($p < 0.001$, $df = 49$ or $2,499$) for all four (symmetric) experimental cases, as illustrated in Table C.2. The original analysis method has a larger absolute t-statistic for the revenue change due to OI and OI + HF/FA, but a smaller absolute t-statistic for the revenue change due to HF/FA and AMD/AMR.

C.3 Conclusions

Although demand forecasting introduces a dependency between the bookings received in one sample and the booking limits set in another sample, and therefore potentially induces

Table C.1: Network A1ONE simulation results paired by sample or by trial for baseline and experimental cases with 50 trials and 50 unburned samples per trial (medium demand, 100% simultaneous passengers, $50 ancillary price, equally appealing disutility scenario).

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
Ticket Revenue	$22,057	$22,018	$22,347	$22,334	$22,493
Ancillary Revenue	$1,805	$1,807	$1,796	$1,796	$1,786
Total Revenue	$23,862	$23,825	$24,143	$24,130	$24,279
Load Factor	83.4%	83.5%	82.2%	82.3%	82.0%
Total Yield	22.00	21.94	22.60	22.56	22.76
Ancillary Sales Rate	33.3%	33.3%	33.6%	33.6%	33.5%
Change from Baseline					
Ticket Revenue		-0.2%	+1.3%	+1.3%	+2.0%
Ancillary Revenue		+0.1%	-0.5%	-0.5%	-1.1%
Total Revenue		-0.2%	+1.2%	+1.1%	+1.7%
Load Factor		+0.1 pts	-1.3 pts	-1.1 pts	-1.4 pts
Total Yield		-0.3%	+2.7%	+2.5%	+3.5%
Ancillary Sales Rate		-0.0 pts	+0.3 pts	+0.3 pts	+0.2 pts
Significance of Change in Total Revenue from Baseline					
By Sample (Original Analysis Method)					
Observations		2,500	2,500	2,500	2,500
Standard Error		0.0%	0.1%	0.1%	0.1%
t-statistic		-9.27	13.42	13.77	14.70
p-value		< 0.001	< 0.001	< 0.001	< 0.001
By Trial (Alternative Analysis Method)					
Observations		50	50	50	50
Standard Error		0.0%	0.1%	0.1%	0.1%
t-statistic		-7.69	14.82	15.00	15.88
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = Observations - 1$.

Table C.2: Network D6 simulation results paired by sample or by trial for baseline and symmetric experimental cases with 50 trials and 50 unburned samples per trial (100% simultaneous passengers, ancillary price = 40% FC 6 fare, equally appealing disutility scenario).

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
Airline 1					
Ticket Revenue	$1,322,131	$1,318,090	$1,363,385	$1,355,249	$1,373,430
Ancillary Revenue	$115,569	$115,787	$113,013	$113,961	$111,546
Total Revenue	$1,437,700	$1,433,877	$1,476,398	$1,469,210	$1,484,976
Load Factor	83.5%	83.5%	82.1%	82.6%	81.6%
Total Yield	14.03	14.00	14.65	14.50	14.84
Ancillary Sales Rate	42.6%	42.6%	42.3%	42.4%	42.1%
Change from Baseline					
Ticket Revenue		-0.3%	+3.1%	+2.5%	+3.9%
Ancillary Revenue		+0.2%	-2.2%	-1.4%	-3.5%
Total Revenue		-0.3%	+2.7%	+2.2%	+3.3%
Load Factor		-0.0 pts	-1.4 pts	-0.9 pts	-1.9 pts
Total Yield		-0.2%	+4.4%	+3.3%	+5.8%
Ancillary Sales Rate		+0.1 pts	-0.3 pts	-0.1 pts	-0.4 pts
Significance of Change in Total Revenue from Baseline					
By Sample (Original Analysis Method)					
Observations		2,500	2,500	2,500	2,500
Standard Error		0.0%	0.0%	0.0%	0.0%
t-statistic		-46.33	66.34	58.27	77.82
p-value		< 0.001	< 0.001	< 0.001	< 0.001
By Trial (Alternative Analysis Method)					
Observations		50	50	50	50
Standard Error		0.0%	0.0%	0.0%	0.0%
t-statistic		-27.59	68.59	56.92	83.55
p-value		< 0.001	< 0.001	< 0.001	< 0.001

Note: Total yield in cents per mile. Standard error expressed as percentage of baseline total revenue. $df = Observations - 1$.

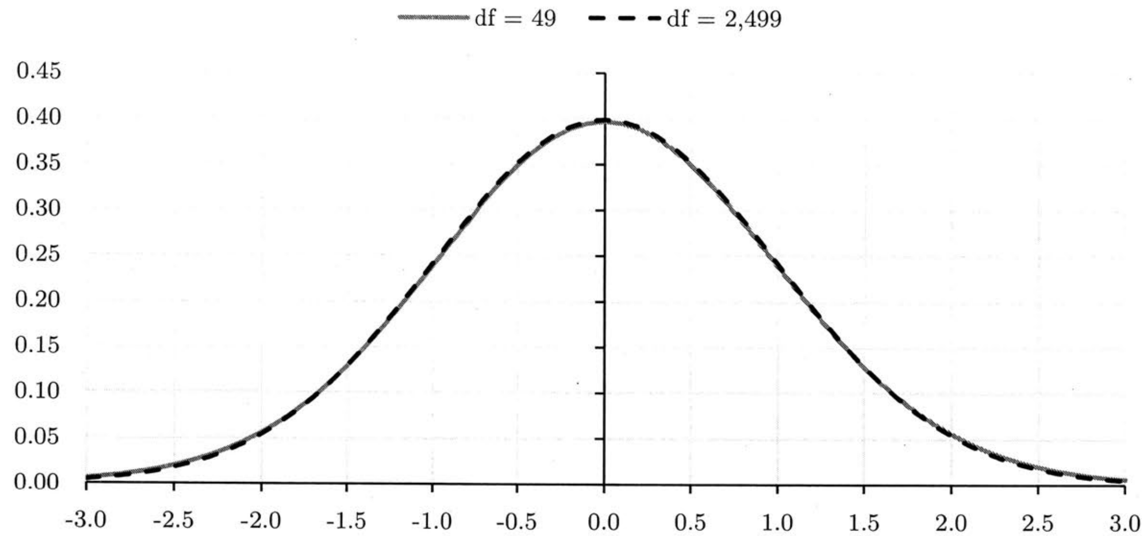

Figure C.4: Probability density function for t distribution with $df = 49$ and $df = 2,499$.

correlations in total revenue changes between different samples, our analysis in this appendix illustrates that (1) these dependencies are weak, with only minor correlation between $\hat{\Delta}_i$ and $\hat{\Delta}_{i-l}$ for the tested values of l (1 to $n_{ob} = 26$) and (2) alternative analysis methods that aggregate revenue by trial, and then perform statistical testing on the outcome of each trial (eliminating any dependencies between different observations used for testing), lead to similar conclusions about the statistical significance of our simulation results: in all tested cases, both analysis approaches reject the null hypothesis that the change in revenue due to the optimizer increment or AMD/AMR (or HF/FA or OI + HF/FA) is zero, in favor of an alternative hypothesis that the change is non-zero. Together, these conclusions validate the statistical analysis methods used in the main body of this thesis.

References

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

Algers, S. and Beser, M. (2001). Modelling choice of flight and booking class—a study using stated preference and revealed preference data. *International Journal of Services, Technology, and Management*, 2(1–2):28–45.

Allon, G., Bassamboo, A., and Lariviere, M. (2011). Would the social planner let bags fly free? Working Paper.

Ancarani, F., Gerstner, E., Posselt, T., and Radic, D. (2009). Could higher fees lead to lower prices? *Journal of Product & Brand Management*, 18(4):297–305.

Azadeh, S. S., Marcotte, P., and Savard, G. (2014). A taxonomy of demand uncensoring methods in revenue management. *Journal of Revenue and Pricing Management*, 13(6):440–456.

Balcombe, K., Fraser, I., and Harris, L. (2009). Consumer willingness to pay for in-flight service and comfort levels: a choice experiment. *Journal of Air Transport Management*, 15:221–226.

Belobaba, P. (1987). *Air Travel Demand and Airline Seat Inventory Management.* Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

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

Belobaba, P. (2015). Course materials for 16.71 The Airline Industry, Airline Pricing Theory and Practice.

Belobaba, P., Odoni, A., and Barnhart, C., editors (2009). *The Global Airline Industry.* John Wiley & Sons.

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

Ben-Akiva, M. (1973). *Structure of Passenger Travel Demand Models.* Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

Ben-Akiva, M. and Gershenfeld, S. (1998). Multi-featured products and services: Analysing pricing and bundling strategies. *Journal of Forecasting*, 17(3–4):175–196.

Ben-Akiva, M. and Lerman, S. (1985). *Discrete Choice Analysis.* MIT Press, Cambridge, MA.

Bertsimas, D. and de Boer, S. (2005). Simulation-based booking limits for airline revenue management. *Operations Research*, 53(1):90–106.

Bockelie, A. (2018). Tests of dynamic offer generation with modified concurrent choice behavior. Presented at the May 2018 PODS Consortium Meeting, Cannes, FR.

Bockelie, A. and Belobaba, P. (2016a). Bid-price behavior in network RM: Forecast decrementing and bid-price bounding. *Journal of Revenue and Pricing Management*, 15(2):95–106.

Bockelie, A. and Belobaba, P. (2016b). Incorporating ancillary services in passenger choice: Revenue and booking impacts. Presented at the 2016 meeting of AGIFORS Revenue Management Study Group, Frankfurt, Germany.

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 Belobaba, P. (2018). Total revenue optimization with the Ancillary Marginal Demand and Ancillary Marginal Revenue transformation heuristics. Presented at the 58th Annual AGIFORS Symposium, Tokyo, Japan.

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

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

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

Brueckner, J. K., Lee, D. N., Picard, P. M., and Singer, E. (2015). Product unbundling in the travel industry: The economics of airline bag fees. *Journal of Economics and Management Strategy*, 24(3):457–484.

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

Bruning, E. R., Hu, M. Y., and Hao, W. (2009). Cross-national segmentation: An application to the NAFTA airline passenger market. *European Journal of Marketing*, 43(11–12):1498—1522.

Carrier, E. (2008). *Modeling the Choice of an Airline Itinerary and Fare Product Using Booking and Seat Availability Data.* Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

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

Chung, J. Y. and Petrick, J. F. (2013). Price fairness of airline ancillary fees: An attributional approach. *Journal of Travel Research*, 52(2):168–181.

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

Conrady, R. (2013). Low cost carriers in western and central europe. In Gross, S. and Michael, L., editors, *The Low Cost Carrier Worldwide.* Routledge.

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.

Cui, Y., Duenyas, I., and Sahin, O. (2016). Unbundling of ancillary service: How does price discrimination of main service matter? Working Paper.

De Marcken, C. (2003). Computational complexity of air travel planning.

Ellison, G. (2005). A model of add-on pricing. *Quarterly Journal of Economics*, 120(2):585–637.

Espino, R., Martín, J. C., and Román, C. (2008). Analyzing the effect of preference heterogeneity on willingness to pay for improving service quality in an airline choice context. *Transportation Research Part E: Logistics and Transportation Review*, 44:593–606.

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):152–170.

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

Fruchter, G. A., Gerstner, E., and Dobson, P. W. (2011). Fee or free? how much to add on for an add-on. *Marketing Letters*, 22(1):65–78.

Gabaix, X. and Laibson, D. (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics*, 121(2):505–540.

Gaeth, G. J., Levin, I. P., Chakraborty, G., and Levin, A. M. (1990). Consumer evaluation of multi-product bundles: An information integration analysis. *Marketing letters*, 2(1):47—57.

Gallego, G. (2013). Single resource revenue management problems with dependent demands. Working Paper.

Gallego, G., Iyengar, G., Phillips, R., and Dubey, A. (2004). Managing flexible products on a network. Technical Report COCR Technical Report TR-2004-01, Department of Industrial Engineering an dOperations Research, Columbia University, New York.

Gallego, G., Li, L., and Ratliff, R. (2009). Choice-based EMSR methods for single-leg revenue management with demand dependencies. *Journal of Revenue and Pricing Management*, 8(2–3):207–240.

Garrow, L., Hotle, S., and Mumbower, S. (2012). Assessment of product debundling trends in the US airline industry: Customer service and public policy implications. *Transportation Research Part A: Policy and Practice*, 46(2):255–268.

Garrow, L., Jones, S., and Parker, R. (2007). How much airline customers are willing to pay: An analysis of price sensitivity in online distribution channels. *Journal of Revenue and Pricing Management*, 5(4):271–290.

Geng, X. and Shulman, J. D. (2015). How costs and heterogeneous consumer price sensitivity interact with add-on pricing. *Production and Operations Management*, 24(12):1870–1882.

Gorin, T. (2004). *Assessing low-fare entry in airline markets: impacts of revenue management and network flows.* Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

Greenleaf, E. A., Johnson, E. J., Morwitz, V. G., and Shalev, E. (2016). The price does not include additional taxes, fees, and surcharges: A review of research on partitioned pricing. *Journal of Consumer Psychology*, 26(1):105–124.

Guo, J. C. (2008). Estimation of sell-up potential in airline revenue management systems. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

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

Harvey, G. (1987). Airport choice in a multiple airport region. *Transportation Research Part A: Policy and Practice*, 21A(6):439–499.

Hess, S. and Adler, T. (2011). An analysis of trends in air travel behavior using four related SP data collected between 2000 and 2009. *Journal of Air Transport Management*, 17:244–248.

Hopperstad, C. (2005). Passenger O/D simulator PODS v.10 technical specifications.

Hopperstad, C. (2008). Q forecasting with fare families – methodology. Presented at the May 2008 PODS Consortium Meeting, Auckland, NZ.

Hopperstad, C. (2009). The un-bucketed DP (UDP): LDP using network displacement costs. Presented at the Jan 2009 PODS Consortium Meeting, Los Angeles, CA.

Hopperstad, C. and Belobaba, P. (2004). Alternative RM algorithms for unrestricted fare structures. In *Proceedings of the 2004 AGIFORS RM Study Group Meeting.*

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.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292.

Kothari, A., Madireddy, M., and Sundararajan, R. (2016). Discovering patterns in traveler behaviour using segmentation. *Journal of Revenue and Pricing Management*, 15(5):334—351.

Lal, R. and Matutes, C. (1994). Retail pricing and advertising strategies. *Journal of Business*, 67(3):345–370.

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

Lee, A. O. (1990). Airline reservations forecasting: Probabilistic and statistical models of the booking process. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

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

L'Heureux, E. (1986). A new twist in forecasting short-term passenger pickup. In *Proceedings of the 26th AGIFORS Symposium.*

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

Liu, Q. and Van Ryzin, G. (2008). On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management*, 10(2):288–310.

Loewenstein, G. (1996). Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes*, 65(3):272–292.

McFadden, D. (1999). Rationality for economists? *Journal of Risk and Uncertainty*, 19:73–105.

McFadden, D. and Train, K. (2000). Mixed MNL models for discrete response. *Journal of Applied Economics*, 15(5):447–470.

Metters, R., Queenan, C., Ferguson, M., Harrison, L., Higbie, J., Ward, S., Barfield, B., Farley, T., Kuyumcu, H. A., and Duggasani, A. (2008). The "killer application" of revenue management: Harrah's Cherokee Casino & Hotel. *Interfaces*, 38(3):161–175.

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.

Mumbower, S., Garrow, L. A., and Newman, J. P. (2015). Investigating airline customers' premium coach seat purchases and implications for optimal pricing strategies. *Transportation Research Part A: Policy and Practice*, 73:53–69.

Oi, W. Y. (1971). A Disneyland dilemma: Two-part tariffs for a Mickey Mouse monopoly. *The Quarterly Journal of Economics*, 85(1):77–96.

Open Allies (2014). Comments of Open Allies for Airfare Transparency. United States Department of Transportation Docket DOT-OST-2014-0056.

Phillips, R. (2005). *Pricing and Revenue Optimization.* Stanford Business Books, Stanford, CA.

Queenan, C. C., Ferguson, M., Higbie, J., and Kapoor, R. (2007). A comparison of unconstraining methods to improve revenue management systems. *Production and Operations Management*, 16(6):729–746.

Richter, H. (1982). The differential revenue method to determine optimal seat allotments by fare type. In *Proceedings of the 22nd AGIFORS Symposium.*

Schiffman, H. R. (1990). Chapter 19: The perception of time. In *Sensation and Perception: An Integrated Approach.* John Wiley & Sons.

Schwieterman, J. P. (1985). Fare is fair in airline deregulation: The decline of price discrimination. *Regulation*, 9:32–38.

Scotti, D. and Dresner, M. (2015). The impact of baggage fees on passenger demand on US air routes. *Transport Policy*, 43:4–10.

Shugan, S. M., Moon, J., Shi, Q., and Kumar, N. S. (2017). Product line bundling: Why airlines bundle high-end while hotels bundle low-end. *Marketing Science*, 36(1):124–139.

Shulman, J. D. and Geng, X. (2013). Add-on pricing by asymmetric firms. *Management Science*, 59(4):899–917.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1):99–118.

Simon, H. A. (1989). Cognitive architectures and rational analysis: Comment. In *21st Symposium on Cognition*. Department of Psychology.

Skwarek, D. K. (1996). Competitive impacts of yield management system components: Forecasting and sell-up models. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Smith, B., Darrow, R., Elieson, J., Guenther, D., and Rao, V. (2007). Travelocity becomes a travel retailer. *Interfaces*, 37(1):68–81.

Smith, B. and Penn, C. (1988). Analysis of alternate origin-destination control strategies. In *Proceedings of the 28th AGIFORS Symposium*.

Stremersch, S. and Tellis, G. J. (2002). Strategic bundling of products and prices: A new synthesis for marketing. *Journal of Marketing*, 66(1):55–72.

Swait, J. and Ben-Akiva, M. (1987). Incorporating random constraints in discrete models of choice set generation. *Transportation Research Part B: Methodological*, 21B(2):91–102.

Talluri, K. and van Ryzin, G. (1998). An analysis of bid-price controls for network revenue management. *Management Science*, 44(11):1577–1593.

Talluri, K. and van Ryzin, G. (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.

Teichert, T., Shehu, E., and von Wartburg, I. (2008). Customer segmentation revisited: The case of the airline industry. *Transportation Research Part A: Policy and Practice*, 42(1):227—242.

Thaler, R. (1985). Mental accounting and consumer choice. *Marketing science*, 4(3):199–214.

Theis, G. W. (2011). *Incorporating Attitudes in Airline Itinerary Choice: Modeling the Impact of Elapsed Time.* Ph.D. thesis, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA.

tnooz (2015). Airline system look-to-book ratios soar, expected to go 10x higher.

Train, K. (2009). *Discrete Choice Methods with Simulations.* Cambridge University Press, New York, NY, 2 edition.

Tuzovic, S., Simpson, M. C., Kuppelwieser, V. G., and Finsterwalder, J. (2014). From 'free' to fee: Acceptability of airline ancillary fees and the effects on customer behavior. *Journal of Retailing and Consumer Services*, 21(2):98–107.

van Ryzin, G. and Vulcano, G. (2008a). Computing virtual nesting controls for network revenue management under customer choice behavior. *Manufacturing & Service Operations Management*, 10(3):448–467.

van Ryzin, G. and Vulcano, G. (2008b). Simulation-based optimization of virtual nesting controls for network revenue management. *Operations Research*, 56(4):865–880.

van Ryzin, G. and Vulcano, G. (2017). Technical note—an expectation-maximization method to estimate a rank-based choice model of demand. *Operations Research*, 65(2):396–407.

Vinod, B. and Moore, K. (2009). Promoting branded fare families and ancillary services: Merchandising and its impacts on the travel value chain. *Journal of Revenue and Pricing Management*, 8(2–3):174–186.

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.

Walker, J. and Ben-Akiva, M. (2002). Generalized random utility model. *Mathematical Social Sciences*, 43(3):303–343.

Weatherford, L. (2016). Tests of revenue management performance under different demand correlation assumptions. *Journal of Revenue and Pricing Management*, 15(5):399–416.

Weatherford, L. and Pölt, S. (2002). Better unconstraining of airline demand data in revenue management systems for improved forecast accuracy and greater revenues. *Journal of Revenue and Pricing Management*, 1(3):234–254.

Wickham, R. (1995). Evaluation of forecasting techniques for short-term demand of air transportation. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Williamson, E. (1988). Comparison of optimization techniques for origin-destination seat inventory control. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Wittman, M. (2018). *Dynamic Pricing Mechanisms for Airline Revenue Management: Theory, Heuristics, and Implications.* Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

Wittman, M. D. and Belobaba, P. P. (2017). Dynamic availability of fare products with knowledge of customer characteristics. *Journal of Revenue and Pricing Management*, 16(2):201–217.

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.

Wollmer, R. (1992). An airline seat management model for a single leg route when lower fare classes book first. *Operations Research*, 40(1):26–37.

Xia, L. and Monroe, K. B. (2004). Price partitioning on the internet. *Journal of Interactive Marketing*, 18(4):63–73.

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

Zhuang, W. and Li, M. Z. F. (2012). Dynamic pricing with two revenue streams. *Operations Research Letters*, 40(1):46–51.

Zou, L., Yu, C., Rhoades, D., and Waguespack, B. (2017). The pricing responses of non-bag fee airlines to the use of bag fees in the US air travel market. *Journal of Air Transport Management*, 65:209–219.