Dynamic Pricing Mechanisms for Airline Revenue Management: Theory, Heuristics, and Implications

by

Michael D. Wittman

S.M., Transportation, Massachusetts Institute of Technology, 2014
B.S., Mathematics & Economics, American University, 2012

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

Doctor of Philosophy in Air Transportation Systems

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2018

Signature redacted

Author ____

Department of Aeronautics and Astronautics
February 1, 2018

Signature redacted

Certified by ____

Peter P. Belobaba
Principal Research Scientist of Aeronautics and Astronautics
Thesis Supervisor

Signature redacted

Certified by ____

Cynthia Barnhart
Chancellor and Ford Professor of Engineering
Committee Member

Signature redacted

Certified by ____

Hamsa Balakrishnan
Associate Professor of Aeronautics and Astronautics
Committee Member

Signature redacted

Accepted by ____

Hamsa Balakrishnan
Associate Professor of Aeronautics and Astronautics
Chair, Graduate Program Committee

Abstract

Even as the distribution and sale of commercial airline tickets has shifted in recent years from physical reservation offices to the Internet, many airline commercial processes remain highly reliant on pre-Internet technologies and standards. This legacy infrastructure compels airlines to publish a discrete set of prices in each market they serve, and to select prices for each itinerary from among only this limited set of possible price points.

Recent advancements in distribution technology, such as the New Distribution Capability (NDC), offer airlines the chance to break away from these constraints. These new standards enable the creation of customized offers with prices that could be generated dynamically in real time. While airlines have shown interest in these new technologies, practical methods for integrating dynamic pricing into existing airline revenue management (RM) and distribution systems have yet to be defined and evaluated by academics or practitioners.

In this work, we propose the first mechanisms for dynamic pricing designed specifically for use in the airline industry. By selectively providing increments or discounts based on demand segmentation and estimates of willingness-to-pay (WTP), our mechanisms can increase airline revenues by stimulating new bookings from price-sensitive travelers while encouraging more price-inelastic travelers to buy up to higher price points. Moreover, the methods are compatible with the pricing, RM, and distribution systems currently used by airlines today.

Our dynamic pricing heuristics emerge from the development of a novel theoretical model of customer choice. Using the model, we introduce a new concept called "conditional WTP" to describe how a customer's willingness-to-pay for an itinerary can change depending on the other alternatives available in his choice set. We show how assuming an unchanging maximum WTP for air travel, as in past work on dynamic pricing, can lead to overestimation of WTP in competitive environments, and describe how an airline's estimates of conditional WTP play an integral role in our dynamic pricing mechanisms.

We test our dynamic pricing methods in the Passenger Origin-Destination Simulator (PODS): a robust agent-based booking simulation that models the interactions between passengers and airlines. In a complex, competitive network, we find that our heuristics can increase airline revenues by up to $1-4\%$ from traditional pricing and RM alone. Incrementing prices can result in revenue gains through an increase in yield, and discounting can lead to higher revenues through demand stimulation and share shift from other airlines. In both cases, we identify a phenomenon we call "forecast spiral-up" which increases yield by protecting more seats for higher-value fare classes. We also develop a variant of the heuristic in which multiple substitutable flights are priced simultaneously, leading to additional revenue gains.

Finally, we provide the first in-depth assessment of the practical implications of dynamic pricing for the airline industry. We focus on airline concerns that dynamic pricing could lead to price wars, excessive discounting, and a race to the bottom. We also evaluate some of the potential legal implications and customer reactions that could emerge as dynamic pricing becomes more commonplace. These analyses provide new insight on how airline competition could potentially change as dynamic pricing is integrated into traditional airline processes.

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

Acknowledgments

It is difficult to overstate the gratitude I have for my doctoral advisor, Peter Belobaba, and his ceaseless support and mentorship throughout this journey. Peter's trust in me and my ideas gave me the confidence to thrive at MIT, and his pedagogy and encouragement have undoubtedly made me a better researcher, presenter, and teacher. I am so fortunate to have had the chance to work with Peter, and I will always be thankful for everything he has taught me and for making my graduate school experience such a positive one.

I am also very thankful to my Ph.D. committee, who are not only experts in their field, but also among the kindest people I have encountered in academia. Cindy Barnhart generously agreed to take time away from her demanding "other job" as Chancellor to be a part of my committee, and Hamsa Balakrishnan has given me encouragement throughout my time at MIT. Bill Brunger's industry perspective led to many interesting conversations and inspired Chapter 6 of this dissertation, and Larry Weatherford put his fine editing skills to work in his cover-to-cover reviews of the document. And while he wasn't a member of this committee, I will always be grateful to Bill Swelbar for his continuous mentorship and support.

This dissertation would not have been possible without the programming and development skills of Craig Hopperstad and Matt Berge. It is the ultimate luxury for a graduate student to have a professional team of software developers on his side, and I am very thankful for the many hours Craig and Matt spent helping me get my ideas implemented into PODS.

I am indebted to the members of the MIT PODS Consortium for their financial support and for their many helpful suggestions and comments as this work evolved. Particularly, my conversations with Thomas Fiig, Dan Zawack, Eric Ruhlin, Jonas Rauch, Luc Choubert, Marcial Lapp, and Praveen Bala challenged my preconceived notions and improved this dissertation. I am also grateful to Thomas Fiig, Jan Vilhelmsen, Robin Adelving, and Jean-Michel Sauvage for their hospitality during my research visit to Amadeus in August, 2016.

I was very fortunate to have taken this journey alongside Adam Bockelie, whose skills and diligence in PODS and RM far surpass mine. There are only so many people in this world who can talk about FRAT5s, Bellman equations, efficient sets, and detruncators, and I was lucky to have one of them in my office. Many other friendships were forged in 35-217, including Michael Abramovich, Tamás Kolos-Lakatos, Eric Hao, Daniel Fry, Alex Jacquillat, Oren Petraru, Germán Escovar Álvarez, Sathya Silva, Luke Jensen, Matthieu de Vergnes, Alex Bachwich, and J.-P. García. Airline 1 thanks you for your dedicated service.

My friends outside of MIT kept me balanced and let me talk about other things than airlines once in a while. I was so lucky to have met David and Jess Cohen-Tanugi, who have been a constant ray of positive light ever since I moved to Cambridge. I was fortunate to live with some excellent housemates at 8 Chatham: Kezia Ofiesh, Matt Chastain, Kenny Wong, and Catherine Rosenberg. And of course, my friends from Bedford—Dave Hursh, Ethan Jorgenson, Tony D'Andragora, and Corey Prescott—thanks for over fifteen years of Fridays.

Finally, to Leslie and Paul Wittman, who are probably as surprised as I am that their only child would end up jetting around the world talking about how to price airline tickets for a living, thank you so much for your unending love, support, and interest in whatever I end up doing. I am so grateful to have you both as my parents.

Contents

List of Figures

List of Tables

1 Introduction

"A high-street shop would never raise the price of an item during the short period of time that I've got it in my hands and am umm-ing and ahh-ing over whether to buy it. So why do airlines think they can do so?"
(Collinson, 2010)

1.1 Motivating dynamic pricing in the airline industry

With a dizzying and ever-changing array of itineraries, product options, and prices, the marketplace for air transportation is one of the most complex shopping experiences that an average consumer is likely to face in his or her daily life. Customers searching for air travel are often presented with dozens of possible alternatives with prices that seem to change at a moment's notice. It is no surprise that customers often believe that airlines currently use complex "dynamic pricing" methods to track their past behavior and create customized fare quotes based on their shopping history, search patterns or other characteristics (Collinson, 2010; McGee, 2013).

In fact, the price displayed for an itinerary is the result of the interaction between two separate and distinct airline functions: *pricing*, which defines a finite number of pre-priced *fare products* based on the competitive landscape of each air transportation market, and *revenue management* (RM), which determines which of these fare products to make available for sale at any given time. These actions are typically made sequentially. While prices for an itinerary may indeed change over time as fare classes become unavailable, these decisions are not typically tied to the characteristics of a specific customer or shopping session.

As information technology improves and as the number of bookings made through online channels increases, airlines are starting to gather more information about the preferences and behaviors of their customers. Moreover, as airlines begin to sell and distribute more complex products, the practice of filing a limited number of pre-defined fare products may limit an airline's ability to implement its commercial strategies.

Legacy technological infrastructure has historically limited the possibility for innovation in pricing and revenue management. Airlines still rely on distribution technology that was created prior to the invention of the Internet. The messages and standards used to distribute and sell airline tickets are rigid and inflexible, and contain only limited information about customers and shopping sessions. These messages are explicitly tied to alphabetic reservation booking designators (RBDs), meaning that airlines can only offer a maximum of 26 price points in a market across all classes of service at any given time.

The industry has slowly begun to modernize distribution capabilities to allow for more robust pricing and revenue management practices. An advanced distribution standard called the

New Distribution Capability (NDC) was first proposed in 2012 by the International Air Transport Association (IATA) and is slowly making its way into the airline industry.

NDC offers several advantages over existing distribution standards. It is much more flexible, allowing for the inclusion of additional information in distribution messages besides simply schedules and fares. NDC also enables the creation of "offers," where customers could be shown potentially unique bundles of fare products and add-on services (Hoyles, 2015).

With NDC, prices would no longer need to be tied to a reservation booking designator. This means that airlines would no longer be limited to choosing among a relatively small, finite number of price points to offer in a market at any given time. NDC has caused some practitioners to speculate that the airline industry is heading for a future without traditional filed fares, in which products and prices are generated "on-the-fly" for each shopping session (Westermann, 2006; Isler and D'Souza, 2009; Westermann, 2013; Bala, 2014; Fiig et al., 2015).

However, airlines are not yet scientifically or technologically ready to move away from filed fares. Many internal airline commercial processes, such as pricing, revenue management, revenue accounting, and revenue integrity, are highly reliant on filed fare classes. In the notoriously risk-averse airline industry, airlines are unlikely to fully abandon decades of pricing and revenue management knowledge and experience in favor of an entirely new dynamic pricing approach.

The development of NDC has allowed airlines to begin to imagine a world in which dynamic pricing mechanisms are *integrated* with traditional revenue management systems. This integrated approach would retain pre-priced fare classes that could be filed and distributed through either legacy distribution architecture or NDC. But it would also allow for *dynamic price adjustments* to be made to pre-priced fare products in certain situations.

For instance, an airline may decide to dynamically discount the price of a certain itinerary for a customer who is a member of the airline's loyalty program, or increase the price of a full flight for customers who are likely traveling for business. While this type of targeted pricing and marketing is commonplace in other industries, it would represent a significant step forward for airlines' commercial capabilities.

While industry working groups have started to consider the technological infrastructure that would be required to integrate dynamic price adjustments into traditional RM systems, the scientific algorithms that would drive dynamic pricing decisions have yet to be designed. It is also not a foregone conclusion that dynamic pricing will increase airline revenues. The revenue performance of these new methods would need to be compared to that of traditional airline RM methods for dynamic pricing to gain footing in a risk-averse industry. And, since dynamic pricing could significantly change the way airlines compete in the marketplace, a thorough examination of the practical implications of dynamic price adjustments is also an urgent need.

This dissertation aims to answer these questions by providing the first in-depth examination of how dynamic pricing mechanisms could be integrated with traditional airline revenue management systems. The primary contribution of the dissertation is a dynamic pricing method called Probabilistic Fare-Based Dynamic Adjustment (PFDynA) that increments or discounts the prices of pre-filed fare products in certain situations. Unlike other approaches in the literature that would require replacing traditional RM systems, PFDynA is designed to be used in conjunction with common RM methods that are currently in use in the industry. This means that airlines would not need to fully abandon well-understood and proven RM approaches to begin practicing dynamic pricing.

Through simulations of PFDynA in a variety of single-carrier and competitive scenarios, we find that dynamic pricing typically increases revenues over traditional revenue management alone. Unlike any past work, our review of the performance of PFDynA closely considers revenue management outcomes such as fare class mixes, load factors, yields, and booking curves, which are important for airline managers to understand the ways in which dynamic pricing could change commercial patterns. We also consider some concerns with dynamic pricing that have emerged in the industry, including the possibility that dynamic pricing mechanisms could spur a "race to the bottom" on price that would be damaging to industry profitability.

In the remainder of this introductory chapter, we first discuss in more detail how emerging technologies could soon allow airlines to dynamically adjust the prices of pre-determined fare products for specific shopping sessions. We then discuss the contributions of the dissertation in more detail, followed by an outline of the document.

1.2 Current processes of airline pricing, RM, and distribution

As shown in Figure 1.1, the commercial process at an airline currently consists of many interrelated but sequential functions—pricing, revenue management, inventory, fare quoting, and distribution. Each of these functions could be performed by different teams at the airline, which may have limited interaction.

First, the airline's pricing department typically files a set of pre-designed *fare products* with a central distribution agency, such as the Airline Tariff Publishing Company (ATPCO). Each fare product typically contains the features shown in Box 1 below (Vinod, 2010).

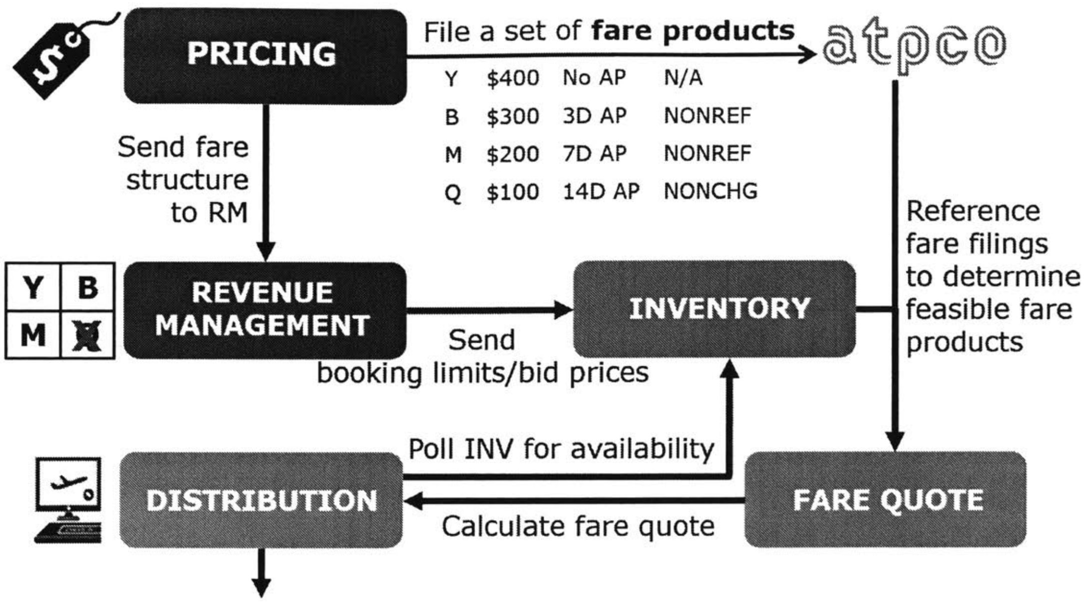

Figure 1.1: Current processes of airline pricing, RM, and distribution

Box 1: Features of fare products

- A *fare class* designator, which is represented by an alphabetic character from A – Z. This is sometimes called a *reservation booking designator* (RBD).
- A *price* for the fare product.
- A *validity date*, which describes the dates over which the fare product is allowed to be sold.
- An *advance purchase requirement*, which describes the number of days prior to departure for which the fare product is allowed to be sold. Common advance purchase requirements include 21 days, 14 days, and seven days prior to departure.
- A number of additional *fare restrictions* or *fare rules*. These restrictions impose additional conditions on when the fare can be sold or whether changes can be made after purchase. Some common fare restrictions include:
  - Round trip purchase requirements
  - Minimum stay requirements
  - Cancellation or reservation change fees
  - Non-refundability conditions

Fare products are typically created for each origin-destination market in which the airline sells tickets. A collection of multiple fare products in a single market is called a *fare structure*. Table 1.1 provides a simple example of a restricted four-class fare structure.

Fare Class	Fare	Adv. Purch.	Min Stay	Refundable?	Change Fee?
Y	$400	N/A	N/A	Yes	N/A
B	$300	3 days	N/A	No	$200
M	$200	7 days	N/A	No	$200
Q	$100	14 days	Sat. Night	No	Not Changeable

Table 1.1: A simple four-class restricted fare structure

The fare structures created by the pricing department are typically visible publicly to all other airlines after they are filed. In practice, airlines often match each others' fare structures exactly in the markets they serve (Lua, 2006; Vinod, 2010). Airline pricing analysts are alerted to changes in other airlines' fare structures by pricing systems, allowing them to respond by matching changes or, at times, choosing not to match for competitive reasons.

After the fare structures are filed, it is the job of the revenue management department and its systems to determine which fare products to make available for purchase at a given time for each future date and flight departure. The number of seats that are made available in each class for a specific flight departure is a function of the airline's forecast of remaining demand, the remaining time before departure, and the capacity of the aircraft, among other factors (Talluri and van Ryzin, 2005; Belobaba et al., 2016).

The output of the revenue management system is a series of *availability decisions*. These decisions describe how many seats are available for purchase in each of the fare classes in the fare structure. Fare classes with at least one seat available for purchase are referred to as *open*. If a fare class has no seats available for purchase on a flight, it is referred to as *closed*.

There are two ways in which a given fare class may switch from open to closed. First, all of the available seats in that fare class may be purchased. In the fare structure in Table 1.1, if there is a single seat in Q-class available for purchase, Q-class would switch from open to closed if a customer purchases that seat. The displayed price of the itinerary would then change from $100 to $200. Second, airlines periodically reoptimize the availability controls for each fare product. During a periodic reoptimization, the RM system could decide to close a class if, for instance, the flight is booking more quickly than was forecast.

The information from the revenue management system about which fare classes are open and which are closed is held in the airline's inventory (Gunther et al., 2010). In response to a shopping request, the airline's distribution system will poll the inventory to see which classes in the fare structure are available. It will also reference the fare structure to determine which classes have rules and restrictions that are *feasible* for the shopping request. For instance, a fare product that requires a Saturday night stay will not be feasible for a two-day Monday – Wednesday shopping request.

Once the set of available, feasible fare products is generated, the distribution system will then typically query the pricing server to determine the final price for the product, including any relevant taxes or carrier-imposed fees. The least-expensive available fare product that meets the conditions of the request is then typically displayed to the customer for purchase, although some airlines may also display other products with different characteristics or included features.

1.2.1 The role of distribution standards: EDIFACT and NDC

Airlines sell and distribute their products through two main channels. The *direct channel* refers to an airline's own website or call center, and provides the airline with the most control over the content displayed in response to a search for air travel. The *indirect channel* involves the substantial portion of bookings that are made through online travel agents (OTAs), brick and mortar retailers, corporate travel services, or metasearch websites (Gunther et al., 2010; Vinod, 2015). The technological backbone of the indirect channel is comprised of the global distribution systems (GDSs), which serve to exchange information, such as schedules and fare product availability, between airlines, travel agents, and passengers. In 2015, the GDSs serviced approximately half of global bookings (Taubmann, 2016).

The indirect channel provides airlines much less flexibility in the ways they can price and sell their products than the direct channel. This is due primarily to technological limitations on the type and volume of information that can be contained in the messages sent between airlines and the GDS. Airlines have historically communicated to travel agents and passengers through the GDS by the means of the EDIFACT standard. EDIFACT is a pre-Internet "language" by which airlines can send information between their computerized reservation systems (CRSs) and the GDS (Niketic and Mules, 1993; Vinod, 2015). Information is exchanged through a series of pre-defined and highly standardized request/response messages.

A stylized example of an EDIFACT PAOREQ message that is requesting fare class availability for a certain flight itinerary (BOS-NRT-BOS on flight JL7/JL8 from 18 July - 27 July) is shown in Figure 1.2. Upon receiving this PAOREQ request message, the airline would respond with a message that details the current seat availability on the requested flights. A stylized PAORES response message in the EDIFACT standard is shown in Figure 1.3.

```
UNB+IATB:1+1APPC+ZE0AV+270616:1617+091331300005C1'
UNH+1+PAOREQ:93:2:IA+091331300005C1'
MSG+1:46'
ORG+1A:BOS+99999992:USD:EN+A0001AAGS'
ODI+BOS+NRT+BOS'
TVL+180716+BOS+NRT+JL+0007+1++ P
TVL+270716+NRT+BOS+JL+0008+1++P'
UNT+6+1'
UNZ+1+091331300005C1'
```

Figure 1.2: Stylized UN/EDIFACT PAOREQ availability request message (Lacroix (1996))

```
UNB+IATB:1+ZE0AV+1APPC+270616:1618+9253C7FD0001+091331300005C1'
UNH+1+PAORES:93:2:IA'
MSG+1:47'
ODI+BOS+NRT+BOS'
TVL+180716+BOS+NRT+JL+0007+1++ P
PDI++Y:9+K:9+Q:9+T:9+V:9+X:4+H:0+L:0
TVL+270716+NRT+BOS+JL+0008+1++P'
PDI++Y:9+K:9+Q:9+T:5+V:4+X:2+H:0+L:0 '
UNT+6+1'
UNZ+1+9253C7FD0001'
```

Figure 1.3: Stylized UN/EDIFACT PAORES availability response (Lacroix (1996))

This message communicates the number of seats on flights JL7 and JL8 that are available for purchase in various fare classes. For example, on flight JL7 on July 18th, at least 9 seats are available in Y, K, Q, T, and V classes and 4 seats are available in X class. Classes are typically listed in descending order in order of price. Since X class seats are available on both JL7 and JL8, X class would be the least-expensive class that would be available for sale on the BOS-NRT-BOS itinerary in the example above.

Note that this EDIFACT request/response framework communicates only a limited amount of information. Only information about fare class availability is displayed, and a separate query needs to be made to request the price of the X class itinerary. No information about ancillary products, such as checked baggage or premium seats, is displayed in this message. While the identity of the travel agent making the search may be communicated through EDIFACT, it is relatively blind to the identity of the customer making the request.

As airline products have become more complex and more as customer information has become available, airlines, vendors, and industry groups have begun to develop new infrastructure for disseminating information through the indirect channel. The result of this effort was the International Air Transport Association (IATA)'s New Distribution Capability (NDC), which launched in 2012 and was approved by the U.S. Department of Transportation in 2014. NDC is an Extensible Markup Language (XML)-based standard that was developed to allow more information to be passed between airlines and the GDS during the shopping process. This could include details regarding ancillary services, rich media (such as pictures of an extra legroom seat), and optional personal information about the customer making the booking request, including frequent flyer numbers (Hoyles, 2015).

When communicating via NDC, a customer can, if she chooses, provide identifying information to the airline that the airline can use when making a fare product offer. This information could include a frequent flyer number, point-of-origin, trip purpose, or other characteristics. Airlines could use this information to create a customized offer—perhaps including a mix of ancillary products—targeted to the passenger making the shopping request (International Air Transport Association, 2016). A highly stylized example of the types of fields that could be present in an NDC product shopping request is shown in Figure 1.4.

```
<ShoppingRequest>
        <RequestDetails>
                <OriginDestination>
                        <Orig> BOS </Orig>
                        <Dest> NRT </Dest>
                </OriginDestination>
                <Dates>
                        <Depart>18JUL2016</Depart>
                        <Return>27JUL2016</Return>
                </Dates>
                <Traveler>
                        <Name>Michael Wittman</Name>
                        <FFAirline1>UA</FFAirline2>
                                <FF#1>LW6xxxxx</FF#1>
                        <FFAirline2>DL</FFAirline2>
                                <FF#2>91xxxxxxx</FF#2>
                        <Point-of-sale>USA</Point-of-sale>
                </Traveler>
                <Preferences>
                        <Checked-Bags> 1 </Checked-Bags>
                </Preferences>
        </RequestDetails>
</ShoppingRequest>
```

Figure 1.4: Highly stylized example of XML-based shopping request

Following a search for air travel, the NDC message would be sent to an aggregator, which then communicates with airlines to retrieve a set of offers for the customer. Along with a flight itinerary, these customized offers could include a bundle of ancillary services, associated fare rules, and an offer-specific price. The offers are then returned to the customer, who can choose to purchase an offer before a certain offer-expiry date (International Air Transport Association, 2016).

Critically, note that the NDC message does not require any information about fare classes to be transmitted, in contrast to the EDIFACT framework. This means that airlines would not necessarily need to pre-file fares with a central agency like ATPCO, nor decide on a fixed number of price points in advance. The offer generation capability of NDC allows airlines much more flexibility in the ways they can generate and price products, but also increasing complexity in terms of how these dynamically priced products would interact with existing airline systems for revenue management, revenue accounting, finance, and marketing.

1.2.2 Standards for Dynamic Pricing Engines

As NDC would require many changes to long-standing processes, airlines have generally proceeded slowly with NDC implementation. In 2015, a working group convened by ATPCO started to consider dynamic pricing mechanisms that would still be compatible with legacy standards. The result was the development of standards for "dynamic pricing engines" (DPEs), which apply price adjustments to pre-filed ATPCO fares (Ratliff, 2017).

A general schematic for dynamic pricing engines as proposed by the ATPCO Working Group is shown in Figure 1.5. During the shopping process, the DPE would be queried to determine if the fare that would normally be offered in response to the shopping request should be adjusted up or down, or whether no adjustment should apply. The DPE-adjusted fare (if applicable) is then sent to the customer. During the pricing phase, the DPE is queried again to ensure that the quoted price adjustment is still valid.

An airline using a DPE would have a rules engine that decides when and how the DPE will adjust the prices of filed fare products. This engine could be based on pre-defined business rules designed by analysts or managers, or by scientific algorithms to compute price adjustments based on session-specific information, demand forecasts, or customer segmentation.

To ensure that DPE-adjusted fares can be tracked and verified, the ATPCO Working Group has proposed filing all DPE adjustments as private fares with ATPCO. Adjusted fares could then be validated in downline processes such as revenue accounting. It is not yet clear whether DPE-adjusted fares will be visible to other airlines, or if they will remain opaque.

The DPE standards proposed by the ATPCO Working Group could clearly allow for the practical deployment of a dynamic pricing mechanism within traditional airline revenue management systems. Development of DPEs has moved to a pilot phase with airlines and vendors working together to determine specifications and data requirements, and others

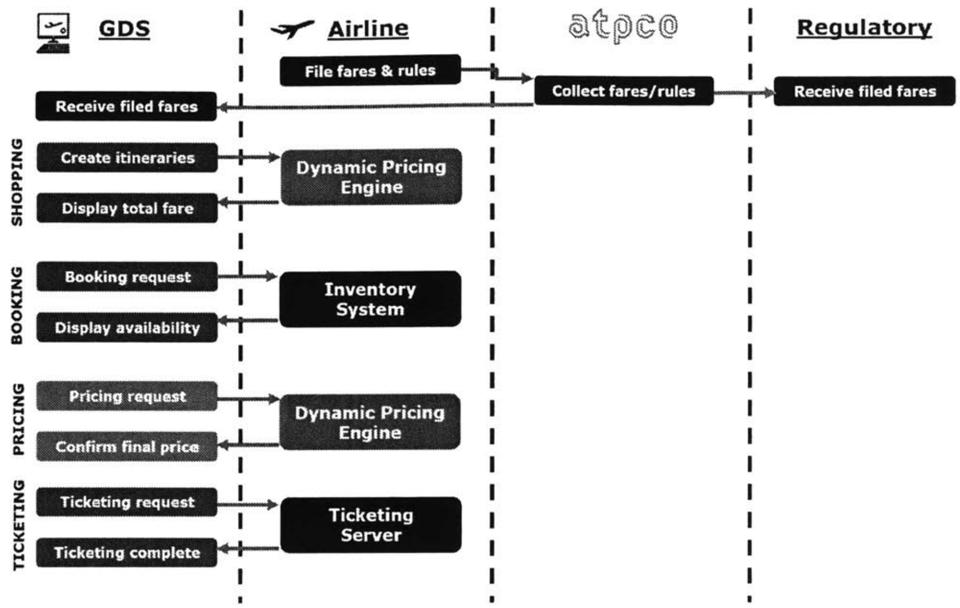

Figure 1.5: ATPCO Dynamic Pricing Engine Schematic (Source: ATPCO Dynamic Pricing Working Group)

are also starting their own development efforts for dynamic pricing mechanisms to use in conjunction with traditional RM (Fiig et al., 2016; Kumar and Walczak, 2016; Ratliff, 2017). These new industry standards for DPEs provide a practical pathway for implementation of the PFDynA dynamic pricing methods developed in this dissertation.

1.3 Contributions of the dissertation

This dissertation provides the first comprehensive overview of the integration of next-generation dynamic pricing mechanisms into traditional airline RM systems. A variety of models, heuristics, and techniques related to dynamic pricing in the context of airline RM are motivated, generated, discussed, and simulated. The following topics represent the core contributions of the dissertation:

- To help guide our discussion of the existing literature in the field, we first present a new **definitional framework** of next-generation pricing mechanisms. The framework draws distinctions between current airline pricing and revenue management methods and next-generation dynamic pricing mechanisms that incorporate session-specific information to make pricing or availability decisions. Within the context of this framework, relevant dynamic pricing research from the operations research and economics literature is reviewed.

- We then develop a novel **customer choice model** that describes our assumptions of how customers make decisions amongst different itineraries. The choice model is inspired by the field of rational choice theory. Unlike common choice-based revenue management approaches based on discrete choice theory or the multinomial logit model, which assumes *probabilistic* choice among alternatives, rational choice theory assumes that passengers make *deterministic* choices among alternatives based on their idiosyncratic tastes and preferences.

- The choice model reveals two important concepts that drive customer decision making: **maximum willingness-to-pay (WTP)**, which is the highest price at which a customer will purchase a given product if there were no other choices available, and **conditional WTP**, which is the the most a customer is willing to pay for a product among an assortment of other alternatives, each with their own prices and attributes.

  Conditional WTP is a new development of this dissertation, as it acknowledges that a customer's reservation price for a product is dependent on the other products from which he can choose. This concept is particularly important for dynamic pricing. Using maximum WTP to compute optimal prices, as has been done in some past work (e.g. Kambour (2014); Kumar et al. (2017)), could overestimate the customer's ability to pay in a competitive environment and result in a suboptimal choice of price for that customer or request.

- The choice model described above supports the development of the first **dynamic price adjustment** mechanism that is designed specifically for use in the airline industry and with airline revenue management. Uniquely among dynamic pricing models in revenue management, our model decouples the dynamic pricing decision from the airline revenue management problem, allowing price adjustments to be made after optimization by a traditional RM system. This dynamic pricing model is best described as a *dynamic price adjustment* mechanism for airline revenue management.

  This feature of the mechanism allows airlines to continue to use the revenue management techniques that have been developed over the past three decades to generate initial assortments of fare products available for purchase. This makes our dynamic pricing model significantly more practical than other approaches, which would require abandoning current fare filings, fare structures, distribution messages, and revenue management technologies.

- From the general dynamic pricing model, we propose several practical **dynamic pricing heuristics** that could be used by airlines after making basic assumptions on conditional WTP and marginal costs. Generally, the heuristics work by providing discounts in certain situations when the customer is likely to have a lower willingness-to-pay, while incrementing prices in other situations when the customer is likely to have a higher willingness-to-pay.

The heuristics take as inputs information from existing airline revenue management systems, such as leg bid prices, when computing their dynamic price adjustments. We allow for increments or discounts to be provided to customers based on an estimate of conditional WTP, which could change from request to request. We also introduce a new mechanism to allow airlines to generate a number of conditional WTP distributions using a single input parameter.

- The dynamic pricing heuristics are **tested in the Passenger Origin-Destination Simulator (PODS)**; a complex simulation of the interactions between passenger purchasing decisions and airline revenue management systems. The heuristics are tested in a number of competitive airline networks, and the revenue results are compared to a base case with no dynamic pricing to estimate the range of potential revenue gains. We also run simulations in which all airlines in the network use dynamic pricing, to observe whether revenue gains are still possible in competition. For most mechanisms we test and in most networks, we find positive revenue performance from dynamic pricing when used by a single airline alone or by all airlines in a competitive environment.

- We also **extend the dynamic pricing heuristics** to price multiple flights simultaneously using a new approach to modeling choice among substitutable flights. The dynamic pricing problem with multiple substitutable flights is much more complex than the single-flight dynamic pricing problem. Using a Hotelling model of product differentiation (Hotelling, 1929), we find the optimal dynamic prices for two flights that are horizontally differentiated by departure time and also vertically differentiated by schedule quality. With knowledge of the underlying departure time preference distribution, we find through simulation results that simultaneously pricing both flights marginally increases the revenue performance of the dynamic pricing heuristics. More complex schemes involving multiple product characteristics or more than two flights are also explored.

- Finally, this dissertation provides the first thorough discussion of the **practical implications** of switching from a quantity-based (Cournot) method of competition in the airline industry, where firms are price-takers and set only the quantity of seats available for sale for each flight at a given time, to a price-based (Bertrand) method of competition in which prices are generated dynamically. Through a review of the related economics and industrial organization literature, we discuss whether a "race-to-the-bottom" could result, which could lead to lower revenues for the airlines. We review some key results regarding the so-called "Bertrand paradox" in several situations, including where products are differentiated, when firms interact repeatedly, and when the prices offered by firms are either visible or opaque to other competitors. We also discuss the potential legal implications of dynamic pricing, as well as how customers may react to these new airline pricing practices.

1.4 Dissertation structure

The remainder of the dissertation is structured as follows. In Chapter 2, we review the relevant literature in dynamic pricing, assortment optimization, and airline revenue management through the lenses of economics and operations research. While each of these disciplines has contributed to the development and analysis of dynamic pricing strategies, the existence of an optimal or heuristic approach to dynamic pricing that fits the unique data availability, history, and technological limitations of the airline industry has remained elusive. In this chapter, we discuss the highlights of each of these streams of literature and propose how they could be fused together.

Chapter 3 presents the theoretical underpinnings of the airline dynamic pricing mechanisms developed in this dissertation. We start by developing a customer choice model for airline itineraries based on the classical building blocks of rational choice theory. We introduce two important concepts: a customer's maximum willingness-to-pay (WTP) for a product and her conditional WTP for the same product in the context of other alternatives.

Using this choice model, we then develop a dynamic price adjustment model, called Probabilistic Fare-Based Dynamic Adjustment (PFDynA), for airline itinerary products. PFDynA works by computing discounts or fare increments for individual itinerary products, based on an estimate of a passenger's conditional WTP and (optionally) characteristics of the customer or the request. The PFDynA heuristic is parameterized using a concept called a "Q-multiplier," which represents the airline's estimate of mean conditional WTP in that market. We close the chapter by proposing several variants of the PFDynA heuristic, depending on the amount of information available about the customer and whether the airline desires to apply increments, discounts, or both.

In Chapter 4, we test the dynamic pricing heuristics developed in Chapter 3 in the Passenger Origin-Destination Simulator (PODS), an airline revenue management simulator that models the interaction between airline revenue management systems and passenger choice in competitive, complex aviation networks. The PFDynA heuristics are tested in a wide variety of monopolistic, duopolistic, and oligopolistic environments.

The results generally show that dynamic price adjustment is revenue-positive for airlines in both single-airline and competitive scenarios. In most cases, the heuristics result in revenue gains over the base case even if all of the airlines in the simulated world are practicing dynamic pricing. We discuss the mechanisms by which dynamic pricing appears to result in simulated revenue gains, including forecast spiral-up, stimulation of demand, and capture of passengers from competitors.

In Chapter 5, we develop and test an extension to the model, called Simultaneous PFDynA, that prices multiple substitutable itineraries at the same time. By pricing two flights at simultaneously, airlines can incentivize price-sensitive travelers to purchase a less-attractive itinerary option, allowing more seats to be sold to high-value, schedule-sensitive travelers on a more-attractive itinerary. This model is tested in PODS for a case of two differentiated flights, and the results show incremental revenue gains over flight-by-flight PFDynA. However, these incremental gains are relatively small compared to the revenue performance of flight-by-flight dynamic pricing alone.

Chapter 6 discusses some practical implications of dynamic pricing for the airline industry. Dynamic pricing would represent a significant change in the way airlines price and sell itineraries, and worries of a "race-to-the-bottom" in which airlines repeatedly undercut each other on price has raised concerns that dynamic pricing could lead to revenue losses in the real world. We discuss some theoretical results from the economics and game theory literature to investigate whether dynamic pricing could lead to lower prices and revenues in practice. We also discuss issues related to the potential legal implications of dynamic pricing, as well as potential customer reactions to these new pricing practices.

Finally, Chapter 7 concludes the dissertation by reviewing the results of the previous chapters and discussing opportunities for future work on dynamic pricing in the airline industry.

2 Definitional Framework and Literature Review

In this chapter, we examine relevant literature for airline pricing and revenue management through the lenses of economics and operations research. We begin by introducing a definitional framework that describes different ways that firms can select prices for a given product. This framework is then used to describe how airlines are beginning to move from traditional revenue management practices towards more dynamic, real-time continuous pricing approaches.

2.1 A definitional framework for dynamic pricing mechanisms

Firms practice *dynamic pricing* when they charge different customers different prices for the same product, based on an *observable state of nature* (Wittman and Belobaba, 2017c). The observable state of nature could include remaining product inventory, time remaining in the selling period, characteristics of the customer, forecasts of demand-to-come, competitor offerings, or other factors. Changes in the observable state of nature could result in the selection of a different price for the same product.

Dynamic pricing need not result in different prices for each individual customer. Some firms may change their prices infrequently, such as on a daily or weekly basis. At the other extreme, firms practice *transactional dynamic pricing* when they select prices for each shopping request based on the characteristics of the request or, at the limit, of the individual customer making the request.

Firms practicing dynamic pricing can use different mechanisms to adjust prices from transaction to transaction. In this framework, we discuss three such mechanisms: assortment optimization, dynamic price adjustment, and continuous pricing.

2.1.1 Assortment optimization

One common way for firms to select prices for a given product is to first pre-define a finite set of possible price points, and then select which of those prices to display to each customer based on the observable state of nature. We call this approach *assortment optimization*, since the firm is selecting one or more price points from a pre-set assortment of possible options.

The selection of prices from the pre-defined menu of price points could be governed by business rules or restrictions. For instance, certain price points may only be available when purchased at least 21 days in advance, or if the customer is a senior citizen. Firms practicing assortment optimization may also vary in how often prices are selected from the menu. Some firms may select prices infrequently (for instance, on a monthly or weekly basis), while others may perform this selection based on the characteristics of each transaction. However, a key feature of assortment optimization is that the underlying set of price points consists of a finite number of possible options, and does not change from transaction to transaction.

The following definition formalizes the concept of assortment optimization.

Definition: Assortment Optimization
Firms practice *assortment optimization* when they first pre-define a finite set of possible price points, then select one or more of those price points to display to customers. Prices may be selected from the set infrequently (e.g., daily or weekly), or on a transaction-by-transaction basis. However, the menu of possible price points does not frequently change, and customers may not be offered prices or products that are not included in the set.

Traditional airline pricing and revenue management is clearly an assortment optimization problem. In most airline RM models, airlines decide which of a series of fare products (for instance, fare classes $\{Y, M, B, Q, X\}$) to offer to a customer at a specific moment in time. Fare product availability could be based on rules or restrictions that govern when each product can be sold, or based on the outcome of a mathematical optimization model. After determining the availability of each fare product, the least-expensive feasible price point is typically displayed to the customer. However, the underlying set of possible price points is updated only infrequently, and this fare structure is typically taken as fixed by airline revenue management systems.

2.1.2 Dynamic price adjustment

Firms practicing assortment optimization may sometimes wish to offer a price for a product that is not included in their pre-set menu of possible price points. Firms could do so by adding to or updating their set of possible price points. This may not always be the best approach—for instance, if the desired price point is only valid for particular customers, or only applicable for a limited period of time.

In such cases, firms may wish to *dynamically adjust* the price points in their pre-defined menu. For instance, a firm may wish to offer a discount off of a pre-defined price point for customers that are members of the firm's loyalty program, or increment the price in situations when market demand is likely to be higher. All of these adjustments are made in reference to a pre-defined price point, and some customers may receive no adjustment.

Definition: Dynamic Price Adjustment
Firms practicing *dynamic price adjustment* begin by pre-defining a set of possible price points, as in assortment optimization. The firm then selects one or more of these price points to make available to a customer. In certain situations, the firm can then apply an adjustment (either up or down) to the selected price point. These adjustments can be determined occasionally (e.g., daily or weekly), or on a transaction-by-transaction basis. Some customers may receive no adjustment to the pre-defined price point.

One example of a firm practicing dynamic price adjustment is the online retailer Amazon. Amazon frequently adjusts the prices of books by marking down from a manufacturer list price. Amazon's prices are displayed in reference to this list price, and the amount of the markdown can change based on demand or remaining inventory. While Amazon practices dynamic price adjustments via markdowns, the on-demand transportation app Uber's "surge pricing" mechanism uses dynamic markups to adjust prices higher in response to high demand or a limited supply of drivers.

By dynamically adjusting prices up or down, firms can effectively fill in the gaps between the pre-determined price points they define in assortment optimization. However, any adjustments are made in reference to one of these price points. Dynamic price adjustment can thus be seen as an intermediate step between assortment optimization and more complex continuous pricing schemes, described next. The heuristics that we develop in the remainder of this dissertation are best defined as dynamic price adjustment mechanisms.

2.1.3 Continuous pricing

With *continuous pricing*, firms are not limited to a finite, pre-defined set of possible price points, Instead, firms can freely select prices from a continuous range of values. As with the previous mechanisms, this price selection process can occur infrequently, or (at the limit) individually for each transaction.

We call this latter approach *transactional continuous pricing*, which represents the most advanced pricing practice currently used by firms today. While continuous pricing gives firms the most flexibility in the prices they can offer to their customers for a given product, it can also be more difficult to implement than simpler assortment optimization approaches that are based on a set of fixed, pre-determined price points.

Definition: Continuous Pricing
Firms practicing *continuous pricing* select prices for their products from a continuous range of possible values. Unlike assortment optimization and dynamic price adjustment, there is no finite, pre-defined set of possible price points. Prices could be selected infrequently or (at the limit) on a transaction-by-transaction basis (transactional continuous pricing).

In recent years, Uber has moved from a dynamic price adjustment model to a transactional continuous pricing mechanism to dynamically compute the price for each ride. The prices charged by Uber's "market-based pricing" system are not based on a fixed tariff schedule, but are determined instead by the characteristics of the ride and current demand and supply conditions (Newcomer, 2017). In this way, two rides that cover the same distance in the same traffic conditions could be charged different prices based on the market conditions in place at the time, as opposed to Uber's previous pricing system in which prices were determined by distance and time alone.

2.1.4 Summing up the definitional framework

The following box summarizes the three pricing mechanisms discussed in this section for selecting prices for a given product:

- Firms practicing ***assortment optimization*** select prices from among a finite, pre-determined set of possible price points.
- Firms practicing ***dynamic price adjustment*** first use assortment optimization to select a price point from a pre-determined set of possibilities. Then, firms can increment or decrement this price point for certain customers or in certain situations.
- Firms practicing ***continuous pricing*** select prices freely from a continuous range of possible values. Unlike assortment optimization and dynamic price adjustment, there is no pre-defined set of possible price points.

While these definitions describe the mechanisms by which firms can set prices for a given product, they do not discuss the factors that determine how one price point is selected over another, nor the scientific models that are used to compute prices. Academic literature has provided countless examples of mathematical optimization models, heuristics, and rules of thumb for dynamically selecting prices. These models vary depending on the mechanism used by firms and whether a set of pre-defined price points is taken as an exogenous input.

Furthermore, the interactions between the state of nature and dynamic pricing mechanisms are treated differently in different fields of study. For instance, the literature in the field of economics focuses on selecting prices using principles of price discrimination to segment customers into different groups. In contrast, the operations research literature typically focuses on optimal solutions to capacity-constrained pricing problems, often using network optimization approaches to select which price points to make available. In the next sections, we discuss literature related to pricing and revenue management mechanisms in each of these fields, which will inspire our later development of a dynamic price adjustment mechanism for airline revenue management.

2.2 Economics literature: Price discrimination

We begin by examining classic pricing mechanisms within the economics literature of industrial organization and the theory of the firm. We will begin with a simple example of a monopolist pricing a single product.

In the simplest textbook examples of monopolistic pricing, firms face a demand function $Q(p)$ as a function of a price p (Tirole, 1988). A monopolist's static pricing decision is to find the price p^* that maximizes her expected profits Π as a function of the demand function and the firm's marginal cost (MC) for the product:

$$p^* = \arg\max_p \Pi(p) = (p - MC) \cdot Q(p) \tag{1}$$

At a constant marginal cost, the firm would decide to charge the same profit-maximizing price to all customers. This makes sense if all of the customers in the market are homogeneous and react to changes in price in the same way. However, in most markets, customers are heterogeneous, having different valuations for the same good. It may therefore be desirable in some situations to charge different prices to different segments of customers, if possible. Economists refer to this concept as *price discrimination.* There are generally three categories or "degrees" of economic price discrimination:

- First-degree (or perfect) price discrimination occurs when each customer is charged a unique price exactly equal to his valuation for the product (Tirole, 1988).[1] This assumes that customer valuations are public knowledge that are visible to the firm, and that customers can be perfectly targeted and charged a unique price. First-degree price discrimination could be thought of as transactional continuous pricing, using the framework discussed in the previous section.

[1]We will explore the concept of customer valuation or willingness-to-pay in more detail in Chapter 3.

- Second-degree price discrimination occurs when customers are charged different prices depending on how many units of the good they choose to purchase. Second-degree price discrimination often occurs using mechanisms like two-part tariffs, in which customers pay a fixed "entry" fee for the right to purchase the product and then a variable tariff for each unit of the good they consume. One classic example of second-degree price discrimination is amusement parks, in which a fixed admission fee is often charged followed by an additional charge for each ride.
- Third-degree price discrimination occurs when firms can segment customers into a fixed number of categories or groups and provide a unique price to each group. Third-degree price discrimination can be seen as an incomplete version of first-degree price discrimination; it is not possible to identify each customer individually and charge him a price equal to his valuation. Instead, segmentation happens on a broader level. Some examples of third-degree price discrimination are senior discounts, deep-discount leisure fares for air travel, and student rates for subscriptions to magazines.

The traditional airline pricing and revenue management problem is most closely related to third-degree price discrimination. By creating restricted fare structures with onerous conditions for purchasing the cheapest product, airlines are attempting to segment customers into groups of price-sensitive leisure customers and restriction-averse business customers. Third-degree price discrimination can increase revenues for the monopolist over a uniform-pricing scenario, but will result in different prices being charged to different groups of customers, which may or may not be efficient from a social perspective (Tirole, 1988).

An alternative approach to price discrimination is quality discrimination, where firms indirectly segment customers into groups by offering goods of different quality. Theorists have found that monopolists practicing quality discrimination will produce a wide range of goods at different price points (Mussa and Rosen, 1978; Tirole, 1988). In some cases, firms may even choose to "damage" some of their goods to further increase their ability to practice price discrimination (Deneckere and McAfee, 1996). This could explain the recent introduction of "basic economy" fare products by large legacy carriers in the United States (Sumers, 2017).

Basic economy products are low-priced fare products that offer few or no benefits like seat selection, complementary baggage, early boarding, or mileage accrual that are associated with traditional fare products. The purpose of these basic economy products is to segment extremely price-sensitive passengers away from traditional leisure fares that include these standard benefits. Indeed, airline executives have discussed the success of basic economy in terms of how many passengers avoid it by instead buying up to a more expensive product (Jansen, 2017).

2.2.1 Competitive price discrimination

Price discrimination typically increases the monopolist's revenue by allowing it to capture more of the area underneath the demand curve. However, increasing the number of firms in the market can lead to different outcomes depending on the degree of price discrimination that is possible. With multiple firms, price discrimination can sometimes lead to *lower* revenues compared to static pricing mechanisms, even when the products are differentiated.

As an example, consider the classic "horizontal differentiation" model of Hotelling (1929), as shown in Figure 2.1. Suppose customers possess locational preferences θ that are uniformly distributed over a line from 0 to 1.[2] There are two firms, one located at location 0 and one at location 1 on the line. Customers face a transportation cost t to move towards either firm. For instance, a customer located at position θ will incur a cost θt to purchase from firm 0.

Figure 2.1: A "Hotelling line" of spatial differentiation

Customers may decide to purchase from a firm located farther away in exchange for a low price. It is possible to show that the equilibrium of this game is to set prices as a function of the firms' marginal costs marked up as a function of the transportation cost t. Hotelling (1929) also shows that when firms are allowed to choose where to locate on the line, "maximum differentiation" (locations 0 and 1) will also emerge as an equilibrium.

Price discrimination changes the outcome of the Hotelling (1929) game. In a classic extension, Thisse and Vives (1988) consider a situation when firms can practice first-degree price discrimination. That is, firms have perfect knowledge of each customer's preference θ and can price accordingly. Thisse and Vives (1988) show that in this scenario, all customers will be charged prices that are lower than in the base Hotelling case.

This is because the presence of personalized information intensifies the competition for each customer; it allows firms to target discounts towards customers that would prefer the other firm's products to try to encourage those customers to switch. As Fudenberg and Villas-Boas (2006) state (p. 68): "competition with customized prices becomes like competition with no differentiation, in which at the equilibrium prices, an infinitesimal small price cut attracts all the demand." Thisse and Vives (1988) show that when no additional customers decide

[2]The Hotelling line could represent any horizontal difference between firms, such as product quality or geographic location relative to a customer's home or work.

to purchase at the lower prices, first-degree price discrimination can lead to lower revenues for both firms, but higher consumer surplus.

2.2.2 Behavior-based price discrimination

Inspired by the work of Thisse and Vives (1988), a rich subset of the economics literature has developed to examine so-called "behavior-based price discrimination" (BBPD). In BBPD, the observable state of nature includes signals of customers' past actions or behaviors that provide some information about their preferences. Firms then use this information to provide targeted discounts to capture rivals' customers while increasing prices for their own captive customers. This type of mechanism could be best defined as dynamic price adjustment, using the framework from the previous section.

Villas-Boas (1999) provides a canonical investigation of BBPD. In the Villas-Boas (1999) model, firms can observe whether a customer has shopped with the firm before or whether she is a new customer. Different prices can then be offered to each type of customer depending on her purchase history. Like many models in the behavior-based price discrimination literature, this type of pricing in a duopolistic setting leads to intense competition, as firms compete for their rival's customers by lowering prices. Patient consumers can then benefit from lower prices as firms are pitted against each other. This outcome generally agrees with the Thisse and Vives (1988) result.

Shaffer and Zhang (2002) also consider behavior-based price discrimination when customers can be perfectly identified. In the Shaffer and Zhang (2002) model, firms are able to offer targeted coupons to specific customers. They show that firms will choose to offer the coupons in such a way to incentivize their competitor's customers to switch to their firm. As a result, as in Thisse and Vives (1988) and Villas-Boas (1999), price discrimination will lead to a race to the bottom, as extensive discounting leads to lower revenues for the firms (assuming that no new customers buy as a result of the lower prices). Choudhary et al. (2005) find that personalized pricing can hurt firms when they can also decide to change qualities of products along with prices. Specifically, high-quality firms are hurt by the ability of low-quality firms to price discriminate, causing both firms to lower their qualities and prices.

However, a series of papers have shown that personalized or customized pricing need not lead to a race to the bottom in all situations. Chen et al. (2001) examine a scenario in which individualized pricing is possible but imperfect. That is, firms will sometimes make mistakes when identifying their loyal customers and customers of the other firm. Counterintuitively, this imperfect targetability actually can help *improve* revenues as it provides a counterincentive against reducing prices. Chen et al. (2001) find that the race to the bottom result is a function of high targetability, and that with low levels of targetability, increases in revenues for both firms are possible. Chen and Iyer (2002) find a similar result when identification is perfect but addressability (the ability to apply a discount) is limited.

The counterintuitive finding that less-accurate targetability can actually improve the performance of price discrimination strategies in competition has continued to be verified by researchers in theoretical environments. Esteves (2014) considers a situation when customer brand preferences are private, but customers provide a noisy signal of their preferences to firms. Esteves (2014) finds that as the information becomes *more* accurate, industry revenues fall and approach the Thisse and Vives (1988) result. In these models, "loyal" customers (whose brand preferences are closer to that of the firm) typically pay a higher price; as more information becomes commonly available about who is loyal and who is not, other firms can target these customers with discounts in an effort to poach them away. This leads to firms partially losing their ability to price discriminate for customers that prefer their products.

Shy and Stenbacka (2015) also find that a moderate amount of customer information and targetability leads to better revenues than high or low levels. In their models, firms are allowed to collect information about their customers, and can either keep it to themselves (weak privacy) or can be mandated to share it with other firms (no privacy). In contrast, firms may be banned from collecting information on preferences, but can still keep a database about who has bought from them before (strong privacy). Shy and Stenbacka (2015) find that the weak privacy case leads to the best outcomes for the firms. This reinforces the result that full and open information about each customer's preference may lead to lower overall revenues for the firms. Their analysis agrees with an earlier result by Esteves (2009), who also finds that partial price discrimination can lead to higher revenues than perfect or no price discrimination.

In sum, the economics literature gives us two counterintuitive results about customization and dynamic pricing. First, we see that even when products are differentiated, personalized or customized pricing in competition may lead to lower revenues for the firms, as competition is intensified for each customer (Thisse and Vives, 1988; Shaffer and Zhang, 2002; Villas-Boas, 1999; Fudenberg and Villas-Boas, 2006). This is the case when customers can be perfectly identified and targeted.

However, we see that if the targetability of the firms is not precise, revenues can actually increase with personalized pricing, because the uncertainty removes some incentives for firms to discount to avoid accidentally charging low prices to loyal or high-value customers (Chen et al., 2001; Chen and Iyer, 2002; Esteves, 2009, 2014; Shy and Stenbacka, 2015). These two counterintuitive results highlight the need to test any dynamic pricing strategy or heuristic in a competitive environment, and to carefully consider the role that perfect or imperfect information plays in the competitive equilibrium from targeted pricing. We consider these effects in our later discussion of competitive implications of dynamic price adjustment in the airline industry in Chapter 6.

While the economics literature benefits from an intuitive approach to dynamic pricing that produces easily-understandable results, it also possesses some deficiencies. Most of the models discussed in this section do not consider competition with capacity constraints, as is the case in the airline industry. Capacity-constrained dynamic pricing problems are discussed in more detail in the operations research section of this literature review.

Furthermore, the economics literature suffers at times from over-simplification. For instance, the Hotelling line is a toy model of the world that produces interesting insights, but changes to the model's assumptions can produce different outcomes. There are few economics papers in this area that specifically consider the airline scenario (e.g. McAfee and te Velde (2007)), and these papers typically do not consider airline revenue management. Finally, economics papers are typically concerned with assessing overall market conditions, not creating actionable heuristics or strategies for individual firms. In contrast, the operations research literature discussed in the next section is more focused on evaluating specific strategies, but the evaluation of market outcomes is often less developed than in the economics literature.

2.3 Operations research literature: Dynamic pricing and assortment optimization

Dynamically pricing products with limited inventories has been a staple of the operations research literature for decades, dating back to an early evaluation of the problem by Kincaid and Darling (1963). Since then, researchers have considered increasingly complex theoretical formulations of pricing selection mechanisms and have proven many general properties. In recent years, the focus has turned towards *customized* pricing and assortment optimization. In a recent review of literature, Chen and Chen (2015) opined that "dynamic personalized pricing over multiple time periods and with [an] inventory consideration is an exciting future research topic for the [operations research] community" (p. 726).

However, the operations research literature on dynamic pricing has remained highly disjoint from the economics literature. Unlike the economics literature, the operations research (OR) literature typically looks at single-firm models, with only select papers considering competition. The OR literature also remains highly theoretical, and putting into practice the results and heuristics can be difficult, even for those papers designed specifically for the airline revenue management problem. It can thus be difficult to compare the performance of advanced pricing mechanisms to tried-and-true, simpler approaches. In this section, we discuss some of the classic and emerging OR literature surrounding dynamic pricing and customized assortment optimization and pricing.

2.3.1 Single-firm dynamic pricing

The main difference between the economics literature and the OR literature is that the latter typically focuses on *capacity-constrained* pricing problems. That is, a seller possesses a finite quantity of a resource which may also have a limited selling season. The seller's goal is to set a price in each period to maximize his total revenue over the entire selling season. The most classic example of a model of this type is described in a paper by Gallego and van Ryzin (1994). The authors model arriving demand as a Poisson process as a function of the price charged. The firm aims to find the optimal pricing path to maximize its revenue over the selling period, constrained by the quantity of the resource available for sale. This model is a clear example of a continuous pricing mechanism, since there is no pre-defined set of possible price points.

Gallego and van Ryzin solve this problem using the classic OR technique of constrained optimization, and are able to find optimal solutions for a certain class of exponential demand functions. For more general demand functions, they provide a deterministic relaxation of the problem. Their results show that the optimal prices decrease as the remaining stock of the good increases and the time remaining in the selling season decreases. That is, products with large inventories or whose selling seasons are about to end will have comparatively lower prices. Using a similar methodology, in Gallego and van Ryzin (1997) the authors expand their work to a network yield management problem, in which each product uses multiple resources (e.g., multiple legs in a flight network).

These two papers have served as the basis for many extensions, mostly focusing on situations when the demand function is more complex or customers behave in more interesting ways (Elmaghraby and Keskinocak, 2003). Bitran and Mondschein (1997) and Zhao and Zheng (2000) consider optimal dynamic pricing strategies when the consumer demand is Poisson but nonhomogeneous—when the customers' valuation of the products change over time. As in the Gallego and van Ryzin papers, Zhao and Zheng (2000) show that prices generally decrease when the stock of available units increases or the time in the selling season decreases. Walczak and Brumelle (2003) complete a similar analysis with a semi-Markov demand model, and Su (2007) and Levina et al. (2009) examine situations when customers are strategic and may decide to wait for a future period in the hope that a product will be discounted.

2.3.2 Personalized dynamic pricing and assortment optimization

While these extensions to the classic Gallego and van Ryzin (1994) model consider more complex and varied models of customer demand, they typically do not assume that customers can be charged customized or personalized prices. That is, the papers do not model the economic concept of third-degree price discrimination.

Aydın and Ziya (2009) provide one of the first operations research applications of price discrimination in dynamic pricing for a problem with limited inventories. In their model, customers belong to one of two types: high-value (H) and low value (L). Customers provide a noisy signal of their type to the firm, which selects its price as a function of its interpretation of this signal and the remaining capacity and time remaining in the selling horizon. The Aydın and Ziya (2009) model considers only two customer types in a monopolistic environment, and does not model the effect of competition.

Following the work by Aydın and Ziya (2009), a number of papers emerged in a parallel stream of literature examining the customized assortment optimization problem. The addition of customization to assortment optimization came after many increasingly complex papers considering choice-based, capacity-constrained assortment optimization (Mahajan and van Ryzin, 2001; Rusmevichientong et al., 2010).

Golrezaei et al. (2014) proposed one of the first models of personalized assortment optimization. Similar to Aydın and Ziya (2009), in the Golrezaei et al. (2014) model customers make a signal to the firm that contains some (potentially noisy) information about that customer's type. The firms can then use this information to customize their selection of products or price points to that specific customer. Their model allows for stockouts, but assumes a fixed capacity of each product or price point that is exogenously set in advance (which is not the case in the airline revenue management environment).

Recent papers have advanced the work on the customized assortment optimization problem with customer type as part of the observable state of nature. Bernstein et al. (2015) consider customer-specific assortment optimization, but the products in their model are of the same price. Chen et al. (2015a) consider a related pricing problem in which inventories are limited, but customer types do not vary. Abeliuk et al. (2016) also consider assortment optimization in contexts where customers are of various types. Chen et al. (2015b) expand the work of Golrezaei et al. (2014) to create an algorithm to construct decision policies based on past observations of customer arrivals and customer choice. They provide an airline example, but only for an ancillary product (priority seating) and not for a series of products in a fare structure as in the traditional airline revenue management problem.

Gallego et al. (2016) present a model of customized assortment optimization based on the choice-based revenue management frameworks of Talluri and van Ryzin (2004) and Gallego et al. (2004). Upon observing a customer's type, the airline uses an assortment selection algorithm to compute an assortment of products or price points as a function of the time and capacity remaining. As in Golrezaei et al. (2014), the Gallego et al. (2016) model generates probabilities that a customer is shown a particular assortment when arriving at a given time. However, the model requires a number of inputs that would be difficult to estimate in reality. They include the probability that a customer of a given type arrives at a certain time, and a detailed choice model as a function of each customer type, product, and assortment combination. The authors suggest that their model "could be very difficult to solve in practice" (Gallego et al. (2016), p. 19) and propose a number of heuristics to generate the assortment of products or price points.

Finally, Besbes and Sauré (2016) provide an exploration of a competitive model where firms compete in a duopoly by selecting products and prices simultaneously. They find Nash equilibrium outcomes in the cases when the sets of products or price points from the two firms do not overlap. Differences in customer choice functions by type are considered. The paper is one of the first to look at competition in assortment optimization, but it is not clear how their theoretical results would extend to the airline RM problem. Other recent work has begun to explore the simultaneous product and price selection problem in other non-airline contexts (Federgruen and Hu, 2015; Li et al., 2015; Alptekinoğlu and Semple, 2016; Jagabathula and Rusmevichientong, 2016).

2.3.3 Dynamic pricing in competition

In contrast to the economics literature, few papers in the operations research literature consider the effect of dynamic pricing strategies in competition. There are, however, some exceptions. Netessine and Shumsky (2005) create revenue management games in competition and assess equilibria, but these papers are static in nature and do not consider pricing decisions. Martínez-de-Albéniz and Talluri (2011) examine an oligopolistic scenario with dynamic pricing and stochastic demand. They find that when demand is uncertain, the race to the bottom outcome occurs as an equilibrium, wherein the seller with the lower reservation price (or marginal cost) will sell at a price just below the reservation price of the next lowest-cost firm. This agrees with an earlier result from Isler and Imhof (2008), who find the race-to-the-bottom in a one shot game with homogeneous products and perfect information. Unlike the behavior-based price discrimination literature, however, these papers do not consider customized pricing for individual passengers, nor do they consider differentiated products.

Gallego and Hu (2014) consider dynamic pricing of differentiated products with more general choice models, and reduce the problem to the solution of shadow prices in the current time. Lin and Sibdari (2009) also find a Nash equilibrium in a dynamic pricing game with capacity constraints and demand specified by the multinomial logit (MNL) choice function. Finally, Grauberger and Kimms (2016) construct games of simultaneous airline price and quantity competition without the Poisson demand assumption used by Martínez-de-Albéniz and Talluri (2011) and Gallego and Hu (2014). Their game assumes complete information and independent fare class demand.

Unlike past results, Grauberger and Kimms (2016) find that price competition can lead to higher prices and revenues than quantity competition alone. However, the games considered in this paper are one-shot games using linear demand functions, and the performance of dynamic pricing is not compared to traditional revenue management systems.

2.3.4 Comparing dynamic pricing to airline revenue management methods

Despite the volume of papers considering various flavors of dynamic pricing, very few papers aim to answer a single important question: how do the outcomes of dynamic pricing compare to traditional airline revenue management, both when practiced by a single airline and in a competitive scenario? As the literature suggests, the performance of a pricing mechanism is highly dependent on the context in which it is practiced and the amount of information available to the airlines about their customers. Filling this gap in the literature is one of the aims of this dissertation.

Among the few papers considering the performance of dynamic pricing relative to other revenue management methods is Zhang and Lu (2013). They construct a version of the nonlinear programming formulation of dynamic pricing for network RM from Gallego and van Ryzin (1997) and solve it using a dynamic programming decomposition approach. They compare the performance of their dynamic pricing heuristic to a choice-based revenue management framework first introduced in Talluri and van Ryzin (2004) and expanded in Liu and van Ryzin (2008).

Zhang and Lu (2013) test both of the methods in a simple hub-and-spoke network, with two airports located on either side of the hub and demand flowing west-to-east across the network. They find that their continuous pricing mechanism results in a revenue gain of 1–6% over static pricing methods or the choice-based availability control of Liu and van Ryzin (2008). While this paper is one of the few to consider the performance of dynamic pricing in an airline network context, it reports only the revenue performance of the methods, and other revenue management outcomes that are important to airlines, such as load factors, booking patterns, and fare class mixes are not reported.

Fiig et al. (2016) compare the performance of a dynamic price adjustment algorithm against traditional airline revenue management, and find revenue gains for the dynamic pricing algorithm in a competitive environment. Their rules-based algorithm controls for the current bid price calculated by the RM system and also allows for the introduction of customer-specific characteristics. In simulation results, the dynamic price adjustment algorithm results in gains of approximately 6% over a baseline case when used by a single airline, and approximately 3% when used by both airlines in the simulation.

Unlike the OR literature on dynamic pricing, Fiig et al. (2016) consider basic revenue management outcomes in their analysis, including load factors, changes in booked passengers, and fare class mix. They also show that use of the dynamic pricing algorithm represents a Nash equilibrium for airlines. However, the simulations of their heuristics assume that the airlines have perfect knowledge of the underlying parameters that drive passenger choice functions. The baseline RM method used in the simulations is also a very simple "threshold" type method that does not reflect state-of-the-art practices in the airline industry.

Finally, in a series of papers, Wittman and Belobaba (2017a,b,c) develop and evaluate a series of dynamic price adjustment mechanisms for airline revenue management. Each of the heuristics presented in these papers assumes that airlines can identify customer characteristics with a certain degree of accuracy and make dynamic adjustments to the availability or prices of fare classes that would normally be output by an airline revenue management system. Like Fiig et al. (2016), these papers consider how some airline RM outcomes change as a result of the dynamic price adjustment heuristics in their simulations. This work serves as a backbone for this dissertation, which motivates and tests a more complex and theoretically justified version of the dynamic pricing heuristics proposed in Wittman and Belobaba (2017c).

2.4 Implications for research

This section has reviewed a number of papers that have examined price selection mechanisms for given products in non-competitive and competitive environments.

In the economics literature, price selection mechanisms are typically based on monopolistic pricing models from the field of industrial organization. Firms that are able to practice first- or third-degree price discrimination can be shown to have higher revenues than firms that are only able to charge a uniform price. However, the literature on behavior-based price discrimination (BBPD) shows that in competition, price discrimination can become much more complicated. Work by Thisse and Vives (1988) and others have shown that personalized pricing can lead to lower revenues in some situations, and Chen et al. (2001) and Esteves (2014), among others, have shown the counterintuitive result that less-accurate targeting may actually improve firm revenues relative to perfect targeting. In Chapter 6, we will also discuss how the availability of information about firms' pricing actions can affect equilibrium outcomes of the repeated dynamic pricing game.

In the operations research literature, the focus of pricing-related papers is usually to examine single-firm examples when capacities are constrained. Many papers focus on evaluating the theoretical properties of models and proving their optimality, often generating heuristics for use when the optimal solution is too computationally intensive to calculate. Indeed, just because a pricing mechanism is optimal does not mean it will be practical; most of the common revenue management techniques used today by airlines are provably non-optimal heuristics (Belobaba, 2016). Few operations research papers consider the market effects of price selection mechanisms in competition, and fewer still compare the performance of various pricing mechanisms to other well-accepted revenue management techniques.

The literature review yields several important implications as we develop the dynamic price adjustment method used in this dissertation. From the economics literature, we learned that competitive environments can yield very different performance of price discrimination methods than non-competitive environments. Therefore, any advanced pricing mechanism must be tested amongst competitors who are (or are not) using those advanced pricing methods themselves. Furthermore, the lack of papers comparing the outcomes of next-generation pricing mechanisms to existing RM methods in language that is familiar to revenue management practitioners represents a significant gap in the literature. This dissertation aims to assess new pricing mechanisms in competitive environments, with a focus on the outcomes of these methods compared to other common RM approaches.

In the next chapter, we begin our development of the Probabilistic Fare-Based Dynamic Adjustment (PFDynA) heuristics that are then simulated and tested in Chapter 4. We start with the development of a customer choice model to guide the pricing decisions that will be made by the heuristic.

3 Models for Customer Choice and Dynamic Price Adjustment

"Man is a deterministic device thrown into a probabilistic Universe. In this match, surprises are expected."
–Amos Tversky (Lewis, 2016)

A customer wishes to make a purchase among a set of alternatives. Each alternative possesses various attributes, including price and quality, which the customer evaluates according to her idiosyncratic tastes. Which alternative, if any, does the customer choose to purchase? Moreover, from the perspective of the firm, which alternatives should be displayed to the customer, and what prices should be set for those alternatives?

The best answer to the questions posed to the firm depends upon how customers are assumed to make choices among alternatives. Once a choice model is assumed and parameterized, a firm can use it to select the prices that maximize expected revenue from each customer.

Countless papers have been written about price selection mechanisms in a variety of environments. These papers typically describe approaches for pricing products as a function of assumed, observable, or inferable information about the way customers make choices. The airline pricing problem, however, poses several additional challenges:

- First, the alternatives that can be made available for sale are limited by the number of seats on each flight in the airline's network.
- Second, the alternatives have limited selling horizons, making departed unsold seats worthless to both the firm and the customer.
- Third, characteristics of air transportation markets often mean that choice behaviors change over the selling period, as early-arriving, lower-budget discretionary travelers are replaced by higher-budget, non-discretionary travelers that book later in the booking process. This runs contrary to common optimization approaches that assign lower prices to products near the end of their selling seasons (Gallego and van Ryzin, 1994).
- Finally, in airline networks with connecting hubs, each alternative may involve the consumption of units (seats) of two or more scarce resources (legs) in the airline's network.

In this chapter, we propose a model of customer choice and dynamic pricing designed specifically for the airline revenue management problem. After developing a generic set of assumptions regarding how customers make decisions amongst alternatives, we show how airlines could dynamically adjust the prices of fare products to maximize expected revenues from each customer, given available information. We also discuss how these heuristics can be integrated with the RM methods used by airlines today. The dynamic price adjustment methods developed in this chapter are tested in simulated environments in Chapter 4.

3.1 A brief introduction: Rational choice theory and discrete choice analysis

The customer choice model presented in this section is designed in the spirit of rational choice theory (RCT) (Elster, 1986; Allingham, 1999; Levin and Milgrom, 2004). RCT is a classic framework in the field of economics that describes how people make decisions amongst alternatives. Largely inspired by the field of logic, RCT underpins many canonical economic models in which rational economic agents make decisions (Becker, 1976).

Rational agents in RCT are endowed with a *preference relation* $\succ$, which allows the agents to rank-order alternatives in order of attractiveness. The preference relation then serves as a *reason* for a *choice rule* $C(S)$, which describes how agents select one or more alternatives out of a set S (Allingham, 1999). RCT assumes that choice is deterministic—once a preference relation has been specified, the agent will select the alternative that he prefers the most.

Preference relations are often assumed to follow a number of logical axioms (Samuelson, 1938; Allingham, 1999; Levin and Milgrom, 2004). Some common axioms include:

- **Completeness:** A preference relation $\succ$ is complete if for any alternatives x and y either $x \succ y$ or $y \succ x$ or $x \sim y$ (x and y are equivalent).
- **Transitivity:** A preference relation $\succ$ is transitive if $x \succ y$ and $y \succ z$ implies that $x \succ z$.
- **Expansion:** Suppose $x \in S$. If $x \in C(\{x, y\})$ for all $y \in S$, then $x \in C(S)$.
- **Contraction:** If $x \in S \subset T$ and $x \in C(T)$, then $x \in C(S)$.
- **Sen's Extension Axiom:** If $x, y \in C(S)$ and $S \subset T$, then if $x \in C(T), y \in C(T)$.
- **Weak Axiom of Revealed Preference:** If $x, y \in S$ and $x \in C(S)$ then if $x, y \in T$ and $y \in C(T)$, $x \in C(T)$.

In particular, if a preference relation $\succ$ follows the transitive axiom, it can be used to construct a real-valued *utility function* $u(\cdot; \succ)$, which assigns a numerical value to each alternative (Allingham, 1999; Levin and Milgrom, 2004). Utility functions can also serve as a reason for choice rules—for instance, the option with the highest utility value could be selected out of a set.

RCT is not without its critics, particularly among psychologists who have observed that people do not always make choices according to rational rules. The seminal work of Daniel Kahneman and Amos Tversky, for which Kahneman won the Nobel Prize in Economics in 2002, found that the choices of participants could be influenced by the framing of the alternatives and their relationship to a reference state (Tversky and Kahneman, 1981; Lewis, 2016). Developments in *bounded rationality* suggest that people possess limited ability to fully parameterize and inform their preference relations, and may "satisfice" by choosing sub-optimal outcomes to avoid costly computation or search (March, 1978; Ariely, 2010).

RCT also differs from the field of discrete choice theory, which assumes that choice is not deterministic but instead probabilistic (Ben-Akiva and Lerman, 1985; Anderson et al., 1992). That is, given the same set of stimuli, the decision maker may choose different decisions on different occasions due to "consumer optimization errors" or attribute measurement errors on the part of the researcher or the customer (Ben-Akiva and Lerman, 1985). Common discrete choice models such as the multinomial logit (MNL) possess this central assumption of probabilistic choice.

3.2 The choice function and preference relation

We begin by introducing some definitions of concepts that will be useful as we expand our choice framework. First, let $\mathcal{L} = \{1, \ldots, l, \ldots, L\}$ represent the flight legs in an airline's network, and $\mathcal{C} = \{c_1, \ldots, c_l, \ldots, c_L\}$ represent the capacities of those flight legs. We then construct a set $\mathcal{F}$ of origin-destination *itinerary products*[3] utilizing these flight legs.

Definition (Itinerary Product): An ***itinerary product*** $x \in \mathcal{F}$ consists of:

- A vector of one or more flight legs $l \in \mathcal{L}$ utilized for an origin-destination journey.
- A *price* $f_x \in \mathbb{R}^+$
- A vector of *attributes* $\mathbf{a}_x$. These attributes could include characteristics of the *itinerary*, including elapsed time, scheduled departure time, number of connections, or characteristics of the *fare product*, including fare restrictions or ancillary fees.[a]

[a]We assume that each attribute can be represented *cardinally*, as opposed to ordinally; for instance, the number of elapsed minutes for an itinerary, or the presence (=1) or not (=0) of a fare restriction. In other words, "the attractiveness of an alternative expressed by a vector of attribute values is reducible to a scalar." (Ben-Akiva and Lerman (1985), p.37)

Customers desire to travel in a given origin-destination market. In response to a search request for air transportation, airlines display subsets of available itinerary products to customers. We refer to a nonempty set $S \subseteq \mathcal{F}$ of itinerary products as an *assortment.*

Among the itinerary products available in an assortment S, a customer evaluates each option and chooses which product, if any, to purchase. The next definition formalizes that action.

[3]Note that there are typically multiple *itinerary products* associated with a given airline *itinerary.* For instance, a nonstop itinerary on a nonstop flight from Boston to Paris could be sold in a number of different fare classes (for instance, classes Y, M, B, and Q), each of which could be associated with a set of fare restrictions and a unique price. In this example, each of classes Y, M, B, and Q for that non-stop itinerary from Boston to Paris would be a unique itinerary product x with a unique price f_x and a set of attributes $\mathbf{a}_x$.

Definition (Choice Function): The ***choice function*** $C(S)$, $\mathcal{F} \to \mathcal{F}$ returns a subset of itinerary products in assortment S that are chosen for purchase by the customer.

The choice function takes as an argument the assortment of itinerary products offered to the customer, and returns a subset (which could be the empty set $\emptyset$) of itinerary products which are chosen for purchase. For instance, if $S = \{x, y, z\}$, $C(S) = \{x\}$ indicates that the customer has chosen to purchase itinerary product x out of the options x, y, and z.[4]

This definition of the choice function does not contain any information about *how* customers make decisions among available options. That is, what causes a customer to select x out of the options x, y, and z? To model this choice itself, we need one additional concept that describes how customers prefer one alternative to another: the ***preference relation*** $\succ$.

Definition (Preference Relation): Given itinerary products $x, y \in \mathcal{F}$, we say that $x \succ y$ if and only if $x \in C(S)$ and $y \notin C(S)$ for all $S \subseteq \mathcal{F}$ such that $x, y \in S$.

A customer prefers x to y ($x \succ y$) when, given any assortment S containing both x and y, the customer always decides to purchase x and not y. Note that this preference relation is generated by evaluating all of the components of each itinerary product, including its price, schedule, and product attributes.

In what situations will $x \succ y$? The answer depends on the customer's idiosyncratic evaluation of the prices and attributes of the two itinerary products. For instance, suppose that itinerary products x and y are identical in all ways except price. That is, $\mathbf{a}_x = \mathbf{a}_y$. Given these otherwise identical itinerary options, we assume that the customer will prefer the one with the lower price:

$$\text{If } \mathbf{a}_x = \mathbf{a}_y, x \succ y \text{ if and only if } f_x < f_y$$

$$\text{If } f_x = f_y, C(S) = \begin{cases} \{x\} & \text{with prob. } \frac{1}{2} \\ \{y\} & \text{with prob. } \frac{1}{2} \end{cases}$$

In most cases, however, there will be differences among the attributes of the two itinerary products. Perhaps one itinerary product has a lower price, but departs early in the morning, or that one itinerary product requires a connection compared to another that is non-stop. In these cases, we assume that customers will consider not only the price of an itinerary, but also their relative valuations of each attribute when assigning preferences to flights. Combining the price and attribute valuations together yields the customer's *perceived price* for the itinerary product.

[4]Note that we do not explicitly limit $C(S)$ to contain no more than one element. While in the airline itinerary choice example, it is likely that a customer will select only a single itinerary product (that is, $x \in C(S)$ if and only if $C(S) = \{x\}$), the choice framework discussed here also holds in more general cases.

Definition (Perceived Price):	A customer's ***perceived price*** for itinerary product x is equal to the quantity $f_x + \boldsymbol{\beta a}_x$.

Here, $\boldsymbol{\beta}$ is a vector of preferences in willingness-to-pay space for each itinerary attribute. For instance, suppose a customer is willing to pay \$150 more for a non-stop itinerary that does not require a connection. This \$150 value would then represent a *disutility* for accepting a flight with a connection. Suppose the price of the non-stop itinerary is \$200 and the price of the connecting itinerary is \$150; then, the perceived price of the non-stop itinerary would be \$200 and the perceived price of the connecting itinerary would be \$300.

3.3 The budget constraint: maximum and conditional willingness-to-pay

3.3.1 Maximum willingness-to-pay

The concept of perceived price allows us to numerically rank and compare itinerary products. However, customers often decide not to purchase any of the options presented to them, either because they are too expensive or because they possess undesirable attributes. For instance, a business customer who needs to arrive in time for an 11am meeting will never purchase an itinerary product for a two-hour flight that departs at 6pm on the same day.

Consider the following scenarios:

Scenario 1:	$S = \{x\}, f_x = \$200$ $C(S) = \{x\}$
Scenario 2:	$S = \{x\}, f_x = \$201$ $C(S) = \emptyset$

In Scenario 1, the firm charges \$200 for itinerary product x, and the customer, facing the choice of purchasing this alternative at \$200 or purchasing nothing, decides to purchase x. In Scenario 2, the same alternative x is offered to the customer, but at a price of \$201. In this scenario, the customer decides not to purchase anything and leaves empty handed.

If prices were set only in whole-dollar integer values, there would clearly be something special about the price \$200. This is the maximum price at which the customer is willing to purchase the product, given there are no other alternatives available in the choice set S. We refer to this value $\theta_x = \$200$ as the customer's *maximum willingness to pay (WTP)* for product x.

Definition (Maximum WTP): Itinerary product x has a **maximum WTP** $\theta_x \in \mathbb{R}^+$ if and only if when $S = \{x\}, x \in C(S)$ if $f_x = \theta_x$ and $x \notin C(S)$ if $f_x = \theta_x + \epsilon\ \forall \epsilon > 0$.

The maximum WTP for a given itinerary product can and will vary from customer to customer as a function of each customer's idiosyncratic tastes. For instance, from our earlier example, an itinerary product x that departs at 6pm may be worthless to a business traveler that needs to arrive at 11am that day. For that customer, θ_x could equal \$0. Yet for another leisure customer returning from a vacation, the 6pm departure may perfectly meet her needs, and θ_x could equal a higher value.

If we aggregate the maximum WTP values across all customers for a given itinerary product x, we can construct a probability distribution Θ_x of maximum WTPs for that itinerary product. Then, for a random customer arriving to book, the maximum WTP θ_x for that itinerary product can be assumed to be drawn from the distribution Θ_x. This probability distribution Θ_x of maximum WTP will become important later when we determine which prices to charge for each itinerary product.

3.3.2 Affordability and preference rules

We say that an itinerary product is *affordable* for a customer if the price does not exceed that customer's maximum WTP for that product.

Definition (Affordability): Itinerary product x is *affordable* if $f_x \leq \theta_x$.

When the customer has no affordable options in S, she will choose to no-go and will not purchase any product. Therefore, the customer's choice from any assortment S is equivalent to her choice from a potentially smaller choice set S' that consists only of affordable products in S. These statements are expressed mathematically in Remarks 1 and 2.

Remark 1: If $f_x > \theta_x$ for all $x \in S$, then $C(S) = \emptyset$.
Remark 2: Let $S' = \{x \in S : \theta_x \geq f_x\}$. Then $C(S') = C(S)$.

With the concept of affordability, we can construct some rules for the preference relation $\succ$:

Preference Rule 1: If x and y are both affordable,
$x \succ y$ if and only if $f_x + \boldsymbol{\beta a}_x < f_y + \boldsymbol{\beta a}_y$.

An additional preference rule is also necessary for the case when one or both of the two products are unaffordable.

Preference Rule 2: Let $S' = \{x \in S : \theta_x \geq f_x\}$. If $x \in S'$ and $y \notin S'$, then $x \succ y$. If $x \notin S'$ and $y \notin S'$, then $x \sim y$ (x and y are equivalent).

Finally, given a preference relation $\succ$ and maximum WTPs for each itinerary product, we can construct the choice function for any set S:

Preference Rule 3: $x \in C(S)$ if and only if $x \in S' = \{x \in S : f_x \leq \theta_x\}$ and $x \succ y$ for all $y \neq x \in S'$

That is, a customer chooses alternative x out of S if x is the most-preferred affordable itinerary in S. Furthermore, if $x \succ y$ for all other affordable itinerary products $y \in S$, it follows from Preference Rule 1 that $f_x + \boldsymbol{\beta a}_x < f_y + \boldsymbol{\beta a}_y$ for any y. In other words, customers choose to purchase the itinerary with the lowest perceived price among the affordable alternatives presented to them.

From Preference Rules 1, 2, and 3, it is also possible to show that the preference relation $\succ$ is transitive. As a result, since $\succ$ is a reason for choice rule $C(S)$, the choice rule is rational (Allingham, 1999).

Proposition 1 (Transitivity): Let $x, y, z \in S$. If $x \succ y$ and $y \succ z$, then $x \succ z$.

Proof: Since $x \succ y$ and $y \succ z$, both x and y are affordable. If z is not affordable, then $z \notin S' = \{x \in S : \theta_x \geq f_x\}$. Therefore, by Preference Rule 2, since x is affordable, it must be the case that $x \succ z$.

If z is also affordable, then by Preference Rule 1, we have $f_x + \boldsymbol{\beta a}_x < f_y + \boldsymbol{\beta a}_y$ and $f_y + \boldsymbol{\beta a}_y < f_z + \boldsymbol{\beta a}_z$. Therefore, $f_x + \boldsymbol{\beta a}_x < f_z + \boldsymbol{\beta a}_z$, and hence $x \succ z$ by Preference Rule 1. □

3.3.3 Conditional willingness-to-pay

When more than one alternative is available in the choice set, the willingness-to-pay of a customer for an itinerary product may change. Consider the following two scenarios:

Scenario 3:	$S = \{x\}$
Maximum WTP	$\theta_x = \$200$
Price	$f_x = \$200$
Choice	$C(S) = \{x\}$
Scenario 4:	$S = \{x, y\}$
Maximum WTP	$\theta_x = \$200, \theta_y = \$200$
Prices	$f_x = \$200, f_y = \$150$
Attributes	$\mathbf{a_x} = \mathbf{a_y}$
Choice	$C(S) = \{y\}$

In Scenario 3, the assortment is $S = \{x\}$ and the customer's maximum WTP for itinerary product x is \$200. When the price of that product $f_x = \$200 \leq \theta_x$, the customer decides to purchase x. This fits our earlier definitions of affordability and maximum WTP.

However, in Scenario 4, we introduce an additional alternative y to the assortment. Suppose that this alternative is identical to x in all ways except price (that is, $\mathbf{a}_x = \mathbf{a}_y$ and hence $\theta_x = \theta_y$, but $f_y < f_x$). Now, however, even though $f_x \leq \theta_x$, the customer does not decide to purchase x. This is because, according to Preference Rule 1, $y \succ x$ since $f_y < f_x$.

The introduction of an additional alternative into the assortment has effectively reduced the customer's willingness-to-pay for itinerary product x. The customer is no longer willing to purchase alternative x at a price $f_x = \theta_x$. Indeed, note that the customer will not purchase itinerary x at any price above $\$150(= f_y)$. To describe this situation, we define an additional WTP concept called *conditional WTP*.

Definition (Conditional WTP): Itinerary product x has a **conditional WTP** $\widetilde{\theta}_{x|S}$ in assortment S if and only if $x \in C(S)$ if $f_x = \widetilde{\theta}_{x|S}$ and $x \notin C(S)$ if $f_x = \widetilde{\theta}_{x|S} + \epsilon \ \forall \epsilon > 0$.

Unlike maximum WTP θ_x, which is a function of only the price and attributes of itinerary product x, conditional WTP $\widetilde{\theta}_{x|S}$ is a function of both the itinerary product and the assortment S in which it is presented. Conditional WTP is thus a function of both the maximum WTP for itinerary product x and the prices and attributes of other alternatives in S.

If only one alternative is available in the assortment, the conditional WTP is equal to the maximum WTP, since there are no other products in the assortment from which to choose. Furthermore, the conditional WTP for an itinerary product is never greater than its maximum WTP, and may often be less, as shown in Scenario 4. These properties are formulated in Remarks 3 and 4:

Remark 3:	If $S = \{x\}$, $\theta_x = \widetilde{\theta}_{x\|S}$
Remark 4:	$\widetilde{\theta}_{x\|S} \leq \theta_x$ for any assortment S for which $x \in S$.

It need not be the case that adding an additional product into assortment S will change the conditional WTP of an itinerary product. The introduction of an irrelevant or unaffordable alternative may not change conditional WTP, as shown in Scenario 5.

<u>Scenario 5:</u>	$S = \{x, y\}$
Maximum WTP	$\theta_x = \$200, \theta_y = \$0$
Prices	$f_x = \$200, f_y = \$150$
Choice	$C(S) = \{x\}$

In Scenario 5, an irrelevant alternative y has been introduced into the assortment, but the customer still selects to purchase x. In this case, conditional WTP $\widetilde{\theta}_{x|S} = \$200 = \theta_x$.

Using Preference Rule 3, we are able to define conditional WTP $\widetilde{\theta}_{x|S}$ precisely in terms of perceived price for any itinerary product in S:

Remark 5:	Suppose $C(S) = \{x\}$. Let $S'' = S \setminus \{x\}$. Then if $C(S'') = \{z\}, \widetilde{\theta}_{x\|S} = \min(\theta_x, f_z + \boldsymbol{\beta}(\mathbf{a}_z - \mathbf{a}_x))$ and if $C(S'') = \emptyset, \widetilde{\theta}_{x\|S} = \theta_x$
Remark 6:	Suppose $C(S) = \{x\}$. Then for any $y \in S$, $y \neq x$, $\widetilde{\theta}_{y\|S} = \min(\theta_y, f_x + \boldsymbol{\beta}(\mathbf{a}_x - \mathbf{a}_y))$
Remark 7:	If $C(S) = \emptyset$, $\widetilde{\theta}_{x\|S} = \theta_x$ for all $x \in S$.

Remarks 5, 6, and 7 provide an expression for conditional WTP for each itinerary product in the assortment S:

- *(Remark 5)* When x is chosen from S, then its conditional WTP $\widetilde{\theta}_{x|S} = f_z + \boldsymbol{\beta}(\mathbf{a}_z - \mathbf{a}_x)$, where z is the next-best affordable option (the option that would have been selected if x was not included in S).
- *(Remark 6)* When x is chosen from S, the conditional WTP for a product $y \in S, y \neq x$ that is not chosen is equal to $\widetilde{\theta}_{y|S} = f_x + \boldsymbol{\beta}(\mathbf{a}_x - \mathbf{a}_y)$.
- *(Remark 7)* When no item is chosen from S, all of the products are unaffordable. That means the prices of all products in S exceed the customer's maximum WTP for those products. Therefore, $\widetilde{\theta}_{x|S} = \theta_x$.

Scenario 6 presents an example that summarizes the concepts presented in this section.

Scenario 6:	$S = \{x, y\}$
Maximum WTP	$\theta_x = \$150, \theta_y = \$120$
Attribute disutilities	$\boldsymbol{\beta}\mathbf{a}_x = \$40, \boldsymbol{\beta}\mathbf{a}_y = \$70$
Prices	$f_x = \$100, f_y = \$80$
Choice	$C(S) = \{x\}$
Conditional WTP	$\widetilde{\theta}_{x\|S} = \$110, \widetilde{\theta}_{y\|S} = \$70$

Some observations about Scenario 6 follow:

- The customer selects itinerary x, because its perceived price is less than the perceived price of itinerary y:

$$f_x + \boldsymbol{\beta}\mathbf{a}_x = \$100 + \$40 < \$80 + \$70 = f_y + \boldsymbol{\beta}\mathbf{a}_y$$

  Equivalently, itinerary x is chosen since x and y are affordable and $x \succ y$.

- The conditional WTP of itinerary x is $\widetilde{\theta}_{x|S} = f_y + \boldsymbol{\beta}(\mathbf{a}_y - \mathbf{a}_x) = \$80 + \$30 = \$110$, and the conditional WTP of itinerary y is $\widetilde{\theta}_{y|S} = f_x + \boldsymbol{\beta}(\mathbf{a}_x - \mathbf{a}_y) = \$100 - \$30 = \$70$.

- Thus, the customer would still select itinerary x at any price up to \$110, and that the customer would switch to purchasing itinerary y if the price f_y were reduced to \$70 or less or the price of itinerary x exceeded \$110.

We close this section with a theorem that links together itinerary choice with the concept of conditional WTP.

Theorem 1: Under preference relation $\succ$, $\text{Prob}(x \in C(S)) = \text{Prob}(\widetilde{\theta}_{x|S} \geq f_x)$

Proof: We proceed by showing the statements $x \in C(S)$ and $\widetilde{\theta}_{x|S} \geq f_x$ are equivalent. To show that $x \in C(S)$ implies $\widetilde{\theta}_{x|S} \geq f_x$, suppose in order to obtain a contradiction that $x \in C(S)$ and $\widetilde{\theta}_{x|S} < f_x$. Note that if $S = \{x\}$, Remark 3 states that $\widetilde{\theta}_{x|S} = \theta_x$, so $\theta_x < f_x$. But this contradicts the definition of maximum WTP, where $x \notin C(S)$ if $f_x > \theta_x$.

Therefore, there exists at least one alternative $y \in S$ such that $y \neq x$. If y is unaffordable, then by Remark 2, $C(S) = C(S')$ where $S' = \{x\}$, and we obtain the same contradiction as in the previous paragraph.

If y is affordable, by Preference Rule 3, $x \succ y$ since $x \in C(S)$. Then, by Preference Rule 1, $f_x + \boldsymbol{\beta}\mathbf{a}_x < f_y + \boldsymbol{\beta}\mathbf{a}_y$. Therefore, $f_x < f_y + \boldsymbol{\beta}(\mathbf{a}_y - \mathbf{a}_x) = \widetilde{\theta}_{x|S}$. But this contradicts our assumption that $f_x > \widetilde{\theta}_{x|S}$. Hence, $x \in C(S)$ implies $\widetilde{\theta}_{x|S} \geq f_x$.

To show that $\widetilde{\theta}_{x|S} \geq f_x$ implies $x \in C(S)$, first suppose $\widetilde{\theta}_{x|S} < f_x$. From the definition of conditional WTP, $x \notin C(S)$ if $f_x > \widetilde{\theta}_{x|S}$. Furthermore, if $\widetilde{\theta}_{x|S} = f_x$, $x \in C(S)$, also from the definition of conditional WTP.

Now, suppose in order to obtain a contradiction that $\widetilde{\theta}_{x|S} > f_x$ and $x \notin C(S)$. If $C(S) \neq \emptyset$, then there exists another itinerary product $y \in S, y \neq x$ such that $y \in C(S)$. Then, y is affordable and $y \succ x$. Therefore, by Preference Rule 3, $f_y + \boldsymbol{\beta}\mathbf{a}_y < f_x + \boldsymbol{\beta}\mathbf{a}_x$ and hence $f_x > f_y + \boldsymbol{\beta}(\mathbf{a}_y - \mathbf{a}_x)$. From Remark 5, this means that $f_x > \widetilde{\theta}_{x|S}$. But this contradicts our earlier assumption that $f_x < \widetilde{\theta}_{x|S}$.

If $C(S) = \emptyset$, then $f_x > \theta_x$ from the definition of maximum WTP. But then we have $\theta_x < f_x < \widetilde{\theta}_{x|S}$, which contradicts Remark 4 that $\widetilde{\theta}_{x|S} \leq \theta_x$.

Thus, $x \in C(S)$ and $\widetilde{\theta}_{x|S} \geq f_x$ are equivalent statements. Therefore, we can conclude that $\text{Prob}(x \in C(S)) = \text{Prob}(\widetilde{\theta}_{x|S} \geq f_x)$. □

Theorem 1 suggests that to estimate the probability that a random customer will select a given alternative, we need to estimate the probability that his conditional WTP for that alternative given assortment S will exceed the price f_x. This task is no small feat, and the estimation of a customer's conditional or maximum WTP could be difficult or impossible in practice. In the following sections, we will describe some simplifying assumptions that could be used when the conditional WTP for each alternative-assortment pair cannot be identified exactly.

3.4 Finding the optimal fare

Suppose a monopolist offers an assortment with a single itinerary product $S = \{x\}$. What price should the monopolist charge for this product to maximize revenue from a customer arriving to book? From microeconomic principles of the theory of the firm, the monopolist faces the following expected revenue $\mathbb{E}(R)$ from the customer:

$$\mathbb{E}(R) = (f_x - MC_x) \cdot Prob(x \in C(S)) \tag{2}$$

In Equation (2), f_x represents the price charged for itinerary product x, MC_x is the marginal cost of itinerary product x (which will be discussed in detail shortly), and $\text{Prob}(x \in C(S))$ is the probability that the customer decides to purchase x as opposed to purchasing nothing.

From Theorem 1 and Remark 3, we can replace $\text{Prob}(x \in C(S))$ with $\text{Prob}(\theta_x \geq f_x)$ in the single-product case. That is, the customer will decide to purchase itinerary product x if the price does not exceed his maximum WTP; otherwise, he will no-go.

$$\mathbb{E}(R) = (f_x - MC_x) \cdot Prob(\theta_x \geq f_x) \tag{3}$$

If the maximum WTP distribution Θ_x is known, we can compute the optimal fare for the itinerary product by differentiating Equation (3) with respect to f_x and setting it equal to zero. The answer will depend on the assumed distribution of Θ_x.

3.4.1 Examples with uniformly-distributed maximum WTP

Suppose that $\Theta_x = U[\$100, \$200]$; that is, that the maximum WTP for product x is uniformly distributed between \$100 and \$200. In this case,

$$Prob(x \in C(S)) = Prob(\theta_x \geq f_x) = 1 - \Theta_x(f_x) = 1 - \frac{f_x - 100}{200 - 100}$$

Then the optimal price is:

$$\frac{\partial}{\partial f_x}(f_x - MC_x) \cdot \left(1 - \frac{f_x - 100}{100}\right) = 2 - \frac{2f_x}{100} + \frac{MC_x}{100} = 0$$
$$\Rightarrow f_x^* = 100 + \frac{MC_x}{2}$$

If the marginal cost of the itinerary product is \$0, the optimal fare is \$100; as the marginal cost increases, the optimal prices increases at a rate of half of the marginal cost.

Now, consider the situation when an additional product is introduced into the assortment S. Suppose that an itinerary product y is introduced into S that is identical to itinerary product x in all attributes except price ($\mathbf{a}_x = \mathbf{a}_y$ and $\theta_x = \theta_y$). In this case, the effects of this itinerary product on the conditional WTP $\widetilde{\theta}_{x|S}$ of product x depends on the price f_y of product y.

For instance, if $f_y = \$250$, y is unaffordable since $\theta_y < f_y$. This means the set of affordable choices is $S' = \{x\}$. Therefore, the conditional WTP $\widetilde{\theta}_{x|S} = \theta_x = \$200$, as in the previous example. However, if $f_y = \$150$, $y \succ x$ if $f_x > \$150$. Therefore, we need to consider the distribution of conditional WTP $\widetilde{\Theta}_{x|S}$ as a function of both f_y and the maximum WTP distribution Θ_x for product x.

- If $\theta_x \leq \$150$ (which occurs with probability $\frac{1}{2}$), then by Remark 5, $\widetilde{\theta}_{x|S} = \min(\theta_x, f_y)$. Since $f_y = \$150$, this means that $\widetilde{\theta}_{x|S} = \theta_x$.
- If $\theta_x > \$150$ (which occurs with probability $\frac{1}{2}$), by Remark 5, $\widetilde{\theta}_{x|S} = \min(\theta_x, f_y)$. Therefore, $\widetilde{\theta}_{x|S} = \$150$.

Hence, the distribution of conditional WTP $\widetilde{\Theta}_{x|S}$ for product x consists of a uniform distribution from \$100 to \$150, and a mass with probability 0.5 at \$150. The mean of this distribution is:

$$\mathbb{E}(\widetilde{\theta}_{x|S}) = 0.5 \cdot \frac{150 - 100}{2} + 0.5 \cdot 150 = \$137.50 < \$150 = \mathbb{E}(\Theta_x)$$

In this example, we have seen that the availability of an attractive alternative in the assortment can change the conditional WTP for product x. Specifically, the mean of the conditional WTP distribution decreases from \$150 in the single-product case to \$137.50 in the two-product case above.

It is worth noting that even though the conditional WTP distribution $\widetilde{\Theta}_{x|S}$ has changed with the introduction of product y in S, the optimal price f_x^* may not change:

- With $MC_x < \$100, f_x^* = \$100 + \frac{MC_x}{2}$. This is identical to the single-product case when $S = \{x\}$.
- With $MC_x > \$100, f_x^* = \$150$. This is in contrast to the single-product case, where $f_x^* = \$100 + \frac{MC_x}{2}$. For instance, with $MC_x = \$120$, f_x^* in the single-product case would be \$160, and f_x^* in the multi-product case would be \$150.

3.4.2 Examples with normally-distributed maximum WTP

It is perhaps too simplistic to represent passenger WTP using a uniform distribution. It is unlikely that passengers will be equally likely to value an itinerary at a high-price than a low price; an alternative would be to represent maximum WTP using a distribution that provides more weight around a middle-ground valuation. One example of such a distribution is the Normal (or Gaussian) distribution. The Normal distribution is commonly used in airline revenue management to represent passenger demand (Talluri and van Ryzin, 2005; Belobaba et al., 2016) and has also been used to model passenger willingness-to-pay in air transportation markets (Kambour, 2014; Kumar et al., 2017).

Suppose that $\Theta_x = N(\mu_x, \sigma_x)$. Then, Prob$(\theta_x \geq f_x) = 1$ - $\Phi_x^{-1}(f_x)$, where Φ_x^{-1} is the inverse cumulative distribution function for the Normal distribution with mean μ_x and standard deviation σ_x. The dynamic pricing problem for a monopolist with assortment $S = \{x\}$ is:

$$f_x^* = \arg\max_{f_x} \mathbb{E}(R) = (f_x - MC_x) \cdot (1 - \Phi_x^{-1}(f_x)) \tag{4}$$

Equation (4) cannot be expressed in closed form, but can be solved using commercial solving packages such as Excel Solver. For instance, with $\Theta_x = N(\$150, \$50)$ and $MC_x = \$0$,

$$f_x^* = \$116.76$$

Note that $\mathbb{E}(\Theta_x) = \$150$, just as in the uniform distribution case, but the optimal fare is higher when a normal distribution assumption is used for the maximum WTP relative to a uniform distribution assumption.

Now, consider adding an additional alternative y to the assortment with $\mathbf{a}_y = \mathbf{a}_x$ and $f_y =$ \$150. From Remark 5,

$$\widetilde{\theta}_{x|S} = \begin{cases} \theta_x & \text{if } \theta_x < \$150 \\ \$150 & \text{if } \theta_x \geq \$150 \end{cases}$$

The cumulative distribution function $\widetilde{\Theta}_{x|S}$ consists of a half-normal distribution[5] with a point-mass at \$150 with probability 0.5. The expected value of this distribution is:

$$\mathbb{E}(\widetilde{\Theta}_{x|S}) = \frac{150 + (\mu_x - \sigma_x\sqrt{\frac{2}{\pi}})}{2} = \$130.05$$

With this distributional assumption, the expected value of the conditional WTP distribution decreases from \$150 in the single-product case to \$130.05 in the case when an additional product is added into the choice set.

These examples help to highlight the finding that if airlines were to dynamically price itinerary products by estimating their customer's *maximum* WTP for that product, they would likely overestimate the optimal price if there are other attractive and affordable alternatives in the choice set. While the maximum WTP may be close to the conditional WTP when other products in the assortment are unaffordable, irrelevant, or possess unattractive attributes, airlines practicing dynamic pricing should consider their customers' *conditional* WTP distributions $\widetilde{\Theta}_x$ as opposed to *maximum* WTP Θ_x as has been considered in some past work (Kambour, 2014; Kumar et al., 2017).

[5] A half-normal distribution is a normal distribution folded at its mean. It has mean $\mu - \sigma\sqrt{\frac{2}{\pi}}$, where μ and σ are the mean and standard deviation of the original normal distribution.

3.5 From theory to heuristics: Integrating dynamic pricing with airline RM

To recap, in the previous sections we have specified a model of customer choice among available itinerary products. Using this model and microeconomic principles, we can describe the expected revenue for a firm pricing a single itinerary product:

$$\mathbb{E}(R) = (f_x - MC_x) \cdot \text{Prob}(\widetilde{\theta}_{x|S} \geq f_x) \tag{5}$$

To find the revenue-maximizing optimal price for itinerary product x, we can find the fare that maximizes the expected revenue in Equation (5).

$$f_x^* = \arg\max_{f_x}(f_x - MC_x) \cdot \text{Prob}(\widetilde{\theta}_{x|S} \geq f_x) \tag{6}$$

While the f_x^* from Equation (6) represents the theoretical optimal fare for a single itinerary product, it is not immediately clear how (1) this price can be generated in real-life situations, and (2) how this fare could interact with the complicated machinery that drives traditional airline revenue management systems. To the first point, we have not yet discussed how the marginal cost MC_x for each itinerary product would be generated for use in Equation (6). Since the airline RM problem is capacity constrained, this marginal cost should reflect the opportunity cost for consuming a unit of capacity on each flight leg used by the itinerary product. Furthermore, while it is clear that a conditional WTP distribution $\widetilde{\Theta}_{x|S}$ will be needed for Equation (6), the specification and estimation of this distribution have also remained unaddressed.

Furthermore, airlines have invested for decades in complex forecasting and optimization technologies for revenue management that they are unlikely to fully abandon for a new dynamic price adjustment approach. Therefore, a practical airline dynamic pricing heuristic should be compatible with the models and methods used currently in existing RM. As discussed in Chapter 1, pricing mechanisms should also be compatible with current industry standards, as well as new technologies such as the New Distribution Capability that will govern the ways in which these fares could be priced, distributed, and sold in the future.

To convert the theoretical discussion in this chapter into a practical heuristic, we need to address the following questions:

1. How can the optimal fare generated from Equation (6) be integrated with existing airline pricing, revenue management, and distribution systems, and with future technologies like the New Distribution Capability (NDC)?
2. How is the marginal cost MC_x computed in Equation (6)?

3. How is the conditional WTP distribution $\widetilde{\Theta}_{x|S}$ estimated in Equation (6)?

After addressing these questions, we describe a heuristic, Probabilistic Fare-Based Dynamic Adjustment (PFDynA), that airlines could use to dynamically adjust the prices of itinerary products within the framework of existing methods and systems.

First, recall from Section 1.2 that the airline shopping process is currently highly siloed. The pricing department and revenue management departments make their decisions separately and often in isolation. Pricing departments design *fare structures* with a number of pre-priced *fare products* k, each of which has a price f_k and a number of restrictions that determines under what conditions the fare product can be sold. The revenue management department then takes these fare structures as exogenous inputs and determine the availability levels for each fare product at each point in time before departure.

The definition below relates the concept of the fare structure to the itinerary product concept discussed in the previous section on airline customer choice.

Definition (Fare Structure): A *fare structure* $\{1, \ldots, k, \ldots, n\}$ where $f_i > f_j$ for all $i < j$ is a collection of all itinerary products for sale for a single itinerary in a single origin-destination market.

Fare structures are conventionally ordered from the most expensive (highest class) to the least expensive (lowest class). Therefore, Class 1 represents the most-expensive product, which is typically the least-restricted, and Class n represents the least-expensive product filed in the market, which is typically the most-restricted.[6]

Suppose that an airline revenue management system has closed[7] classes $\{k+1, \ldots, n\}$ and opened classes $\{1, \ldots, k\}$ for purchase. The itinerary product k with price f_k is then the *lowest available fare product* for the itinerary. In the following discussion, we will typically use the letter k to refer to this lowest-available fare product.

Definition (Lowest Available Fare Product): An itinerary product k is the *lowest available fare product* for an itinerary if fare class k is open and $f_k < f_j$ for all other open classes $j \neq k$ for the itinerary.

The optimal price f_k^* from Equation (6) can be reframed as a *dynamic price adjustment* Δ_k^* relative to the filed fare f_k of the lowest available itinerary product.

[6]Note that all of these itinerary products in the fare structure are being sold for the exact same itinerary.
[7]Recall that a fare class is *closed* if the RM system determines there are no seats available for purchase.

Dynamic Price Adjustment Equation:

$$\Delta_k^* = \arg \max_{\Delta_k \in [\ell_k, u_k]} ((f_k + \Delta_k) - MC_k) \cdot \text{Prob}(\widetilde{\theta}_{k|S} \geq (f_k + \Delta_k)) \tag{7}$$

After the dynamic price adjustment Δ_k^* has been computed using Equation (7), the modified fare for itinerary product k is:

$$f_k^* = f_k + \Delta_k^* \tag{8}$$

Note that Δ_k^* could be positive or negative. If Δ_k^* is positive, we say that the price of itinerary product k has been *incremented.* If Δ_k^* is negative, we say the price of itinerary product k has been *discounted.*

The Role of Bounds

In Equation (7), bounds ℓ_k and u_k have been imposed on the dynamic price adjustment Δ_k^*. Typically, we will set the bounds as follows: $\ell_k = f_{k+1} - f_k$ and $u_k = f_{k-1} - f_k$. In other words, the fare modifications are bounded by the gaps between the fare of the lowest available class f_k and the adjacent classes above or below it.[8]

These bounds serve several purposes. First, they help to prevent a situation known as a *fare inversion* (Vinod, 2015). Suppose that the fare of the lowest available fare product k was \$200, and the next higher product $k-1$ in the fare structure was \$250. If the dynamic price adjustment was set to $\Delta^* = \$100$, then the fare of product k would be \$300. In this case, no one would buy product k; the cheaper (and likely less-restricted) fare product $k-1$ would be more attractive. Therefore, we typically bound the adjusted fare above by the next highest class in the fare structure.[9]

More importantly, the bounds allow the existing airline revenue management system to continue to influence the dynamic price adjustments. If an airline's RM system closes classes $k+1$ through n, it does so based on its forecast of demand-to-come for each fare product and the capacity remaining on each flight leg. By passing this information through to the dynamic pricing algorithm, the fare modifications can be made incorporating information from existing revenue management processes, and will not completely override the recommendations provided by the RM system.

The bounds also provide a mechanism for temporal variations in dynamic price adjustments. Note that the cumulative WTP distribution $\widetilde{\Theta}_{x|S}$ is not necessarily assumed to be time-

[8]In the case that $k = n$, ℓ_k is typically assumed to be $f_n - f_{n+1}$. Similarly, if $k = 1, u_k = f_1 - f_2$.

[9]This assumes that availability of products does not change along with dynamic pricing. In an alternative heuristic, the fare of product $k-1$ could be set to \$300 and class k could be dynamically closed. Such a heuristic could represent a combination of dynamic pricing with assortment optimization.

dependent. That means that no matter when a customer arrives to book, he would be given the same "optimal" dynamic price adjustment (as a function of the bid price and the conditional WTP distribution). This does not align with current pricing and revenue management practices, which typically increase prices closer to flight departure. Incorporating adjacent fare-class bounds allows the revenue management system and existing pricing structure to dictate the range in which the dynamic price adjustment can occur. However, we also investigate the effect of removing fare class bounds in our analysis in Chapter 4.

3.5.1 Bid prices and marginal cost

Airlines are businesses with high fixed costs and low marginal costs. The marginal cost of accepting an additional booking is relatively small—perhaps limited to a negligible reservation processing cost, a printed boarding pass from a kiosk, and an additional can of soda once on board.

However, the seat in which that passenger is sitting could have a much higher marginal cost. For instance, a connecting passenger paying \$300 to travel from Boston to Chicago to Omaha may occupy seats that could have been sold to a local Boston - Chicago passenger for \$200 and another local Chicago-Omaha passenger for \$250. In this case, the connecting passenger paying \$300 is displacing \$450 worth of potential revenue.

In network revenue management, airlines account for these displacement costs through a concept known as a *leg bid price*. The leg bid price can be thought of as the opportunity cost or marginal value of capacity for a flight leg; that is, the expected additional total network revenue that would be obtained if an additional seat was added to the flight leg (Williamson, 1992). In the language of linear programming, the bid price is the *shadow price* from relaxing the capacity constraint on the flight leg by one unit.

There are a number of ways to obtain a vector of leg bid prices for each leg in $\mathcal{L}$. Here, we will focus on two methods: the deterministic linear programming approach (DLP) and the dynamic programming approach (DP).

Obtaining bid prices from a deterministic linear program

In the DLP approach, the leg bid prices are obtained from the dual solution to a deterministic linear program, which attempts to maximize the expected revenue from allocating seats to itinerary products over a flight network. There are several formulations of DLPs for airline revenue management; the following approach is adapted from Bertsimas and de Boer (2005). Let f_k represent the revenue from itinerary product $k \in \mathcal{F}$, c_l represent the capacity on leg $l \in \mathcal{L}$, and D_k represent the mean *deterministic* demand forecast for itinerary product k.

We then solve the following linear program:

$$\max \sum_{k \in \mathcal{F}} f_k x_k$$

Subject to:

$$x_k \leq D_k$$

$$\sum_{k:l \in k} x_k \leq c_l \qquad l = 1, \ldots, L$$

$$x_k \geq 0$$

Here, the decision variables x_k indicate how many of each itinerary product k are allocated for sale. Yet these decision variables themselves are not useful for allocating capacity; the assumption that demand is deterministic and the non-nested nature of the model means that the allocation levels suggested by the LP will not perform well when facing more realistic stochastic demand (Williamson, 1992; Bertsimas and de Boer, 2005; Belobaba, 2016).

However, the shadow prices π_l associated with the constraints $\sum_{k:l \in k} x_k \leq c_l$, which limit the number of seats sold on each flight leg to not exceed the leg capacity, are indeed useful. These shadow prices π_l represent the estimated additional revenue that could be gained if the capacity of leg l was increased by one unit. Note that if there is excess capacity on a leg (i.e., the constraint $\sum_{k:l \in k} x_k \leq c_l$ is not binding), then the bid price for that leg would be zero.

Obtaining bid prices from a dynamic program

Alternatively, a dynamic programming approach could be used to compute leg bid prices. As opposed to the deterministic linear program, which is computed for a single point in time, a dynamic program considers the time remaining before the flight's departure. For instance, the Lautenbacher and Stidham (1999) dynamic programming approach to network RM uses the following time-dependent leg value function $V_{l,t}$, also called a Bellman equation.

$$V_{l,t}(x_l) = \sum_{k \in \mathcal{F}: l \in k} \lambda_{k,t} \cdot \max\{f_k + V_{l,t+1}(x_l + 1), V_{l,t+1}(x_l)\} + \lambda_{0,t} V_{l,t+1}(x_l) \tag{9}$$

In the value equation, x_l represents the current number of bookings on leg l, $\lambda_{k,t}$ is the forecast demand for itinerary product k at time t,[10] and $\lambda_{0,t}$ represents the probability of no sale at time t.

The value function is solved recursively starting at the last period (i.e., just prior to flight

[10] If the time periods t were short enough to allow at most one arrival, the quantity $\lambda_{k,t}$ could represent an arrival rate for demand for itinerary product k.

departure) and working backward in time. At any given moment, the value function represents the expectation of value-to-come at time t for a marginal seat on leg l. At a given time t, we can compute the leg bid price by considering the change in the value function when the remaining capacity decreases by 1:

$$\pi_{l,t}(x_l) = V_{l,t}(x_l) - V_{l,t}(x_l + 1)$$

Uses of bid prices for availability control and dynamic pricing

After leg bid prices are obtained from either the DLP or DP methods, they can be used to control availability for itinerary products. One common network availability heuristic is called *additive bid price control.* In additive bid price control, an itinerary product k is made available for sale if and only if its fare f_k exceeds the sum of the bid prices $\sum_{l \in k} \pi_l$ for the legs used by the itinerary (Williamson, 1992; Talluri and van Ryzin, 1998). While additive bid price control is not always optimal (Bertsimas and de Boer, 2005; Talluri and van Ryzin, 2005), it remains a common heuristic used by airlines (Belobaba, 2016).

Another common availability method for network RM availability control is called *displacement-adjusted virtual nesting (DAVN)* (Smith and Penn, 1988). In DAVN, displacement-adjusted leg fares $f'_{k,l} = f_k - \sum_{l' \in k : l' \neq l} \pi_l$ are computed for each itinerary product/leg combination, by subtracting the displacement costs for other legs the itinerary consumes in the network. These adjusted fares are then clustered into buckets with other adjusted fares of similar value, and a leg-based RM method such as EMSRb is used to determine itinerary product availability (Williamson, 1992; Bertsimas and de Boer, 2005).

For our purposes in dynamic pricing, we will use the bid prices obtained from the RM system to approximate the marginal cost for each itinerary product in Equation (7).

Assumption 1: In Equation (7), the marginal cost of itinerary product x can be approximated by the sum of the bid prices π_l of legs $l \in x$. That is, $MC_x = \sum_{l \in x} \pi_l$.

Since the bid price represents the marginal value of a unit of capacity on each leg, subtracting the sum of the bid prices from the fare f_k will account for the displacement costs that are associated with selling a single seat on each leg used by itinerary product k. Furthermore, it will ensure that the dynamic pricing formulation will never set a fare below the bid price, which would be in conflict with availability control methods such as additive bid price control.

3.5.2 Assumptions on the conditional WTP distribution: The Q-multiplier

So far, we have introduced the concept of dynamic price adjustments, and discussed how bid prices from the airline's revenue management system could be used as marginal costs in the dynamic fare modification equation. To conclude, we need to address the conditional WTP distribution $\widetilde{\Theta}_{x|S}$ used to compute the probabilities $\text{Prob}(\widetilde{\theta}_{x|S} \geq f_x)$.

The estimation of customer willingness to pay as been the subject of considerable academic and industry attention. One approach focuses on estimating the *price elasticity* for air travel (Gillen et al., 2003; Njegovan, 2006). This is particularly useful for macro-scale or network planning purposes, where airlines would wish to determine how much a change in price could stimulate or dampen demand. Other approaches have focused on trying to determine the willingness-to-pay for ancillary services, like extra leg room or a checked bag (Espino et al., 2008; Balcombe et al., 2009).

Traditional airline revenue management models, however, do not typically use price elasticity or willingness-to-pay estimates as inputs. Instead, many models use historical booking data to construct a forecast of demand-to-come for each itinerary product (Talluri and van Ryzin, 2005; Belobaba et al., 2016). In more basic models, these demands may be fare class specific: that is, a certain amount of demand is expected to request in Q class on a given flight from Boston to Paris departing in four weeks. This demand is usually assumed to be stochastic, typically with a Normal or Poisson distribution (Belobaba, 2016). More recent forecasting models have considered ways to account for buy-up and buy-down behavior to generate conditional estimates of demand (Belobaba and Hopperstad, 2004; Fiig et al., 2010).

There is no single accepted method that is used to estimate or compute customer willingness to pay or price elasticity in the airline industry. Indeed, estimating a customer WTP distribution from data is beyond the scope of this dissertation—we are more interested with developing the theoretical building blocks that airlines could use to estimate the WTP distributions specific to their customer segments and their markets.

As such, we will make a distributional assumption on the shape of the customer's conditional WTP distribution for use in our dynamic pricing heuristic. We will use a Gaussian (Normal) distribution to portray the range of conditional WTP values for a given itinerary product in a given assortment. There is precedent to use such an assumption; the use of Normal distributions is common in airline revenue management models, such as EMSRb (Belobaba, 1992), to measure passenger demand (McGill and van Ryzin, 1999). Furthermore, in several recent publications, researchers have also started using Normal distributions to model (maximum) passenger WTP (Kambour, 2014; Kumar et al., 2017). A literature review by McGill and van Ryzin (1999) also states that the "Normal probability distribution gives a good continuous approximation to aggregate airline demand distributions" (p. 237) as a result of the Central Limit Theorem.

Assumption 2 formalizes the use of a Normal distribution to model conditional WTP:

Assumption 2:	The conditional WTP distribution $\widetilde{\Theta}_{x\|S}$ for itinerary product x in assortment S is Normally distributed with mean $\mu_{x\|S}$ and standard deviation $\sigma_{x\|S} = \gamma \cdot \mu_{x\|S}$, with coefficient of variation γ.

Note that the Normal distribution used to approximate the conditional WTP distribution is parameterized by a mean $\mu_{x|S}$ and a coefficient of variation γ. However, the means of the conditional WTP distributions could vary significantly from market to market.

For instance, a long-haul international market would likely have a very different conditional WTP distribution than a short-haul, one-hour flight. Similarly, the mean conditional WTP of markets that are highly leisure-oriented may be less than for markets frequented by less-discretionary business travelers. And since conditional WTP is a function not only of the market and fare product, but also of the assortment in which the product is displayed, highly competitive markets with many available alternatives will likely lead to conditional WTP distributions with smaller means than those served only infrequently or by a single carrier.

It would be difficult or impossible for an airline to estimate the means of the thousands or millions of individual market, passenger, assortment, and product combinations needed to produce conditional WTP distributions for every environment in practice. To operationalize the heuristic, we need a mechanism for airlines to easily specify mean WTP parameters for a wide variety of conditional WTP distributions. Assumption 3 introduces this parameter, which we call a "Q-multiplier."

Assumption 3:	For any itinerary product x and assortment S, the mean of the conditional WTP distribution $\widetilde{\Theta}_{x\|S}$ is equal to $QMULT \cdot f_n$, where $QMULT$ is an input Q-multiplier and f_n is the lowest filed fare in the itinerary's fare structure, regardless of availability.

The Q-multiplier is essentially a ratio between the mean of the conditional WTP distribution and the lowest filed fare in the fare structure. For instance, if the lowest filed fare f_n is \$100, selecting a Q-multiplier of 1.5 and a coefficient of variation $\gamma = 0.3$ would imply that the conditional WTP distribution for an itinerary product x in assortment S would be Normally distributed with a mean of \$150 (\$100 $\cdot$ 1.5) and a standard deviation of \$45 (\$150 $\cdot$ 0.3). By changing the Q-multiplier with a fixed coefficient of variation, the airline can adjust both the mean of the conditional WTP and the standard deviation simultaneously. The Q-multiplier thus serves as a mechanism by which airlines could generate a large number of conditional WTP distributions without having to specify WTP individually for thousands of market/assortment combinations.

Note that Assumption 3 does not specify subscript(s) for the Q-multiplier estimate. At the limit, an unique Q-multiplier could be generated for each booking request, taking into account passenger characteristics, the products in the assortment, fare restrictions, ancillary fees and amenities, and itinerary qualities. Such fine-tuned estimation is likely to be out of reach in the short term. Instead, airlines will likely generate a limited number of representative Q-multipliers. For instance, Q-multipliers could be generated for leisure-oriented, highly-competitive markets, as well as business-oriented markets with only a limited number of itinerary products available. These Q-multipliers could then be applied as needed to compute the dynamic price adjustment for any given scenario.

In general, airlines should use relatively higher Q-multipliers when:

- The customer's maximum WTP is likely to be higher (for instance, for non-discretionary, business passengers);
- The itinerary product has relatively few fare restrictions;
- There are relatively few alternatives in the assortment S;
- The itinerary product has attractive attributes (for instance, a short elapsed flight time, nonstop flight, or an attractive departure time) relative to other alternatives in the assortment S; and/or
- There is limited time remaining before flight departure (since higher WTP passengers are more likely to arrive closer to flight departure).

If passengers can be segmented into types, either based on their past purchasing behavior or the characteristics of their requests, it would be possible to specify Q-multiplier estimates for each passenger type. For instance, if passenger requests could be segmented into discretionary leisure and non-discretionary business, and it is assumed that the less-discretionary business travelers would have a higher WTP, a higher Q-multiplier could be used for business requests relative to leisure requests. We will further describe this possibility for passenger type segmentation in Section 3.7.

3.6 The Probabilistic Fare-Based Dynamic Adjustment (PFDynA) heuristic

By incorporating the customer choice model developed earlier, Theorem 1, Equation (7), and Assumptions 1, 2, and 3, we arrive at the following heuristic: Probabilistic Fare-Based Dynamic Adjustment, or PFDynA:

Heuristic 1: Probabilistic Fare-Based Dynamic Adjustment (PFDynA):
Suppose k is the lowest available fare product for an itinerary, and n is the lowest filed fare in the fare structure, regardless of availability. Let $\ell_k = f_{k+1} - f_k$ (or 0 if $k = n$), and $u_k = f_{k-1} - f_k$ (or $f_1 - f_2$ if $k = 1$). The PFDynA price adjustment Δ_k^* is:

$$\Delta_k^* = \arg \max_{\Delta_k \in [\ell_k, u_k]} \left((f_k + \Delta_k) - \sum_{l \in k} BP_l \right) \cdot \left(1 - \Phi^{-1} \left(\frac{(f_k + \Delta_k) - \mu_k}{\sigma_k} \right) \right) \quad (10)$$

where $\mu_k = QMULT \cdot f_n$ and $\sigma_k = \gamma_k \cdot \mu_k$. The offered price for product k is then:

$$f_k^* = f_k + \Delta_k^*$$

The fare modification Δ_k^* could be computed by any solver that is capable of computing probabilities from Normal cumulative distribution functions. The following example illustrates in more detail how the PFDynA dynamic price adjustment is calculated:

Example 1: Suppose the fare structure for a given itinerary consists of six classes $\{1, 2, 3, 4, 5, 6\}$ priced as follows: $\{\$600, \$500, \$400, \$300, \$200, \$100\}$. At a given point in time, suppose that the itinerary's bid price is equal to \$150. Using additive bid price control, any filed fare above \$150 would be available for purchase. Given the fare structure above, class $k = 5$ would be the lowest available fare product and the displayed price of this itinerary would be $f_k = \$200$.

Suppose that for a given customer and assortment, a Q-multiplier estimate of 1.5 has been chosen with a coefficient of variation of $\gamma = 0.3$. Then the conditional WTP distribution $\widetilde{\Theta}_{k|S}$ is assumed to be Normally distributed with a mean of $\mu_{k|S} = QMULT \cdot f_n = 1.5 \cdot \$100 = \$150$ and a standard deviation of $\gamma \cdot \mu_{k|S} = 0.3 \cdot \$150 = \$45$. Then, the dynamic price adjustment is computed from Equation (10) as follows:

$$\Delta_k^* = \arg \max_{\Delta_k \in [-\$100, \$100]} ((\$200 + \Delta_k) - \$150) \cdot \left(1 - \Phi^{-1} \frac{(\$200 + \Delta_k) - \$150}{\$45} \right) = -\$16.17$$

Here, the PFDynA heuristic suggests a \$16.17 discount for this customer. The PFDynA fare offered for Class 5 would be $f_5 - \$16.17 = \$183.83$.

3.6.1 Mechanisms by which PFDynA could affect airline revenues

Recall that according to our customer choice model, a customer will choose to purchase itinerary product k as long as $f_k \leq \widetilde{\theta}_{k|S}$. As a result, dynamically adjusting the filed fare f_k using Equation (10) could lead to several potential benefits:

- If the filed fare f_k is significantly below the customer's conditional WTP $\widetilde{\theta}_{k|S}$, incrementing the fare by an amount less than $f_k - \widetilde{\theta}_{k|S}$ would still result in the customer purchasing the itinerary product, but at a higher fare. In this case, incrementing fares can increase the airline's revenue by converting some of the customer's consumer surplus into revenue. The increase in revenue would be equal to the fare modification Δ_k^*.
- If the filed fare f_k is above the customer's conditional WTP $\widetilde{\theta}_{k|S}$, discounting the fare to a level below $\widetilde{\theta}_{k|S}$ could lead the customer to purchase itinerary product k when she otherwise may have chosen not to travel or booked an alternative itinerary product, perhaps with a competitor. The increase in revenue in this case would be equal to the modified fare $(f_k + \Delta_k^*) \leq f_k$, where Δ_k^* is negative.

Since Δ_k^* is bounded by the gap in adjacent class fares, it will generally be small relative to f_k. As a result, discounting could lead to a greater potential gain in revenue than incrementing.

We also need to consider the risks of dynamically modifying fares. In these cases, suppose that a customer would have ordinarily purchased itinerary product k at the original fare f_k.

- If the customer's fare is incremented above her conditional WTP $\widetilde{\theta}_{k|S}$, she will choose not to travel or to purchase another itinerary product, perhaps with a competitor. The loss of revenue in this situation would equal the entire original filed fare f_k.
- If the customer receives an unnecessary discount, she will still choose to book itinerary product k, but at a price lower than f_k. The loss of revenue to the airline in this case would equal the dynamic discount $\Delta_k^* < 0$.

In this case, the potential loss in revenue is generally greater for incrementing than discounting. This makes intuitive sense—incrementing a fare above a passenger's conditional WTP will result in the loss of the entire booking, as opposed to a small loss in revenue from providing an unnecessary discount. Therefore, discounting appears to produce a greater potential reward with less potential risk than incrementing. In our simulation results, we will examine this relationship between the rewards and risks of incrementing versus discounting in various situations.

3.7 Passenger segmentation in dynamic pricing: Methods for PFDynA

One of the proposed benefits of the New Distribution Capability is that it will allow targeted offers to be generated for specific search requests. At the limit, this could mean that each individual passenger could have a customized offer generated and priced specifically for their attributes and characteristics. That is, if the conditional WTP $\widetilde{\theta}_{x|S}$ was known or could be estimated for each passenger, NDC provides a mechanism to set the price equal to $\widetilde{\theta}_{x|S}$ for each passenger. This mechanism can be seen in economic terms as perfect first-degree price discrimination.

However, for many reasons, individualized customized pricing is unlikely to take hold in the airline industry for many years. IATA Resolution 787, which provided a framework for NDC, was only ratified by the U.S. Department of Transportation under the condition that anonymous shopping would always be allowed (Rice, 2015). Therefore, there will always be a portion of shopping requests for which the airlines will be unable to gather customized information. Furthermore, the use of personalized information in airline marketing and IT infrastructure today has been limited, and a full ramp-up to customization could take time. It is also unclear whether customization would be permitted from a regulatory perspective.

Customization aside, it is reasonable to expect that airlines might be able to *segment* customers or requests into different categories. For instance, based on past travel patterns, search histories, or characteristics of the request, an airline may be able to segment a request into discretionary leisure or non-discretionary business categories. We would expect that passengers segmented into the non-discretionary business category to arrive later in the booking process and have higher maximum WTP values.

Definition (Customer Types): Let W represent a set of customer types. Then for any itinerary product x and assortment S, the conditional WTP $\widetilde{\theta}_{x|S}$ is drawn from the same distribution $\widetilde{\Theta}^{w}_{x|S}$ for all customers of type $w \in W$.

The segmentation of passengers into types allows for more fine control over PFDynA actions. First, since the conditional WTP distributions are different for passengers of different w-types, it would be reasonable to use different Q-multipliers as estimates of conditional WTP in the PFDynA heuristic. For instance, the Q-multiplier used for non-discretionary business travelers should likely be higher than the one used for discretionary leisure travelers, since the former group is likely to have a higher maximum WTP on average.

Second, if passengers can be segmented into types, it would be possible to target increments or discounts towards specific groups of passengers. It may be unwise to provide highly price-sensitive leisure passengers with incremented prices, since these customers may choose to book with a competitor even when facing only a small increase in price. Through segmentation, various variations on the PFDynA heuristic could be created, in which different groups would be eligible for different types of fare modifications.

Consider a world in which W = {Business, Leisure}, and business passengers are assumed to have higher conditional WTPs on average than leisure passengers. The box below shows several possible variants of PFDynA for use in this situation.

- **Increments-Only PFDynA:** For customers belonging to the "business" segment, compute a PFDynA increment Δ_k^* using Equation (10), bounded below by 0 and above by $f_{k-1} - f_k$ (or $f_1 - f_2$ if $k = 1$). Do not adjust prices for leisure customers.
- **Discounts-Only PFDynA:** For customers belonging to the "leisure" segment, compute a PFDynA discount Δ_k^* using Equation (10), bounded above by 0 and below by $f_{k+1} - f_k$ (or 0 if $k = n$). Do not adjust prices for business customers.
- **Two-Way PFDynA:** Compute both PFDynA increments for business customers and PFDynA discounts for leisure customers, as described above. However, business customers are not eligible for discounts, and vice versa.
- **Unsegmented PFDynA:** In conjunction with any of the methods described above, the same Q-multiplier is used for customers of both types.

Consider our previous Example 1, in which an airline used the following six-class filed fare structure: {\$600, \$500, \$400, \$300, \$200, \$100} for an market. At a certain point in time, the bid prices for the legs in the itinerary sum to \$150. In this case, the lowest class that would be open for sale would be Class 5 at a price of \$200. Suppose that the Q-multiplier for business passengers in this market is estimated at 2.5, and the Q-multiplier for leisure passengers is estimated at 1.5. Table 3.1 provides examples of the dynamically modified fares for both passenger types using the heuristics described above.

Method	Passenger Type	
	Business (Q = 2.5)	Leisure (Q = 1.5)
Itinerary Bid Price	$150	
No PFDynA (Standard RM)	$200	$200
PFDynA Optimal Price (Equation (10))	*$247*	*$184*
Increments-Only PFDynA	**$247**	$200
Discounts-Only PFDynA	$200	**$184**
Two-Way PFDynA	**$247**	**$184**
Unsegmented Two-Way PFDynA (Q = 2.0)	**$213**	**$213**

Table 3.1: Offered fares for various segmented and unsegmented PFDynA heuristics

Using Q-multiplier estimates of 2.5 and 1.5, the optimal offered fares from Equation (10) are \$247 and \$184, respectively. However, as shown in Table 3.1, various PFDynA methods provide different combinations of the PFDynA optimal fares and the normal fares. For instance, when unsegmented two-way PFDynA is used with a single Q-multiplier estimate of 2.0, both passenger types would be offered a fare of \$213.

Depending on how customers could be segmented, and on the regulatory and competitive permissibility of discounts and increments, airlines may select to practice any of the modified PFDynA heuristics listed above. In practice, booking requests could be segmented based on characteristics of the *customer*, such as frequent flyer status, past travel history, search patterns, or on characteristics of the *request*, such as how many days prior to departure the request is made, the length of stay, the destination, or other features. Airlines currently use request characteristics to segment incoming demand, and there is a growing literature on passenger segmentation in the airline industry. Readers interested in segmentation approaches can refer to Teichert et al. (2008); Brunger (2010); Kothari et al. (2016); and Nicolae et al. (2016).

In the next chapter, we also test situations in which segmentation is not possible with 100% accuracy, as well as the unsegmented heuristics discussed above wherein passenger segmentation is not possible. As discussed in Chapter 2, the economics literature has suggested that imperfect targetability may actually improve the performance of customized or personalized pricing methods in competitive environments (Chen et al., 2001; Esteves, 2014).

3.8 Summary

In this chapter, we proposed a complete framework for customer choice and dynamic price adjustment of airline fare products. First, we presented a model of customer choice based on rational choice theory. Customers in this model choose to purchase the itinerary product with the lowest perceived price among affordable alternatives. Using a concept we called conditional WTP, we linked together customer choice in this model to the fare and attributes of the product as well as the other alternatives offered to the customer.

Then, we discussed how microeconomic principles could allow us to compute the optimal price for an itinerary product. Unlike other approaches that set prices based on customers' *maximum* WTP, our optimal price was a function of the distributional assumption on *conditional* WTP. We provided examples using uniform and Normal distributions to show how the optimal price could vary when alternatives are added into the choice set.

Finally, we converted this theoretical discussion into an actionable heuristic that would interact with existing airline pricing, revenue management, and distribution processes. Using the filed fare structure as a base, we discussed a Normal distribution assumption for conditional WTP, as well as a parameter—the Q-multiplier—that could be used to specify this distribution. We also discussed how bid prices could be used as a proxy for marginal cost in our dynamic pricing equations.

The resulting heuristic, Probabilistic Fare-Based Dynamic Adjustment (PFDynA), dynamically computes a fare modification to the lowest available filed fare as a function of the itinerary product's bid price and the assumed conditional WTP distribution. We presented several segmented variants of the PFDynA base heuristic, including Increments-Only and Discounts-Only PFDynA, in which only certain segments of customers are eligible for increments or discounts.

PFDynA could increase revenue by encouraging passengers with higher WTP to buy up to higher price points, or by stimulating new bookings by providing a discount. However, depending on the customer's exact conditional WTP, PFDynA fare modifications could lead to unnecessary discounts or increments that could lead the customer to no-go or book with a competitor. To see which of these effects dominates, and to estimate the revenue performance of the heuristic, PFDynA needs to be tested in a robust simulation with passenger choice, competition, and revenue management. In the next chapter, we test the PFDynA heuristics described above in the Passenger Origin-Destination Simulator, or PODS, to measure the effects of dynamic price adjustments in realistic airline environments.

4 Simulations of Dynamic Pricing with Airline Revenue Management

In Chapter 3, we developed a methodology—Probabilistic Fare-Based Dynamic Adjustment (PFDynA)—to dynamically adjust the prices that would ordinarily be offered by an airline's revenue management system. In this chapter, we test the performance of PFDynA in a realistic revenue management simulation environment using the Passenger Origin-Destination Simulator (PODS).

PODS is a complex revenue management simulator that has been in development for over two decades. Originally developed by engineers at Boeing Commercial Airplanes in the early 1990s, PODS has become one of the most advanced simulators currently in use to test the performance of revenue management methods for the airline industry. An implementation of PFDynA was programmed in PODS, allowing the heuristic to be tested in a variety of monopolistic and competitive network scenarios and to be compared to traditional airline revenue management methods.

This chapter is structured as follows: first, we briefly describe some of the core functionalities of PODS, including the passenger generation and choice process as well as some of the revenue management and forecasting methods commonly used in PODS. We also describe how PFDynA was implemented in PODS.

We then present the results from a number of PODS simulation tests of PFDynA in various network environments. We test PFDynA in simple one-airline and two-airline networks, as well as in complicated multi-airline networks with hundreds of flight legs and origin-destination markets. The results of these simulations suggest that PFDynA has the potential to increase airline revenues by increasing yield, stimulating new demand, and improving fare class mix. We close by reviewing some of the common patterns that emerge when an airline uses PFDynA in PODS.

4.1 The Passenger Origin-Destination Simulator (PODS)

PODS and its underlying mechanics have been described in detail in many past works (e.g. Gorin (2000); Carrier (2003); Cléaz-Savoyen (2005)). Here, we focus mostly on reviewing the features of PODS that are most relevant for our tests of PFDynA.

4.1.1 An overview of PODS

The Passenger Origin-Destination Simulator is an agent-based model that repeatedly simulates the booking process for flights on a single departure day. A single PODS *simulation run* consists of two or more *trials*, each of which generally consists of 600 *samples*. Each sample can be seen as a simulation of the booking process leading up to a single departure day. In PODS, the booking process typically starts 63 days before the departure day.

During each sample, PODS models the interactions between two different types of agents: passengers and airlines.

- *Passengers* are generated by PODS with a desire to travel in a single air transportation market. Each passenger possesses randomly-selected idiosyncratic attributes, including a maximum budget for air travel, a time-of-day departure preference, and disutilities associated with various fare restrictions. When a passenger arrives to book, he is shown an assortment of zero or more available itinerary products[11] for his desired market. Each passenger chooses to purchase exactly one itinerary product, or decides to *no-go* (not travel). The details of the passenger choice process are discussed in Section 4.1.2.
- *Airlines* offer service across a network of flight legs. These flight legs are combined into various itineraries that are sold in a variety of origin-destination markets. Airlines use revenue management systems to determine which itinerary products to make available during the booking process. Over the course of the trial, airlines use information about bookings in previous samples to update their revenue management forecasts, but have no information about the underlying passenger generation process.

Figure 4.1 shows a schematic of the interactions between passengers and airlines in PODS. During each sample, passengers are generated and arrive to book one at a time. Each passenger constructs a choice set of affordable itinerary products (if any), and uses a customer choice model to select one of the products in the choice set for purchase. If the passenger chooses to purchase an itinerary product, the airline records the booking in a database.

The airline uses this database, which includes observations of bookings in previous samples, to construct a demand forecast for each itinerary product on future departure days (or samples). The demand forecast is then used by the airline's revenue management optimizer to determine which itinerary products to make available. This loop continues throughout the trial, with the airline learning from its booking history in previous samples to update its forecasts and refine its itinerary product availability decisions in future samples.

[11] In the PODS literature, itinerary products are sometimes referred to as "path/classes." We retain the phrase "itinerary product" in this chapter for consistency with Chapter 3.

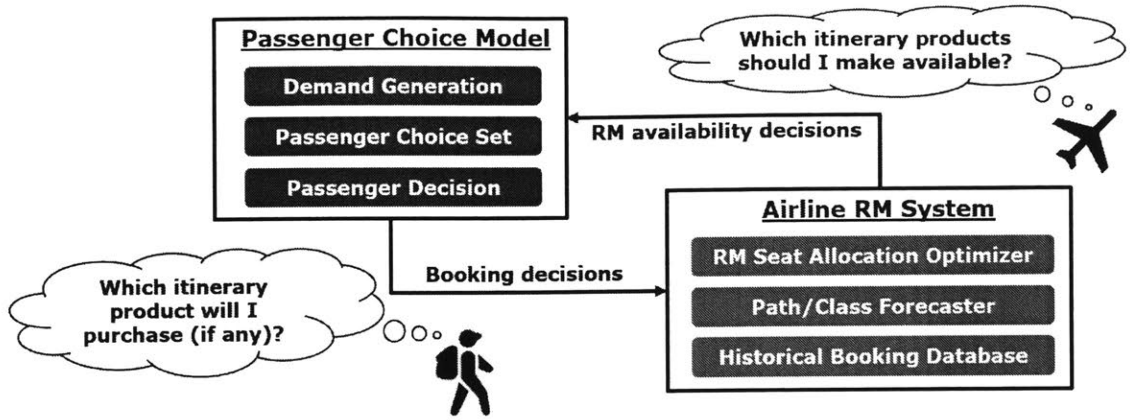

Figure 4.1: Schematic of interactions between passengers and airlines in PODS

Airlines in PODS do not possess any underlying information about passenger behavior and, like real-world airlines, must construct their forecasts and availability decisions based on historical bookings. When reporting results, we typically discard the first 200 samples of each trial as a "burn-in period." This allows the airline's revenue management forecaster to build up a robust history of bookings in previous departure days, leading to better revenue management optimization decisions. We then average together the airline's revenue from each sample after the burn-in period to generate the revenue results of the simulation run. For a typical simulation run with two trials, the revenue results reported are the average of 800 individual samples (400 post-burn samples for each trial).[12]

4.1.2 The passenger generation process and customer choice model in PODS

During each sample, PODS generates hundreds or thousands of individual passengers which arrive to book one at a time. Each passenger possesses a number of randomly generated attributes that govern their choice behavior. These attributes include:

- *Passenger type.* Each PODS passenger is generated with a single passenger type, which is generally interpreted as reflecting whether the customer is traveling for business or leisure. Leisure-type customers are more likely to have a lower maximum WTP (discussed below) and arrive earlier in the booking process. Business-type customers are more likely to have a higher maximum WTP and arrive closer to flight departure. Figure 4.2 shows typical arrival rates for business and leisure customers in PODS.

[12]With 800 samples, T-tests have generally found that revenue changes between two simulation runs of 0.05% or greater are statistically significant at a 95% level.

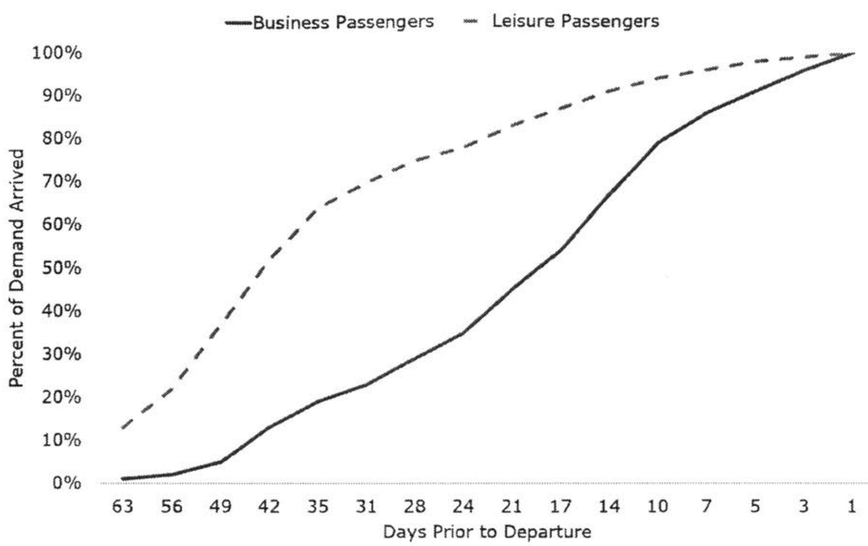

Figure 4.2: Example arrival curves by passenger type in PODS

- *Maximum Willingness-to-Pay.* Each PODS passengers has a randomly-drawn maximum WTP, which describes the maximum amount the customer will pay for air transportation in the market in which they desire to travel. This value is analogous to the concept of maximum WTP described in Chapter 3.

  Each passenger's maximum WTP is drawn from a negative exponential distribution which varies by passenger type and by market. The probability that a type w passenger's maximum WTP θ in some market m is greater than some fare f is:

$$\text{Prob}(\theta > f) = \min\left[1, e^{\frac{-\ln(2)\cdot(f - basefare_{w,m})}{(emult_w - 1)\cdot basefare_{w,m}}}\right]$$

  Equivalently, the negative exponential distribution $\Theta_{w,m}$ used to generate maximum WTP for type w passengers in market m has a mean of

$$\mu_{w,m} = basefare_{w,m} \cdot \left(1 + \frac{(emult_w - 1)}{\ln(2)}\right)$$

  and a standard deviation of

$$\sigma_{w,m} = basefare_{w,m} \cdot \left(\frac{(emult_w - 1)}{\ln(2)}\right)$$

  The distribution $\Theta_{w,m}$ of maximum WTP is parameterized by two values:

  - $basefare_{w,m}$ is a fare value that 100% of the mean demand of type w passengers in market m are assumed to be able to afford. As an example, in some PODS networks $basefare_{w,m}$ is set equal to the lowest filed fare in the market's fare structure for leisure passengers, and 2.5 times the lowest filed fare for business passengers.

– $emult_w$ is a price elasticity multiplier for passenger type w. which is typically greater than 1. Multiplying the *emult* parameter by the basefare yields a fare value that 50% of type w passengers are able to afford. For example, with a basefare of \$100 and an *emult* of 1.5, exactly 50% of passengers would be able to afford a fare of \$150 or greater.

Figure 4.3 shows an example of the distribution of maximum WTP θ for passengers with $basefare$ = \$100 and $emult$ = 1.5.

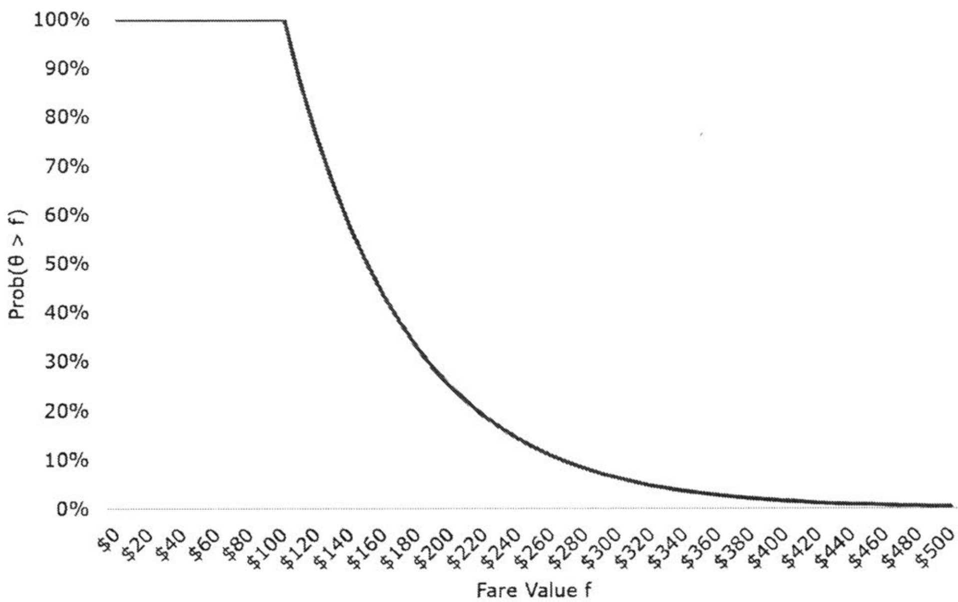

Figure 4.3: Example of negative exponential distribution for maximum WTP in PODS (basefare = \$100, emult = 1.5)

- *Time-of-day departure preference.* Each PODS passenger is generated with a time-of-day departure preference, which we will call ω. This preference is drawn from a distribution that depends on the duration of a hypothetical nonstop journey between the origin and destination. The passenger uses their time-of-day departure preference to create a *decision window* around their preferred departure time ω. The size of this decision window also depends on the length of the journey.

  Passengers prefer itineraries that fit entirely within their decision window; itineraries that fall either partially or entirely outside the decision window incur a "replanning disutility," which is considered when the passenger computes the perceived price of each affordable itinerary.

- *Fare restriction disutility costs.* Each PODS passenger incurs a *disutility cost* for purchasing a itinerary product that includes a fare restriction. In the tests shown in this chapter, itinerary products can have up to three restrictions: R1, R2, and R3. R1 is assumed to be the most onerous restriction, and R3 is the least onerous.

  Each passenger has random disutility values for each fare restriction which are drawn from Normal distributions. The means of each distribution are higher for business passengers than leisure passengers, reflecting an assumption that business passengers will be more affected by fare restrictions like Saturday-night minimum stays or nonrefundability. The passenger's disutility costs are used during the passenger choice process when a customer is computing the perceived price of each itinerary.

After a passenger is generated, he is shown an assortment S of itinerary products in his desired OD market, depending on which itinerary products are made available by each airline's revenue management system. The passenger then decides which itinerary product (if any) in the assortment to purchase, using the following methodology:

- First, the passenger prunes the assortment S by removing any unaffordable products. Recall from Chapter 3 that a product is unaffordable if its price exceeds the passenger's maximum WTP. In that chapter, we referred to the set of affordable products $\{x \in S : f_x \leq \theta\}$ as S'. Note that disutility costs are not considered when evaluating affordable products.

- If $S' = \emptyset$, either the passenger cannot afford any of the available itinerary products, or all of the itinerary products in that OD market are unavailable. In either case, the passenger chooses to *no-go* and purchases nothing. The passenger then leaves the simulation; he is not able to re-shop later in the booking process.

- If S' is nonempty, the passenger can afford at least one available itinerary product and will definitely choose to travel. To decide which itinerary product to purchase, the passenger computes his perceived price for each itinerary product. Recall from Chapter 3 that the perceived price PP_x for itinerary product x is $PP_x = f_x + \boldsymbol{\beta}\mathbf{a}_x$, where f_x is the price of itinerary product x, $\mathbf{a}_x$ is a vector of attributes of the itinerary product, and $\boldsymbol{\beta}$ is a vector of disutility costs in willingness-to-pay space.

  In PODS, $\mathbf{a}_x$ is a vector of [0,1] values that indicates whether the itinerary product contains any fare restrictions and/or is associated with a replanning penalty, and $\boldsymbol{\beta}$ is a vector of disutility costs for each fare restriction and the replanning disutility, which are generated as described above.

- The customer deterministically purchases the itinerary product in S' with the lowest perceived price. If a customer is indifferent between two or more itinerary products (because the itinerary products have the same perceived price), he will select randomly among them.

The passenger choice process in PODS matches the general choice model described in Chapter 3. Passengers first create a set of affordable itinerary products, and then deterministically select the itinerary product with the lowest perceived price among this set of affordable options. Because these customer choice processes are fundamentally similar, the results from Chapter 3 (e.g. Theorem 1, which found that the probability that a customer selects an itinerary product from an assortment equals the probability that his conditional WTP for that product exceeds its price) will also hold in PODS.

4.1.3 Revenue management optimization and forecasting methods in PODS

Airlines in PODS use revenue management systems to determine which itinerary products to make available at each *time frame*, or data collection point (DCP).[13] PODS has been programmed with many of the leg-based and network optimization and forecasting methods that are commonly used by airlines in the real world. This section describes some of the RM methods that are used in our simulations.

Leg-based Revenue Management

In simple, point-to-point flight networks, *leg-based* revenue management methods can be used to determine fare class availability. Leg-based RM methods manage the inventory of seats on each flight leg individually and independently. The outputs of a leg-based RM method are *booking limits*—the number of seats to make available on each flight leg for each fare class in an airline's fare structure.

The Expected Marginal Seat Revenue (EMSRb; Belobaba (1992)) heuristic is a common leg-based revenue management method that is used by many airlines due to its robustness, relative simplicity, and good revenue performance. Given a fare structure and a forecast of demand for each fare class, EMSRb computes the *expected marginal seat revenue* (EMSR) for *protecting* a certain number n of seats for a given fare class and all fare classes above it. This value is equal to the probability of selling at least n seats in that class (and all classes above it) times the average fare value for that class (and all classes above it).

[13] We use 16 time frames in each sample in PODS, spanning a booking period of 63 days prior to the flight's departure.

The probability of selling n or more seats in a fare class depends on the RM system's forecast of future demand for each fare class on each flight leg. While there are many types of forecasting methods used in airline revenue management, one common approach is *pick-up forecasting*. Pick-up forecasting works by calculating an average number of bookings that were made on previous flight departures between the current time period and departure, and adding that estimate to the number of bookings already in hand (Gorin, 2000). This approach is typically called *standard forecasting*, and is often paired with leg-based RM methods like EMSRb.

While EMSRb and standard forecasting have been found to produce strong and robust revenue results (Belobaba, 2016), there are several cases in which the assumptions of leg-based RM are ill-suited for real-world applications. Since it controls availability on a flight leg-level, leg-based RM is not best suited for airlines with hub-and-spoke connecting networks. In these complex networks, airlines may wish to control availability on an origin-destination (OD) level, as opposed to a flight leg level. For instance, the airline may wish to open a lower fare class on a given flight leg only for passengers that are connecting to high-revenue international flights. In the 1990s, a set of methods for *network revenue management* emerged to allow airlines to control availability for each combination of OD and fare class.

EMSRb and standard forecasting typically assume that demand is independent between fare classes and follows a Gaussian distribution. This assumption of independent class demand is best suited for restricted fare structures, in which fare products are highly differentiated from class to class. In less-restricted fare structures where multiple fare classes may have exactly the same fare restrictions, demand for classes will be dependent on which classes are open at a given time. This violates EMSRb's assumption of independent class demand, and could lead to a phenomenon called *forecast spiral down* where the airline does not appropriately account for demand in more expensive fare classes (Cooper et al., 2006). Advanced forecasting and optimization methods which take into account the possibility of spiral down have also been developed to improve the performance of RM systems with less-restricted fare structures.

Network Revenue Management

Airlines can use network revenue management (also called *OD control*) to determine the availability of each fare class in each origin-destination market. One common heuristic for network RM is displacement-adjusted virtual nesting (DAVN) (Smith and Penn, 1988). While it is simpler than other network RM methods based on dynamic programming, DAVN is well-understood, robust, and remains in use in many real-world RM systems.

As described in Chapter 3, the main idea of DAVN is to generate a *displacement-adjusted revenue value* for each itinerary product/leg combination, taking into account the displacement costs of all other flight legs used by the itinerary product. To obtain these values, a linear program (LP) is used to optimize the allocation of seats for each itinerary product across the airline's network, based on a deterministic forecast of demand. Since this linear program assumes that demand is deterministic rather than stochastic, the allocation levels that result from the LP are not particularly useful in practice (Belobaba, 2016). Rather, the dual of the LP is used to produce shadow prices π_l (also called *bid prices*) which represent the marginal revenue value of the last seat on each flight leg l.[14]

For each flight leg, displacement-adjusted revenues are calculated for the itinerary products that use that leg. The revenue value of the itinerary product is adjusted by subtracting the bid price(s) of each other flight leg utilized by that product. If an itinerary product uses only a single flight leg, its displacement-adjusted revenue value is equal to its fare.

The displacement-adjusted itinerary products are then grouped into "virtual buckets" which are used for inventory management. For instance, the airline could decide to manage a virtual bucket containing all itinerary products with displacement-adjusted revenues between $100 and $200, another bucket for products with displacement-adjusted revenues between $200 and $300, and so forth. The availability of each virtual bucket on a flight leg (and the itinerary products therein) is then controlled using a leg-based RM heuristic like EMSRb.

With DAVN, the bid prices used to compute displacement-adjusted revenues generally come from a deterministic network linear program, and availability decisions for each itinerary product are based on a leg-based RM heuristic. There are also other network revenue management approaches that use a method called *additive bid price control* to determine itinerary product availability. With additive bid price control, the fare value of each itinerary product is compared with the sum of the bid prices on all of the itinerary's legs. If the fare f_k exceeds the sum of the bid prices $\sum_{l \in k} \pi_l$, the itinerary product is made available; otherwise, the itinerary product is not available.

Some RM methods generate leg bid prices using more complex optimization techniques than DAVN. For instance, the Probabilistic Bid Price (ProBP) method (Bratu, 1998) computes bid prices through an iterative algorithm that converges across the airline's flight network, and a method called unbucketed dynamic programming (UDP) computes bid prices using a dynamic program on individual legs (Belobaba, 2016). With any of these approaches, the vector of flight leg bid prices can be used to determine itinerary product availability using additive bid price control. In later sections, we will test the sensitivity of the performance of PFDynA when the airline uses each of these common network RM optimization methods.

[14] Note that if a flight is forecast to depart with empty seats, the marginal revenue of the last seat on the plane will be $0. In practice, it is common for many of an airline's flight legs to have a bid price of $0.

Forecasting and Optimization Methods for Less-restricted Fare Structures

As mentioned earlier, the assumption of independent class demand inherent in EMSRb and standard forecasting breaks down in environments in which fare structures are semi-restricted or unrestricted. In these environments, passengers will generally purchase the lowest-available itinerary product; there is no incentive to choose a more-expensive product to avoid an onerous fare restriction like a required Saturday-night stay. Using an RM method that assumes independent class demand will ignore the possibility of buy-down (or sell-up) in these fare structures.

Various next-generation forecasting and optimization methods have been proposed which are more appropriate for less-restricted fare structures. One approach, called *hybrid forecasting* (Boyd and Kallesen, 2004), works by separating bookings into two categories depending on whether the passenger booked in the lowest-available fare class. Customers that purchase higher-value fare classes when a less-expensive class is available are classified as *product-oriented demand*. The assumption of independent class demand is more appropriate for these customers since they appear to make their decisions based on the attributes of each itinerary product. Standard forecasting approaches can be used to account for these bookings.

For customers that book in the lowest available class (called *price-oriented demand*), there is a possibility that these customers would have sold up to a more-expensive class if the less-expensive class had been closed. Forecasts for this type of demand are adjusted to account for the possibility of sell-up. This adjustment process is based on the airline's estimates of the probability of sell-up from each class to the next. Adding together the product-oriented and price-oriented demand forecasts yields the hybrid forecast for that class.

Hybrid forecasting is often paired with an optimization technique called *fare adjustment* (Fiig et al., 2010) that is appropriate for less-restricted fare structures. Fare adjustment also acknowledges that customers may sell-up and buy-down amongst various classes in a less-restricted fare structure. Opening a less-expensive class may make it more likely that a customer will make a purchase, but it also could lead to buy-down by customers that would also have been willing to buy at a higher price point.

Fare adjustment uses assumed sell-up rates to compute the change in marginal revenue associated from opening each fare class and to adjust the fare values associated with each class in the RM optimizer. Fiig et al. (2010) and others have shown that this marginal revenue transformation can preserve the assumption of independent class demand that is made by traditional RM optimization methods, even with less-restricted fare structures.

Each of the revenue management optimization and forecasting methods described in this section have been programmed and tested in PODS simulations over the past two decades.

PODS simulations have found that a combination of hybrid forecasting and fare adjustment (HF/FA) with network RM heuristics like DAVN and ProBP can result in revenue gains of about 1 - 2 % over leg-based RM methods with standard forecasting. When combined with hybrid forecasting and fare adjustment, revenue gains of up to 3 - 4% are possible over traditional leg-based methods with standard forecasting (Belobaba, 2016).

4.1.4 PODS networks and fare structures

PODS can be configured to simulate a variety of airline networks and competitive environments. Here, we discuss the networks that will be used for the simulations in this chapter.

Network A1ONE

Network A1ONE is a simple network that consists of one origin-destination market (AAA to BBB) and one airline (AL1). The airline operates a single nonstop flight in the market, which is shown in Figure 4.4.

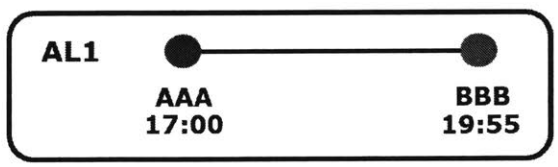

Figure 4.4: Structure of PODS Network A1ONE

Class	Fare	Restrictions			
		Adv. Purch.	R1	R2	R3
1	$500	N/A	No	No	No
2	$390	3 days	No	No	Yes
3	$295	7 days	No	Yes	Yes
4	$200	10 days	Yes	No	Yes
5	$160	14 days	Yes	Yes	No
6	$125	21 days	Yes	Yes	Yes

Table 4.1: Restricted fare structure for PODS Network A1ONE

The airline in Network A1ONE uses a traditional, restricted fare structure with six fare classes.[15] This fare structure is shown in Table 4.1, and is typical of the types of fare structures used by legacy carriers in U.S. domestic markets before recent periods of fare simplification. With this fare structure, the network is calibrated such that the airline obtains

[15] Note that since there is only a single nonstop flight in this network, "fare class" is synonymous with "itinerary product" for this network.

an average load factor of 83.3% when the airline uses EMSRb as its revenue management method. We call this scenario "Medium Demand." In this base case, the majority of passengers that book with the airline are leisure passengers (60.2%) and the remainder are business passengers (39.8%).

For some tests, a less-restricted version of the fare structure is tested, as shown in Table 4.2. The less-restricted fare structure removes advance purchase requirements for all fare classes, as well as the onerous restriction R1. Some fare classes retain restrictions R2 and R3. However, note that in terms of restrictions, fare classes 3 through 6 are identical. This fare structure is similar to the those used by many carriers in Europe, as well as increasingly by low-cost and ultra-low-cost carriers in the United States.

Class	Fare	Restrictions			
		Adv. Purch.	R1	R2	R3
1	$500	N/A	No	No	No
2	$390	N/A	No	Yes	No
3	$295	N/A	No	Yes	Yes
4	$200	N/A	No	Yes	Yes
5	$160	N/A	No	Yes	Yes
6	$125	N/A	No	Yes	Yes

Table 4.2: Less-restricted fare structure for PODS Network A1ONE

Network A2TWO

Network A2TWO is a variant of A1ONE that adds an additional airline. As with A1ONE, A2TWO has a single origin-destination market from AAA to BBB. Two airlines serve this nonstop market with a single, identical non-stop flight each, as shown in Figure 4.5. A2TWO can be configured with either restricted or less-restricted fare structures that are identical to Network A1ONE. In a Medium Demand base case with restricted fare structures, the airlines in Network A2TWO have average load factors of about 83.3% when both airlines use EMSRb and standard forecasting, with a similar split of leisure and business passengers as Network A1ONE.

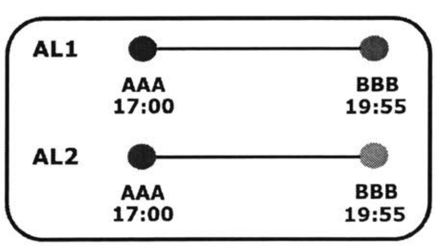

Figure 4.5: Structure of PODS Network A2TWO

Network U10

Network U10 is among the most complex networks that are commonly run in PODS. It contains four airlines, each of which operates a hub-and-spoke network structure with a single connecting hub. Each airline operates more than 100 non-stop flights each day to and from the connecting hub, and may also operate a select number of non-stop flights. All together, the four airlines in Network U10 operate 442 flight legs each day, serving a total of 572 origin-destination markets. A map of the network structure of Network U10 is shown in Figure 4.6.

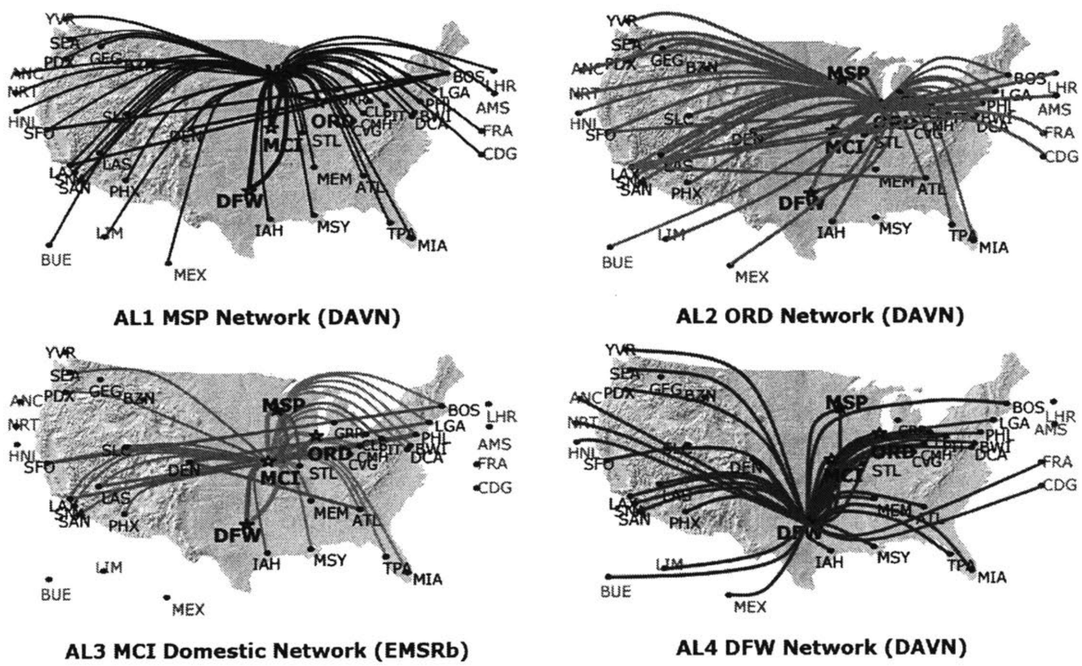

Figure 4.6: Structure of PODS Network U10

Airlines 1, 2, and 4 (AL1, AL2, and AL4) in Network U10 are designed to represent legacy carriers operating a domestic and international connecting hub operation through a single hub in the middle of the country.[16] Airline 3 (AL3) is designed as a low-cost domestic carrier, operating a mix of connecting hub service and point-to-point flights. Due to the hub-and-spoke design of their networks, AL1, AL2, and AL4 typically use a network revenue management system like DAVN. The low-cost carrier AL3 typically uses a leg-based EMSRb heuristic. This mix of revenue management methods is common in the real world, as some carriers retain leg-based RM systems while others have transitioned to OD control.

[16]Although the hubs used in Network U10 may overlap with hubs used by actual U.S. carriers, the network is designed generically and does not represent the network or fare structure of any particular carrier.

Demand in Network U10 is configured to flow from west to east across the network. For instance, a passenger may desire to travel from Denver (DEN) to Miami (MIA). Each of the airlines in Network U10 operate DEN – MIA connecting service through their respective hubs (DFW, MSP, MCI, and ORD). The DEN– MIA passenger will have a choice of up to 11 itineraries spread throughout the day from which to choose.

For each individual market, all of the airlines in Network U10 use identical fare structures, although the design of the fare structures varies from market to market. One of three sets of fare restrictions is used for the fare structures in each market, depending on whether the market is short-haul domestic or long-haul international, and on whether the low-cost carrier operates service in the market. Table 4.3 shows an example of the fare structure used by all four airlines operating connecting service in the DEN – MIA OD market.

Class	Fare	Restrictions			
		Adv. Purch.	R1	R2	R3
1	\$481	0 days	No	No	No
2	\$307	0 days	No	Yes	No
3	\$256	3 days	No	Yes	No
4	\$204	7 days	No	Yes	No
5	\$183	7 days	No	Yes	No
6	\$161	7 days	No	Yes	No
7	\$153	14 days	No	Yes	No
8	\$145	14 days	No	Yes	No
9	\$138	14 days	No	Yes	No
10	\$130	21 days	No	Yes	No

Table 4.3: Fare structure for DEN – MIA market in PODS Network U10

While Network U10 is closer to the complex competitive networks that airlines face in the real-world, the simplicity of Networks A1ONE and A2TWO makes it easier to identify the effects of dynamic pricing on revenues and booking patterns. By testing PFDynA heuristics in each of these three distinct networks, we can start to gain insights into the performance of dynamic pricing in both simple and complex environments, with and without competition.

4.2 Implementation of PFDynA in PODS

Before proceeding to the simulation results, we first describe how the PFDynA dynamic price adjustment heuristic was implemented in the PODS simulator. First, recall the general PFDynA heuristic that was introduced in Chapter 3 as Equation (10), which we reproduce below as Equation (11).

Heuristic 1: Probabilistic Fare-Based Dynamic Adjustment (PFDynA):

Suppose k is the lowest available fare product for an itinerary, and n is the lowest filed fare in the fare structure, regardless of availability. Let $\ell_k = f_{k+1} - f_k$ (or 0 if $k = n$), and $u_k = f_{k-1} - f_k$ (or $f_1 - f_2$ if $k = 1$). The PFDynA price adjustment Δ_k^* is:

$$\Delta_k^* = \arg \max_{\Delta_k \in [\ell_k, u_k]} \left((f_k + \Delta_k) - \sum_{l \in k} BP_l \right) \cdot \left(1 - \Phi^{-1} \left(\frac{(f_k + \Delta_k) - \mu_k}{\sigma_k} \right) \right) \tag{11}$$

where $\mu_k = QMULT \cdot f_n$ and $\sigma_k = \gamma_k \cdot \mu_k$. The offered price for product k is then:

$$f_k^* = f_k + \Delta_k^*$$

For the lowest-available fare product k of each itinerary, PFDynA computes a price adjustment Δ_k^* and a modified price f_k^*. If Δ_k^* is positive, $f_k^* > f_k$ and we say the price has been *incremented.* If Δ_k^* is negative, $f_k^* < f_k$ and we say the price has been *discounted.* If $\Delta_k^* = 0$, either because of a bound or because a particular fare class or customer type is not eligible for a price adjustment, the price of the itinerary product remains unchanged.

The implementation of PFDynA in PODS requires the following simulation inputs:

- A *filed fare structure* for each market, which describes the set of itinerary products that can be made available by the airline's revenue management system.
- A choice of *RM optimization and forecasting methods,* which are used to compute itinerary product availability.
- An *input Q-multiplier* $QMULT_w$ for each passenger type w. As discussed in Chapter 3, the Q-multiplier is the airline's estimate of the mean of the conditional WTP distribution for a given passenger type, expressed as a ratio of the lowest filed fare in the market. The conditional WTP distribution is assumed to follow a Normal distribution for each passenger type.[17]

  For example, if the lowest filed fare in the market $f_n = \$100$, and an airline estimates the mean of the conditional WTP distribution is \$150, the airline should use an input Q-multiplier of $QMULT = 1.5$. Q-multipliers will vary by passenger type; since business passengers are likely to have higher conditional WTPs than leisure passengers, a higher input Q-multiplier is used for business passengers than leisure passengers. The input Q-multiplier may also vary over the course of the booking process.

[17] Recall that the underlying distribution used to compute *maximum* WTP in PODS follows a negative exponential distribution. Therefore, the Normal distribution assumption of conditional WTP will always produce an inexact approximation of true WTP. This is a desired outcome, because airlines in the real world would also not have access to the exact underlying distribution of passenger WTP.

- A *input coefficient of variation* γ_w for each passenger type w, which is used to construct the standard deviation of the Gaussian (Normal) distribution of conditional WTP.
- The *customer types* (leisure; business) that the airline decides will be eligible for increments and/or discounts.
- A *customer type identification accuracy* α. With $\alpha = 1$, the airline can perfectly identify all requests by the customer type that made the request. With $\alpha < 1$, the airline will correctly segment the request with probability α and incorrectly segment the request (i.e., as the opposite type) with probability $1 - \alpha$.

With these parameters selected, PFDynA is implemented in PODS as follows:

1. The RM system selects the assortment S of itinerary products to make available.
2. When a customer arrives to book, the airline correctly perceives the customer's type with probability α.
3. If the customer's perceived type is eligible for a price adjustment, Δ_k^* is computed[18] as described above for the lowest-available fare class for each itinerary. The price modification is a function of the input Q-multiplier for the identified passenger type, the input coefficient of variation for that passenger type, the price of the lowest filed fare in the market, and the sum of the bid prices of all legs used by the itinerary product.
4. The price of the lowest-available fare class for each itinerary is changed to $f_k^* = f_k + \Delta_k^*$.
5. The customer constructs a set of affordable itinerary products $S' = \{x \in S : f_x \leq \theta\}$, where θ is the customer's maximum WTP. The customer's choice set creation comes after PFDynA makes any applicable adjustments to the prices of itinerary products.
6. The customer deterministically selects the itinerary product in S' with the lowest perceived price. If S' is empty, there are no affordable itinerary products and the passenger decides to no-go.
7. If the customer purchases an itinerary product, the airline records the booking in its historical booking database. The booking is recorded in the itinerary product's original fare class, even if PFDynA has made a price adjustment.
8. If PFDynA adjustments have been made to any itinerary products, the airline resets the prices of those itinerary products to their filed fares for the next customer arrival.

[18] In PODS, a line search "squeezer" method is used to find the arg max of the Δ_k^* expression. To decrease the computational burden, the price adjustment is calculated for each passenger type for each itinerary (path) at the beginning of each PODS time frame and held constant for the entire time frame.

This process repeats for each customer arrival: the RM system decides which products to make available; the airline segments the request as leisure or business; PFDynA decides whether to increment or decrement the fare of the least-expensive itinerary product for each itinerary; the customer forms a choice set and decides whether to make a booking; and the airline records any applicable booking and resets the prices of each itinerary product to their filed fares. Note that the value of the price adjustment for business and leisure passengers are computed once at the beginning of each time frame, but the determination of whether or not to adjust the price occurs separately for each customer arrival.

4.2.1 RM system awareness of PFDynA

One of the benefits of PFDynA compared to other dynamic pricing approaches is that it can be directly integrated into existing revenue management systems. This is because PFDynA provides a price adjustment to the pre-filed prices of the products that would ordinarily be made available by the airline's RM system. Ideally, the RM system should be made aware of any adjustments that PFDynA makes to the pre-filed fares. Adjustments in input fares and fare ratios could lead to changes in the fare class availability decisions made by the RM system, as well as any bid prices exported by the RM system that are used to determine availability and compute PFDynA price adjustments.

In PODS, *RM system awareness* of PFDynA was implemented by maintaining a database of the average price paid for each itinerary product over the previous 26 samples. These average prices, which incorporate any PFDynA adjustments, were then fed into the RM optimizer to compute itinerary product availability in the next sample. This is similar to how an implementation of dynamic pricing would likely occur in the real world—a revenue accounting department would compile data on actual prices paid for each itinerary product and pass that information to the RM team to update their systems. Unless otherwise noted, tests of PFDynA in this chapter were conducted with RM awareness enabled.[19]

4.3 PFDynA in single-carrier Network A1ONE

In this section, we test the performance of the PFDynA heuristic in the single-airline, single-flight Network A1ONE. We begin by assuming airlines are able to identify the passenger type of each booking request with accuracy $\alpha = 1$. As discussed in Chapter 3, this segmentation could be made based on identifiable characteristics of the request (e.g. point-of-sale, search

[19]Simulations have shown that in most cases, the presence of RM awareness makes little difference in the performance of PFDynA heuristics. This is because, as we will see, relatively few passengers receive price adjustments when PFDynA is used, so the average fares paid for each itinerary product are usually close to the product's pre-filed fare. Since the difference between average fares paid and filed fares is small, RM awareness has little effect on availability decisions or bid prices.

parameters, duration and schedule of the trip, origin/destination of the trip, or other characteristics) or characteristics of the passenger (e.g. her frequent-flyer number, past travel history, past search history, demographics, or other identifiable information).

As introduced in Chapter 3, we will be testing three different variants of the PFDynA heuristic which target increments or discounts for particular customer segments:

- **Increments-Only PFDynA:** For customers belonging to the "business" segment, compute a PFDynA increment Δ_k^* using Equation (11), bounded below by 0 and above by $f_{k-1} - f_k$ (or $f_1 - f_2$ if $k = 1$). Do not adjust prices for leisure customers.
- **Discounts-Only PFDynA:** For customers belonging to the "leisure" segment, compute a PFDynA discount Δ_k^* using Equation (11), bounded above by 0 and below by $f_{k+1} - f_k$ (or 0 if $k = n$). Do not adjust prices for business customers.
- **Two-Way PFDynA:** Compute both PFDynA increments for business customers and PFDynA discounts for leisure customers, as described above. However, business customers are not eligible for discounts, and vice versa.

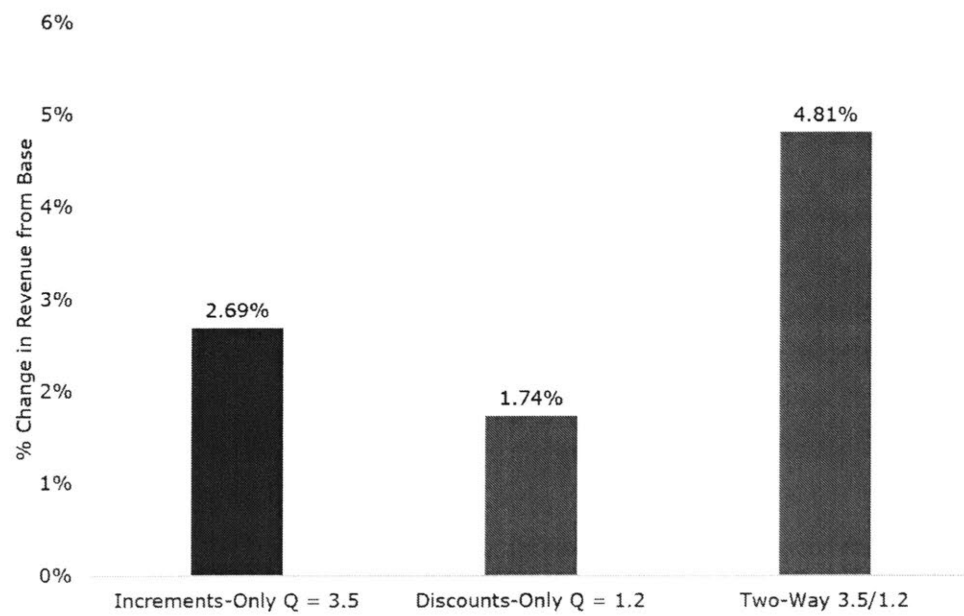

Figure 4.7: Percent change in revenue when AL1 uses PFDynA in Network A1ONE (Medium Demand)

Figure 4.7 presents the percent change in revenue from the base when Airline 1 (AL1) practices Increments-Only, Discounts-Only, or Two-Way PFDynA in Network A1ONE. The tests in Figure 4.7 assume a Medium Demand scenario in which AL1 uses EMSRb[20] with virtual nesting, standard forecasting, a restricted fare structure, input Q-multipliers of 3.5 for business passengers and 1.2 for leisure passengers, an input γ of 0.3, and has 100% passenger type segmentation accuracy. The effects of changes to these base case assumptions on the performance of PFDynA are discussed in detail in the subsequent sections.

With this combination of assumptions in this network, providing increments for business passengers leads to a revenue gain of about 2.7% and giving discounts to leisure passengers leads to a revenue gain of about 1.7%. The Two-Way heuristic, which provides increments to business passengers and discounts to leisure passengers (but not vice versa), increases revenues by about 4.8%. This gain is greater than the sum of the two individual heuristics, suggesting some complementarities from providing both increments and discounts. Next, we explore each of these three heuristics in more detail.

4.3.1 Increments-Only PFDynA

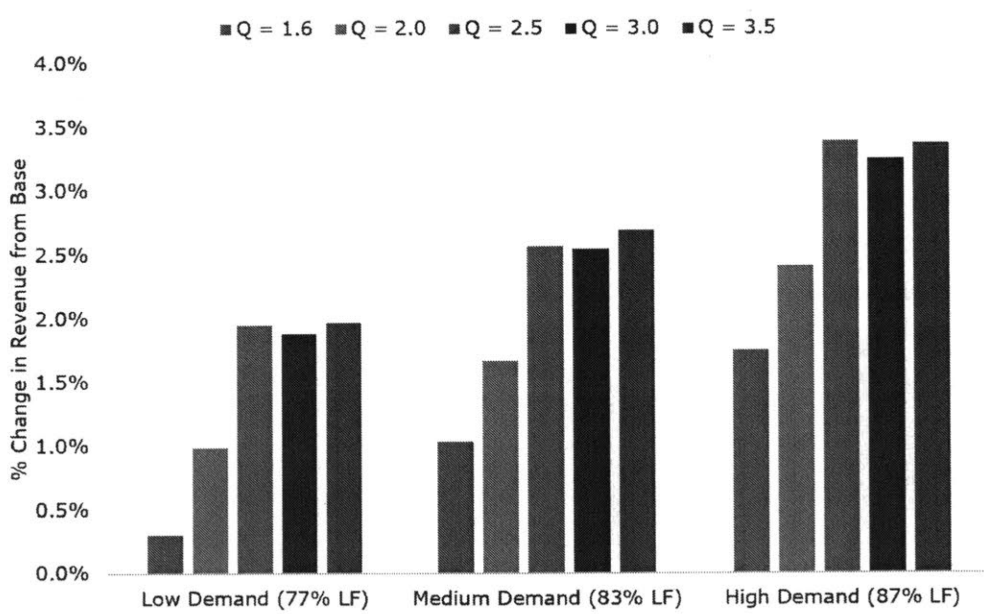

Figure 4.8: Percent change in revenue from base when AL1 uses Increments-Only PFDynA in Network A1ONE

In Figure 4.8, we show the revenue performance of Increments-Only PFDynA in Low,

[20] In practice, the airline in PODS used DAVN with virtual buckets constructed so that each bucket contained exactly one fare class. Since there is only one flight leg in Network A1ONE, this means that the optimization and availability outcomes of DAVN are identical to those of leg-based EMSRb. However, unlike EMSRb, DAVN produces a leg bid price that can be used for PFDynA calculations.

Medium, and High Demand scenarios in Network A1ONE. We also vary the input Q-multiplier from 1.6 to 3.5; within each demand scenario, the input Q-multiplier increases from left to right across the figure. Recall that the input Q-multiplier represents the airline's estimate of the mean conditional WTP for business-type passengers. If an airline uses a higher input Q-multiplier, this means that the airline thinks that business passengers have a higher conditional WTP on average.

In Figure 4.8, the performance of the Increments-Only PFDynA heuristic generally increases as the input Q-multiplier increases. This makes sense in a monopolistic network such as A1ONE. Since there are no other airlines in the network and only a single flight, business passengers' conditional WTP will be higher than in a competitive network where there are many possible itineraries. Higher input Q-multipliers cause PFDynA to apply increments more often than with lower input Q-multipliers, particularly when less-expensive, lower fare classes are open.

Incrementing prices for business passengers has several downstream effects. First, some business passengers facing an incremented price will choose not to fly, because the increment has made travel unaffordable. This will cause the airline to lose bookings and results in lower average load factors than the base case, as shown in the left panel of Figure 4.9. The load factor decrease is intensified when the airline uses a higher input Q-multiplier, because increments will be of greater value and will be provided more often.

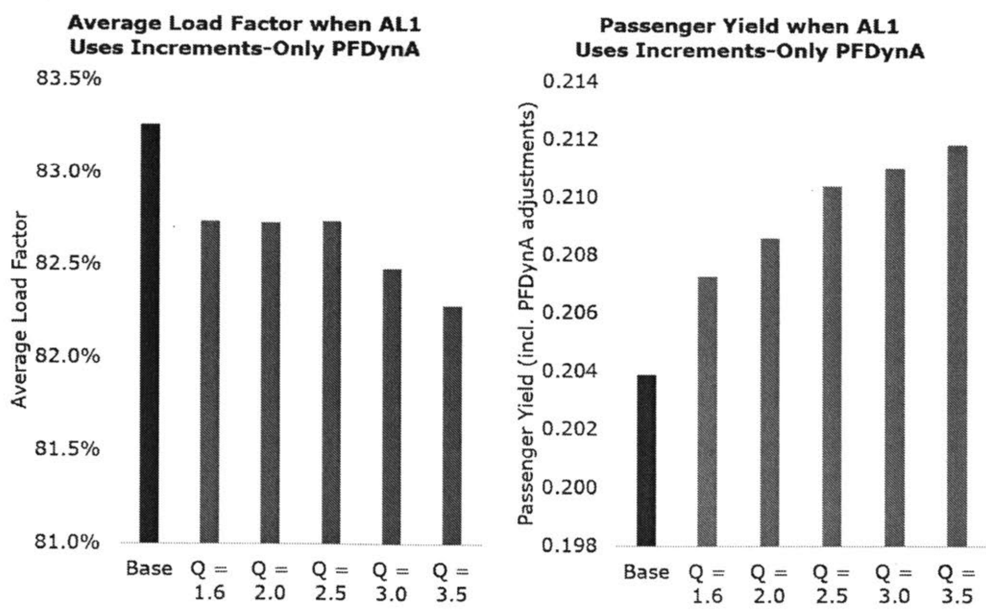

Figure 4.9: AL1 average load factor and yield when AL1 uses Increments-Only PFDynA in Network A1ONE (Medium Demand)

On the other hand, some business passengers will still choose to fly even at the incremented price. Some of these passengers will book the lowest-available class at the incremented price. Other business passengers facing an incremented price in the lowest-available class may choose to buy up to a higher fare class with fewer restrictions. This is particularly common when the airline uses a highly restricted fare structure, since buying up could allow the passenger to avoid onerous fare restrictions. These two behaviors increase the airline's yield (revenue per revenue passenger-mile), as shown in the right panel of Figure 4.9. As the input Q-multiplier increases, the yield also increases, since increments are provided more often.

This highlights the inherent tradeoff in Increments-Only PFDynA. By incrementing prices in certain situations for business passengers, these passengers will pay a higher fare on average, as shown in Figure 4.10. This figure shows the average fare paid by business passengers over the booking period of the flight. When the airline uses Increments-Only PFDynA, the average fare paid by business passengers increases, particularly in the period between 7 and 28 days prior to departure.

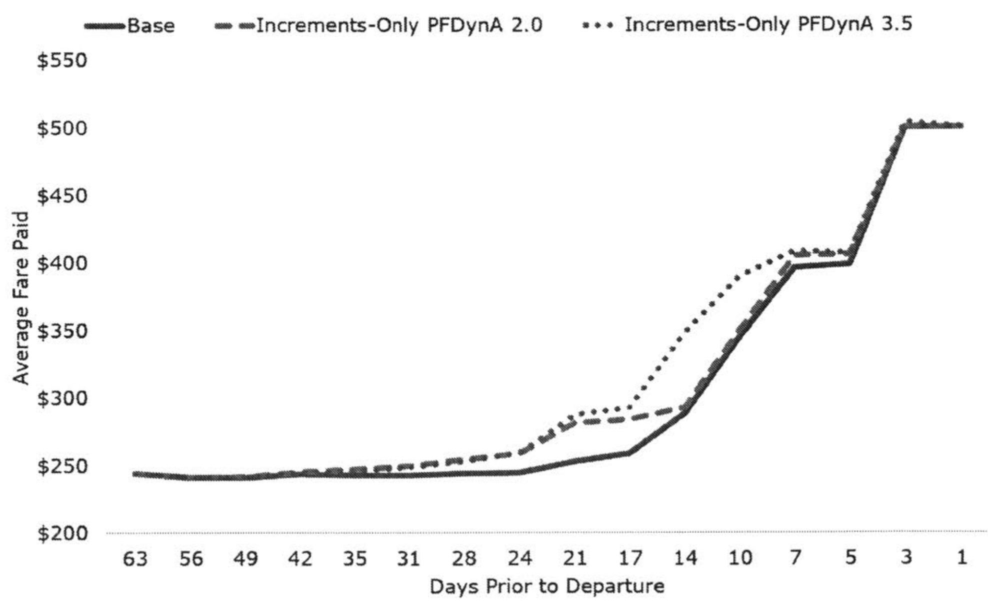

Figure 4.10: Average fare paid by business-type passengers when AL1 uses Increments-Only PFDynA in Network A1ONE with Q = 2.0 and Q = 3.5 (Medium Demand)

These increases in paid fares will cause the airline to see an increase in yield. However, some passengers may choose not to fly as a result of these increments, leading to a decrease in load factor. Since revenue increases overall when the heuristic is used in Network A1ONE, the gain in yield outweighs the loss in load factor in this simple network.

There is also another downstream effect of Increments-Only PFDynA that affects leisure passengers. Recall that some business passengers choose to buy-up to higher fare classes to avoid incremented prices in lower classes. This means that relative to the base, there will be more observations of bookings in higher classes when Increments-Only PFDynA is used. Over time, the RM forecaster will start to incorporate this increase in higher-class bookings into its forecasts. Facing higher forecasts in higher booking classes, the RM optimizer will start to protect more seats for these higher classes, and leave fewer seats remaining for the lower classes that are typically booked by leisure passengers.

This shift in bookings from lower classes to higher classes is shown in Figure 4.11, which displays the average number of bookings by business passengers and leisure passengers when the airline uses Increments-Only PFDynA. Compared to the base, bookings increase in the more expensive classes FC1 and FC2, while bookings decrease in FCs 3 – 6. While this shift in bookings towards higher fare classes is due in part to business passengers buying up to avoid increments, it is also partly due to the revenue management system closing down the least-expensive fare class (FC6) more often. We refer to this pattern as *forecast spiral-up.* As we will see throughout this chapter, forecast spiral-up is an important hallmark of PFDynA, whether it is providing increments or discounts.

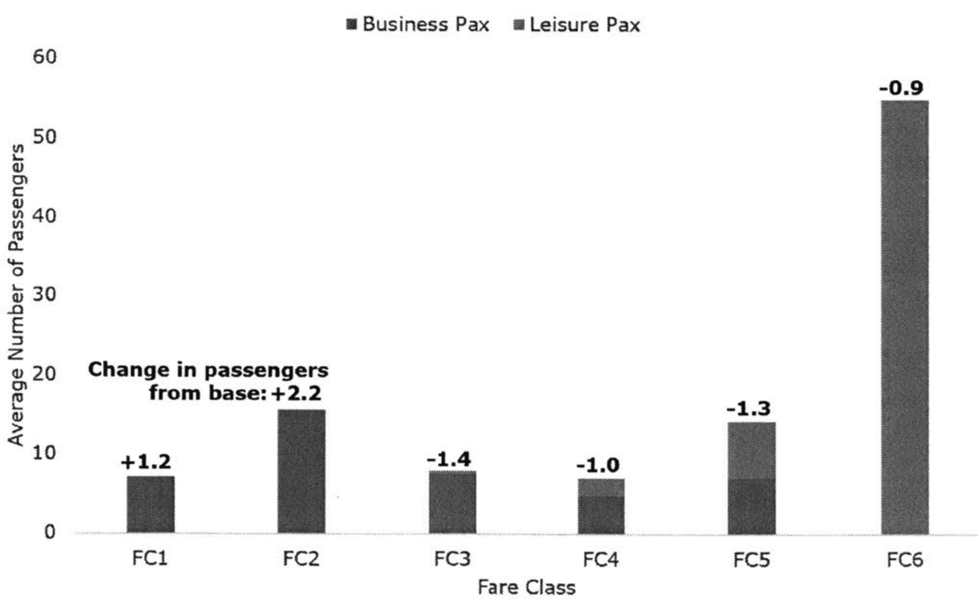

Figure 4.11: Average bookings by fare class and passenger type when AL1 uses Increments-Only PFDynA in Network A1ONE (Q = 3.5, Medium Demand)

4.3.2 Discounts-Only PFDynA

Figure 4.12 shows the results of using the Discounts-Only PFDynA heuristic in Network A1ONE in various demand scenarios and with various Q-multipliers. With this heuristic, leisure-type passengers are eligible for discounts as directed by Equation (11). In contrast to the incrementing heuristic, Discounts-Only PFDynA produces higher revenue results in this network when the airline uses lower input Q-multipliers, rather than higher values. Lower input Q-multipliers suggest that the airline has estimated that leisure passengers have relatively lower conditional WTPs, on average. With lower Q-multiplier values, the heuristic provides discounts more often. Also in contrast with Increments-Only PFDynA, which produced the highest percent revenue increase in the high-demand scenario, Discounts-Only PFDynA performs best in the low-demand scenario when there are more empty seats.

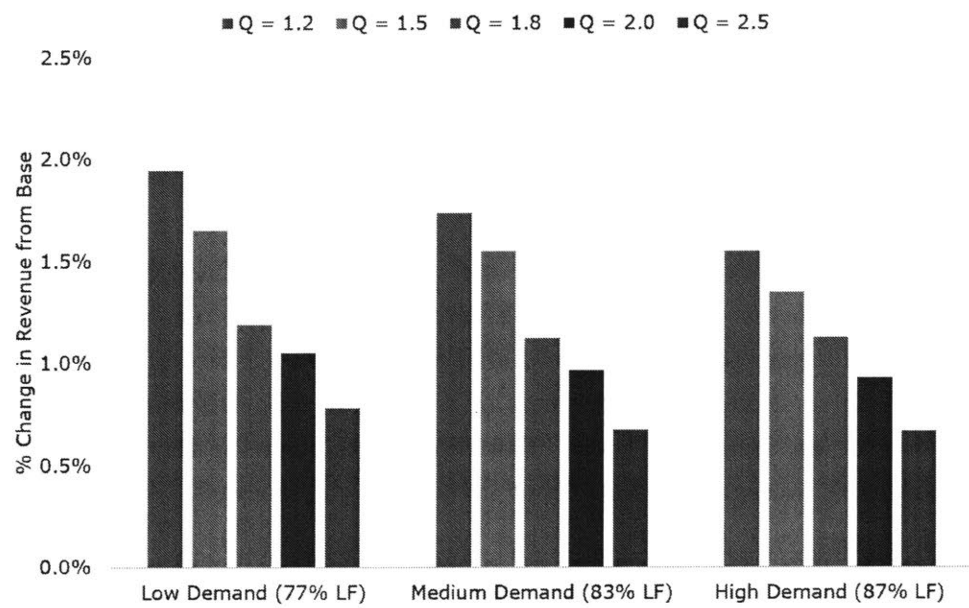

Figure 4.12: Percent change in revenue from base when AL1 uses Discounts-Only PFDynA in Network A1ONE

The primary purpose of providing discounts through the PFDynA heuristic is to stimulate demand from lower-WTP passengers who otherwise would not have traveled. As a result, we would hope to see an increase in the airline's average load factor when discounts are given. This is indeed the case; as shown in the left panel of Figure 4.13, the airline's load factor does increase when the Discounts-Only PFDynA heuristic is used. The increase in load factor is greatest when the airline uses the lowest Q-multiplier, because discounts are provided more often in this case.

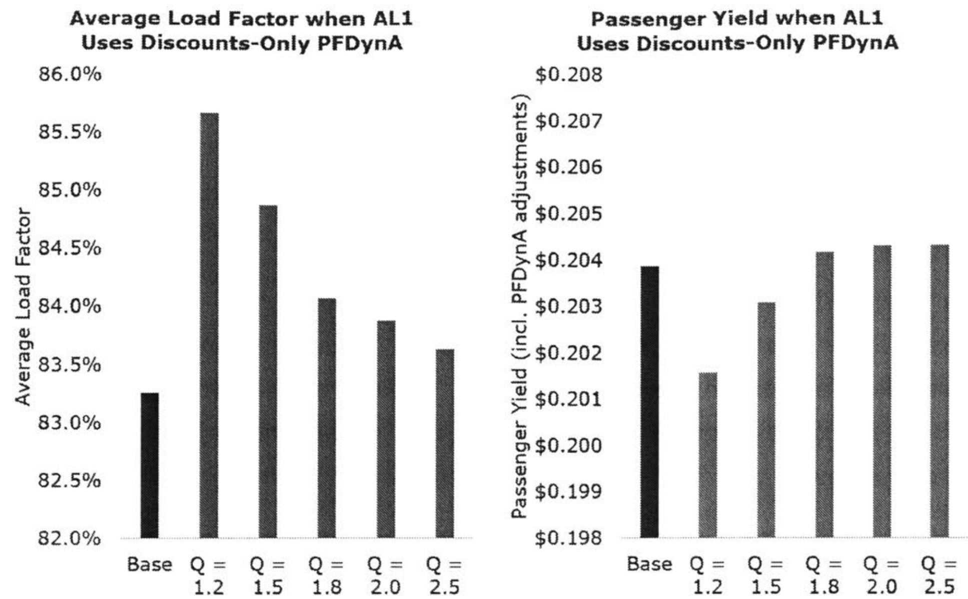

Figure 4.13: AL1 average load factor and yield when AL1 uses Discounts-Only PFDynA in Network A1ONE (Medium Demand)

More surprising is the result from the right panel of Figure 4.13. Since giving incremented prices resulted in an increase in yield (from Figure 4.9), one might reasonably expect that giving discounts to leisure passengers would always lead to a decrease in yield. Yet this is not the case. With some input Q-multipliers (e.g. 1.8, 2.0, 2.5), Discounts-Only PFDynA can actually lead to an *increase* in yield. This discounting heuristic therefore can increase not only the airline's average load factor, but its average yield as well.

The cause of this unexpected outcome is once again forecast spiral-up. When Discounts-Only PFDynA provides discounts to leisure passengers, it does so particularly in situations when the price that would be offered by the RM system is relatively high. This occurs when lower classes have been closed by the RM system, and only higher classes are open. For some leisure passengers traveling in these higher fare classes, the discount makes air travel affordable. Discounts-Only PFDynA therefore leads to an increase in leisure passengers booking in (discounted) higher fare classes.

Just as in the Increments-Only case, the RM forecaster will over time begin to incorporate this increase in bookings into its demand forecast for higher classes. Given higher forecast demand in higher classes, the RM optimizer will again begin to protect more seats for higher classes and reduce availability in lower classes. As shown in Figure 4.14, practicing Discounts-Only PFDynA causes the RM system to close the least-expensive fare class (FC6)

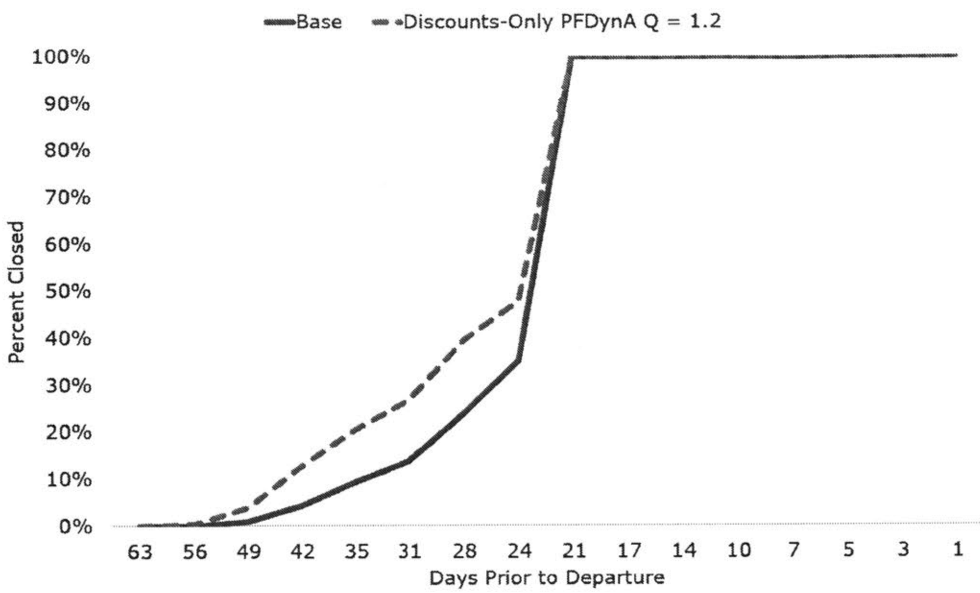

Figure 4.14: Percent of samples in which Fare Class 6 is closed by the RM system when AL1 uses Discounts-Only PFDynA in Network A1ONE (Q = 1.2, Medium Demand)

more often than in the base case.[21]

The higher closure rate of FC6 has several downstream effects on booking patterns. First, leisure passengers arriving early in the booking process may be less likely to find FC6 open. Some of these passengers (who could only afford FC6) may decide not to book, but others may decide to buy up to a higher fare class. Then, in the middle of the booking process when advance purchase requirements are starting to close down lower fare classes, leisure passengers start to receive discounts, leading to demand stimulation and new bookings in higher classes. Finally, at the end of the booking period, there is a slight displacement of business passengers due to the earlier increase in leisure bookings. Figure 4.15 shows this pattern of fewer early (up to 21 days out) leisure bookings due to FC6 closures, more mid-period (days 21 to 5) leisure bookings due to demand stimulation from PFDynA discounts, and slight displacement of late-arriving business passengers (within 5 days of departure).

Therefore, Discounts-Only PFDynA primarily increases revenue in the monopolistic Network A1ONE due to stimulation of demand that otherwise would not have been able to afford to fly. These passengers typically book in relatively high fare classes, leading to forecast spiral-up, a shift of the booking curve from FC6 towards higher fare classes, and increased closures of Fare Class 6.

[21]21 days prior to departure, FC6 is closed automatically due to the advance purchase restriction associated with the class.

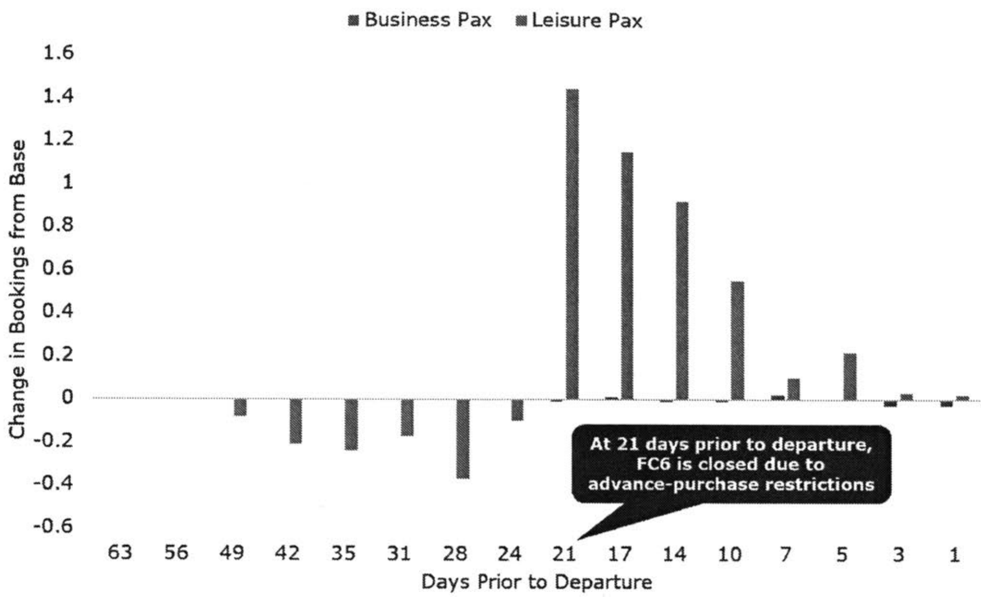

Figure 4.15: Change in bookings from base by passenger type during the booking period when AL1 uses Discounts-Only PFDynA in Network A1ONE (Q = 1.2, Medium Demand)

Note that that the Discounts-Only PFDynA heuristic in Equation 11 does not give discounts below the lowest filed fare in the market (i.e., fare class FC6). As will be discussed in Chapter 6, providing discounts below the lowest filed fare in the market's fare structure could increase the risk of a race to the bottom in a competitive market. Because no discounts are given in the lowest fare class, only a small portion of AL1's leisure passengers book with discounts when PFDynA is used, as shown in Figure 4.16. In total, 16.5% of AL1's leisure passengers book with a PFDynA discount in this network with Q = 1.2, representing 10.2% of AL1's total passengers.

4.3.3 Two-Way PFDynA

Two-Way PFDynA is a combination of the two previous heuristics—business passengers are eligible for increments and leisure passengers are eligible for discounts, but not vice versa.[22] Figure 4.17 shows the revenue gains from practicing Two-Way PFDynA with two different combinations of input Q-multipliers: 2.0 for business/1.5 for leisure (which assumes that conditional WTP between the two passenger types is relatively similar) and 3.5 for business/1.2 for leisure (which assumes that there is a greater gap in conditional WTP between the types).

[22] We will later test an unsegmented version of PFDynA in which it is not possible to identify customer types, and both types are eligible for both increments and discounts.

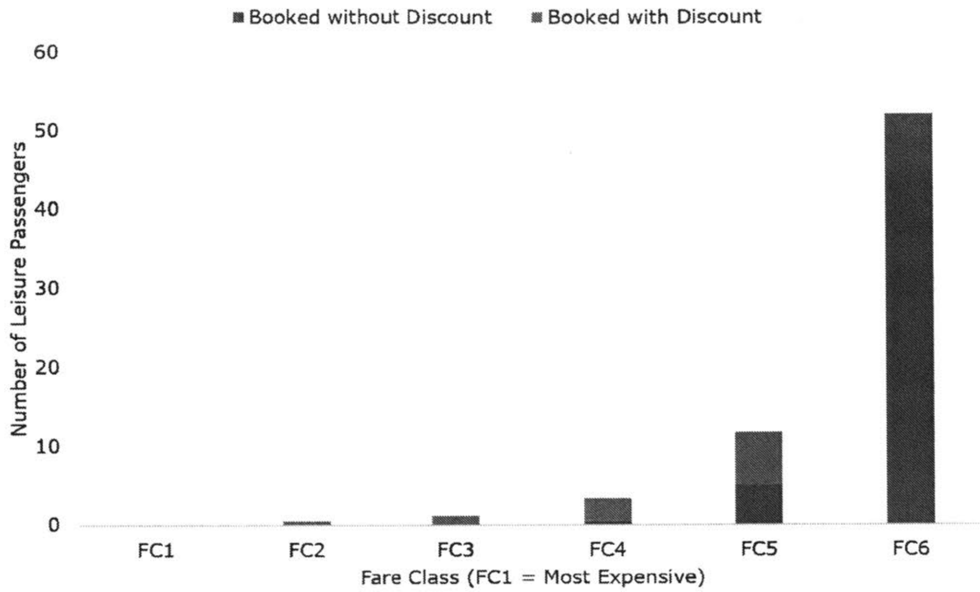

Figure 4.16: AL1 leisure passengers booking with and without PFDynA discounts by fare class when AL1 uses Discounts-Only PFDynA in Network A1ONE (Q = 1.2, Medium Demand)

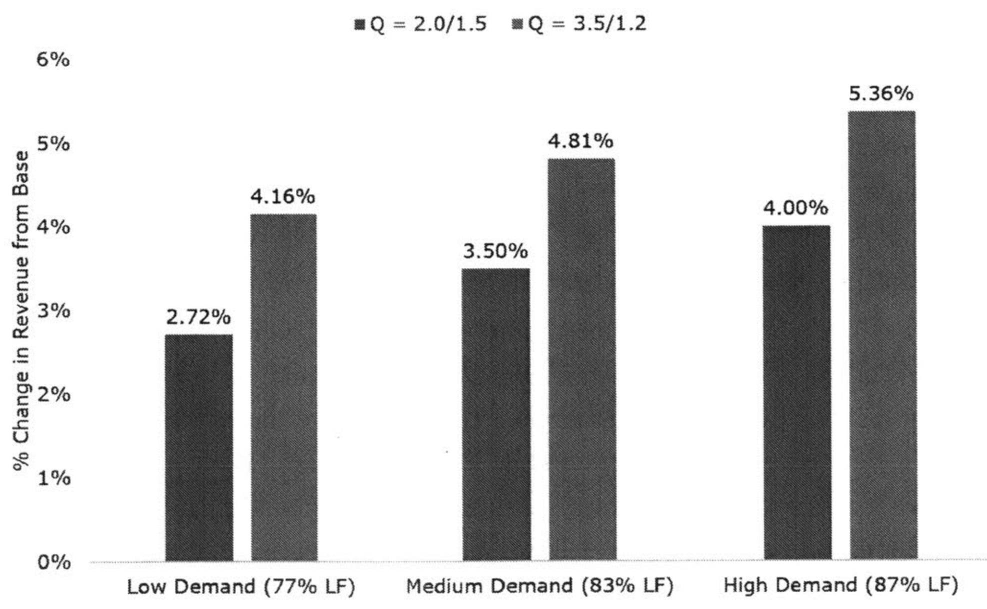

Figure 4.17: Percent change in revenue from base when AL1 uses Two-Way PFDynA in Network A1ONE

In Medium Demand, the 3.5/1.2 input Q-multiplier combination leads to a revenue gain of about 4.8%. This exceeds the sum of the 2.7% gain from the Increments-Only heuristic in A1ONE (Figure 4.8) and the 1.7% gain from the Discounts-Only heuristic (Figure 4.12). This suggests that the forecast spiral-up that each of these methods produce individually can compound when both increments and discounts are provided. Indeed, practicing Two-Way PFDynA leads to greater closure rates of the least-expensive fare class FC6 than if Increments-Only or Discounts-Only PFDynA is used alone, as shown in Figure 4.18.

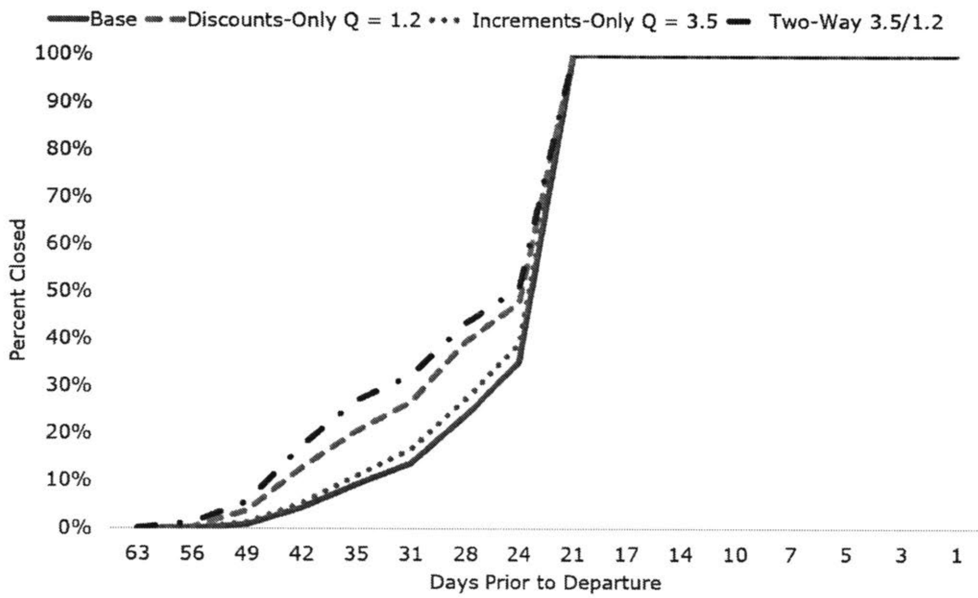

Figure 4.18: Percent of departures in which Fare Class 6 is closed by the RM system when AL1 uses three different PFDynA heuristics in Network A1ONE (Medium Demand)

The prices for both passenger types throughout the booking period, shown in Figure 4.19, also compounds the effects of the two individual heuristics. Business passengers pay higher fares on average than leisure passengers, since they are more likely to choose to purchase a higher fare class to avoid onerous fare restrictions. Using Two-Way PFDynA leads to an increase in average fares paid by business passengers throughout the booking period, as a result of increments, buy-up to higher classes, and forecast spiral-up.

Since leisure passengers receive discounts from Two-Way PFDynA, they pay lower fares closer to departure when the lowest available fare is relatively expensive. However, lower fares also result in an increase in booking volumes for leisure passengers. In Medium Demand, the airline's load factor increases from 83.3% to 84.6% when Two-Way PFDynA 3.5/1.2 is used, and passenger yield also increases from $0.204 per RPM to $0.210 per RPM.

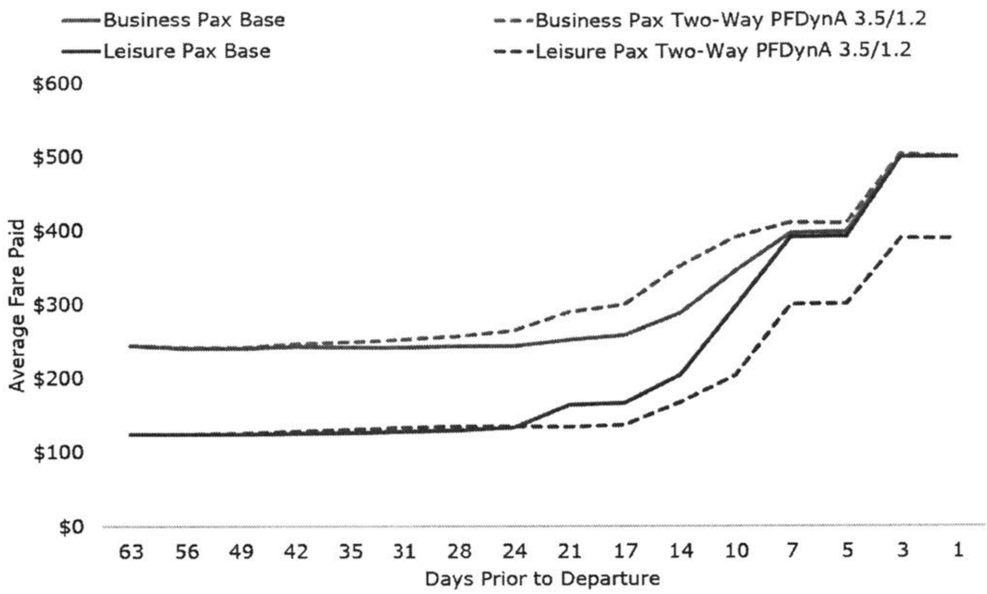

Figure 4.19: Average fare paid by business and leisure passengers when AL1 uses Two-Way PFDynA in Network A1ONE (Medium Demand)

While Two-Way PFDynA can lead to large gains in revenue in this network, relatively few passengers actually book with PFDynA-adjusted fares. With Two-Way PFDynA Q = 3.5/1.2, 4.1% of business passengers book with a PFDynA increment and 18.2% of leisure passengers book with a PFDynA discount. This is because PFDynA only applies increments and discounts in particular situations—namely, increments are provided to business passengers when the lowest-available fare is relatively low, and discounts are provided to leisure passengers when the lowest-available fare is relatively high.

4.3.4 Sensitivity analysis: Segmentation inaccuracy

The tests in the previous section assume that the airline is correctly able to segment the customer types of incoming booking requests with perfect accuracy. In the real world, it is unlikely that any airline would be able to perfectly segment every request, either due to a lack of available information, a lack of technological sophistication, or actions by customers to mask their true customer types. We therefore test the performance of PFDynA when airlines makes errors in segmenting customer types and, separately, when customers cannot be segmented at all.

First, we test the sensitivity of the segmented PFDynA heuristics when the airline can segment customers, but sometimes makes mistakes in its segmentation. These mistakes could be due to the airline's error in identifying characteristics of leisure and business requests, or due to customers "gaming" the system and tricking the airline into thinking they are the opposite type. For example, with a segmentation accuracy level of $\alpha = 0.8$, the airline will think that 20% of leisure requests are actually business passengers, and vice versa. Depending on the PFDynA method used, this will cause the heuristic to provide increments to some leisure customers and/or discounts to some business customers.

A priori, we should expect that lower levels of segmentation accuracy will affect the performance of PFDynA. Misidentifying a leisure customer as a business customer will lead to providing incremented prices to some leisure customers with relatively low WTPs. This will cause some leisure customers to no-go who otherwise would have booked with the airline. Conversely, misidentifying business customers as leisure customers could cause the airline to provide unnecessary discounts to business customers who would have booked at the normal filed fare. This could also reduce the airline's revenue.

Figure 4.20 shows the effects of using PFDynA heuristics with Q = 3.5 for business passengers and Q = 1.2 for leisure passengers with various segmentation accuracies. As expected, lower segmentation accuracy reduces the revenue performance of the heuristic. However, in this network, revenue gains are still possible even with moderate levels of inaccuracy. Revenue gains result from Discounts-Only PFDynA with accuracy levels as low as 80%, and from Increments-Only PFDynA and Two-Way PFDynA with accuracy levels of as low as 70%.

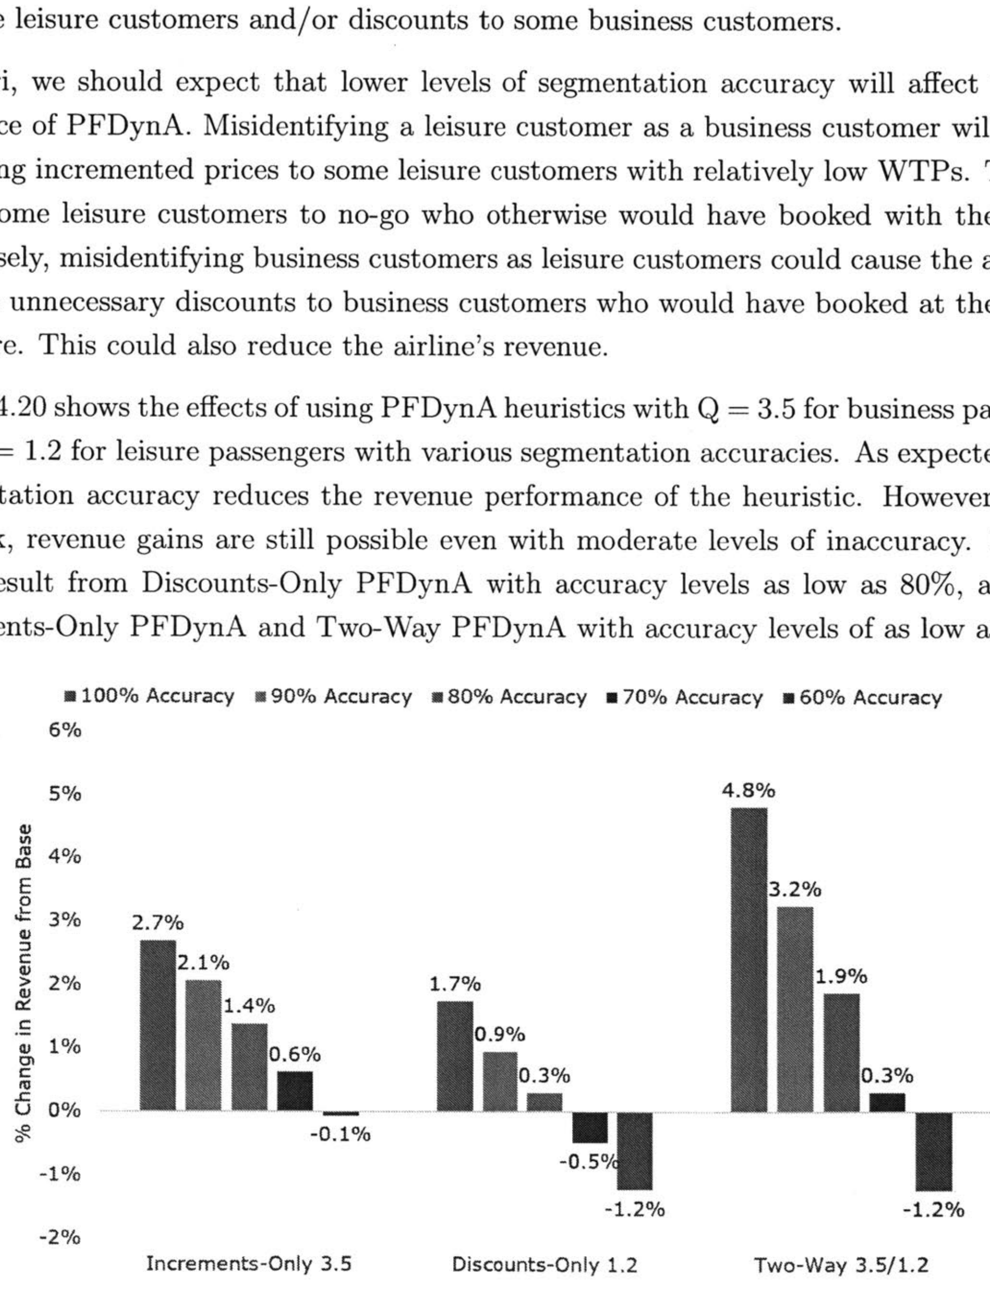

Figure 4.20: Percent change in revenue from base when AL1 uses three different PFDynA heuristics in Network A1ONE with segmentation inaccuracy (Medium Demand)

When segmentation mistakes are made with Increments-Only and Two-Way PFDynA, AL1's load factor decreases since increments are mistakenly given to leisure passengers, some of whom will choose to no-go. However, the airline's yield increases. With Discounts-Only PFDynA, segmentation mistakes lead to higher load factors than the base case but lower yields, since business customers receive unnecessary discounts.

In this monopolistic network, mistaking leisure passengers as business passengers is less costly because there is no other airline that can capture leisure passengers who are given increments. As we will see in tests in more competitive networks, mistaking leisure passengers as business passengers can be more costly than the opposite error, because providing increments to leisure passengers will cause some of them to book with other airlines.

4.3.5 Sensitivity analysis: Unsegmented PFDynA

In the tests shown in Figure 4.20, we assume that airlines are still able to accurately segment booking requests the majority of the time, and are able to provide different prices to requests that meet different segmentation criteria. Yet some airlines may be unable to segment booking requests at all. This could be due to a lack of segmentation technology, a lack of information about booking requests from various distribution channels, or an inability or unwillingness to provide different prices to different demand segments.

Next, we test the performance of unsegmented versions of the PFDynA heuristics. Airlines with Unsegmented PFDynA are unable to differentiate leisure requests from business requests. This means that (1) all passengers are eligible for increments *and* discounts, regardless of their customer type, and (2) a single input Q-multiplier is used for all customer types in each time frame. The Unsegmented PFDynA heuristics are summarized in the box below.

- **Unsegmented Increments-Only PFDynA:** All customers are eligible for increments, which are computed using Equation (11). A single input Q-multiplier, which could change over time, is used for all customers.
- **Unsegmented Discounts-Only PFDynA:** All customers are eligible for discounts, which are computed using Equation (11). A single input Q-multiplier, which could change over time, is used for all customers.
- **Unsegmented Two-Way PFDynA:** All customers are eligible for increments and discounts, as directed by Equation (11). A single input Q-multiplier, which could change over time, is used for all customers.

Since the input Q-multiplier represents the airline's estimate of mean conditional WTP, it seems reasonable that this input parameter should increase over time as changes occur in the booking mix. The beginning of the booking process will mostly be populated by leisure customers, and business customers are more likely to arrive later in the booking process. In our tests of Unsegmented PFDynA, we test both constant and linearly increasing input Q-multiplier schemes.

In Figure 4.21, we show the revenue performance of segmented and unsegmented versions of Increments-Only PFDynA. Recall that in the segmented version, only business passengers are eligible for increments, whereas in the unsegmented version, all customers are eligible for increments. The results show that in Network A1ONE, giving increments to all passenger types leads to significant revenue losses. With higher Q-multiplier values, revenue losses of about 3.5% result from unsegmented Increments-Only PFDynA. A linear Q-multiplier scheme which increases from Q = 1.2 at the beginning of the booking period to Q = 3.5 at the end of the booking period provides some improvement over most constant schemes, but still results in a revenue loss of -1.6%.

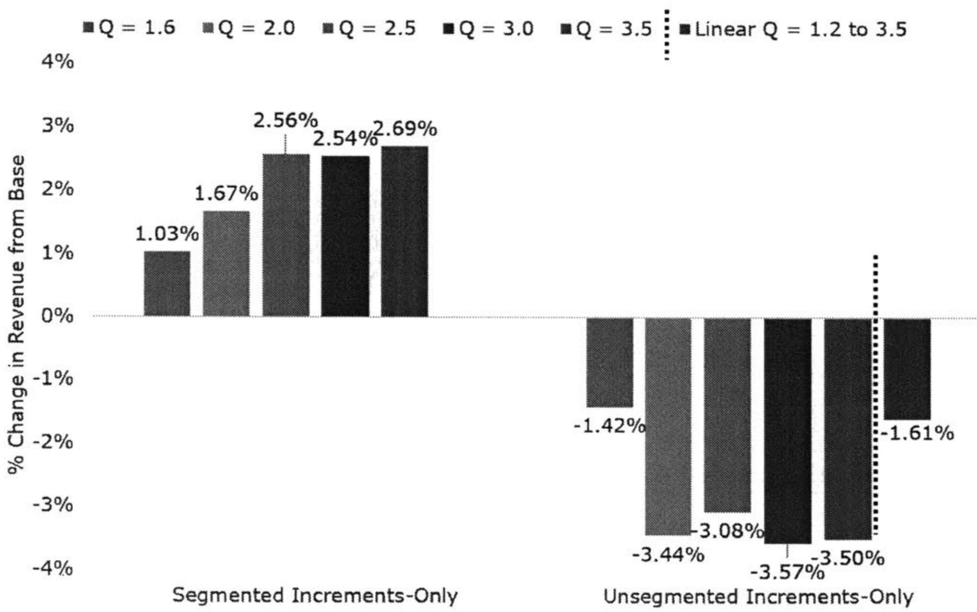

Figure 4.21: Percent change in revenue from base when AL1 uses Segmented and Unsegmented Increments-Only PFDynA in Network A1ONE (Medium Demand)

Unsegmented Increments-Only PFDynA leads to revenue losses because of the increments it provides to leisure passengers. These passengers are willing to pay less on average than business passengers, and are more likely to be unable to afford the flight if the price is incremented. Unsegmented Increments-Only PFDynA leads to a significant loss in load factor as leisure passengers decide to no-go. With Q = 3.5, load factors decrease from 83.3% in the base case to just 65.7% when the unsegmented heuristic is used. While passenger yield increases by over 22% as a result of the increments, this increase is not enough to make up for the loss in passengers who decide to no-go as a result of the increments.

Figure 4.22 shows a similar comparison between segmented and unsegmented versions of Discounts-Only PFDynA. In the segmented version, only leisure passengers are eligible for discounts, while in the unsegmented version all passengers are eligible for discounts. As in the Increments-Only case, the unsegmented version of Discounts-Only PFDynA leads to significant revenue losses of up to 4.2%. The linearly-increasing Q-multiplier scheme again improves on the constant schemes, but even the linear scheme results in a 0.6% revenue loss relative to the base case.

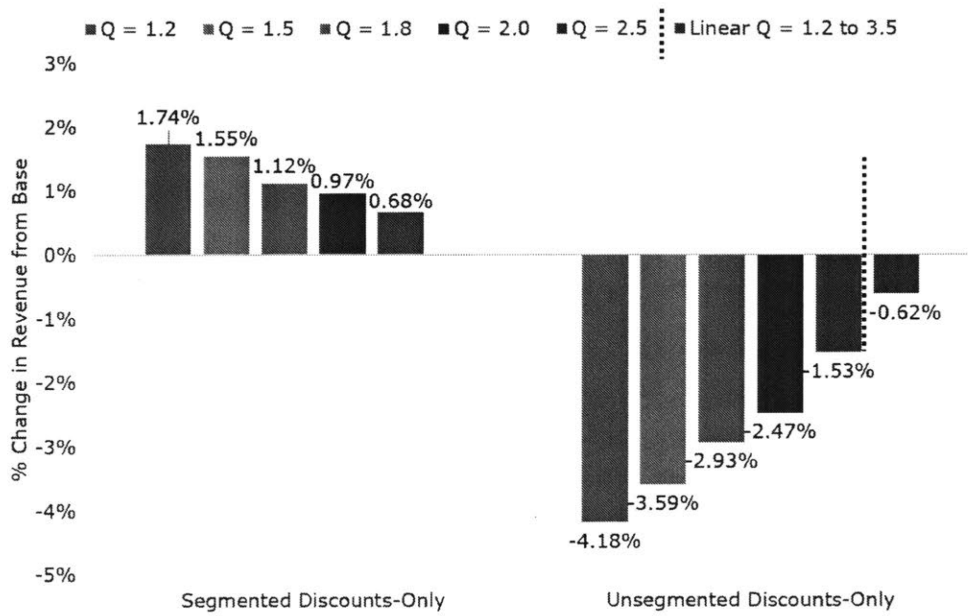

Figure 4.22: Percent change in revenue from base when AL1 uses Segmented and Unsegmented Discounts-Only PFDynA in Network A1ONE (Medium Demand)

Unsegmented Discounts-Only PFDynA leads to revenue losses because it gives too many unnecessary discounts to business passengers, who otherwise would have booked without a discount. Load factors increase over the base case; with a constant Q = 1.2, AL1's average load factor reaches 87.3% compared to a base load factor of 83.3%. This load factor increase as a result of discounting comes with a corresponding 8.6% decrease in yield. In this scenario, 29% of AL1's passengers book with a discount, including 51% of business passengers (who are more likely to book in the higher classes in which discounts are applied). There is not enough demand stimulation to make up for the discounting, causing revenues to fall.

Finally, it is not surprising that since Unsegmented Increments-Only and Discounts-Only PFDynA both resulted in revenue losses, the unsegmented version of Two-Way PFDynA (in which all passenger types are eligible for increments and discounts) also reduces revenue. Figure 4.23 shows the revenue performance of the unsegmented Two-Way PFDynA heuristic. While the linear Q-multiplier scheme works better than the constant schemes, even the linear scheme results in a 2.1% reduction in revenue. The unsegmented Two-Way heuristic provides both unnecessary discounts to business passengers and unnecessary higher prices to leisure passengers, who then decide not to fly more often. This latter effect appears to dominate the former. Except when Q = 1.2, using Unsegmented Two-Way PFDynA leads to lower load factors and higher yields than the base case, as leisure customers receive increments and decide to no-go in greater numbers.

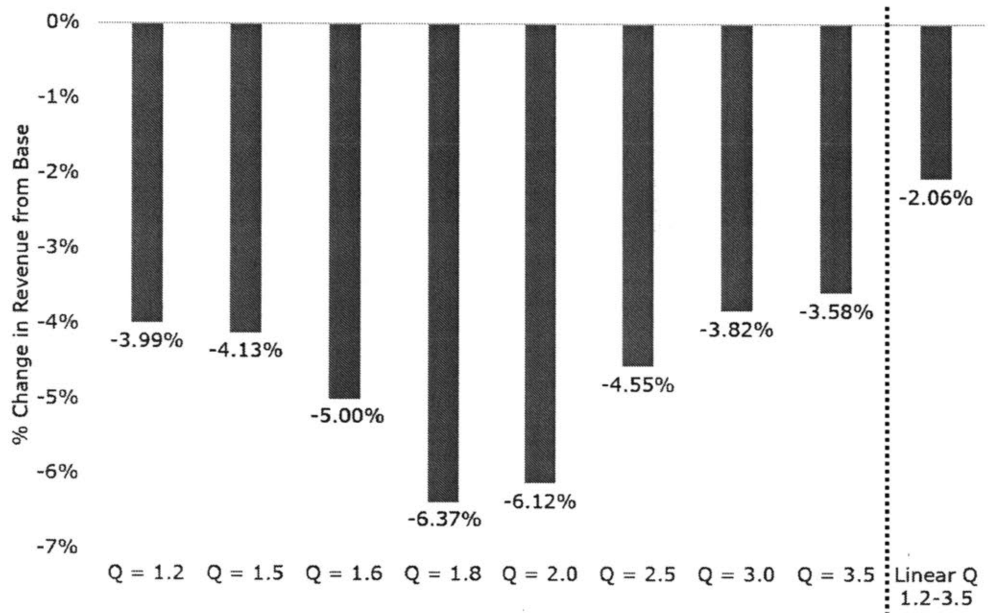

Figure 4.23: Percent change in revenue from base when AL1 uses Unsegmented Two-Way PFDynA in Network A1ONE (Medium Demand)

The results from this section suggest that in Network A1ONE, segmentation is critical to the success of the PFDynA heuristic. The segmentation need not be perfect; even an airline that makes mistakes up to 20-30% of the time can still see revenue gains from PFDynA in this network. Segmentation improves the performance of the heuristic by reducing the chance that a business passenger will receive an unnecessary discount, or that a leisure passenger will decide not to fly as a result of an increment. That is, in the segmented version of the heuristic, price adjustments are made less frequently and targeted to the situations in which they are most likely to be effective—namely, increments for business passengers when prices are relatively low, and discounts for leisure passengers when prices are relatively high.

4.3.6 Sensitivity analysis: PFDynA with a less-restricted fare structure

We close our analysis of PFDynA in Network A1ONE by examining a case when the airline uses a less-restricted fare structure (shown earlier in Table 4.2). This less-restricted fare structure requires advanced forecasting and optimization methods to confront buy-down, even in the base case. We use a combination of hybrid forecasting and fare adjustment[23] to optimize RM availability for this fare structure. The base case in Medium Demand has an average load factor of about 83.2%, but Figure 4.24 shows that the fare class mix has deteriorated compared to the restricted fare structure with more bookings occurring in lower classes and fewer business passengers buying-up to higher classes. The revenue of the less-restricted base case is 7.8% lower than the restricted base case due to this spiral down.

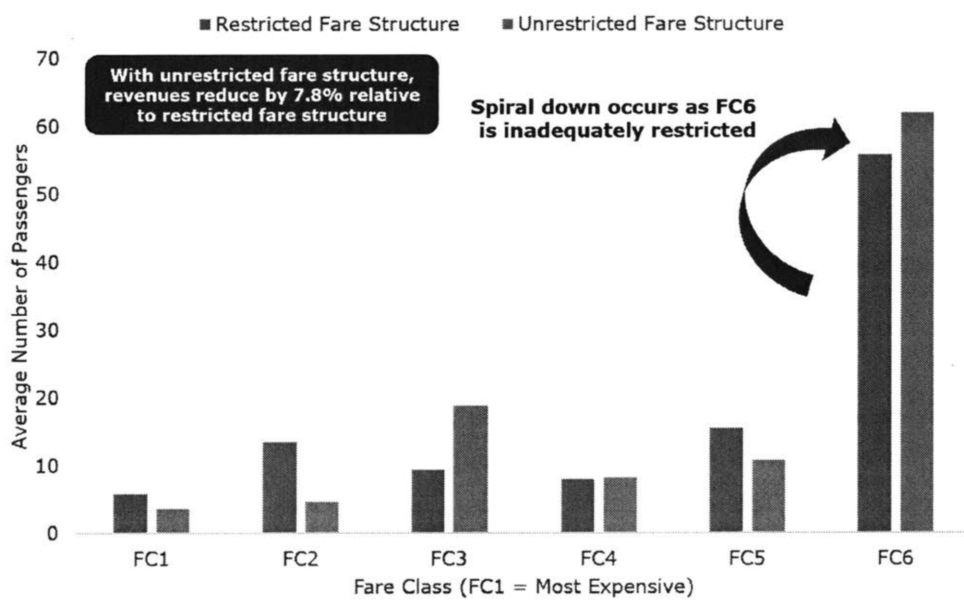

Figure 4.24: Fare class mixes of base Network A1ONE with fully restricted and less-restricted fare structures (Medium Demand)

Figure 4.25 shows that PFDynA produces higher percent revenue gains over the base when a less-restricted fare structure is used, as opposed to a restricted structure. With input Q-multipliers of 3.5 for business passengers and 1.2 for leisure passengers, Increments-Only and Discounts-Only PFDynA increase revenues by 8.6% and 5.4%, respectively. The Two-Way

[23]The "KI" method of fare adjustment is used with a scalar of 0.75, and a PODS FRAT5c curve is used for estimated sell-up probabilities. For more information about these methods and the FRAT5 curve, see Cléaz-Savoyen (2005).

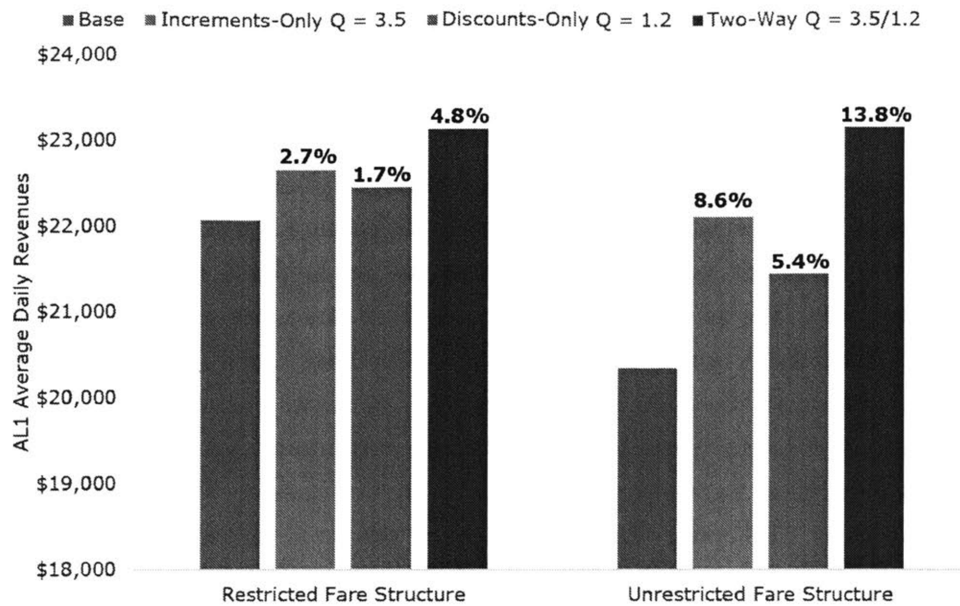

Figure 4.25: Percent change in revenue when AL1 uses three different PFDynA heuristics with restricted and less-restricted fare structures in Network A1ONE (Medium Demand)

PFDynA heuristic increases revenues by nearly 14% over the less-restricted base case and leads to higher total revenues than when Two-Way PFDynA is used in the restricted case.

PFDynA performs so well in less-restricted fare structures because it helps improve the airline's fare class mix. Recall that the three mechanisms that lead to PFDynA revenue gains are: (1) an increase in yield from business passenger buy-up; (2) demand stimulation from leisure passengers in higher fare classes; and (3) forecast spiral-up that causes the RM system to protect more seats for higher booking classes. Each of these three effects leads to an increase in observations in higher-class bookings. As shown in Figure 4.26, PFDynA shifts bookings from FC6 to FC5, and from FCs 3 and 4 to more expensive FCs 1 and 2.

The forecast spiral-up induced by PFDynA causes the RM system to close the least-expensive fare class (FC6) more often near the beginning of the booking process. This leads to an increase in average paid fares at the beginning of the booking process, as shown in Figure 4.27. Average fare levels also increase during the middle of the booking period, as spiral-up leads to higher closure rates for lower classes, and as discounts stimulate new leisure bookings in higher classes. When Two-Way PFDynA is used in less-restricted Network A1ONE, average fares paid increase from the base case throughout the entire booking process. In this single-airline network, this allows Two-Way PFDynA revenues in a less-restricted fare structure to approach those from when a restricted structure is used.

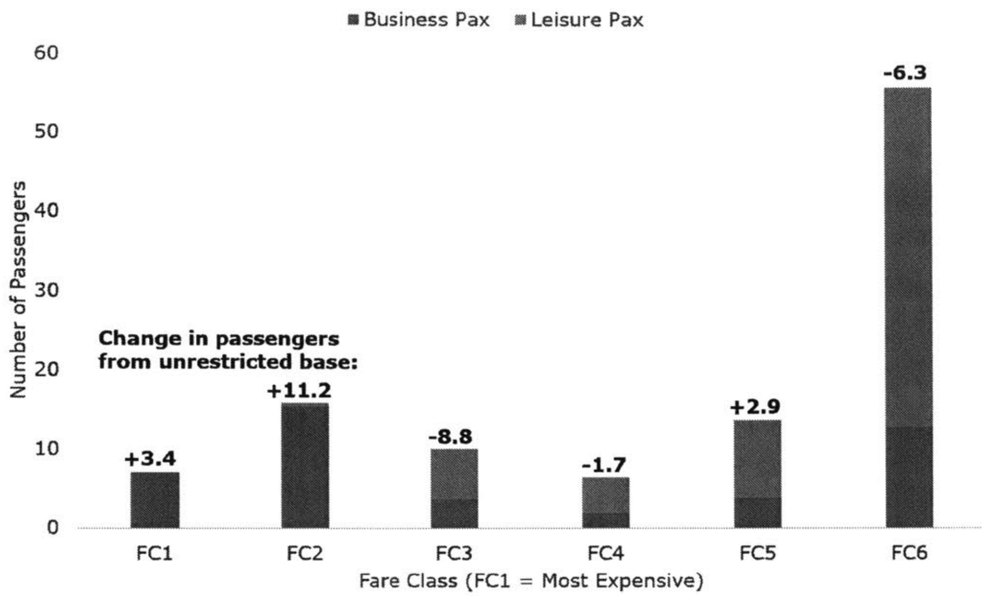

Figure 4.26: Average bookings by fare class when AL1 uses Two-Way PFDynA Q = 3.5/1.2 in less-restricted Network A1ONE (Medium Demand)

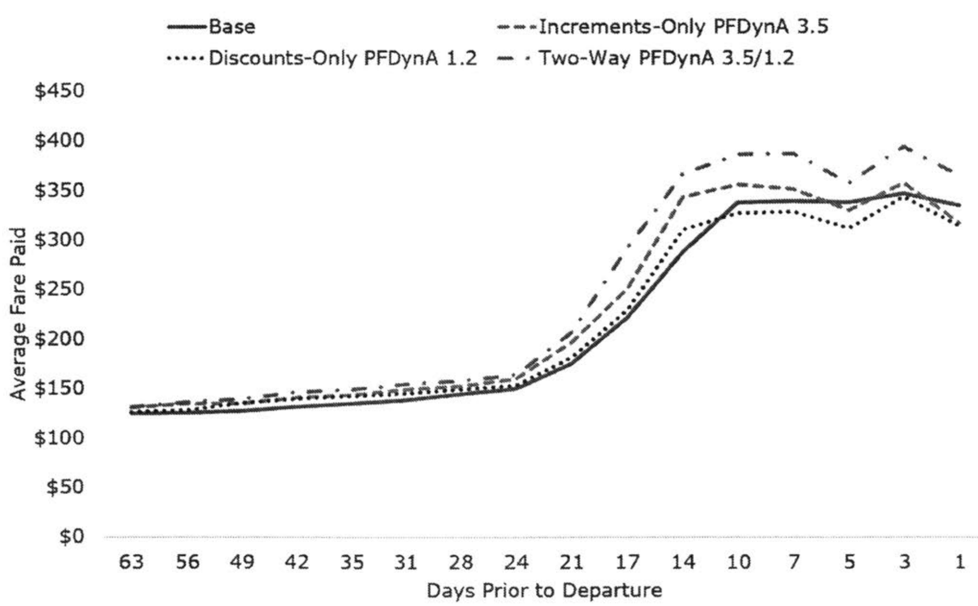

Figure 4.27: Average fares paid when AL1 uses three different PFDynA heuristics in less-restricted Network A1ONE (Medium Demand)

4.3.7 Conclusions from tests of PFDynA in Network A1ONE

In this section, we tested the PFDynA heuristics in a simple network with a single airline operating a single flight. We found that all three PFDynA heuristics—Increments-Only, Discounts-Only, and Two-Way—increased the airline's revenue over the base case in a variety of simulation environments. When Two-Way PFDynA was used by an airline using EMSRb and standard forecasting with input Q-multipliers of 3.5 for business passengers and 1.2 for leisure passengers, revenues increased by 4.8% over a Medium Demand base case. PFDynA produces three main revenue management outcomes that drive the simulated increases in revenue in Network A1ONE:

- Incrementing prices for certain business passengers leads to an increase in average fare levels paid by business passengers. This increase in fare levels is due to some business passengers booking with an increment, and to some passengers choosing to buy up to higher classes to avoid an increment. In both cases, the airline's *yield increases* over the base case. However, incrementing will make the flight unaffordable for some business travelers, who will choose to no-go. This lowers the airline's average load factor.
- Discounting prices for certain leisure passengers *stimulates new bookings* from leisure passengers who otherwise would not have been able to afford to travel. These new bookings increase the airline's load factor. Since discounts can also lead to new leisure passengers booking in higher (more-expensive) classes, these new bookings also increase the airline's average yield over the base case.
- Incrementing business fares and discounting leisure fares both lead to an increased number of bookings in relatively higher fare classes. The RM system's forecaster starts to increase its forecasts of higher-class bookings, which causes the RM system's optimizer to start protecting more seats for these higher classes later in the booking process. We call this process *forecast spiral-up*. Forecast spiral-up results in an increase in yield, a shift in fare class mix towards higher classes, and higher closure rates for the least-expensive fare classes throughout the booking process.

Revenue gains were possible when PFDynA was used with both a restricted and less-restricted fare structures. However, the airline's ability to segment booking requests into leisure and business categories was critical to the performance of PFDynA. This segmentation need not be perfect; positive revenue results were still possible even when the airline made classification errors 20%-30% of the time. An unsegmented version of PFDynA did not produce revenue gains in A1ONE, since it was unable to avoid giving unnecessary discounts to business passengers and mistaken increments that made travel unaffordable for some leisure travelers.

The results in this section also rely on the key assumption that there is only a single carrier in the market. In the presence of competition, PFDynA outcomes could begin to change as passengers shift across airlines. Passengers facing increments could begin to book with other carriers who are not increasing prices, and offering a discount could help an airline attract passengers away from competitors. Given the hyper-competitive nature of the airline industry, tests of PFDynA in competitive environments are warranted. In the next section, we begin to explore the effects of competition on the PFDynA outcomes discussed above.

4.4 Adding a competitor: PFDynA in Network A2TWO

Network A2TWO is a variant of Network A1ONE that introduces an additional airline (AL2) into the market. AL2 operates a single flight in the market with an identical departure time, arrival time, and seat capacity as the incumbent carrier AL1. The underlying demand in the network is scaled up such that each airline maintains an average load factor of about 83.3% in the base case, but all other simulation parameters remain identical to the more simple Network A1ONE. Both airlines in these tests use a fully-restricted fare structure.

In our tests of dynamic pricing in Network A2TWO, we will focus specifically on the effects of competition on the airline(s) using PFDynA. In some of the simulations, only a single airline (AL1) will use PFDynA. In others, both airlines will use PFDynA. We begin by assuming both airlines use a restricted fare structure, EMSRb with standard forecasting, and have 100% passenger type segmentation accuracy.

4.4.1 Increments-Only PFDynA (AL1 only)

Figure 4.28 shows the percent change in revenue from the base case when only AL1 uses Increments-Only PFDynA with various input Q-multipliers, an input γ of 0.3, and 100% passenger segmentation accuracy. In these figures, both airlines use DAVN and standard forecasting. Unless otherwise stated, the results in this section are from a Medium Demand scenario in which base-case load factors for both airlines are about 83.3%.

Figure 4.28 shows that when AL1 alone uses PFDynA in Network A2TWO, it can achieve revenue gains of up to 1.1%. However, there are several differences between the performance of Increments-Only PFDynA in A2TWO as compared to single-carrier Network A1ONE (Figure 4.8). In Network A1ONE, the revenue gains of PFDynA were higher than in A2TWO, and increased along with the input Q-multiplier—a higher Q-multiplier meant increments were applied more often, producing higher yields and more forecast spiral up.

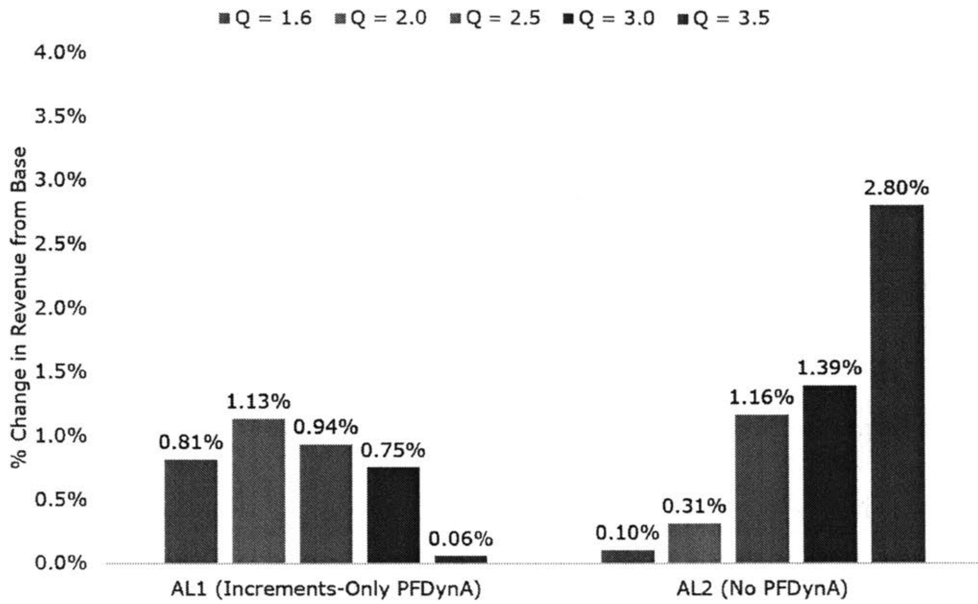

Figure 4.28: Percent change in revenue from base when only AL1 uses Increments-Only PFDynA in Network A2TWO

This is not the case in Network A2TWO. Here, a lower Q-multiplier (Q = 2.0) produces the best revenue results, and these results degrade as the airline continues to increase its input Q-multiplier up to Q = 3.5 (which performed the best in the tests with Network A1ONE).

We noted earlier that conditional WTP would likely decrease when there were more competitors in the market, since there are more options in the customer's choice set with a lower price. This appears to be the case in Network A2TWO. The airline's best choice of input Q-multiplier is lower than in the single-airline case due to the presence of the competitor airline AL2, which is not incrementing its prices.

When AL1 starts incrementing prices for some of its business passengers, these passengers may choose to book with the competitor instead of accepting the higher price. This is why AL2 sees revenue gains as high as 2.8% even when only AL1 is using the dynamic pricing heuristic. If AL1 overestimates the conditional WTP of its business passengers, it sees almost no change in its own revenue while leading to a significant revenue boost for AL2.

As AL1 increases its input Q-multiplier, it sees lower load factors and higher yields, as was the case in single-carrier Network A1ONE. Yet these losses in load factor are exacerbated by the presence of the competing airline. Note from Figure 4.29 that AL2's load factor begins to increase immediately as AL1 practices Increments-Only PFDynA, and even its yield increases as AL1 uses higher input Q-multipliers. This is because AL2 begins to capture some of

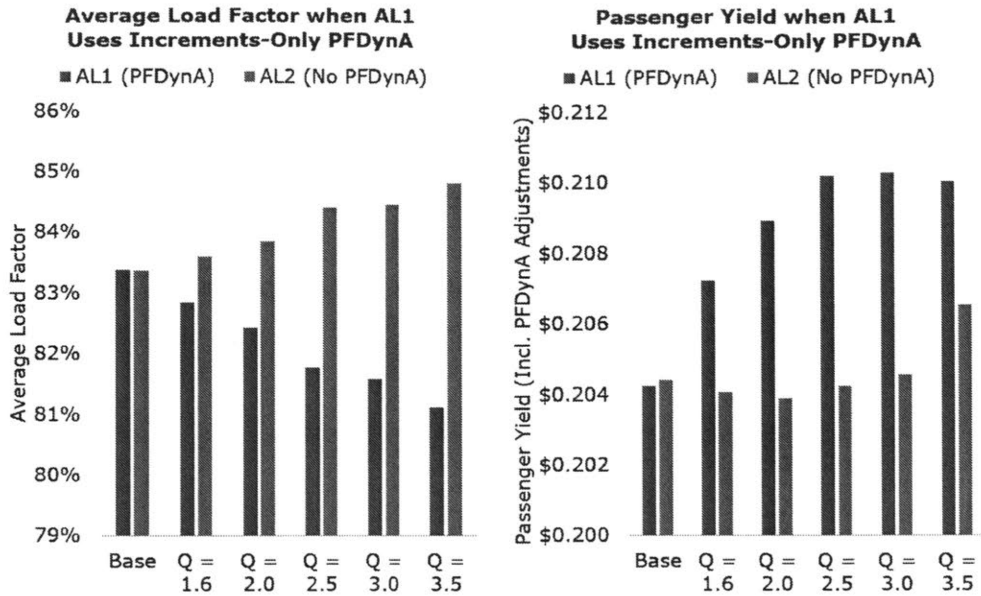

Figure 4.29: Average load factor and passenger yield when only AL1 uses Increments-Only PFDynA in Network A2TWO

AL1's business passengers who face incremented prices. Since forecast spiral-up causes AL1 to close its least-expensive fare class (FC6) more often, AL2 also captures some of AL1's leisure passengers at the beginning of the booking process, as shown in Figure 4.30.

Increments-Only PFDynA causes AL1 to lose bookings at the beginning of the booking process (due both to spiral-up and loss of business passengers to AL2). This early aggressiveness pays off in the last few days before departure. Since AL1's RM system protects more seats for higher classes as a result of PFDynA and forecast spiral-up, it will have more seats available for last-minute business passengers who pay high prices for higher classes. This results in an increase in bookings close to departure, as shown in Figure 4.30. AL1's fare class mix also shifts up, with more bookings in higher, more-expensive fare classes.

4.4.2 Discounts-Only PFDynA (AL1 only)

Discounts-Only PFDynA also shows differences in performance as compared to Network A1ONE. The change in revenue from the base when only AL1 uses Discounts-Only PFDynA to apply discounts to leisure passengers is shown in Figure 4.31. The revenue gains from the heuristic are up to 5% when an input Q-multiplier of Q = 1.5 is used. This is significantly higher than the gains of about 1.7% that were found in Network A1ONE (Figure 4.12). Note that the revenues of Airline 2 fall by as much as 2% when AL1 uses the discounting heuristic.

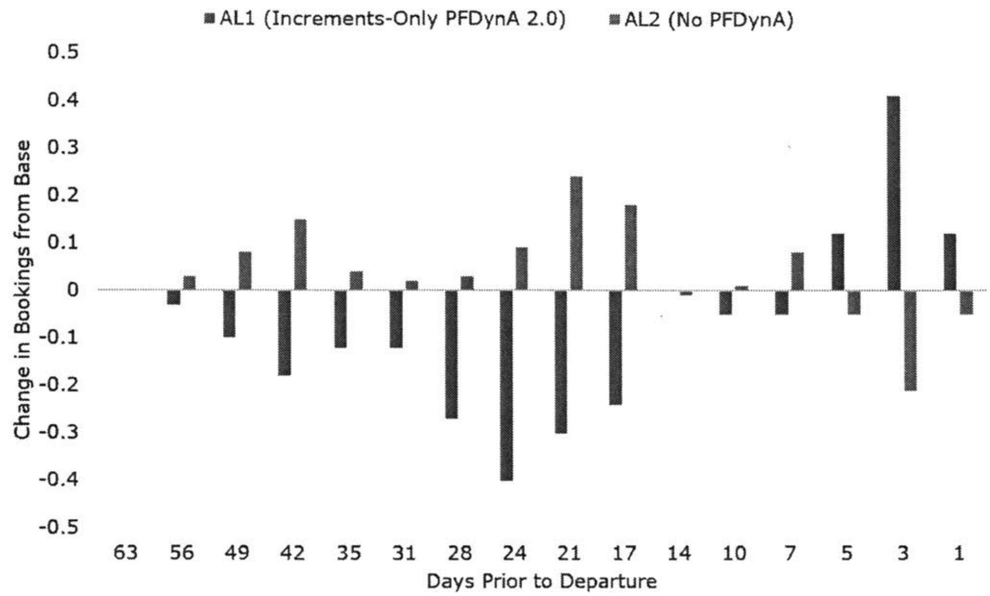

Figure 4.30: Change in bookings over time when only AL1 uses Increments-Only PFDynA in Network A2TWO (Q = 2.0)

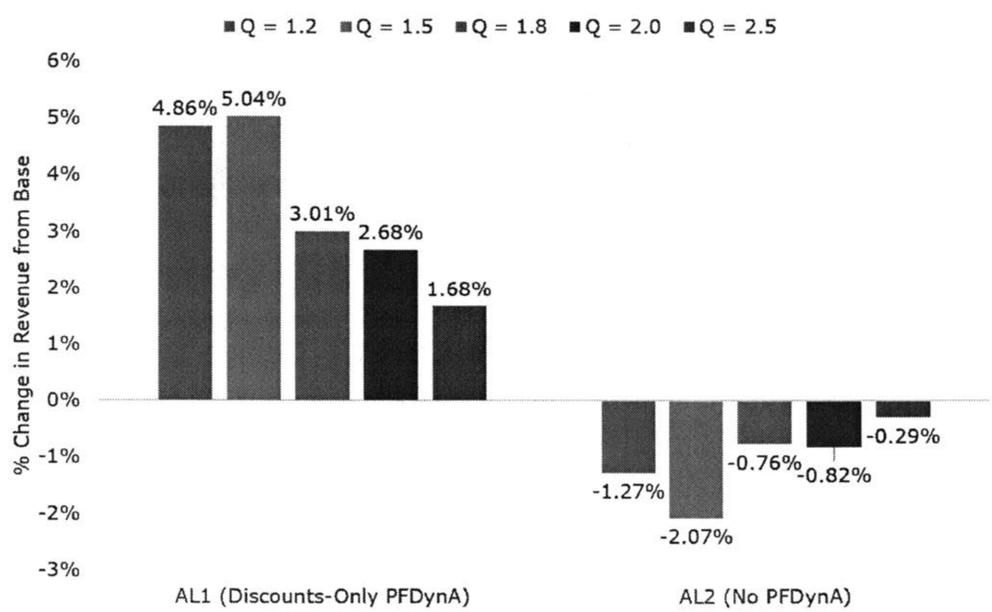

Figure 4.31: Percent change in revenue from base when only AL1 uses Discounts-Only PFDynA in Network A2TWO

Once again, the addition of a competitor changes the ways in which PFDynA leads to revenue gains. Previously, Discounts-Only PFDynA relied purely on demand stimulation and forecast spiral-up to increase revenues. Now, the presence of a competitor that is not practicing PFDynA allows AL1 to begin to capture some passengers from AL2 by enticing them with discounts. As a result, AL2's load factor falls, while AL1's load factor increases from 83.3% to 87.1% (left panel of Figure 4.32) when an input Q-multiplier of Q = 1.5 is used. This increase in load factor can come without a corresponding decrease in yield. As shown in the right-hand panel of Figure 4.32, AL1's yield increases by about 0.6% compared to the base. Even though AL1 is the airline that is discounting, it actually maintains a *higher* average yield than AL2, which is using traditional pricing and RM.

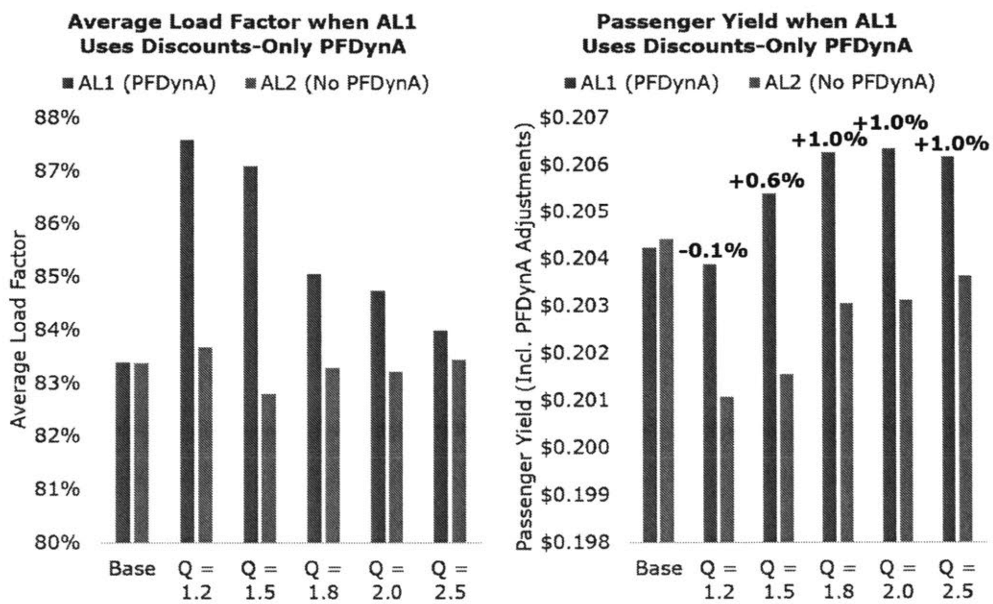

Figure 4.32: Average load factor and passenger yield when only AL1 uses Discounts-Only PFDynA in Network A2TWO

This is again due to forecast spiral-up. Discounts-Only PFDynA stimulates new bookings and captures passengers from AL2 in higher booking classes, causing AL1's RM system to protect more seats for these higher classes. This means AL1 closes FC6 much more often when it uses Discounts-Only PFDynA than AL2, as shown in Figure 4.33. As a result, AL2 will have more availability in lower classes, particularly for early-arriving leisure passengers. This causes AL2 to pick up additional low-paying FC6 passengers, lowering its yield, while AL1 captures more higher-class bookings through strategic discounting (Figure 4.34).

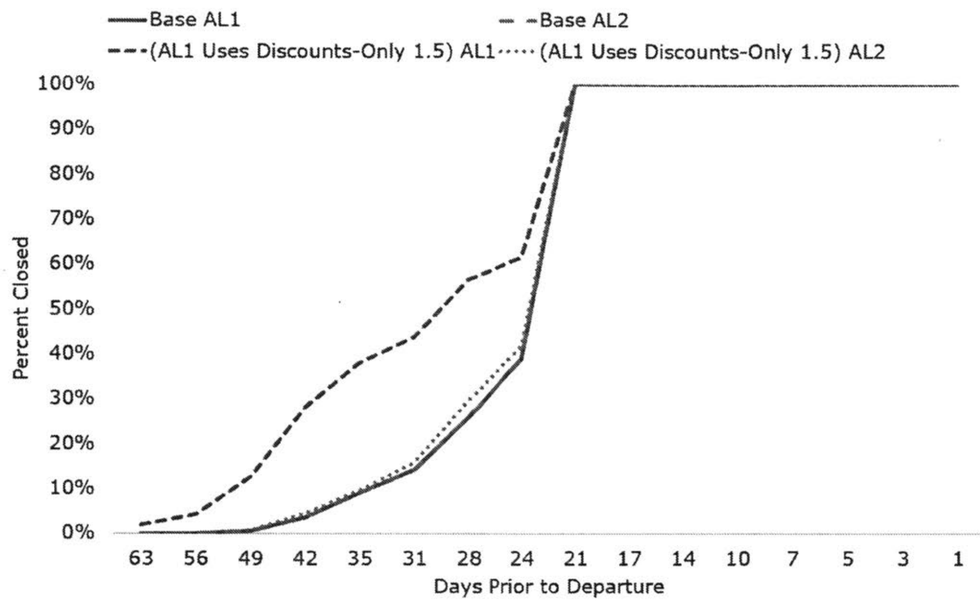

Figure 4.33: FC6 closure rates when only AL1 uses Discounts-Only PFDynA in Network A2TWO (Q = 1.5)

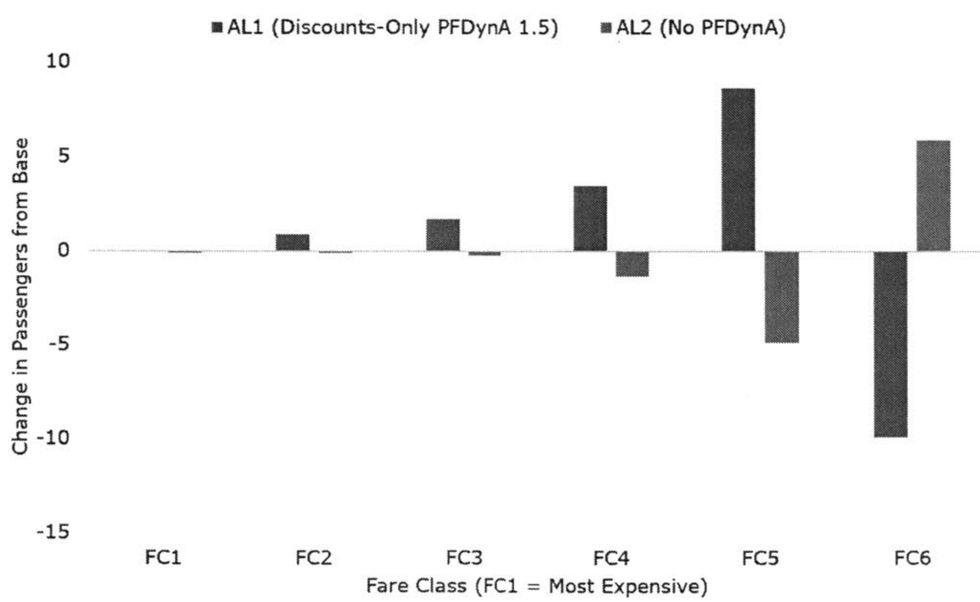

Figure 4.34: Changes in bookings by fare class when only AL1 uses Discounts-Only PFDynA in Network A2TWO (Q = 1.5)

4.4.3 Two-Way PFDynA (AL1 only)

Two-Way PFDynA compounds the revenue gains from the Increments-Only and Discounts-Only heuristics. Figure 4.35 shows that AL1 sees revenue gains of about 6.3% when it uses Two-Way PFDynA with Q = 2.0 for business passengers and Q = 1.5 for leisure passengers. Its load factor increases from a base-case of 83.4% to 86.1% when it uses Two-Way PFDynA.

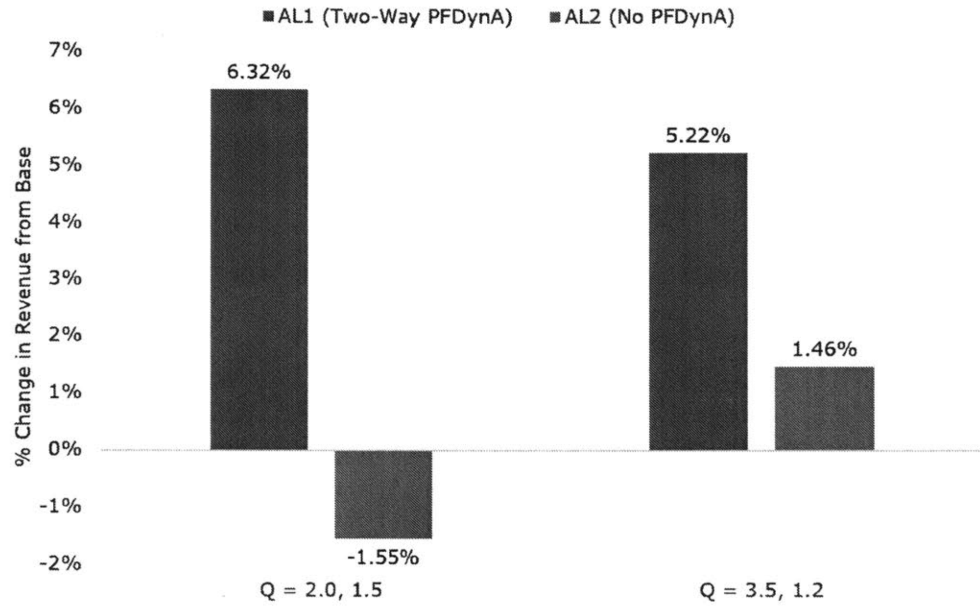

Figure 4.35: Percent change in revenue from base when only AL1 uses Two-Way PFDynA in Network A2TWO

AL1's use of Two-Way PFDynA also captures some passengers from AL2, leading to a slight reduction in AL2's load factor and a revenue loss of over 1.5%. While AL1's yield increases due to forecast spiral-up and the increments it applies to business passengers, AL2 sees a yield decline since it is more open in FC6 than AL1 early in the booking process. Two-Way PFDynA thus increases both load factor and yield for AL1, and decreases both load factor and yield for AL2. If we change the input Q-multipliers to cause more aggressive incrementing and discounting (Q = 3.5 and 1.2), the revenue gains from PFDynA decrease slightly for AL1 (to 5.2%) relative to the Q = 2.0/1.5 case. AL2 now sees a revenue gain, since it is able to capture business passengers from AL1 as a result of more aggressive incrementing.

Returning to input Q-multipliers of 2.0 and 1.5, Two-Way PFDynA also results in changes to the average fares paid over time, as shown in Figure 4.36. As a result of forecast spiral-up and increments for business passengers, AL1's average fares increase early in the booking process relative to the base case. Later in the booking process, when it begins to apply

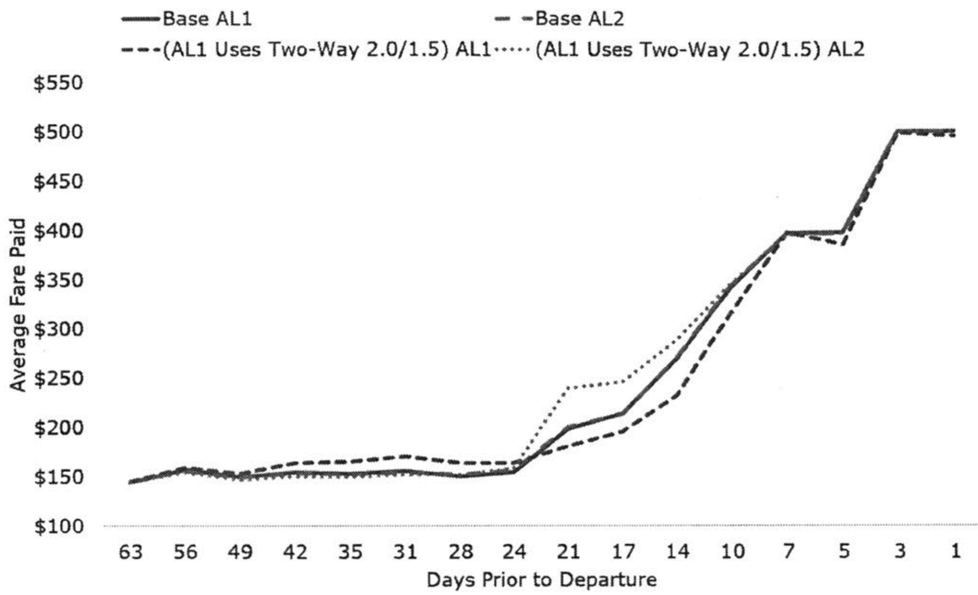

Figure 4.36: Average fares paid over time when only AL1 uses Two-Way PFDynA in Network A2TWO (Q = 2.0/1.5)

discounts for leisure passengers, average fares decrease from the base. However, an increase in late bookings, as shown in Figure 4.37, leads to an overall increase in revenue (6.3%).

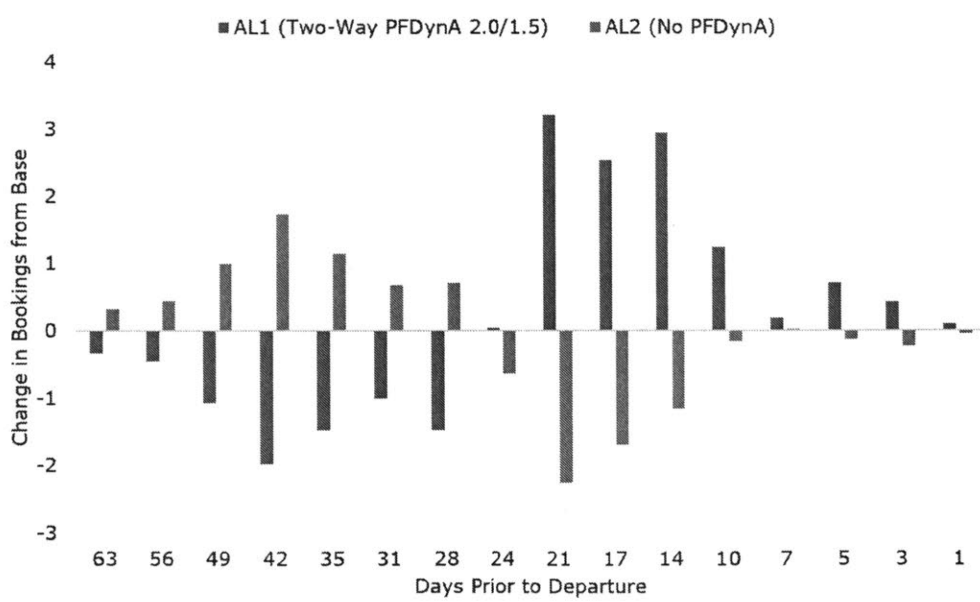

Figure 4.37: Change in bookings from base over time when AL1 uses Two-Way PFDynA in Network A2TWO (Q = 2.0/1.5)

Finally, an examination of the change in fare class mix for both airlines confirms the increase in yield, stimulation of high-class bookings, and forecast spiral-up produced by Two-Way PFDynA. Note from Figure 4.38 that AL1 sees an increase in bookings in all classes higher than FC5 as a result of Two-Way PFDynA. AL2, which is practicing traditional RM, sees a decline in bookings in FCs 4 and 5 (due to AL1 discounting) as well as an increase in bookings in FC6. This increase is not enough to overcome losses in load factor and yield, leading to AL2's revenue decline relative to the base case of 1.6%.

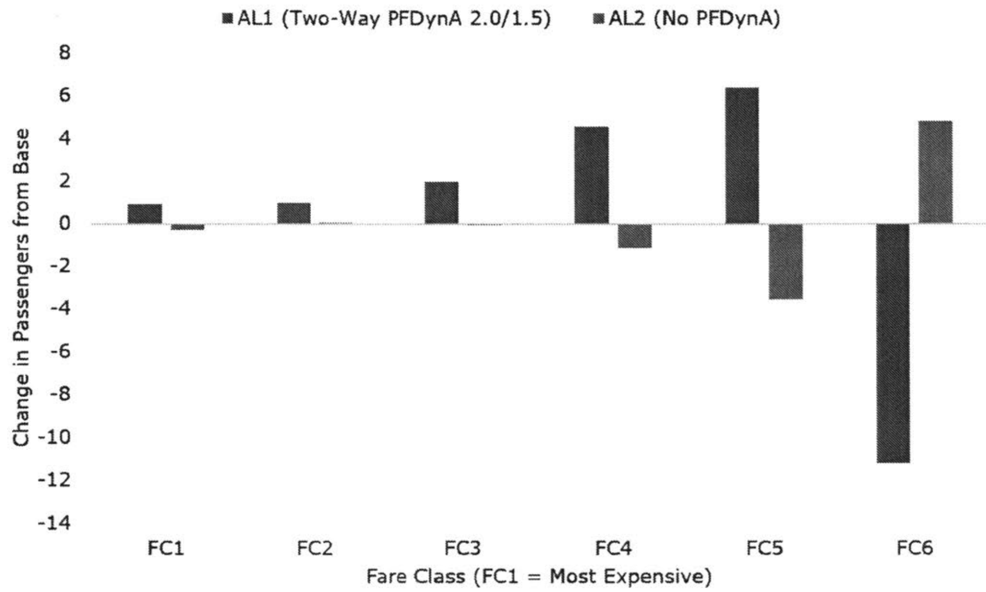

Figure 4.38: Change in bookings from base by fare class when AL1 uses Two-Way PFDynA in Network A2TWO (Q = 2.0/1.5

The addition of a competitor in Network A2TWO altered the performance of PFDynA relative to the single-carrier Network A1ONE. With a competitor, incrementing for business passengers is slightly less successful due to the chance of a passenger deciding to book with a competitor that is not raising prices. Conversely, discounting for leisure passengers can be more successful than the single-airline case, because the airline practicing PFDynA is able to attract customers away from its competitors through targeted discounting. The best-case input Q-multiplier values for business and leisure passengers also changed from the single-airline case.

Overall, Two-Way PFDynA (when used by a single airline) can produce higher revenue gains in the competitive case than in the monopolistic case. The heuristic leads to increases in both yield and load factor for the airline using PFDynA, and decreases in yield and load factor for the airline practicing traditional pricing and revenue management.

In a real-life scenario, it is unlikely that AL2 would stand idly by as its competitor starts practicing dynamic pricing. Instead, AL2 would also likely begin to use dynamic pricing. Next, we investigate what happens when PFDynA is used by both airlines in the simulation.

4.4.4 Both airlines use PFDynA

Figure 4.39 reports the percentage change in revenue when one airline (AL1) or both airlines use the PFDynA heuristics in Network A2TWO with Q = 2.0 for business passengers and Q = 1.5 for leisure passengers. When used by both airlines, each of the three PFDynA heuristics results in revenue gains for both AL1 and AL2 over the base case of traditional RM. The percent revenue gains for AL1 in the all airline case of 1.6% for Increments-Only PFDynA, 1.8% for Discounts-Only PFDynA, and 3.3% for Two-Way PFDynA are very similar to the revenue gains from the single-airline case in Network A1ONE. The same mechanisms that led to PFDynA revenue gains in A1ONE—increases in business yield, stimulation of new demand, and forecast spiral-up—have the same effect when PFDynA is used by both airlines.

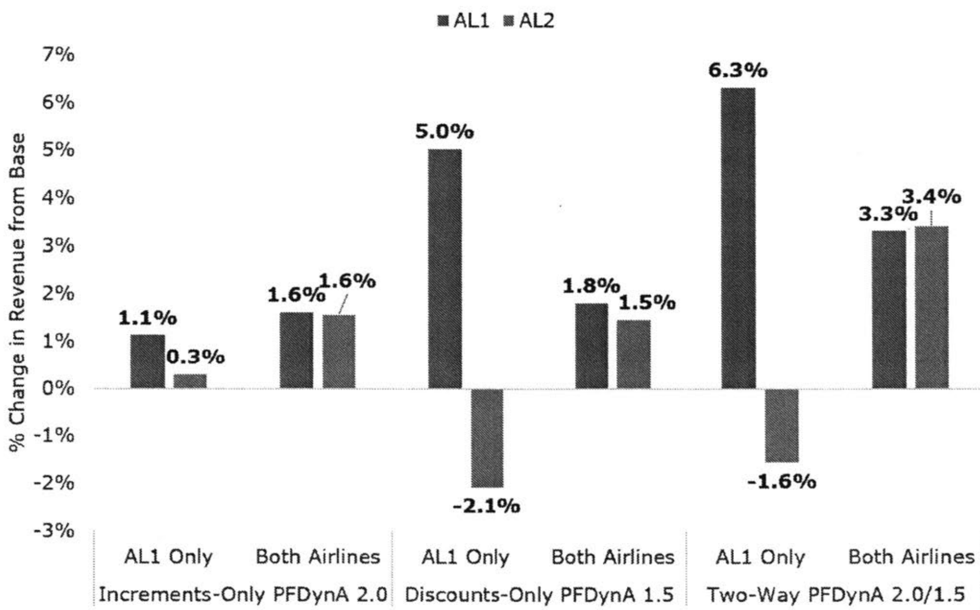

Figure 4.39: Percent change in revenue from base when one or both airlines use PFDynA in Network A2TWO

For each of the three heuristics, using PFDynA is the best response to the use of PFDynA by a competitor. For instance, AL2 sees a -1.6% revenue loss when Two-Way PFDynA (Q = 2.0/1.5) is used by AL1 alone. But if AL2 can itself use Two-Way PFDynA, it will cut AL1's gain from 6.3% to 3.3% and also lead to a 3.4% revenue gain for itself. In the language of economic game theory, "Use PFDynA" is a dominant strategy over "Do Not Use PFDynA" in this network and with these simulation parameters. "Use PFDynA" therefore represents a Nash equilibrium for this one-shot game and leads to revenue gains for both airlines that rival the gains of AL1 in the single-carrier case.

4.4.5 Sensitivity analysis: Segmentation accuracy and fare restrictions

Segmentation inaccuracy in the two-carrier case leads to different effects than in the single-carrier case, as seen in Figure 4.40. Recall that in the single-carrier case, accuracy rates of as low as 70% still resulted in revenue gains when Increments-Only PFDynA was used by AL1. In the multiple-carrier case, however, segmentation inaccuracy with Increments-Only PFDynA is much more costly. This is both because the baseline revenue gains of Increments-Only PFDynA are lower than in the single carrier case, and because AL2 is able to capture leisure passengers who are mistakenly identified as business passengers.

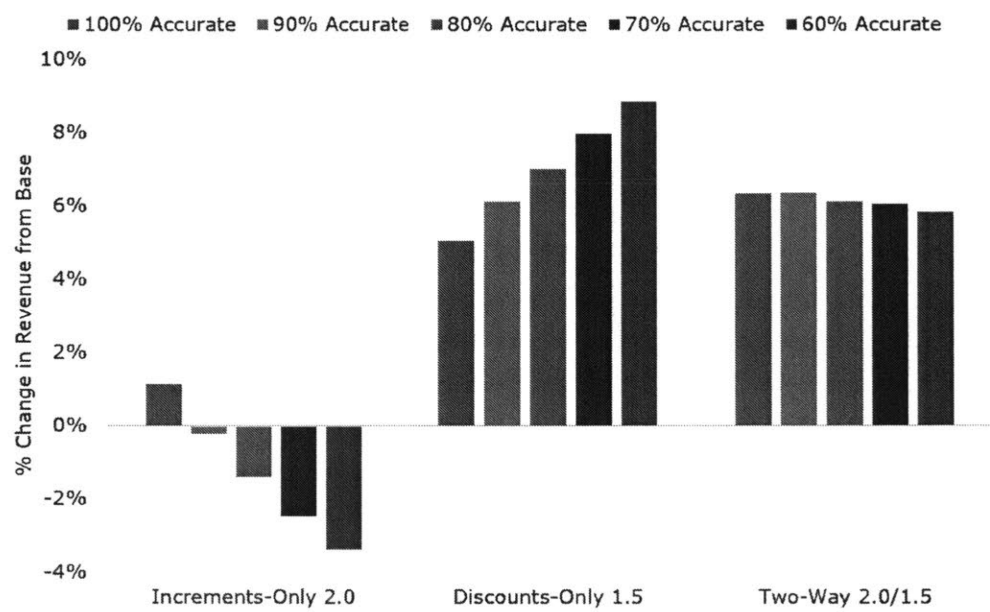

Figure 4.40: Percent change in AL1 revenue from base when only AL1 uses three different PFDynA heuristics in Network A2TWO with segmentation inaccuracy (Medium Demand)

Conversely, identification inaccuracy with Discounts-Only PFDynA actually *increases* its revenue performance when it is used by only AL1 in Network A2TWO. This is because mistakenly giving discounts to business passengers will also lead to new bookings, as business passengers are captured from AL2. The revenue losses from unnecessary discounts due to inaccuracy are made up by this gain in high-value business passengers. Balancing these two effects, Two-Way PFDynA produces approximately the same revenue gains of 5 - 6% when used by AL1 only in Network A2TWO, regardless of AL1's passenger segmentation accuracy.

When PFDynA is used by both airlines with segmentation inaccuracy[24] in Network A2TWO, we see similar patterns as in the single-airline A1ONE case. As shown in Figure 4.41, less accurate segmentation reduces revenue gains for all three heuristics. Increments-Only PFDynA is revenue positive when used by both airlines inaccuracy rates of at least 40%, and inaccuracy rates of 10-20% are still sufficient for revenue gains for Two-Way and Discounts-Only PFDynA, respectively. However, as in Network A1ONE, greater segmentation ability is critical for positive revenue performance when PFDynA is used by both airlines.

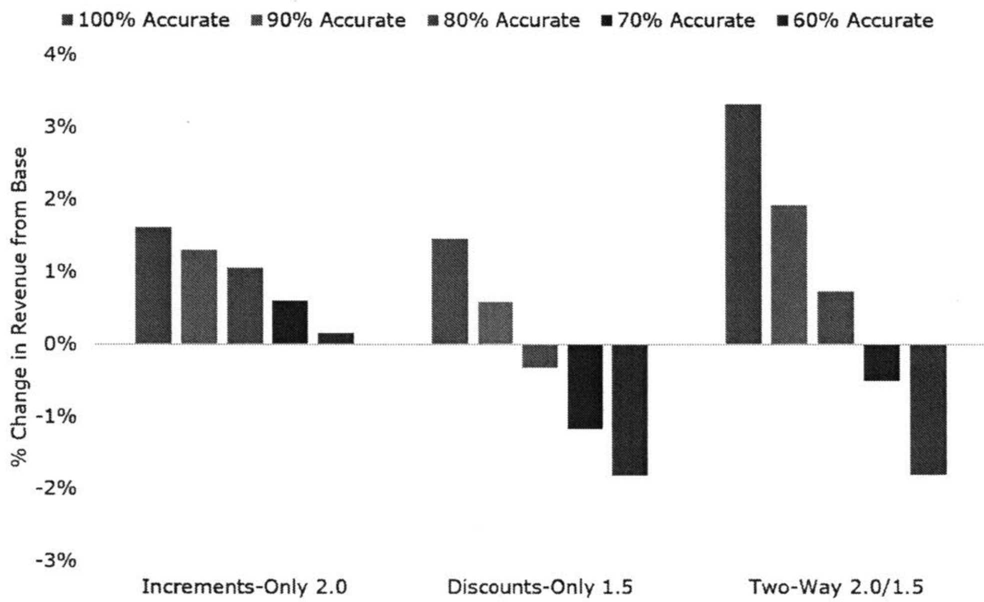

Figure 4.41: Percent change in AL1 revenue from base when both airlines use PFDynA in Network A2TWO with segmentation inaccuracy (Medium Demand)

Finally, we test PFDynA when both airlines in Network A2TWO use the less-restricted fare structure from Table 4.2. Both airlines here use EMSRb with hybrid forecasting (FRAT5c) and fare adjustment (with 0.75 scaling). Without PFDynA, moving from a restricted to a less-restricted fare structure leads to a 7.6% reduction in revenue compared to the base case.

The use of PFDynA by one or both airlines with a less-restricted fare structure leads to revenue gains in Network A2TWO, as shown in Figure 4.42. The mechanisms of the revenue gains are similar to those in Network A1ONE—increased business yield, leisure demand stimulation, forecast spiral-up, as well as shifts of passengers between airlines when PFDynA is used only by a single carrier.

[24]Each airline correctly segments any given passenger with probability α.

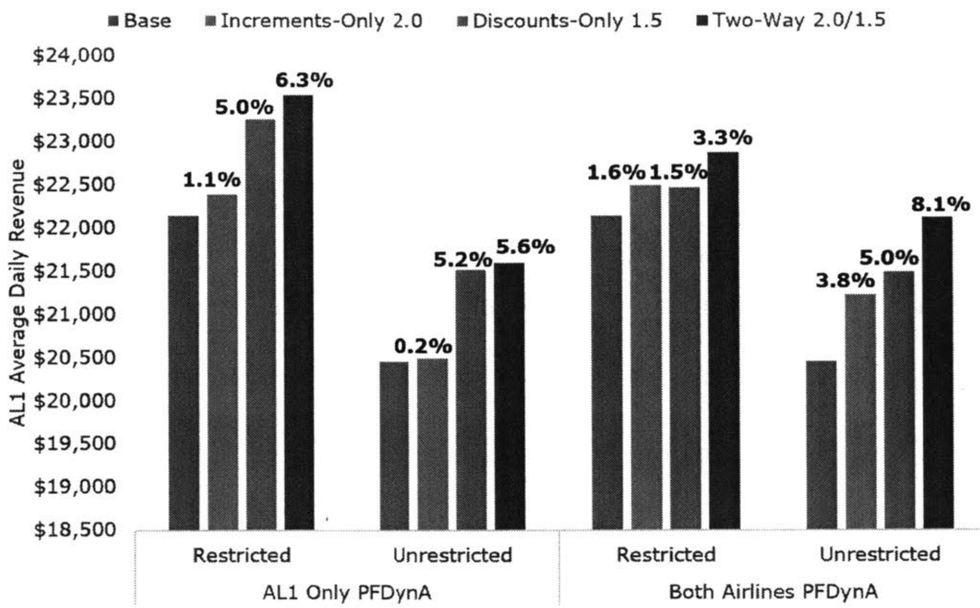

Figure 4.42: AL1 average daily revenue when only AL1 or both airlines use PFDynA in Network A2TWO with restricted or less-restricted fare structures

The revenue gains of 5.6% from Two-Way PFDynA when used by AL1 alone is not enough to recover the 7.6% loss in revenue from moving to a less-restricted fare structure. However, the 8.1% increase in revenue from Two-Way PFDynA when used by both airlines in the less-restricted case exceeds the restricted base case revenue for both airlines.

4.4.6 Conclusions: The effects of competition on PFDynA

In this section, we evaluated the use of PFDynA in a network with two airlines, each of which offers an identical non-stop flight. The simulation results showed that the three PFDynA heuristics can provide revenue gains in this network when used by a single airline or by both airlines in the network. The heuristics result in many of the same revenue management outcomes as in Network A1ONE: namely, increase in business passenger yield, demand stimulation, forecast spiral-up, and more frequent closures of lower-value fare classes.

However, the addition of a competitor airline adds some nuances to the mechanics that drive PFDynA revenue performance. When used by one airline, the Increments-Only heuristic results in lower revenue gains than in the single-airline case, because business passengers facing incremented prices may choose to book with a competitor. Conversely, the Discounts-Only heuristic performs better in a competitive environment if used by a single airline, because the airline can shift demand from the competitor by offering targeted discounts. If both airlines use PFDynA in Network A2TWO, the underlying RM outcome and revenue results are nearly identical to the single-airline Network A1ONE case.

While adding a competitor led to some new dynamics in PFDynA performance, Network A2TWO is still simple, with both airlines offering only a single flight with no hub-and-spoke connections. Next, we test the performance of PFDynA in one of the most complex networks that has been calibrated for PODS: Network U10. These tests in Network U10 are important because this simulation environment most closely parallels to the actual networks operated by real-world airlines.

4.5 PFDynA with four airlines in hub-and-spoke Network U10

Recall from Figure 4.6 that Network U10 consists of four airlines, each of which operates a complex hub-and-spoke operation. The airlines operate a total of 442 flight legs each day, serving a total of 572 origin-destination markets. Airlines AL1, AL2, and AL4, which represent traditional network carriers, use DAVN with standard forecasting as their base revenue management method. AL3, which represents a low-cost carrier with more a simpler point-to-point network, utilizes leg-based EMSRb.[25] Since each airline operates a unique network, baseline load factors and yields will vary from airline to airline. In the Medium Demand scenario, Airline 1 (AL1) has a average load factor of 83.5%.

To begin, we will investigate cases where only a single airline, AL1, uses the PFDynA heuristics in Network U10. Figure 4.43 shows the percent change in revenue from the base when AL1 uses PFDynA with input Q-multipliers of 2.0 for business passengers and 1.5 for leisure passengers in various demand scenarios with 100% passenger segmentation accuracy.

Each of the three PFDynA dynamic pricing heuristics produces revenue gains when used by a single airline in this network. In all three demand scenarios, Two-Way PFDynA produces the highest gains, of about 3.5% - 4.0%. As in Network A2TWO, Discounts-Only PFDynA outperforms Increments-Only PFDynA when used by a single airline in a competitive network. Furthermore, Discounts-Only and Two-Way PFDynA perform best in the low-demand scenario, where there are more likely to be empty seats, and Increments-Only PFDynA performs best in the high-demand scenario where there is enough demand to backfill passengers who are lost to other airlines due to incrementing.

[25] We assume a bid price of $0 for the calculation of PFDynA price adjustments in situations where AL3 practices PFDynA with EMSRb.

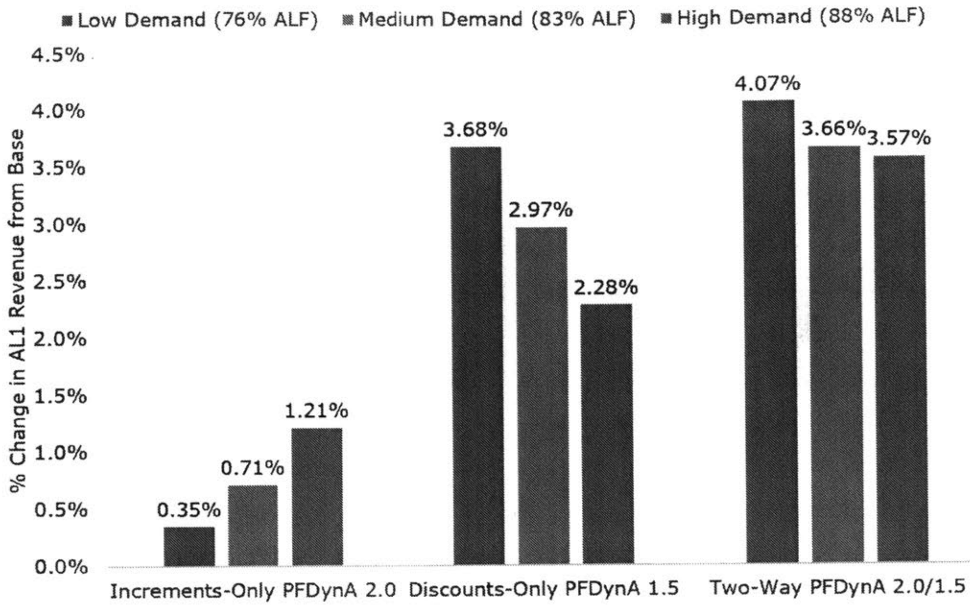

Figure 4.43: Percent change in AL1 revenue from base when AL1 uses three different PF-DynA heuristics in Network U10 (Q = 2.0/1.5, Medium Demand)

4.5.1 Increments-Only PFDynA (AL1 only)

Figure 4.44 shows the change in AL1's revenue when it uses Increments-Only PFDynA with a variety of input Q-multipliers in Network U10 under a Medium Demand scenario.

An increase in input Q-multiplier reflects an increase in the airline's estimated conditional WTP for business-type passengers, and a corresponding increase in the frequency and value of price increments. As in Network A2TWO, overestimating conditional WTP can cause revenue losses, as passengers shift to other airlines as a result of increments. The airline practicing Increments-Only PFDynA sees a decrease in load factor and an increase in yield (not shown), and other airlines gain passengers who decide not to book with AL1 due to the increments. As in Network A2TWO (Figure 4.28), an input Q-multiplier of 2.0 produced the highest revenue gains in Network U10.

Figure 4.45 shows that all other airlines gain revenue even when only AL1 practices Increments-Only PFDynA. The revenue gains of other airlines increase as AL1's estimates of conditional WTP for business passengers increase. In the case when AL1 uses an input Q-multiplier of 3.5 for business passengers, it sees a revenue loss of 0.25% whereas other airlines see gains of up to 1.5%, despite taking no direct dynamic pricing action.

Otherwise, the Increments-Only heuristic performs similarly in Network U10 compared to

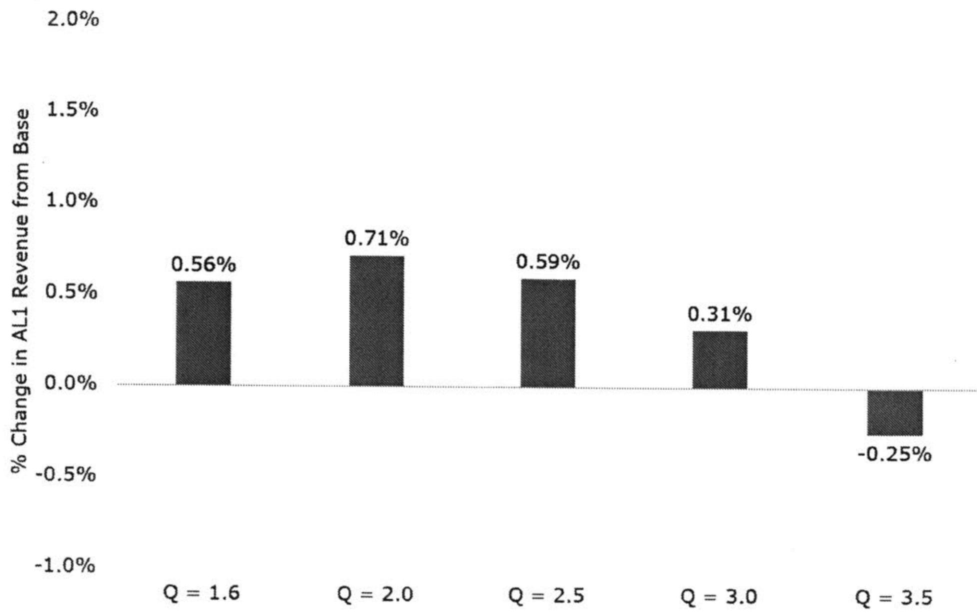

Figure 4.44: Percent change in AL1 revenue from base when AL1 uses Increments-Only PFDynA with a variety of Q-multipliers in Network U10 (Medium Demand)

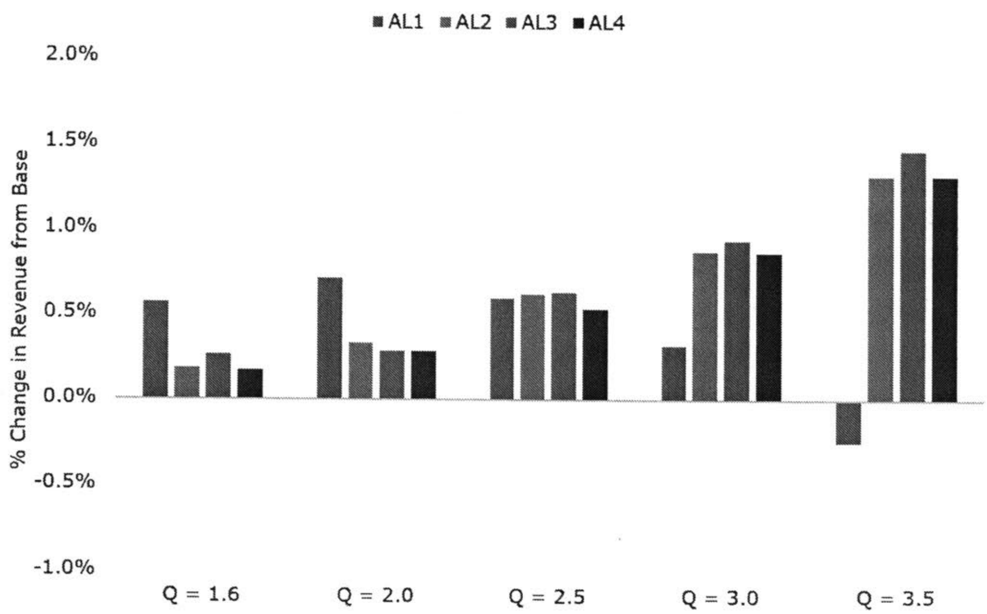

Figure 4.45: Percent change in revenue from base for all four airlines when AL1 uses Increments-Only PFDynA in Network U10 (Q = 2.0, Medium Demand)

the single-airline and two-airline cases discussed earlier. Revenue gains come from an increase in business passenger yield through incrementing. About 23% of the airline's business passengers, which represents 8.6% of the airline's total passengers, book with a PFDynA increment when the airline uses an input Q-multiplier of 2.0. Since U10 uses a mix of restricted and less-restricted fare structures, incrementing a lower fare class will induce some business passengers to buy up to higher fare classes. As in Network A1ONE, this leads to forecast spiral-up, more frequent closures of the least expensive fare class (FC10), and an upward shift in fare class mix.

4.5.2 Discounts-Only PFDynA (AL1 only)

Discounts-Only PFDynA also shows similar performance to Network A2TWO. Figure 4.46 shows the percent change in AL1 revenue when it uses Discounts-Only PFDynA with various input Q-multipliers under Medium Demand. With lower input Q-multipliers, the airline estimates that conditional WTP for leisure passengers is lower, causing the PFDynA heuristics to provide discounts more often. However, the overall level of discounting is still relatively modest. With Q = 1.5, for instance, only about 12% of AL1's leisure passengers (about 8% of its total passengers) end up booking with a PFDynA discount.

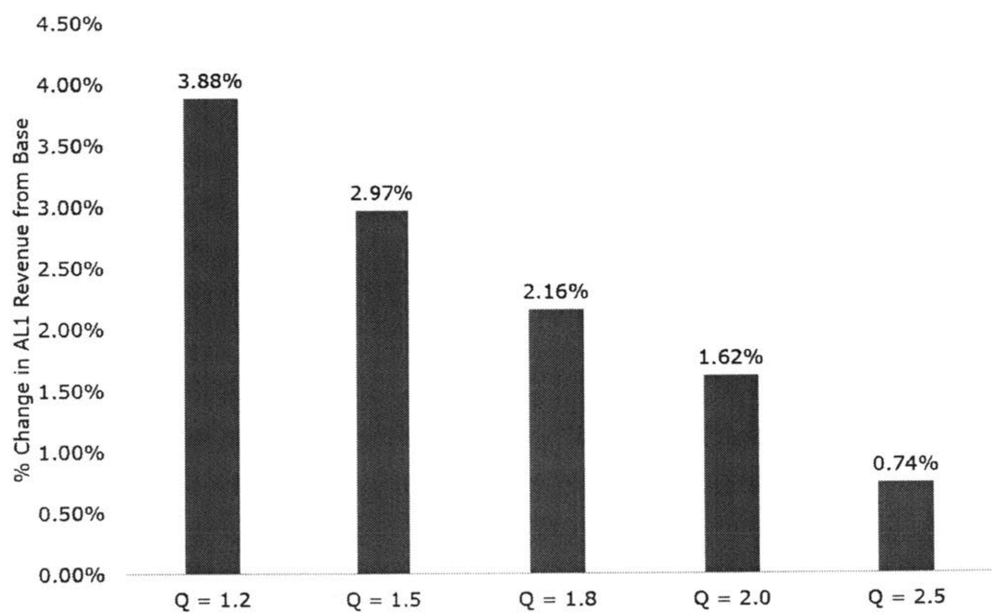

Figure 4.46: Percent change in AL1 revenue from base when AL1 uses Discounts-Only PFDynA with a variety of Q-multipliers in Network U10 (Medium Demand)

When AL1 is the only airline practicing PFDynA in a competitive network, it can capture more and more passengers from other airlines. This can be seen in the revenue performance of other airlines when AL1 uses Discounts-Only PFDynA (Figure 4.47). With only AL1

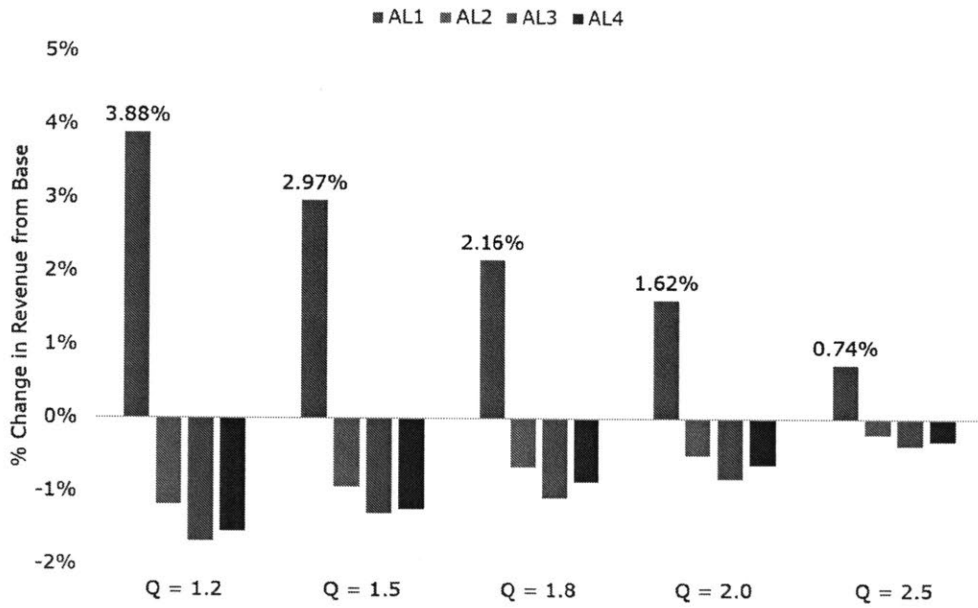

Figure 4.47: Percent change in revenue from base for all airlines when AL1 uses Discounts-Only PFDynA with a variety of Q-multipliers in Network U10 (Q = 1.5, Medium Demand)

using Discounts-Only PFDynA, other airlines see losses in revenue and load factor when AL1 discounts. AL1 itself gains load factor when it uses Discounts-Only PFDynA due to demand stimulation and passenger capture, and with input Q-multipliers of at least Q = 1.5, it also sees an increase in passenger yield (not shown) due to forecast spiral-up.

4.5.3 Two-Way PFDynA (AL1 only)

Finally, Two-Way PFDynA produces revenue gains by compounding the effects of Increments-Only and Discounts-Only PFDynA.

In Medium Demand with Q-multipliers of 2.0 for business passengers and 1.5 for leisure passengers, Two-Way PFDynA produces a revenue gain of about 3.7% when used by AL1 alone. AL1's average load factor increases from 83.5% to 85.2%, and its yield increases by about 1.7% over the base case. Moreover, the other airlines that are not practicing Two-Way PFDynA see a reduction in revenue of between 0.7% and 0.9% when AL1 uses the heuristic, as shown in Figure 4.48.

Along with the shifts of passengers between airlines as a result of increments and decrements, AL1 sees many of the same patterns in PFDynA performance as in the simple Network A1ONE case. First, the combination of increments (which encourage business passenger

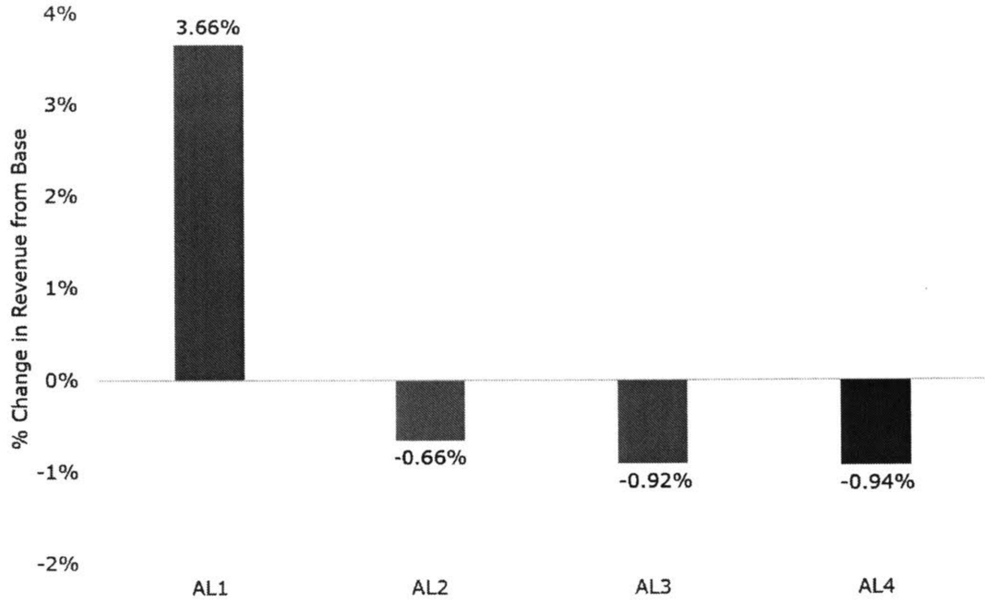

Figure 4.48: Percent change in revenue from base for all airlines when only AL1 uses Two-Way PFDynA in Network U10 (Q = 2.0/1.5, Medium Demand)

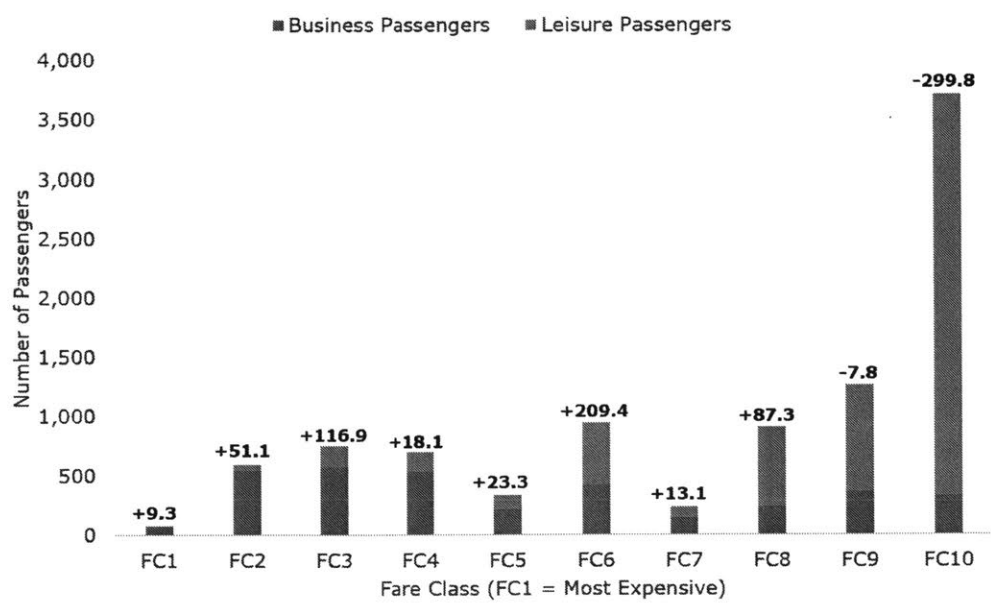

Figure 4.49: AL1 passengers by fare class and passenger type when only AL1 uses Two-Way PFDynA in Network U10 (Q = 2.0/1.5, Medium Demand)

buy-up) and discounts for leisure customers in higher classes causes an upward shift in AL1's fare class mix, as seen in Figure 4.49. The airline sees more bookings in its eight most-expensive booking classes, and fewer bookings in its least-expensive classes 9 and 10. This is due in part to forecast spiral-up closing the least-expensive class (FC10) more often than in the base case.

This change in fare class mix is closely linked to changes in average fares paid over time, which is shown for all three PFDynA methods and the base case in Figure 4.50. Early in the booking process, all three PFDynA methods lead to an increase in average fares paid over the base, either through business passenger increments (Increments-Only PFDynA), forecast spiral-up due to discounts in higher classes (Discounts-Only PFDynA), or a combination of the two (Two-Way PFDynA). While average fare levels remain higher than the base throughout the booking process with Increments-Only PFDynA, fare levels become lower than the base later in the booking process with Discounts-Only and Two-Way PFDynA, as leisure customers receive discounts in high-priced classes (at this point, lower classes have been closed by advance purchase requirements).

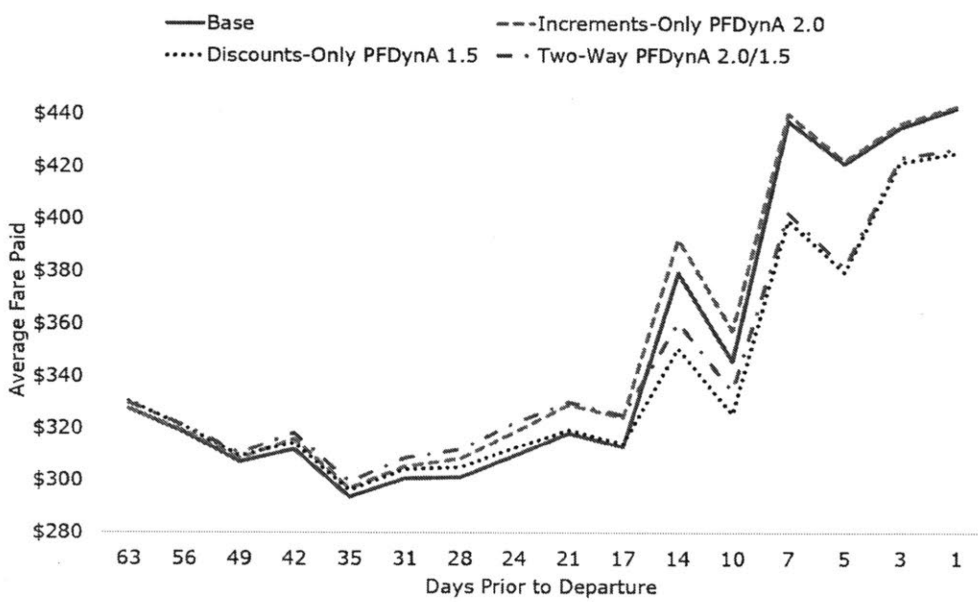

Figure 4.50: AL1 average fares paid when AL1 uses three different PFDynA heuristics in Network U10 (Medium Demand)

Yet these lower fare levels later in the booking process are coupled with higher booking volumes, as shown in Figure 4.51. This chart shows the change in bookings by business and leisure passengers relative to the base throughout the booking process when AL1 uses Two-Way PFDynA. The higher fare levels early in the booking process seen in Figure 4.50 result in slightly fewer early bookings. Later on, when the RM system starts to close down lower classes due to advance purchase requirements, leisure passengers begin receiving discounts,

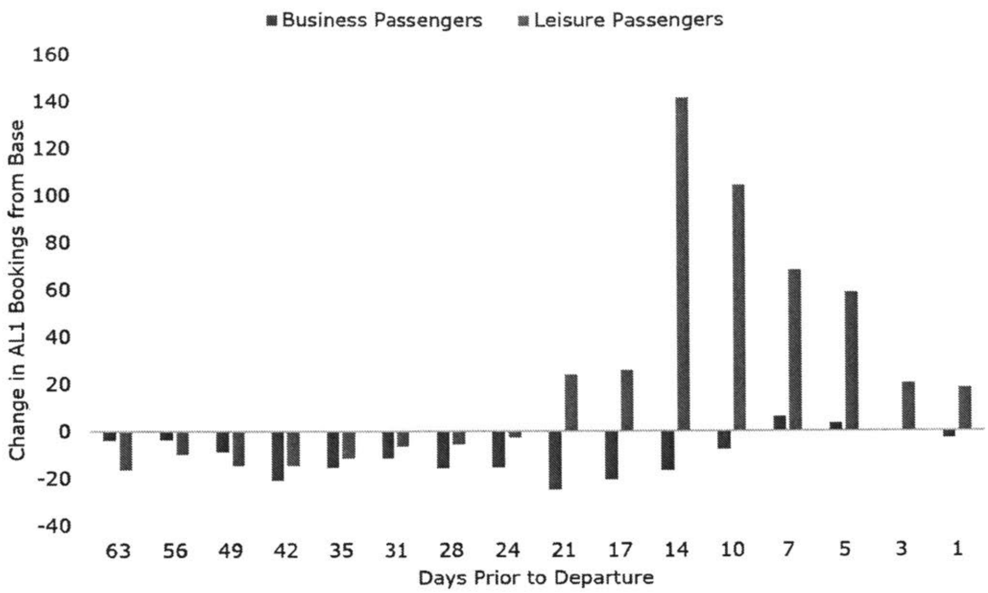

Figure 4.51: AL1 change in bookings from base over time when AL1 uses Two-Way PFDynA in Network U10 (Q = 2.0/1.5, Medium Demand)

causing an increase in leisure bookings in these later time periods. We can also see a slight increase in business bookings in later periods, since forecast spiral-up means that the RM system is protecting more seats for late-arriving passengers (of all types).

Overall, the booking process concludes with about 16% of AL1's passengers booking with a PFDynA price adjustment—8.4% of passengers book with an increment, and 7.8% book with a discount. As shown in Figure 4.52, most of AL1's passengers book at filed fares with no PFDynA adjustment applied. Business passengers are more likely to receive increments in lower fare classes, and leisure passengers are more likely to receive discounts in higher fare classes.

While the complex network structure of Network U10 leads to more opportunities for passenger substitution amongst different carriers as a result of PFDynA increments or discounts, the underlying mechanisms that cause PFDynA's revenue gains are relatively unchanged. When used by one airline, PFDynA increases revenues by increasing yields from business passengers, stimulating demand from leisure passengers, either through new bookings or through demand shift from other airlines, and forecast spiral-up, which reinforces the improvement in fare class mix and results in more frequent closures of the least-expensive class.

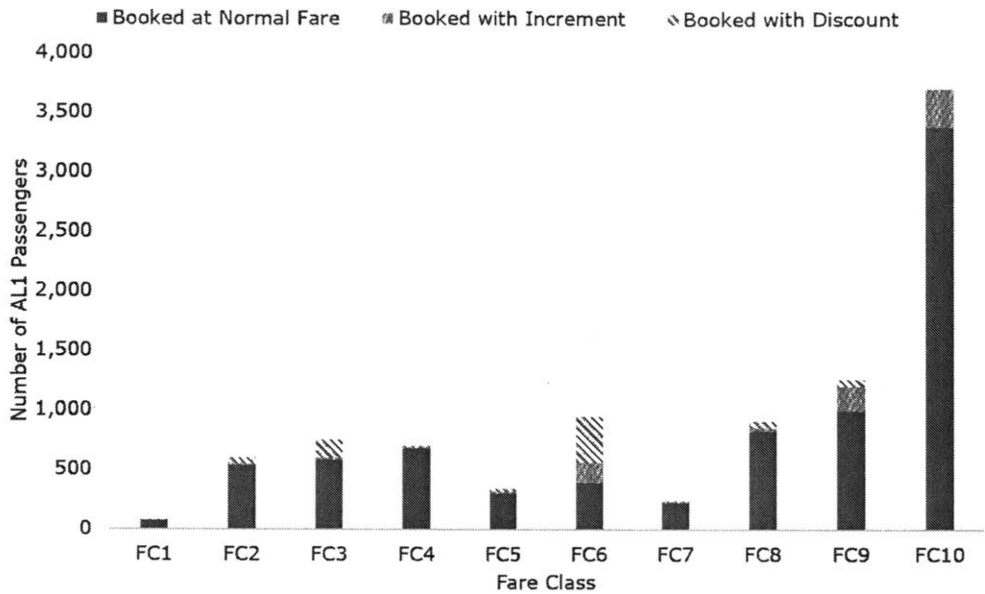

Figure 4.52: AL1 passengers booking with and without PFDynA adjustments by fare class when AL1 uses Two-Way PFDynA in Network U10 (Q = 2.0/1.5, Medium Demand)

4.5.4 Multiple airlines use PFDynA

Figure 4.53 shows the revenue results when all of the airlines in the simulation use one of the three PFDynA heuristics with Q = 2.0 for business passengers and Q = 1.5 for leisure passengers. Note that the use of PFDynA heuristics by all airlines in the simulation is revenue positive for each of the airlines. As in previous tests in Network A2TWO, Increments-Only PFDynA provides the highest revenue gains when used by all airlines, since there is no alternative airline in the network that is not incrementing prices for business passengers.

Conversely, Discounts-Only PFDynA provides lower revenue gains when used by all airlines, since it is relying solely on the effects of demand stimulation and forecast spiral-up to drive its revenue gains, and not on the capture of passengers from other airlines that are not practicing PFDynA. Two-Way PFDynA again combines the effects of Increments-Only and Discounts-Only PFDynA, and produces revenue gains of about 1.7% to 2.4% when used by all of the various airlines in this network.

In Network A2TWO, we found that practicing (as opposed to not practicing) PFDynA was a Nash equilibrium, because no airline had an incentive to discontinue the use of the heuristic to gain additional revenue. This is also the case in Network U10. For instance, suppose one airline (say, AL4) decides to discontinue use of Increments-Only PFDynA. By not applying increments to business passengers, AL4's goal would be to attract business passengers from

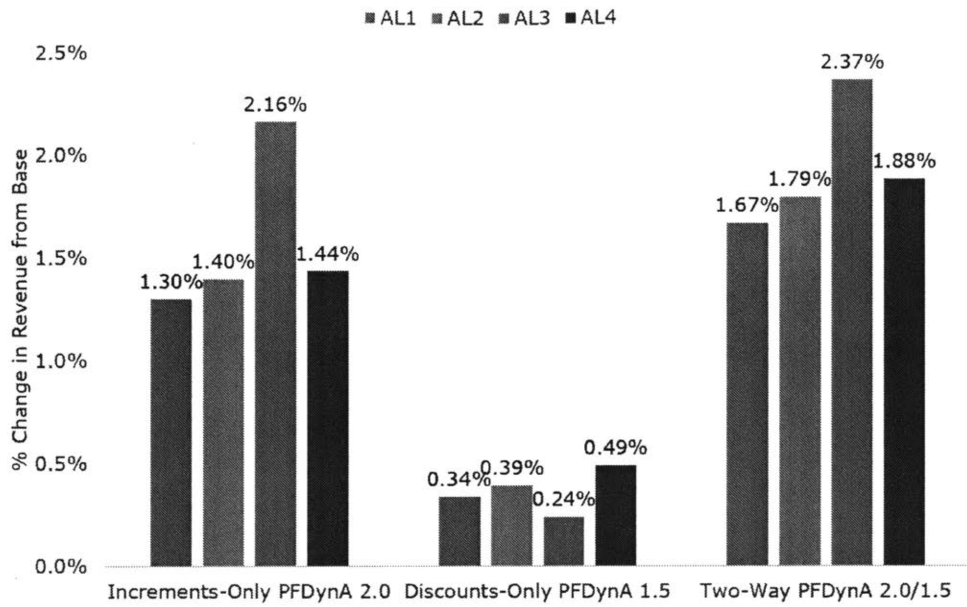

Figure 4.53: Percent change in airline revenue from base when all airlines use three different PFDynA heuristics in Network U10 (Medium Demand)

the three other airlines in the simulation who are incrementing. When AL4 stops using Increments-Only PFDynA, it does indeed see a nearly 7% increase in business passengers. However, this increase in passenger comes at the expense of yield, as AL4 gives up the benefits of price incrementing.

As shown in Figure 4.54, AL4's decision to discontinue the use of Increments-Only PFDynA in this network leads to lower revenues for all airlines compared to the all-airline case. AL4 sees lower revenues because it is no longer increasing its yields by providing increments, and other airlines see lower revenues because some of their business passengers shift to book with AL4, which is not incrementing. Since no airline has the incentive to discontinue use of the Increments-Only heuristic, its use by all airlines is a Nash equilibrium in Network U10.

The use of Discounts-Only and Two-Way PFDynA by all airlines is also a Nash equilibrium in Network U10. Discounts-Only PFDynA presents a particular first-mover advantage—its revenue gains are higher with fewer airlines using PFDynA. This is because the presence of even a single airline that is not discounting prices will allow airlines that are practicing Discounts-Only PFDynA to shift demand from that airline. This revenue benefit increases for AL1 as fewer and fewer airlines use PFDynA, as shown in Figure 4.55.

This suggests that the first airlines that are able to implement a discounting heuristic will likely see higher revenue gains than airlines that wait and practice the heuristic after other

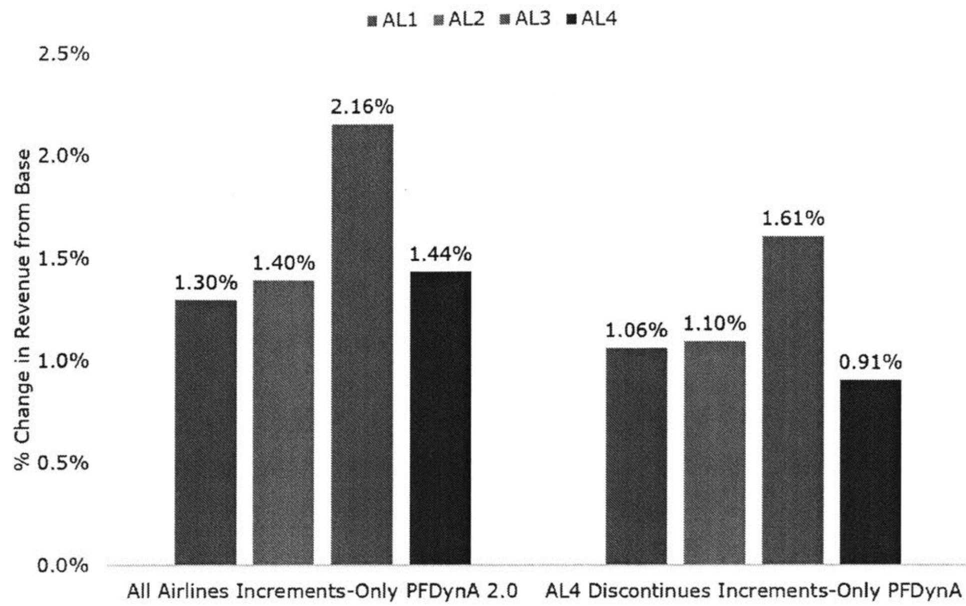

Figure 4.54: Percent change in airline revenue from base when three or four airlines use Increments-Only PFDynA in Network U10 (Medium Demand, Q = 2.0)

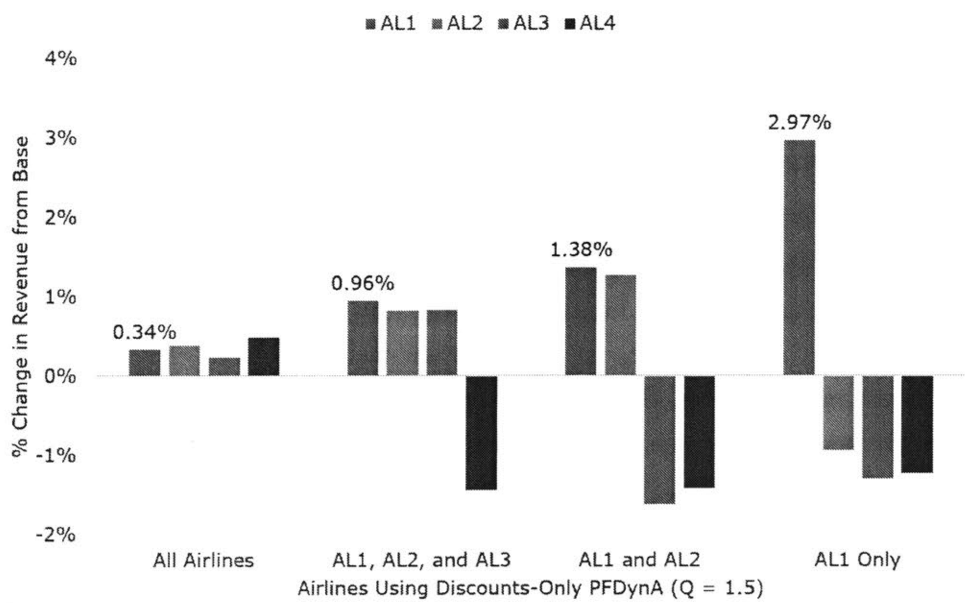

Figure 4.55: Percent change in airline revenue from base when various airlines use Discounts-Only PFDynA in Network U10 (Medium Demand, Q = 1.5)

airlines have already adopted it. Furthermore, airlines that are not practicing Discounts-Only PFDynA see revenue losses when the heuristic is practiced by their competitors, due to the shift of demand induced by discounting.

In this network, practicing PFDynA is the best response to the use of PFDynA by a competitor airline. In Network U10 with these particular simulation parameters, there is no incentive for any airline to discontinue the use of the heuristics in the pursuit of higher revenue gains. Moreover, each of the three PFDynA heuristics produces revenue gains when used by all airlines in the simulation. This outcome is similar to the multiple-airline case in Network A2TWO.

4.5.5 Sensitivity analyses: Segmentation inaccuracy and RM methodology

In Figure 4.56, we find that the choice of revenue management optimization method makes little impact on the performance of PFDynA.When AL1 uses ProBP or UDP instead of DAVN, baseline revenues and load factors show little change relative to the case when it uses DAVN. The PFDynA heuristics have nearly the same effect over the base case regardless of the RM optimization method chosen.

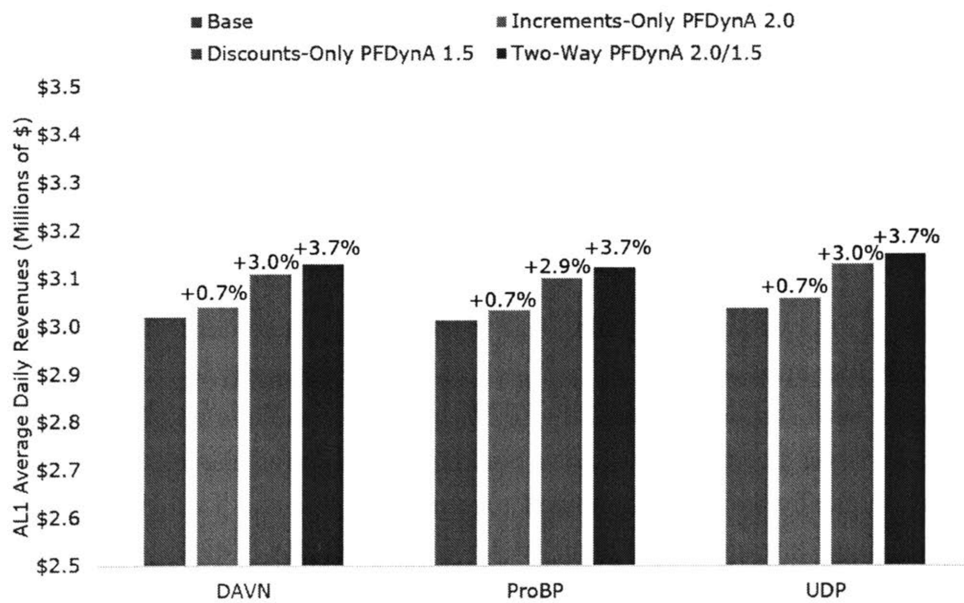

Figure 4.56: Average AL1 daily revenues when only AL1 uses PFDynA with three different RM optimization methods in Network U10 (Medium Demand, Q = 2.0/1.5)

Tests of PFDynA with hybrid forecasting and fare adjustment also lead to similar PFDynA performance as standard forecasting. When AL1 moves to HF/FA in Network U10 (using FRAT5c and a scalar of 0.25), its revenues increase by 1.2% over the standard forecasting base case. These revenue gains come with a decrease in load factor and an increase in yield, as HF/FA results in lower classes being closed more often than standard forecasting. When PFDynA is implemented on top of HF/FA, revenues increase by 0.5%, 3.0%, and 3.6% over the HF/FA base for Increments-Only, Discounts-Only, and Two-Way PFDynA, respectively, as shown in Figure 4.57.

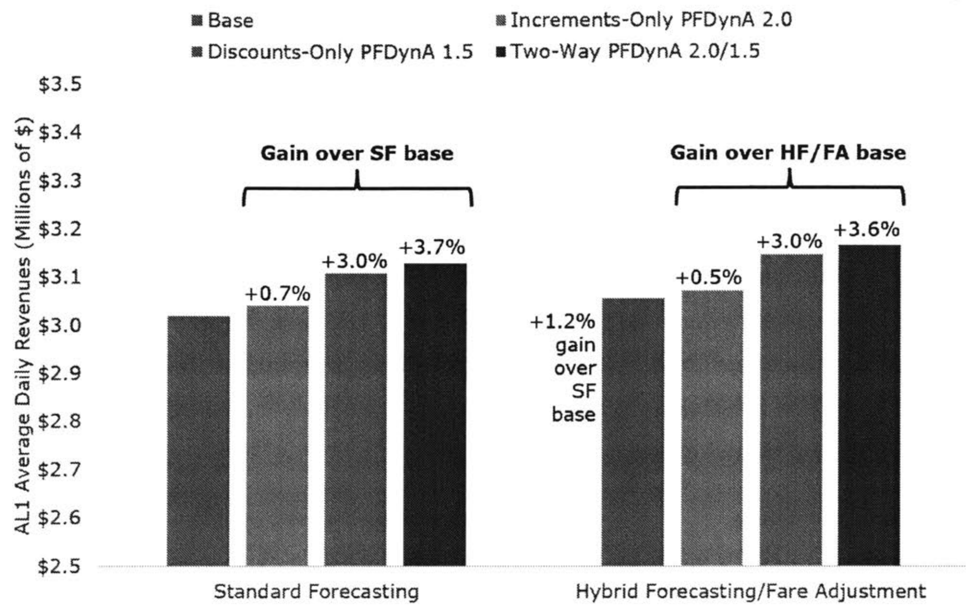

Figure 4.57: Average AL1 daily revenues when only AL1 uses PFDynA with standard forecasting (SF) or hybrid forecasting and fare adjustment (HF/FA) in Network U10 (Medium Demand, DAVN, Q = 2.0/1.5)

These changes in revenue are nearly identical to those under a standard forecasting regime. The exception is Increments-Only PFDynA, which produces slightly lower revenue gains under HF/FA than with standard forecasting (0.5% versus 0.7%, Figure 4.57). This is because HF/FA leads to fewer business passengers booking in the lower classes to which PFDynA increments are typically applied. Combined together, HF/FA and Two-Way PFDynA with Q = 2.0/1.5 produces a revenue gain of about 4.9% over the DAVN with standard forecasting base case when used by one airline in Network U10.

Finally, we test the effects of segmentation errors on PFDynA performance. Figure 4.58 shows the performance of the three different PFDynA heuristics when used by only AL1 with various passenger segmentation accuracy levels. When used by only a single airline, Increments-Only PFDynA quickly becomes revenue-negative with only slight segmentation inaccuracy (20%). This is because other airlines in the network will be able to capture leisure passengers from AL1 that receive mistaken increments. In contrast, the Discounts-Only and Two-Way heuristics are revenue positive at all accuracy levels down to 60%. Mistakenly

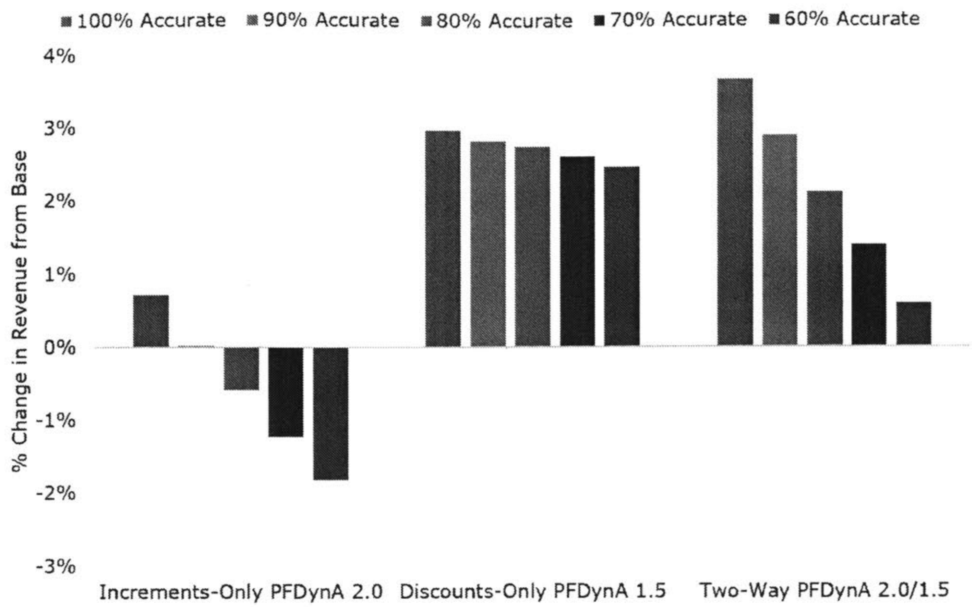

Figure 4.58: Percent change in AL1 revenue from base when only AL1 uses three different PFDynA heuristics with segmentation inaccuracy in Network U10 (Medium Demand, Q = 2.0/1.5, DAVN)

giving lower prices to some business passengers may result in unnecessary discounting, but will also result in the capture of business passengers from other airlines. For all three heuristics, lower segmentation accuracy leads to worse heuristic performance.

When all of the airlines use the heuristic (Figure 4.59), Two-Way PFDynA is still revenue positive for AL1 (and all other airlines) regardless of segmentation accuracy. However, Increments-Only PFDynA is now also revenue positive at all accuracy levels, and Discounts-Only PFDynA becomes revenue-negative below accuracy rates of about 80%. The better performance of Increments-Only PFDynA with inaccuracy in this case is because there is no airline that is not practicing the incrementing strategy that is able to capture the leisure passengers that receive mistaken increments as a result of segmentation inaccuracy. With the Discounts-Only heuristic, giving unnecessary discounts to business passengers will attract relatively few additional bookings when all airlines are using the heuristic, and will cut into the revenue gains from forecast spiral-up and leisure demand stimulation.

Generally, segmentation accuracy matters more for Increments-Only PFDynA when there are relatively few airlines using the heuristic, and more for Discounts-Only PFDynA when there are many airlines using the heuristic. In both the single-airline and all-airline case for Network U10, Two-Way PFDynA was revenue-positive at all accuracy levels tested.

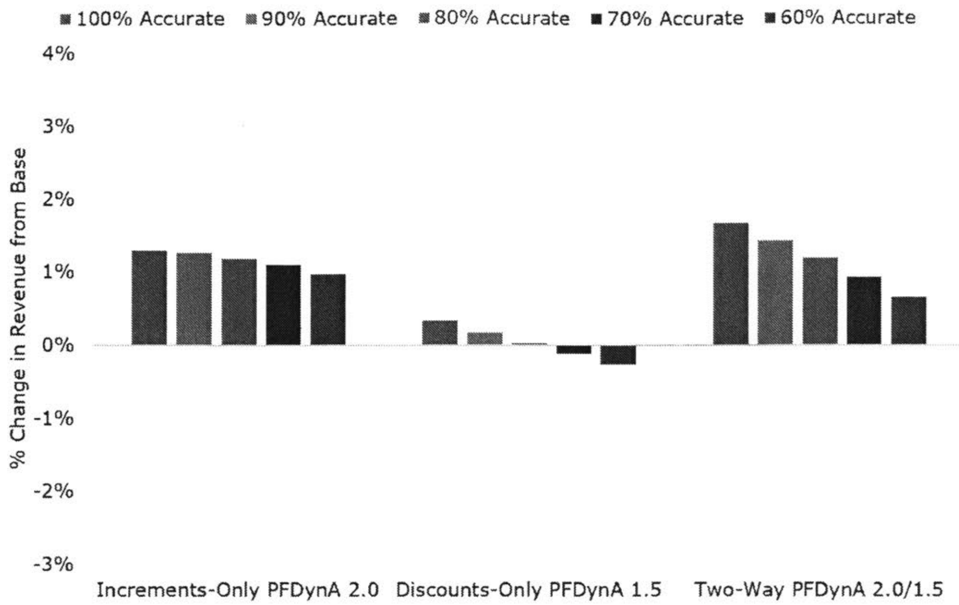

Figure 4.59: Percent change in AL1 revenue from base when all airlines use PFDynA with segmentation inaccuracy in Network U10 (Medium Demand, Q = 2.0/1.5, DAVN)

4.5.6 Sensitivity analysis: Unbounded PFDynA

In previous tests of PFDynA, price adjustments were bounded by the prices of the adjacent classes above and below the lowest available class. These bounds serve several purposes—they help prevent increments from causing fare inversions, and they also cause the modified prices to be fairly close to those that would be offered by traditional pricing and RM. With bounded price adjustments, the increments and discounts provided by PFDynA are relatively small, and relatively few customers book at modified fares. It is worthwhile to investigate the sensitivity of PFDynA to these bounds.

In *Unbounded PFDynA*, the optimal price $f_k^* = f_k + \Delta_k^*$ is computed in the same way as in the bounded version of the heuristic. However, as opposed to bounded PFDynA, which limits the adjusted price to the prices of the classes above or below the lowest available fare class, Unbounded PFDynA almost always sets the price of the lowest available class equal to exactly f_k^*. This should theoretically improve revenue performance, since the prices offered will be closer to the computed f_k^*, but may also have other adverse effects in a competitive environment, or when WTP estimation or customer type segmentation is not precise.

There are several special cases of Unbounded PFDynA that require additional attention. First, with Unbounded Discounts-Only PFDynA, we maintain the bound that stipulates prices are not decreased below the lowest fare in the airline's fare structure. This is because discounting prices below the lowest filed fare in the real world could spark a price war and a race to the bottom that could wipe out any simulated gains in revenue.

For Unbounded Increments-Only PFDynA, we must consider the possibility of fare inversions. Suppose that classes $\{1, 2, 3, 4\}$ are open with prices of $\{\$400, \$300, \$200, \$100\}$, respectively, and that f_k^* has been computed to be \$220. If we were to increment the price of Class 4 to \$220, there would be a fare inversion; a customer could just choose to purchase Class 3 at a price of \$200, ignoring the increment.

To address this, we first identify the highest open class with a fare lower than f_k^*. In our example, that class would be Class 3 at a price of \$200. We then increment the price of this class to equal f_k^*. In this case, a \$20 increment would be applied to Class 3, making its adjusted price equal to \$220. Finally, we then close all classes with fares strictly less than f_k^*. In the example above, this would leave us with classes $\{1, 2, 3\}$ open at prices $\{\$400, \$300, \$220\}$, and Class 4 closed. This method prevents fare inversions while ensuring that the customer will book at a price equal to at least f_k^*.

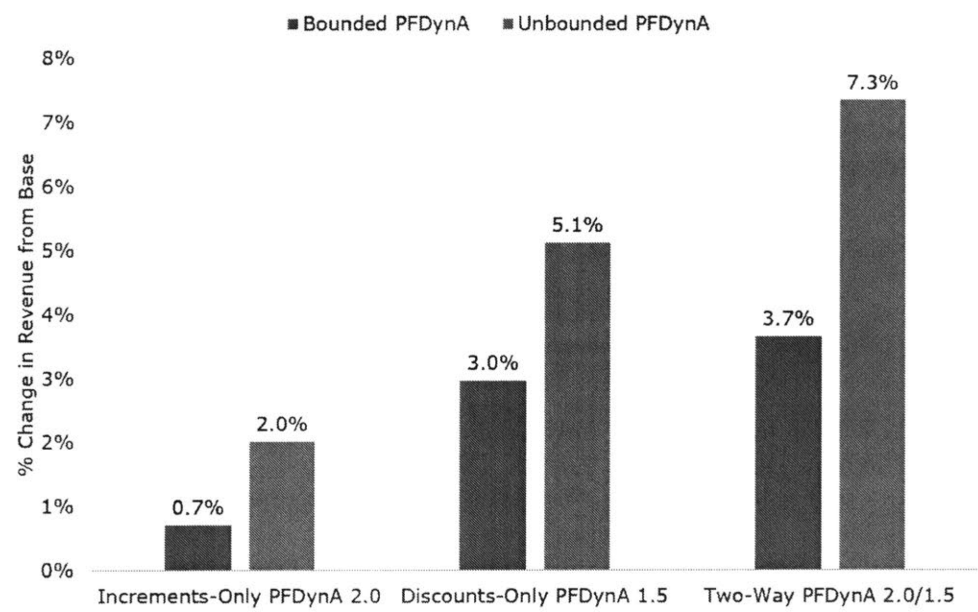

Figure 4.60: Percent change in AL1 revenue from base when only AL1 uses bounded or unbounded PFDynA in Network U10 (Medium Demand, Q = 2.0/1.5, DAVN)

Figure 4.60 shows the performance of Unbounded PFDynA compared to the bounded versions discussed earlier in this chapter. Note that when Unbounded PFDynA is used by a single airline, its performance exceeds that of bounded PFDynA. Both Increments-Only and Discounts-Only PFDynA see an increase in revenue performance of between 1.3% and

2.1%, and the performance of Two-Way PFDynA nearly doubles from 3.7% to 7.3%. While its revenue gains are higher, Unbounded PFDynA still produces revenue gains through the same mechanisms as bounded PFDynA—namely, through an increase in business passenger yield from incrementing, demand stimulation and share shift from discounting, and, in some cases, forecast spiral-up. .

Figures 4.61 and 4.62 show the familiar patterns of lower load factors and higher yields as a result of Increments-Only PFDynA in the unbounded case, as well as higher load factors for Discounts-Only PFDynA. Note that unlike bounded Discounts-Only PFDynA, the unbounded version of the discounting heuristic does not produce an increase in yield.

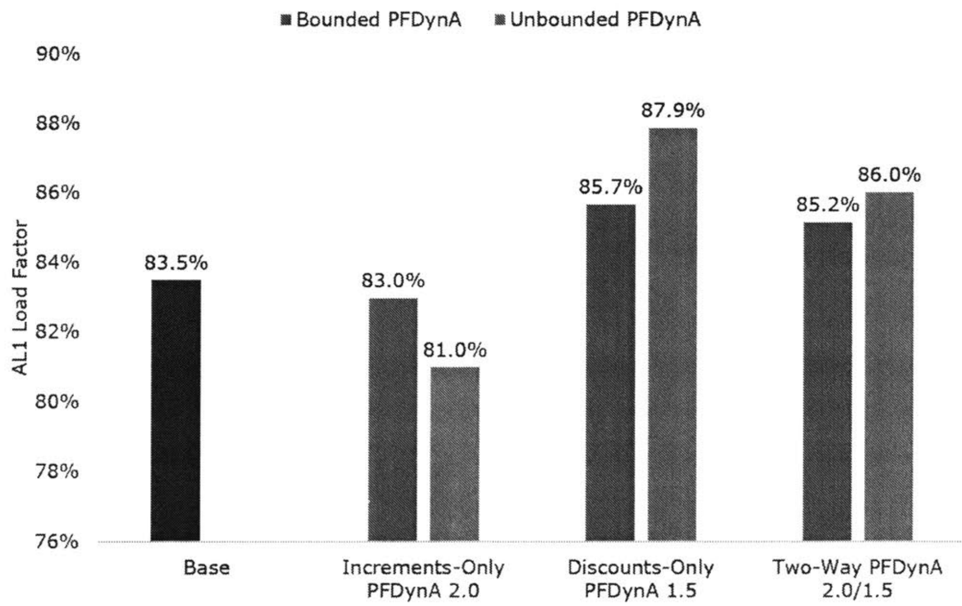

Figure 4.61: AL1 load factors when only AL1 uses bounded or unbounded PFDynA in Network U10 (Medium Demand, Q = 2.0/1.5, DAVN)

Since more passengers book with a discount under Unbounded Discounts-Only PFDynA, and since the discounts are generally for higher amounts, the effects of forecast spiral-up are not enough to overcome the reduction in average fares paid by leisure passengers. Nevertheless, revenues increase by 5.1% over the base as demand is stimulated and passengers shift from other airlines to the airline that is providing the discount.

When Unbounded PFDynA is used, more passengers book with PFDynA increments or discounts than in the bounded case, as shown in Figure 4.63. However, it is worth noting that most passengers still book at filed fares. Even with Unbounded Two-Way PFDynA (with Q-multipliers of 2.0 and 1.5), over 76% of passengers book at filed fares, about 15% receive

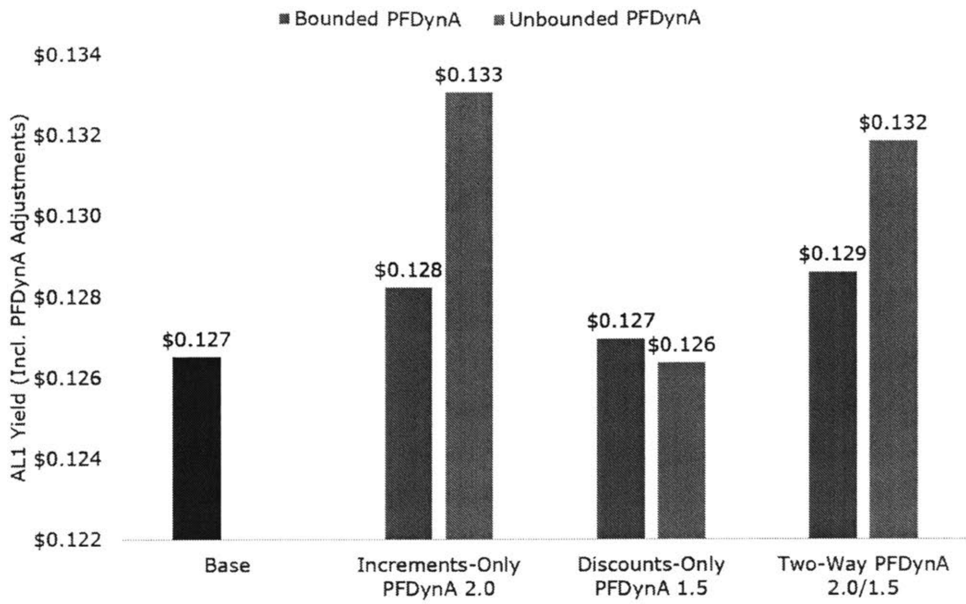

Figure 4.62: AL1 passenger yield (including PFDynA adjustments) when only AL1 uses bounded or unbounded PFDynA in Network U10 (Medium Demand, Q = 2.0/1.5, DAVN)

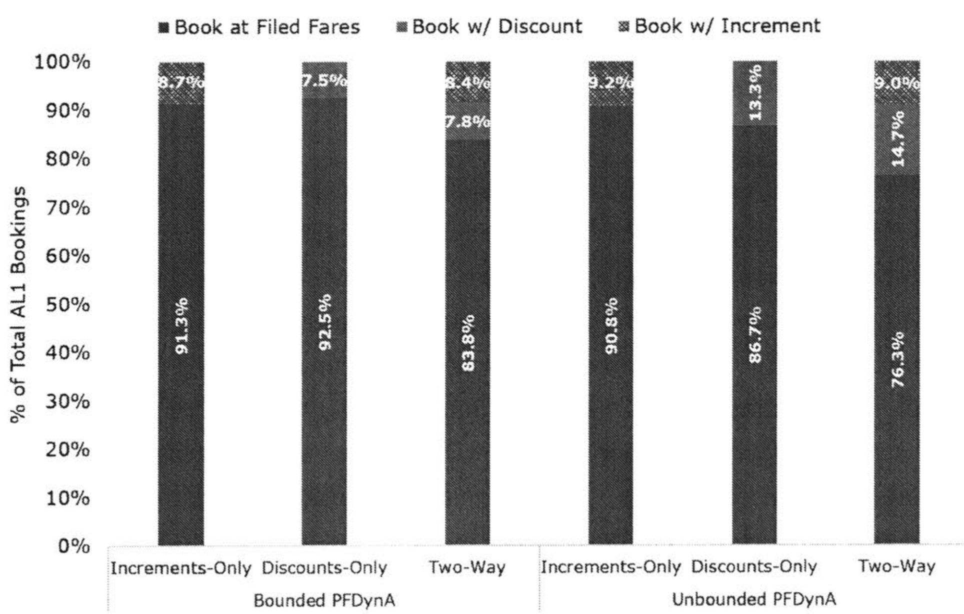

Figure 4.63: Percent of AL1 passengers booking with PFDynA adjustments when only AL1 uses bounded or unbounded PFDynA in Network U10 (Medium Demand, Q = 2.0/1.5, DAVN)

a discount, and about 9% book at an incremented price. The Two-Way PFDynA pattern of early incrementing for business passengers and late discounting for leisure passengers is reflected in the average fares paid by each passenger type in Figure 4.64.

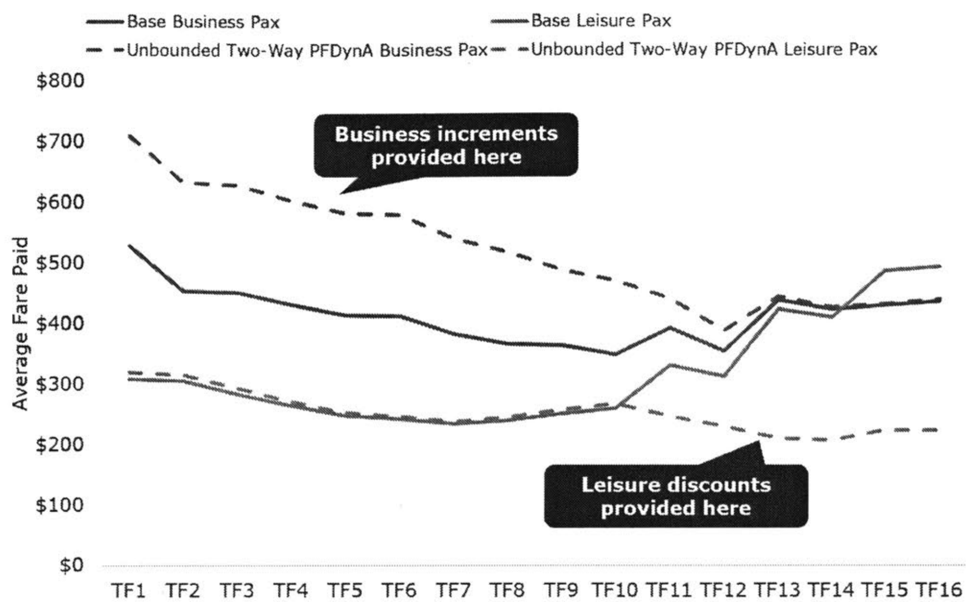

Figure 4.64: Average fare paid by passenger type when only AL1 uses bounded or unbounded Two-Way PFDynA in Network U10 (Medium Demand, Q = 2.0/1.5, DAVN)

When all airlines use Unbounded PFDynA, AL1's revenue gains increase from the bounded case, as shown in Figure 4.65. Most of this increase is driven by the revenue gain from Increments-Only PFDynA. Since all airlines are incrementing prices more frequently than in the bounded case, more revenue will be extracted from business passengers. As in the bounded case, if one airline discontinues Unbounded Increments-Only PFDynA, revenues of all airlines fall. The airline that has stopped incrementing will no longer see the revenue gains from the increments, and other airlines will lose business passengers and revenue to the airline that is not incrementing. Unbounded Discounts-Only PFDynA is also a Nash equilibrium in this network for similar reasons.

Despite the revenue performance of Unbounded PFDynA, there may be reasons why airlines would still prefer to use the bounded versions of the heuristic. First, the bounded heuristic makes smaller adjustments to filed fares than the unbounded heuristic. This may make airlines feel more comfortable in practicing dynamic price adjustment, and could lead to less instability in the marketplace by avoiding wildly changing prices.

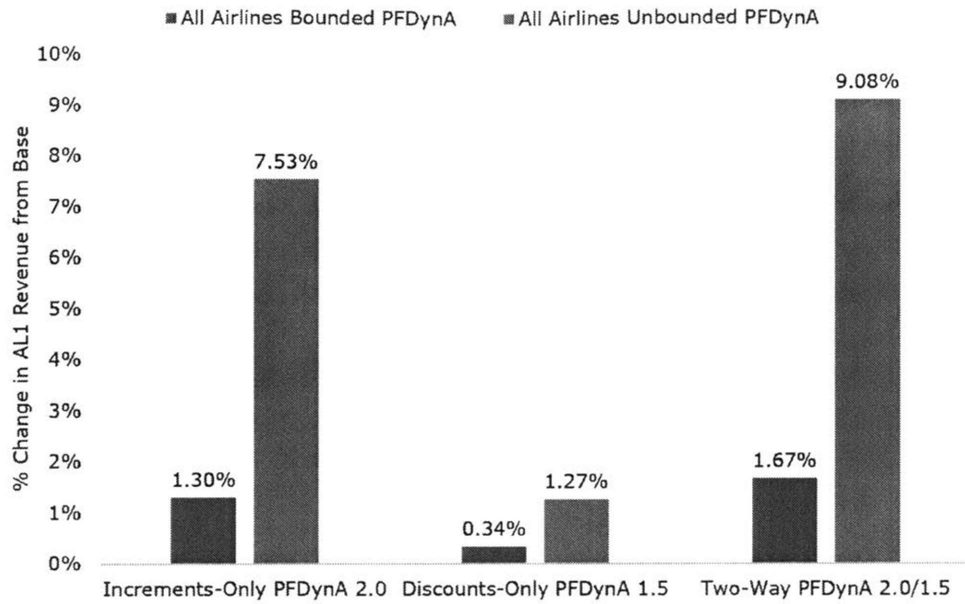

Figure 4.65: Percent change in revenue from base when all airlines use bounded or unbounded PFDynA in Network U10 (Medium Demand, Q = 2.0/1.5, DAVN)

Unbounded PFDynA is also highly sensitive to WTP estimation or passenger type segmentation errors. Figure 4.66 shows the revenue performance of bounded and unbounded Increments-Only PFDynA when passenger segmentation accuracy is not perfect. Note that while the unbounded version of the heuristic outperforms the bounded version when used by a single airline, even a slight degree of segmentation inaccuracy can wipe out the revenue gains and cause a loss in revenue. When the magnitude of the increments and discounts is greater, mistakes in segmentation or estimation can make a greater impact on the bottom line. The bounds thus serve as a hedge against inaccuracies in these processes.

In sum, removing the bounds on PFDynA price adjustments can increase the revenue gains of the heuristics. This makes theoretical sense, since the prices offered will be closer to the f_k^* computed during the PFDynA calculations. These improvements in revenue performance hold both when PFDynA is used by a single airline and in competition. However, the unbounded heuristics are potentially more risky—since their price adjustments are larger, they could be more likely to change well-established pricing patterns, potentially drawing responses from competitors. The unbounded heuristics are also more sensitive to inaccuracies in passenger segmentation accuracy or willingness-to-pay. For these reasons, airlines may wish still to apply fare class bounds in their implementations of dynamic price adjustment as a hedge against this uncertainty and risk.

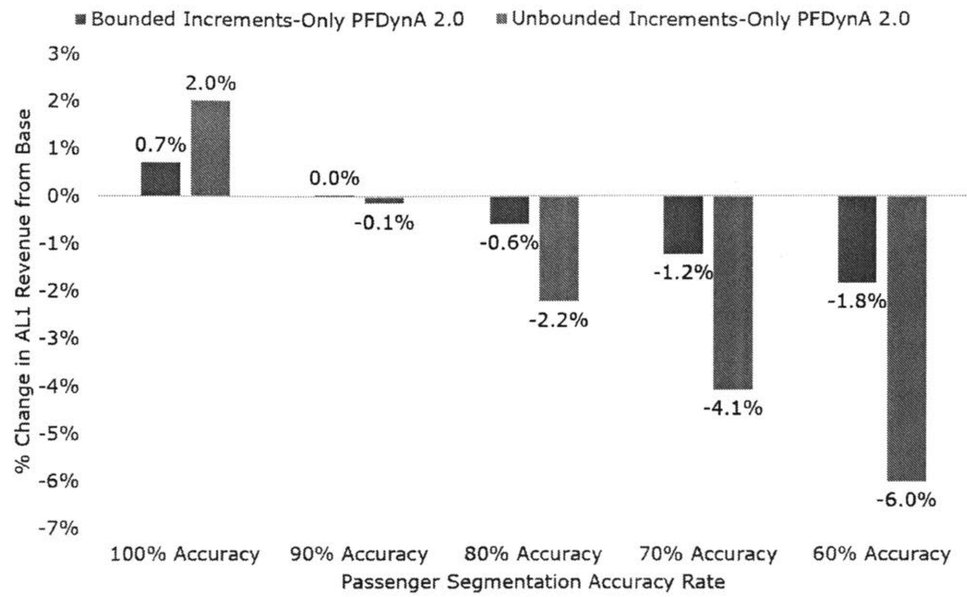

Figure 4.66: Percent change in AL1 revenue from base when only AL1 uses bounded or unbounded Increments-Only PFDynA in Network U10 with segmentation inaccuracy (Medium Demand, Q = 2.0/1.5, DAVN)

4.6 Conclusions from tests of PFDynA in PODS

In this chapter, we performed an extensive battery of tests of the PFDynA heuristics designed in Chapter 3 in the Passenger Origin-Destination Simulator (PODS), a complex agent-based simulation that models the interactions between hypothetical airlines and passengers. We tested versions of the heuristics in which business passengers were eligible for increments (Increments-Only PFDynA); in which leisure passengers were eligible for discounts (Discounts-Only PFDynA); and in which business increments and leisure discounts were both possible (Two-Way PFDynA).

These three heuristics were tested in a variety of PODS networks, from a simple single-airline, single-flight instance (Network A1ONE) to a complex network consisting of four airlines operating hundreds of total flights each departure day (Network U10). We also tested PFDynA under a number of conditions, including different RM optimization methods, forecasting methods, passenger segmentation accuracy, and fare structures.

In this wide range of tests, the PFDynA heuristics produced remarkably consistent revenue management outcomes and, in most cases, resulted in revenue gains for the airlines that used the heuristics. The revenue increases from PFDynA over traditional airline revenue management were in the range of 1 – 6% in our PODS simulations, which closely matches the results reported in other academic studies of dynamic pricing compared to traditional RM (Zhang and Lu, 2013; Fiig et al., 2016; Kumar et al., 2017).

Our simulations show that PFDynA generally produces revenue gains through several mechanisms:

- Providing increments to high-WTP passengers causes an **increase in passenger yield**. Passengers facing higher prices may choose to book with an increment, or may buy-up to a higher class with fewer onerous restrictions to avoid the increment if the airline is using a restricted fare structure. However, incrementing may make air travel unaffordable for some business passengers, or lead a passenger to book with another airline that is not incrementing prices. Increments are typically given early in the booking process when lower classes are open and fares are relatively low.

- Providing discounts for lower-WTP passengers aims to **stimulate demand**. Offering a discount may make air travel affordable for some customers who otherwise would have chosen not to fly. In a competitive network, discounting may also attract passengers from other airlines that are not offering discounts. However, providing too many discounts to passengers who would have booked at normal prices may lead to revenue dilution. PFDynA typically offers discounts to lower-WTP passengers when lower fare classes have been closed by the RM system and the offered fare is relatively high.

- Both incrementing and discounting lead to more bookings in higher fare classes over the base case. This leads to an outcome we call **forecast spiral-up**. When the revenue management forecaster starts to observe more bookings in higher classes, it adjusts upwards its future forecasts for those classes. The RM optimizer then begins to protect more seats for higher classes for future departures, and closes less-expensive classes more often. This can lead to an improvement in fare class mix and an increase in yield, even when discounting heuristics are used.

- The ability to **segment booking requests** into different customer types (and the ability to charge different prices for each of these segments) is critical for the success of PFDynA. Tests of unsegmented PFDynA heuristics did not produce revenue gains in monopolistic settings or in multi-airline networks when all airlines used PFDynA. This is because no single Q-multiplier scheme is able to accomplish the goals of PFDynA—incrementing low prices for high-WTP while keeping prices low for low-WTP customers early in the booking process, and then discounting for low-WTP customers while maintaining high price levels for high-WTP customers later in the booking process.

  Demand segmentation does not need to be complicated or perfectly accurate; our tests considered just two demand segments and made rudimentary guesses about conditional

WTP distributions for each segment, yet found that revenue gains were still possible even if airlines could not perfectly identify each request. In some scenarios (e.g., only AL1 using Discounts-Only PFDynA in Network U10), segmentation inaccuracy of as high as 40% still produced revenue increases from the base case. Depending on the method, PFDynA can sometimes produce revenue losses in some cases when segmentation is inaccurate or not possible.

As tested in this chapter, PFDynA made simple estimates about customers' conditional WTP that did not change as a function of the assortment of products offered to the customer. Yet as we found in Chapter 3, the presence of a more desirable or less-expensive alternative could change a customer's conditional willingness to pay. In the next chapter, we consider an extension of PFDynA in which itinerary attributes are taken into account in the calculation of dynamic price adjustments. We test this variant, which prices multiple flights at the same time, in relation to the Two-Way PFDynA heuristic tested in this chapter.

5 Simultaneous Dynamic Pricing of Multiple Differentiated Flights

5.1 Motivation

In the PFDynA price adjustment methodology motivated and developed in Chapter 3 and tested in the PODS revenue management simulator in Chapter 4, price adjustment decisions were made for each itinerary product individually. A dynamic price adjustment was generated based on the fare that would normally be offered by the airline's revenue management system for that flight. As is the case in most dynamic pricing models, the interactions between the products offered in the assortment are not directly modeled, and are only captured indirectly through the input Q-multiplier.

As discussed in Chapter 3, the presence of additional products in the assortment offered to the customer can change the customer's choice behavior. The presence of a low-cost alternative could cause a customer to select a flight that is not the best fit for her schedule preferences, for example. So far, our heuristics have focused on modifying the prices for each itinerary individually, without considering any possible interactions or substitutability between different itineraries. Dynamic price adjustments that do not explicitly consider differentiated attributes of itineraries could inaccurately assess customer choice and miss out on potential revenue gains.

For instance, a heuristic that simultaneously prices multiple itineraries could lower the price on a less-attractive itinerary for price-sensitive customers, while increasing the price on a more-attractive itinerary for schedule-sensitive customers. To date, the literature on heuristics for airline revenue management has not developed a model for simultaneously pricing multiple itineraries in a way that is consistent with existing revenue management practices. The relative gains of simultaneously pricing multiple itineraries compared to pricing each flight independently have also not been explored.

In this chapter, we develop and evaluate an extension to the PFDynA heuristic—Simultaneous PFDynA—that jointly computes the prices for multiple itinerary products. The itinerary products considered in our exposition are assumed to be *differentiated* along one or more dimensions. For instance, itinerary products may differ in terms of departure time, number of connections, or elapsed time. Instead of pricing itinerary products one at a time, we jointly find a set of price adjustments to the lowest-available fare class for all itineraries available by an individual airline to each customer in a given OD market.

We also modify the dynamic price adjustment equation to explicitly consider the choice probabilities for each itinerary. Unlike other work on pricing multiple differentiated products, many of which use the common multinomial logit (MNL) choice model, we use a locational choice model (Gaur and Honhon, 2006; Ulu et al., 2012) to model customer choice between differentiated itineraries. This allows for our dynamic prices to better take into account differences between itineraries, but also adds complexity to the problem, particularly when more than two itineraries are available.

The remainder of the chapter is structured as follows: we first discuss some of the existing literature regarding dynamic pricing of multiple itineraries. We then present our model for customer choice between horizontally differentiated itineraries, discuss how it is different from past work, and introduce our heuristics for simultaneous dynamic pricing. Using the PODS simulator, we then show results of practicing these heuristics compared to more simple approaches. We then extend the model to vertically differentiated attributes, and to three or more flights.

5.2 Literature review: Simultaneous dynamic pricing of multiple flights

In the past ten years, there have been several operations research papers which have proposed models for dynamically pricing multiple substitutable flights. These papers typically focus on proving the optimality of their models, and largely do not consider how or if these methods could be integrated into existing airline RM systems. Comparisons in performance between simultaneous dynamic pricing and existing RM methods are also not discussed.

Zhang and Cooper (2009) presented one of the first RM models incorporating multiple substitutable itineraries in an airline environment. Their model assumes a Markov decision process to describe how customers make selections between different flights. However, the focus of their paper is the performance of their heuristics relative to a theoretical upper bound, and not the performance of dynamic pricing in relation to existing airline RM methods. It is also unclear if or how their model would interact with current RM systems.

Apart from Zhang and Cooper (2009), several papers in the operations research literature have described other approaches for multiple-flight RM (Ratliff et al., 2008; Akçay et al., 2010; Chen et al., 2010; Gallego and Wang, 2014). These papers typically use a multinomial logit (MNL) choice model to describe customer decision making between various itinerary options. The MNL model has many advantages—it is well studied and understood, and is easy to extend to multiple flights.

However, as discussed in the next section, the MNL model will tend to offer similar prices for flights with similar marginal costs of capacity, which may not be desirable in a situation with highly differentiated flights. Moreover, these papers do not typically consider the integration of dynamic pricing into existing RM systems, nor how the performance of dynamic pricing compares with traditional airline RM methods. Also unlike our approach, these models are not dynamic price adjustment methodologies, as they do not focus on adjusting prices for existing fare products contextually based on the characteristics of each shopping request.

Until this chapter, no work has considered dynamic pricing of multiple itineraries in a framework that is compatible with existing revenue management practices. Previous papers also do not provide a straightforward comparison between simultaneous dynamic pricing and flight-by-flight dynamic pricing. We attempt to fill both of these gaps in the literature with our model, which is introduced in the next section. As opposed to the papers above that use Markov or MNL models to model passenger choice, we use a locational choice model, also called a Hotelling line (Hotelling, 1929). This approach is more common in the economics literature, but relatively rare in the operations research literature to which most papers on dynamic pricing belong.

5.3 Models and heuristics for simultaneous dynamic pricing

5.3.1 Flight-by-flight dynamic pricing

First, consider a simple dynamic pricing problem with multiple flights. Suppose an airline operates two non-stop flights in a single isolated origin-destination market. Assume these flights are identical in every way except departure time: one of the flights (Flight 1) departs at 9am, and the other (Flight 2) departs at 8pm. The flights are shown in Figure 5.1.

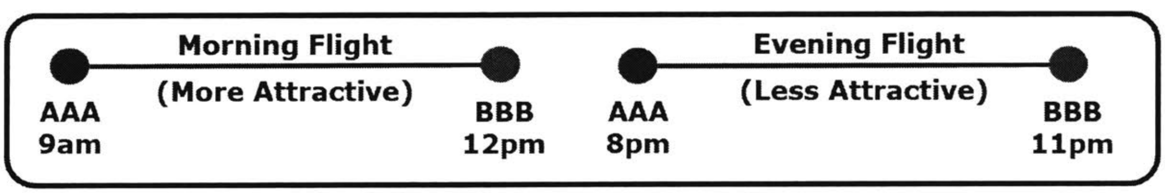

Figure 5.1: A single market with two nonstop flights

The airline wishes to set a price f_i for each flight $i \in \{1, 2\}$ to maximize expected revenue from each shopping session. First, suppose that the airline is pricing each flight individually. For each shopping session, the airline will solve these two equations to maximize expected revenue for each flight separately:

$$f_1^* = \arg\max_{f_1} [(f_1 - BP_1) \cdot Prob(1|f_1)] \tag{12}$$

$$f_2^* = \arg\max_{f_2} [(f_2 - BP_2) \cdot Prob(2|f_2)] \tag{13}$$

In these equations, f_i represents the price for Flight $i \in \{1, 2\}$, BP_i represents the bid price (marginal cost of capacity) for Flight i, and $Prob(i|f_i)$ represents the probability that Flight i is purchased at price f_i. As in Chapter 3, note that the marginal cost of capacity (i.e., the bid price) from the RM system is an input into this dynamic pricing equation, and not the output of a simultaneous optimization of price and availability. This allows dynamic price adjustment to be used in conjunction with any existing airline RM method that outputs a bid price.

We also do not yet specify a functional form for the purchase probability $Prob(i|f_i)$. This purchase probability could take many forms. For instance, in PFDynA in Chapter 3, we assumed that passengers had a conditional willingness-to-pay $\widetilde{\theta}_{x|S}$ for each flight, and that $Prob(i|f_i) = Prob(\widetilde{\theta}_{x|S} > f_i)$. However, this formulation does not explicitly consider the substitutability between the two flights. If we reasonably assume that each customer will only purchase one flight, actual purchase probability $Prob(i|f_i)$ is a function not only of the price of Flight i, but also the price of the other flight, as well as the attributes of both flights.

5.3.2 Simultaneous dynamic pricing of multiple flights

We now consider how the choice function could be modified to explicitly incorporate the presence of substitutable itineraries. If we assume that each customer will purchase at most one flight, the customer's purchase probability $Prob(i|f_i)$ is a function not only of the price of Flight i, but also the price of the other flight, as well as the attributes of both flights.

That is, with simultaneous dynamic pricing, we wish to jointly optimize the prices of both flights as follows:

$$\{f_1^*, f_2^*\} = \arg\max_{f_1, f_2}[(f_1 - BP_1) \cdot Prob(1|f_1, f_2) + (f_2 - BP_2) \cdot Prob(2|f_1, f_2)] \quad (14)$$

In the formulation above, we simultaneously select the prices for both flights to maximize the total expected revenue. This increases the dimensionality of the problem, as shown in Figure 5.2. The flight-by-flight problem can be seen as finding the optimal point on the expected revenue curve for each flight individually, whereas the joint dynamic pricing problem can be seen as finding the optimal point on the entire revenue surface.

The purchase probability $Prob(i|f_1, f_2)$ in the simultaneous pricing model incorporates the prices and attributes of both flights. This means we need to specify a choice model to describe how we assume customers will make choices between flights. One possibility, which is common in the operations research literature on pricing multiple flights (e.g. Dong et

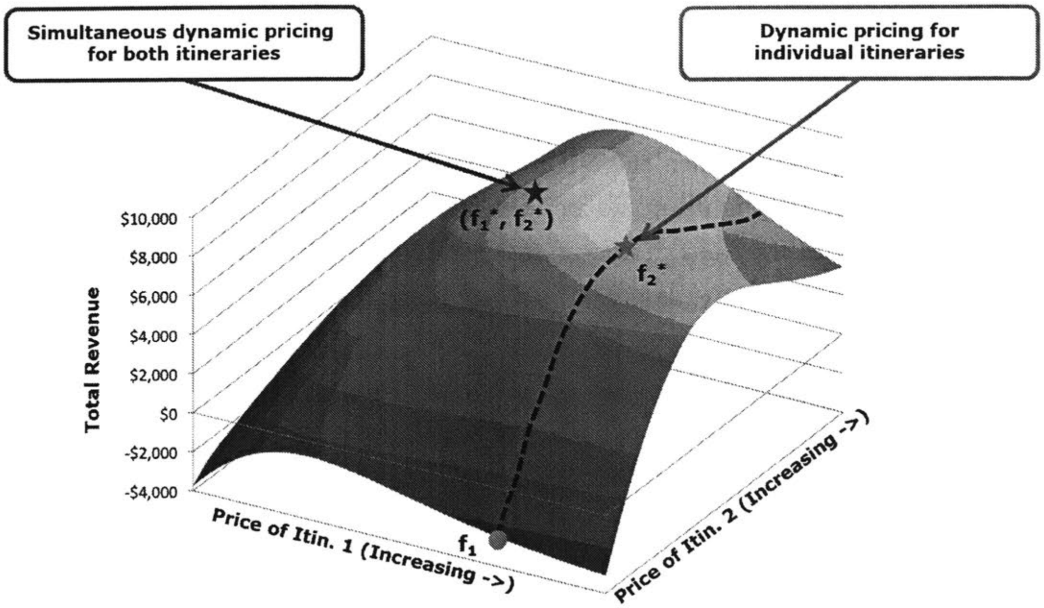

Figure 5.2: Stylized example of the revenue surface for simultaneous dynamic pricing

al. (2009); Chen et al. (2010); Suh and Aydın (2011)), is to use a multinomial logit (MNL) choice model.

In a basic MNL model, customers possess a utility V_i for each flight. For instance, this utility V_i could be a function of the price and some measure of schedule quality: $V_i = -\beta f_i + SQ_i + \epsilon_i$. Here, V_i is the customer's utility for Flight i, β is a measure of price elasticity, SQ_i is a measure of the customer's perception of the schedule quality of Flight i, and ϵ_i is a random error term. The customer's choice probability $Prob(i|f_1, f_2)$ is:

$$Prob(i|f_1, f_2) = \frac{e^{V_i}}{\sum_{j \in \{1,2\}} e^{V_j}} \tag{15}$$

As discussed earlier, the multinomial logit model has some advantages. For instance, adding additional flights into the model is relatively easy; we would just need to add additional terms to the denominator of the probability calculation. However, there are also some disadvantages to the MNL approach for our particular context.

Particularly, Aydın and Ryan (2000); Gallego and Wang (2014) and others have shown variants of the result that at optimal fares, $f_1^* - BP_1 = f_2^* - BP_2$. That is, the markup over the marginal cost of capacity (in our case, the bid price) would be the same for both flights when an MNL model is used to set optimal prices.

This means that if the bid prices for both flights were identical, the optimal prices for both flights would be the same using an MNL model, regardless whether one flight has more desirable attributes than the other. Specifically, if the bid price of both flights were zero, which often occurs in practice when there is more seat capacity than forecast demand on a flight, the prices of the two flights would be the same. As a result, the MNL-based model may not sufficiently capture the differences in schedule quality between flights when the marginal costs of capacity are the same.

We will use a different approach from past literature to construct our choice probabilities. This approach is commonly known as a Hotelling line, after a 1929 paper by economist Harold Hotelling. It is commonly used to model the choice of differentiated products in the economics literature. In the few papers in which the Hotelling model is used in the operations research literature, it is often referred to as a locational choice model (Gaur and Honhon, 2006; Ulu et al., 2012; Alptekinoğlu and Semple, 2016).

In the Hotelling choice model, we represent the attributes of both itineraries on a horizontal line, as shown in Figure 5.3. This is a natural representation of attributes like departure time, which are spread over the course of the day. Let D_i represent the departure time of each flight, and draw the line such that 0:00 is on the left end of the line and 24:00 is on the right end of the line.

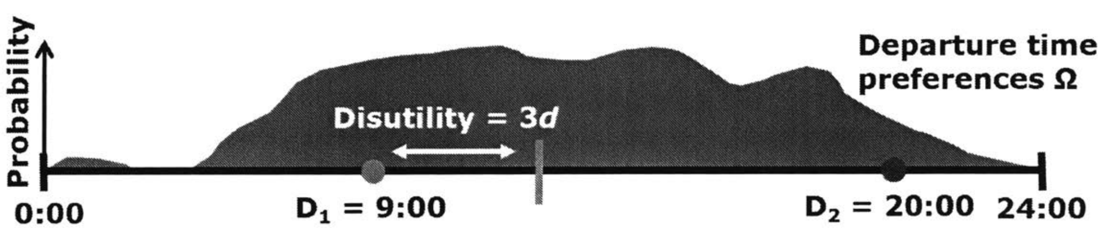

Figure 5.3: A Hotelling line for two horizontally differentiated flights

We assume that customers have random departure time preferences ω that are spread across the day according to some distribution Ω. It makes sense the preferences would vary from customer to customer; perhaps some customers prefer an early flight to make a morning meeting, while others prefer a later flight to return home after a vacation. At the same prices, there is no single departure time that is preferable to another for all customers. We call attributes like departure time *horizontally differentiated attributes*, since customers' preferences are heterogeneously distributed across the horizontal line.

For most customers, there will not be a scheduled flight that departs at exactly their preferred departure time ω. We assume that customers incur a value-of-time disutility d for each hour they move away from their preferred departure time in either direction. For instance, suppose that for a particular customer, $\omega = $ 12pm, $D_1 = $ 9am, and $D_2 = $ 8pm. This customer would face disutility $3d$ for selecting Flight 1 and $8d$ for selecting Flight 2. By adding together this disutility and the price of the flight, we can compute the perceived price (PP) of each flight:

$$PP_i = f_i + d|\omega - D_i| \tag{16}$$

As in Chapter 3, customers in our model also possess a maximum out-of-pocket willingness-to-pay budget θ which is distributed according to some distribution Θ. Θ could be different for different types of customers; for instance, for leisure customers versus those traveling for business. If the price of an itinerary $f_i < \theta$, we say the itinerary is *affordable*. Amongst the itineraries she can afford, the customer deterministically selects the itinerary with the lowest perceived price. Note that if the customer can only afford one itinerary, she will select that itinerary with probability 1. If a customer can afford neither itinerary, she will no-go and purchase nothing.

Then, the probability that a customer will purchase Flight 1 given fares f_1 and f_2 is:

$$Prob(1|f_1, f_2) = Prob(f_1 < \theta < f_2) + Prob(\theta > \max[f_1, f_2]) \cdot Prob(PP_1 < PP_2) \tag{17}$$

The first term in this expression represents the probability that a customer will be able to afford only Flight 1, in which case she will purchase it with probability 1. Note that if $f_2 < f_1$, the first term will equal zero. The second term represents the probability that the customer can afford both flights, and that she will select Flight 1 because it has a lower perceived price.

The probability $Prob(PP_1 < PP_2)$ in the expression above will depend on the distribution of departure time preferences ω. Specifically, it depends on the location of each customer's ω relative to an indifference point ω^*:

$$\omega^* = \frac{D_1 + D_2}{2} + \frac{f_2 - f_1}{2d} \tag{18}$$

If a customer has a departure time preference $\omega = \omega^*$, she will be indifferent between selecting Flight 1 and Flight 2 at given fares f_1 and f_2. Note that if the flights are priced identically ($f_1 = f_2$), the indifference point will be halfway between the departure times of the two flights. Otherwise, ω^* is a function of the difference in fares between the two flights, as well as the value-of-time disutility d.

If a customer has a departure time preference ω to the left of ω^*, she will prefer Flight 1 at fares f_1 and f_2. If her departure time preference is to the right of ω^*, she will prefer Flight 2 at those fares. Note that depending on the position of ω^* and ω, the customer may not always select the flight with the departure time closest to her departure time preference. If her value-of-time disutility is low, and if the difference in fares is large, the customer may decide to pick an itinerary with an unattractive departure time in order to save money.

This gives us enough information to specify the probabilities $Prob(1|f_1, f_2)$ and $Prob(2|f_1, f_2)$.

$$Prob(1|f_1, f_2) = Prob(f_1 < \theta < f_2) + Prob(\theta > \max[f_1, f_2]) \cdot Prob(\omega < \omega^*) \tag{19}$$

$$Prob(2|f_1, f_2) = Prob(f_2 < \theta < f_1) + Prob(\theta > \max[f_1, f_2]) \cdot Prob(\omega > \omega^*) \tag{20}$$

These choice probabilities can be substituted into the simultaneous dynamic pricing equation (14), which is reproduced as Equation (21) below:

$$\{f_1^*, f_2^*\} = \arg\max_{f_1, f_2}[(f_1 - BP_1) \cdot Prob(1|f_1, f_2) + (f_2 - BP_2) \cdot Prob(2|f_1, f_2)] \tag{21}$$

The computation of these probabilities requires that airlines have estimates of three parameters: the distribution of WTP Θ, the distribution of departure time preferences Ω, and the value-of-time disutility d. If these parameters vary by passenger type (i.e. leisure and business), and customers can be segmented by type, then a separate set of parameters would need to be estimated for each passenger type.

5.3.3 Heuristic for simultaneous dynamic pricing: Simultaneous PFDynA

In this section, we discuss a heuristic for simultaneous dynamic pricing of two flights with a single horizontally differentiated attribute (departure time). This heuristic extends the probability fare-based dynamic adjustment (PFDynA) concept introduced in Chapter 3.

In the Simultaneous PFDynA heuristic, leisure passengers are eligible for discounts from the original RM fare, and business passengers are eligible for increments, as in Two-Way PFDynA. To compute the adjusted prices, the pair of prices $\{f_1^*, f_2^*\}$ is computed as in equation (21). We assume airlines have knowledge of the underlying departure time preference distribution Ω and value-of-time disutility d.

For WTP, we follow Chapter 3 and assume that the passenger WTP distribution Θ_w in each market is Normally distributed for each passenger type w. To parameterize these distributions, we again use an input Q-multiplier: a ratio of the lowest filed fare in the market that is used to set the mean of the assumed WTP distribution. That is, if the lowest filed fare in the market is \$100, and the input Q-multiplier was 1.5, the estimated Normal distribution of passenger WTP would have a mean of \$100 * 1.5 = \$150. To complete the specification, we assume a coefficient of variation of $\gamma = 0.3$, as in Chapter 4. The distribution of WTP in this market is assumed to be Normal with a mean of \$150 and a standard deviation of 0.3 * \$150 = \$45.

Hence, we can compute:

$$Prob(\theta > f_i) = 1 - \Phi^{-1}\left(\frac{f_i - QMULT \cdot f_n}{\gamma \cdot QMULT \cdot f_n}\right) \tag{22}$$

where Φ^{-1} here is the inverse cumulative density function of the Normal distribution and f_n is the lowest filed fare in the market.

In practice, a grid search algorithm can be used to find the pair of fares that maximizes expected revenue. As in Chapters 3 and 4, we bound the incremented or discounted fares of the lowest-available fare class by the adjacent higher and lower fare classes in the pre-filed fare structure. This prevents fare inversions (by ensuring the incremented fare does not rise above the unadjusted fare in the next highest class) and ensures that the price adjustments are not too large relative to other price points in the fare structure.

5.4 Simulation results

Simultaneous PFDynA was tested in several PODS networks. Network A1TWO is a network that mimics the example given in Figure 5.1. It contains one market and one airline, which operates two non-stop flights in that market. The flights depart at 9am and 8pm. In the context of PODS, the 9am (morning) flight is seen as the more-attractive flight to all passengers, and the 8pm (evening) flight is seen as the less-attractive flight. Network A2FOUR is an extension of A1TWO with one market and two airlines, each of which operates a 9am and an 8pm departure. This allows us to test the performance of Simultaneous PFDynA in a competitive environment.

In this test of Simultaneous PFDynA, we assume that airlines can segment perfectly between leisure and business passengers. As described above, leisure passengers will be eligible for discounts from the RM fare under Simultaneous PFDynA, and business passengers will be eligible for increments. This matches the implementation of Two-Way PFDynA from Chapter 4.

The airline uses a six-class restricted fare structure, which is shown in Table 5.1. R1, R2, and R3 represent fare restrictions, such as a Saturday night minimum stay, that are onerous to the customers. R1 is the most onerous restriction, and R3 is the least onerous. The network was calibrated in three different scenarios: low demand (78% average load factor), medium demand (83% ALF), and high demand (87% ALF). As in Chapter 4, the airlines use leg-based EMSRb as their RM system, along with standard forecasting.

Class	Fare	Adv. Purch.	R1	R2	R3
1	$500	N/A	N	N	N
2	$400	3 days	N	N	Y
3	$300	7 days	N	Y	Y
4	$200	10 days	Y	N	Y
5	$150	14 days	Y	Y	N
6	$100	21 days	Y	Y	Y

Table 5.1: Fare structure used in Network A1TWO and A2FOUR

As mentioned above, the airlines have knowledge of the underlying departure time preference distribution Ω used by PODS. The value-of-time disutility is also known to the airline, and is set to $40 per hour for business passengers and $20 per hour for leisure passengers.[26] To specify estimates of WTP, the airlines use the Q-multiplier method described in Chapter 3. For these tests, the airlines use an input Q-multiplier of 2.0 for business passengers and 1.5 for leisure passengers, along with a coefficient of variation of 0.3. Note that the airline does not know the actual departure time preference or maximum WTP for any specific customer.

5.4.1 Tests of Simultaneous PFDynA in a single airline network

Figure 5.4 shows the revenue performance of Simultaneous PFDynA over the base case of traditional RM in the low demand, medium demand, and high demand scenarios. Simultaneous PFDynA shows good revenue performance, increasing revenues by 5.4% in the low-demand scenario to 7.4% in the high-demand scenario. This is among the higher end of the revenue gains for dynamic pricing presented in past work (Zhang and Lu, 2013; Fiig et al., 2016; Wittman and Belobaba, 2017c). It should be noted that the airline using Simultaneous PFDynA is assumed to have a great deal of information about the passengers: the actual departure time preference distribution and value-of-time disutility, as well as a good estimate of leisure and business WTP and 100% passenger type identification accuracy.

Simultaneous PFDynA leads to revenue gains because it increases both yields and load factors, as shown in Figure 5.5. By giving discounts to selected leisure passengers, particularly

[26] Other values of time were also tested with similar results. $40 and $20 per hour are similar to the values of time used for internal use by the U.S. Department of Transportation (2016).

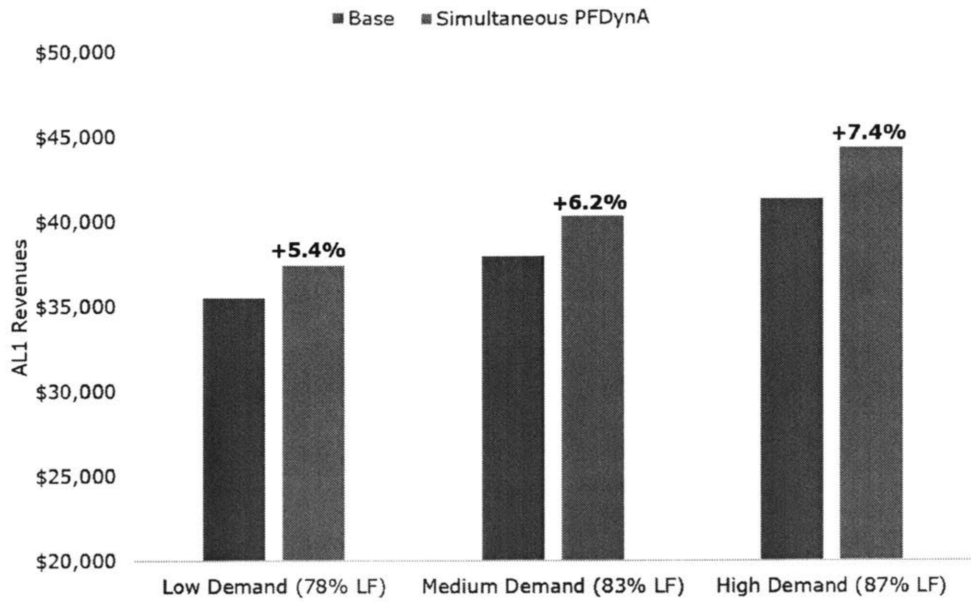

Figure 5.4: Revenue performance of Simultaneous PFDynA in three demand scenarios

those booking in higher fare classes, the airline can stimulate additional demand that otherwise would have chosen to no-go. Also, by incrementing fares for certain business passengers, particularly those booking in lower fare classes, the airline can increase yields as well.

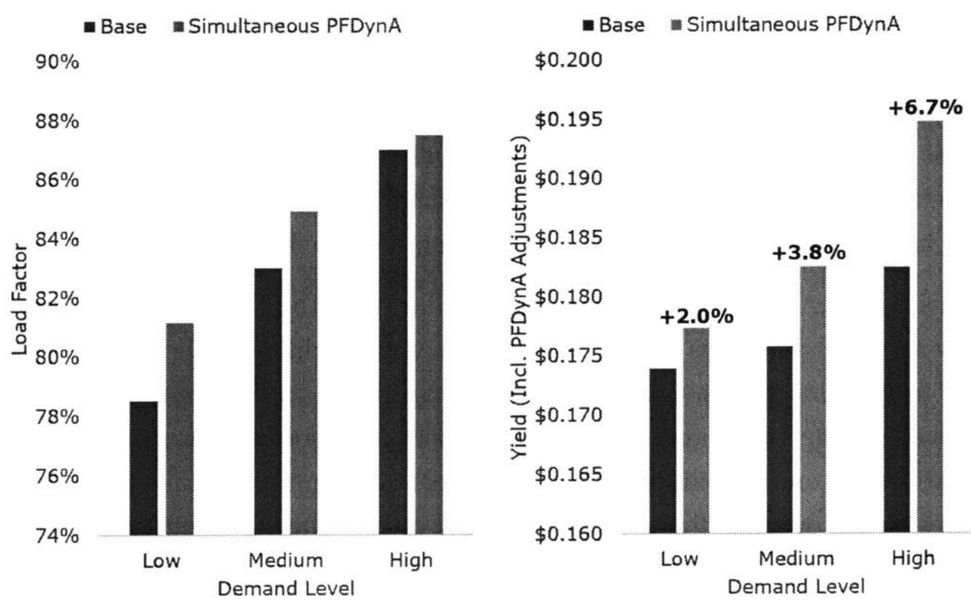

Figure 5.5: Average load factor (left panel) and passenger yield ($ per RPM; right panel) with Simultaneous PFDynA

Most passengers book with neither an increment nor a discount. On the more-attractive morning flight, 6.3% of passengers book with an increment, and 9.3% book with a discount. On the less-attractive evening flight, just 1.4% of passengers book with an increment, and 17.8% book with a discount (not shown). The targeted discounts on the evening flight shift demand from the morning flight to the evening flight, as shown in Figure 5.6.

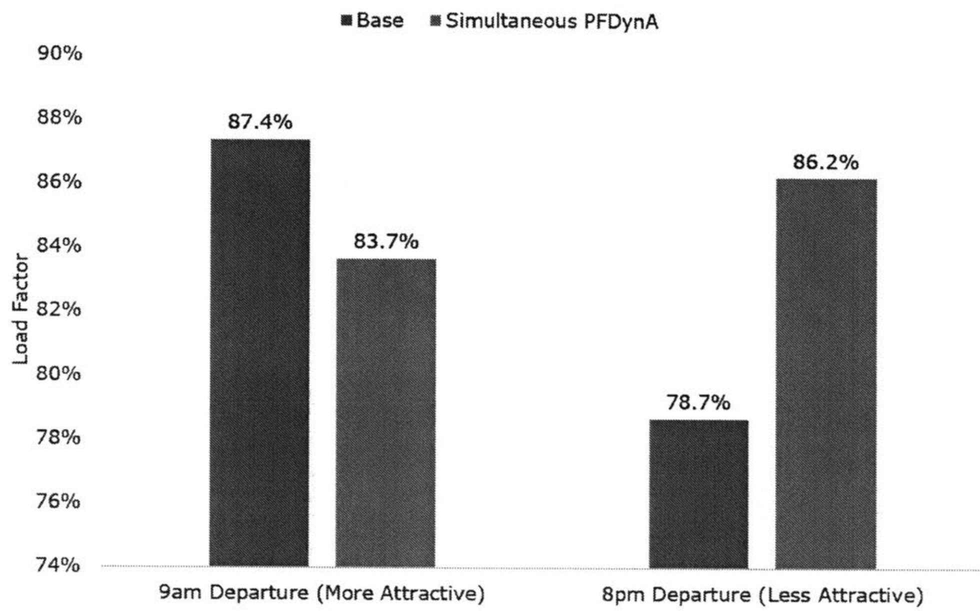

Figure 5.6: Average load factors by flight with Simultaneous PFDynA

In the base case, the more-attractive morning flight had a higher load factor of 87.4%, compared to just 78.7% for the less-attractive evening flight. By discounting the evening flight for certain customers, the load factor of that flight increases to 86.2% when Simultaneous PFDynA is used, and the load factor of the morning flight reduces to 83.7%.

Simultaneous PFDynA also improves the fare class mix on both flights. Figure 5.7 shows that both flights see fewer bookings in the least-expensive fare class (FC6) when Simultaneous PFDynA is used. This makes sense for the morning flight, since price-sensitive demand is shifted to the evening flight and increments encourage business passengers to buy up to higher classes on the morning flight. It may be more surprising to see that the fare class mix also improves on the evening flight, on which 17.8% of customers booked with discounts and only 1.4% booked with increments.

The rationale behind this result is the concept of *forecast spiral-up* which was discussed in Chapter 4. In Simultaneous PFDynA, discounts are provided to leisure customers booking in relatively higher fare classes. This causes more bookings to be recorded in these higher

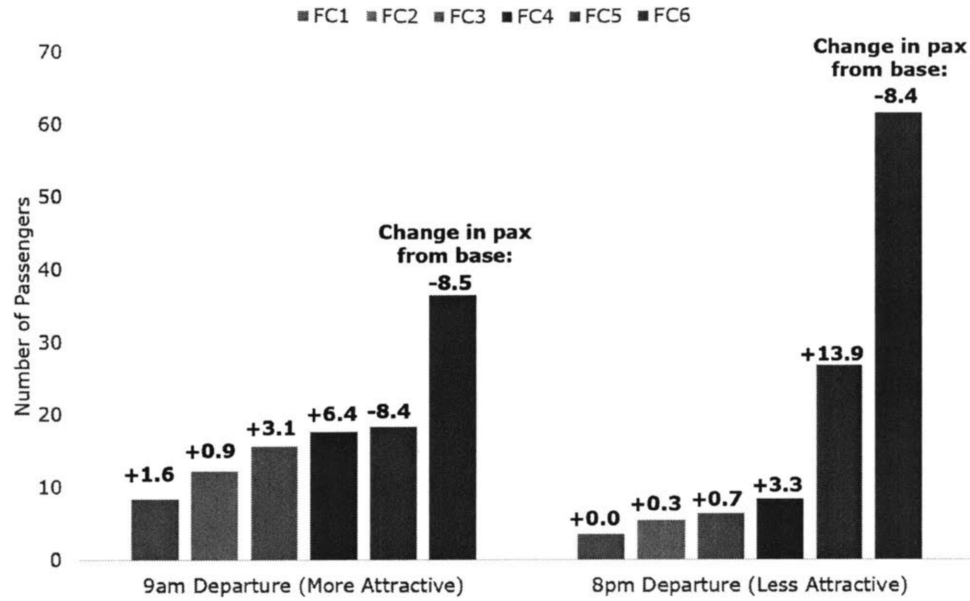

Figure 5.7: Fare class mix by flight when AL1 uses Simultaneous PFDynA

classes. Then, the RM forecaster begins to account for the additional bookings it is seeing in higher classes by increasing demand forecasts for those classes. As a result, the RM optimizer begins to protect more seats for higher classes, saving fewer seats for the lowest FC6.

Due to forecast spiral-up, we see higher bid prices throughout the booking period and FC6 closed more often when the airline is using Simultaneous PFDynA. We also see fewer early bookings (due to reduced availability in the least-expensive fare class relative to the base) and more later bookings (due to discounts being provided to leisure customers in higher fare classes) when Simultaneous PFDynA is used. Both of these effects combine to increase both yields and load factors, leading to revenue increases.

5.4.2 Adding a competitor

For the next series of tests, a competitor is added into the network. The competitor also offers a morning flight at 9am and an evening flight at 8pm. The competitor airline also uses EMSRb, and the demand in the new network (A2FOUR) baseline is recalibrated to an average load factor of 83.5% for both airlines in a medium-demand scenario. We will investigate two cases: only Airline 1 (AL1) using Simultaneous PFDynA, and both airlines using Simultaneous PFDynA.

As shown in Figure 5.8, Simultaneous PFDynA is revenue positive when practiced by one or

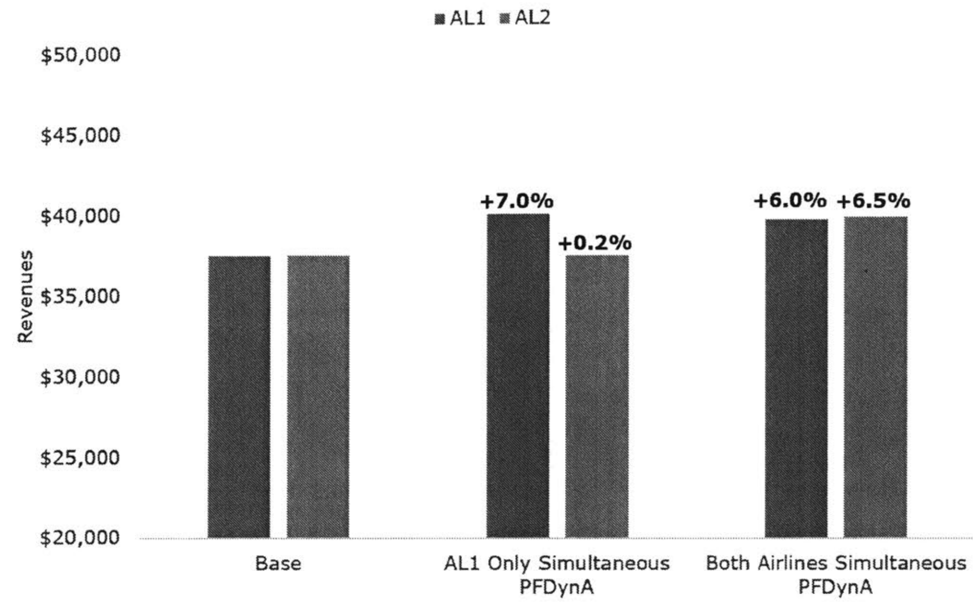

Figure 5.8: Revenues for AL1 and AL2 when one or both airlines use Simultaneous PFDynA

both airlines. If one airline practices Simultaneous PFDynA in this network, it sees a revenue gain of 7%, and its competitor also sees a small revenue gain of 0.2%. If both airlines practice Simultaneous PFDynA, they both see revenue gains between 6 and 6.5%, which is similar to the gains shown in the single-airline case in the medium-demand scenario.

When there are multiple airlines in the network, the use of Simultaneous PFDynA by only one airline can cause passengers to shift between carriers. When Simultaneous PFDynA offers discounts, a passenger that would have ordinarily booked with AL2 may decide to book with AL1 if that airline offers the passenger a discount. Conversely, if a business passenger is quoted an incremented price by an airline practicing Simultaneous PFDynA, she may decide to book at the normal fare with an airline that is not practicing PFDynA.

The net result of this behavior, as shown in Figure 5.9, is an increase in load factor for both AL1 and AL2 when only AL1 uses Simultaneous PFDynA. For AL1, the effect of giving discounts and gaining new bookings seems to outweigh the lost business passengers from incrementing. AL2 is also able to increase its load factor relative to the base, as it recaptures some of AL1's business passengers. AL2 will also tend to have better availability in the least-expensive FC6 than AL1, which shuts down FC6 more often as a result of higher bid prices from forecast spiral up. This means AL2 will gain more bookings in the lowest fare class, lowering its yield relative to the base. When both airlines practice Simultaneous PFDynA, they both see higher yields and load factors, as in the single-airline case.

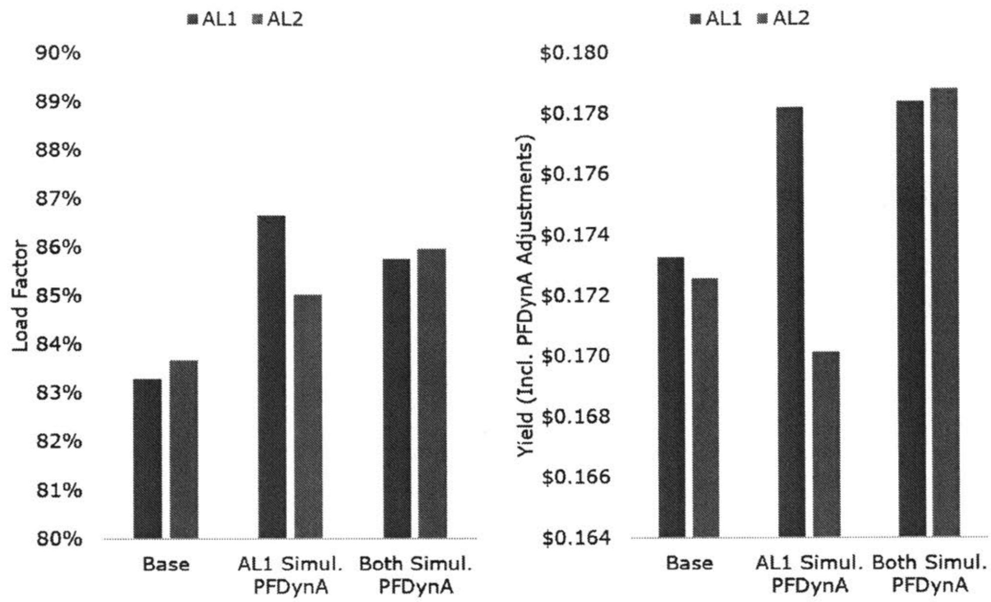

Figure 5.9: Airline load factors (left panel) and yields (right panel) when one or both airlines use Simultaneous PFDynA

5.4.3 Comparison to flight-by-flight dynamic pricing (Two-Way PFDynA)

Finally, we compare the performance of Simultaneous PFDynA to Two-Way PFDynA, as described in Chapters 3 and 4. With Two-Way PFDynA, business passengers are eligible for increments and leisure passengers are eligible for discounts, as in Simultaneous PFDynA. However, instead of the Simultaneous PFDynA choice model in Equation (21), a simpler equation (23) is used:

$$Prob(i|f_i) = f_i \cdot Prob(\theta > f_i) \tag{23}$$

In Two-Way PFDynA, the characteristics or prices of alternative itineraries are not taken into consideration in the choice probability, which simply evaluates the probability that a customer's WTP exceeds the fare. As in Simultaneous PFDynA, $Prob(\theta > f_i)$ is computed assuming a Normal distribution for WTP, with an input Q-multiplier of 2.0 for business passengers and 1.5 for leisure passengers. Despite the fact that it does not explicitly consider substitution between different itineraries, we found that Two-Way PFDynA produced good revenue performance, even in complex networks with many available itineraries.

Figure 5.10 shows the results of using either flight-by-flight dynamic pricing (Two-Way PF-DynA) or Simultaneous PFDynA in the single-carrier Network A1TWO. Note that as with

Simultaneous PFDynA, Two-Way PFDynA produces positive revenue results, from 4.8% over the base case in the low-demand environment and 6.5% over the base case in the high-demand environment. Furthermore, Simultaneous PFDynA improves revenue performance over the flight-by-flight approach in all three demand scenarios. However, this improvement is relatively small relative to the total gain in revenue: between 0.6 and 0.9%.

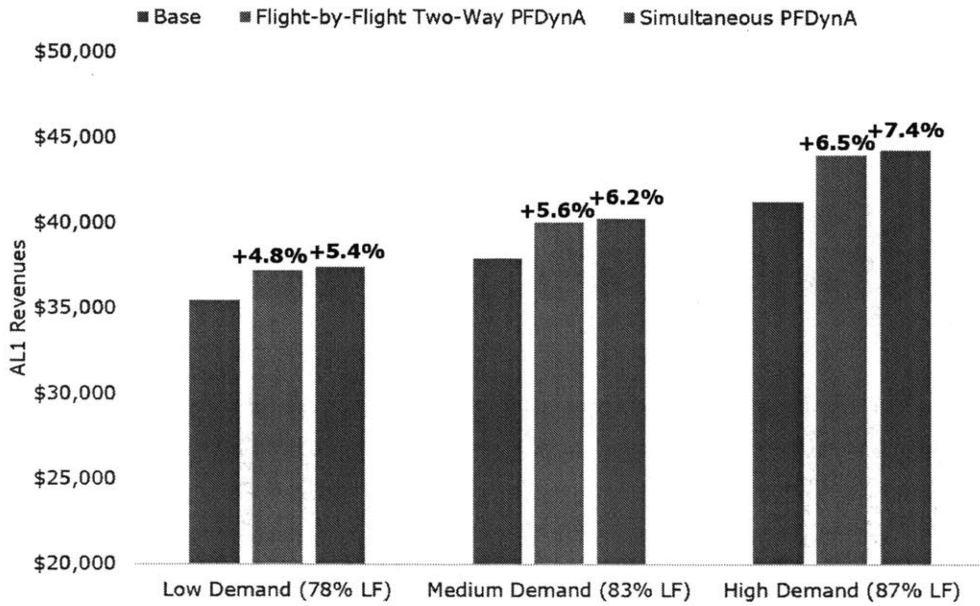

Figure 5.10: Revenues when AL1 uses flight-by-flight dynamic pricing (Two-Way PFDynA or Simultaneous PFDynA) in single-airline Network A1TWO in three demand scenarios

Two-Way PFDynA produces good revenue performance because it too is able to increase both yields and load factors for the airline that uses it. As with Simultaneous PFDynA, Two-Way PFDynA stimulates new demand by targeting discounts to leisure passengers booking in higher fare classes, as well as encouraging buy-up from business passengers booking in lower fare classes. Furthermore, as can be seen in Figure 5.11, Two-Way PFDynA also results in a shift in fare class mix from lower classes to higher classes as a result of forecast spiral-up. However, vis-à-vis Figure 5.7, the shift in fare class mix due to Two-Way PFDynA is not as substantial as the shift from Simultaneous PFDynA.

Simultaneous PFDynA reinforces the desirable shift in demand that it was designed to produce: namely, price-sensitive customers shift from the more-attractive 9am departure to the less-attractive 8pm departure. However, this shift in demand leads to only marginal (although positive) changes in revenue over the case when each flight was priced individually. Figure 5.12 shows that even flight-by-flight Two-Way PFDynA also leads to lower load factors on the morning flight and higher load factors on the evening flight, even without directly

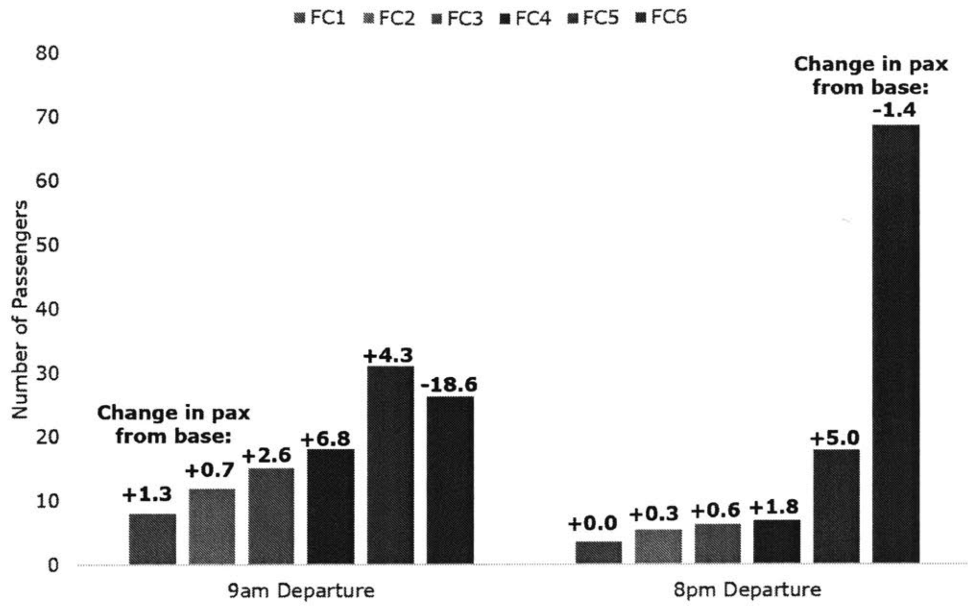

Figure 5.11: Fare class mix by flight when AL1 uses flight-by-flight dynamic pricing (Two-Way PFDynA) (Single-Airline Network A1TWO)

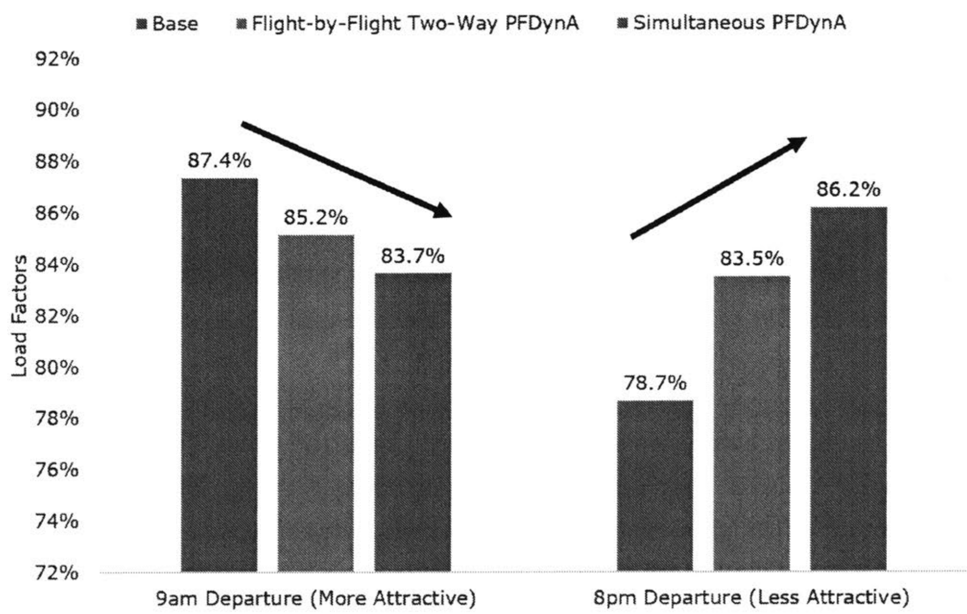

Figure 5.12: Load factors by flight when AL1 uses flight-by-flight dynamic pricing (Two-Way PFDynA) or simultaneous dynamic pricing of both itineraries (Single-Airline Network A1TWO)

considering the substitutability between the two flights in its choice probabilities.

While it is encouraging that simultaneously pricing both flights together leads to higher revenues than pricing each flight separately, recall that the airline using Simultaneous PFDynA was assumed to know a significant amount of additional information: the actual underlying departure time preference distribution, as well as the value of time disutilities for each passenger type. Even under this very optimistic assumption, the additional revenue gain over dynamic pricing of each flight individually was less than one percent.

5.5 Incorporating vertically-differentiated attributes

For a horizontally-differentiated attribute like departure time, customers have heterogeneous preferences that are spread over the course of the day. At equal fares, there is no itinerary that is preferred to another by all customers. Some customers may prefer to depart in the morning, and others may prefer to depart in the evening. However, there are some attributes for which customers have directionally homogeneous preferences. For instance, all customers should prefer an itinerary with a shorter elapsed time than a longer elapsed time, all else equal. Elapsed time is thus an example of a *vertically-differentiated attribute*, by which products are differentiated by quality (Pepall et al., 2011).

We can update the choice model used in Simultaneous PFDynA to incorporate itineraries that are differentiated by both horizontal and vertical attributes. For instance, suppose that v_i is a numerically-valued vertical attribute, and ψ represents the disutility per unit of that attribute in willingness-to-pay space. We can modify the calculation of a customer's perceived price for a particular itinerary i to incorporate this vertically-differentiated attribute:

$$PP_i = f_i + d|\omega - D_i| + \psi v_i \tag{24}$$

If one itinerary has a less attractive vertical attribute than the other, the perceived price for that itinerary will be higher than the other at equal fares $f_1 = f_2$. This would change the indifference point ω^* at which a customer is indifferent between either flight:

$$\omega^* = \frac{(D_1 + D_2)}{2} + \frac{((f_2 + \psi v_2) - (f_1 + \psi v_1))}{2d} \tag{25}$$

With this modified indifference point, Simultaneous PFDynA is more likely to set a lower price for the less attractive flight. This increases the chance that a customer will only be able to afford the less attractive flight. The indifference point ω^* also moves closer to the less attractive flight. This means that if a customer can afford both flights, they will be more likely to prefer the more attractive itinerary. We assume here that the presence of a vertically-differentiated attribute does not change the customer's maximum willingness-to-pay for the flight; the only change made to the heuristic in Equation (21) is to update the calculation of the indifference point with the vertically differentiated attribute(s).

5.5.1 Vertically-differentiated attributes in PODS: The Path Quality Index

To test Simultaneous PFDynA with horizontally- and vertically-differentiated features, we utilize a PODS attribute called the Path Quality Index (PQI). The Path Quality Index is a measure of the attractiveness of each path which typically ranges from a value of 1.0 (most attractive) to 3.0 (least attractive). Each passenger type faces a disutility for each unit of PQI when they evaluate the perceived price of each itinerary. If ψ represents the per-unit disutility for each unit of PQI, an itinerary with a PQI of 1.0 would give the customer a disutility of ψ, and an itinerary with a PQI of 3.0 would give the customer a disutility of 3ψ. In this section, we will initially use a mean PQI disutility ψ of \$30 for business passengers and \$10 for leisure passengers for each unit of PQI.[27]

We introduce vertical differentiation in Network A1TWO by changing the Path Quality Index of the evening flight (Flight 2), which departs at 8pm. We test various PQI values for Flight 2 from 1.0 to 3.0, while the PQI of Flight 1 remains at 1.0. An increase in PQI can be seen as a reduction of schedule quality for Flight 2.

Changing the PQI of Flight 2 leads to increases in base revenues for the airline, even without PFDynA. As shown in Figure 5.13, "damaging" the evening departure by giving it a lower path quality increases AL1's revenue by up to 0.7% in this network over the case when both flights have identical PQI. This is because as Flight 2 becomes less and less attractive, customers that care about path quality are more likely to want to book Flight 1. This improves the fare class mix of Flight 1; as more passengers switch from booking Flight 2, the RM system observes more demand and starts to close down low-fare classes for the more attractive flight. This results in an increase in paid fares for Flight 1 and a decrease for Flight 2, as shown in Figure 5.14.

That is, damaging the path quality of the evening flight already starts to accomplish some of the goals of Simultaneous PFDynA by moving product-sensitive demand to the more attractive morning flight, even without any dynamic pricing. This result agrees with economics literature that suggests that firms may intentionally "damage" their products to induce customers that care more about quality to purchase an "undamaged" product at a higher price point (Deneckere and McAfee, 1996).

[27] As discussed later, various other values of ψ were also tested.

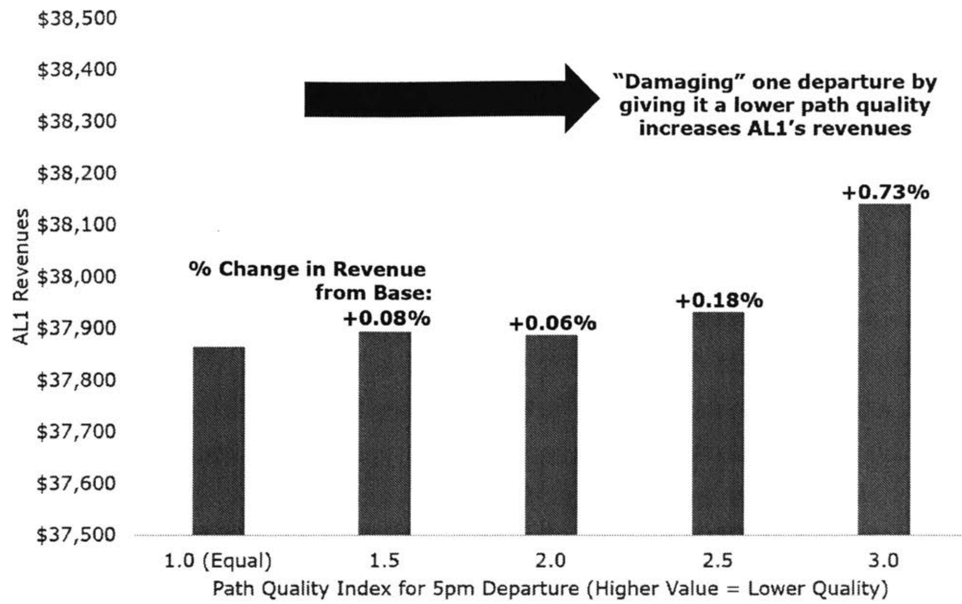

Figure 5.13: AL1 total base revenues in Network A1TWO with various PQI values for 5pm departure

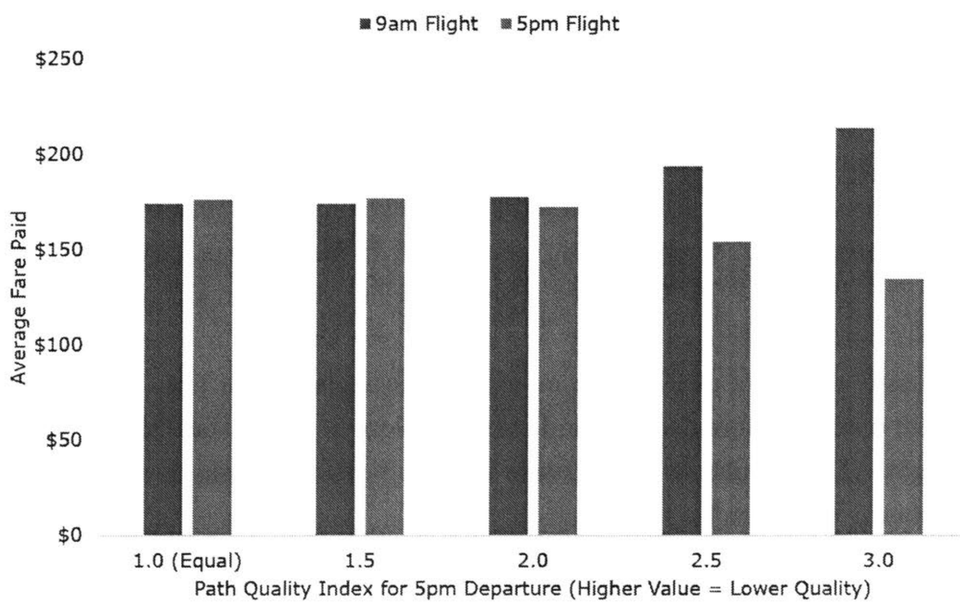

Figure 5.14: AL1 average fare paid in Network A1TWO with various PQI values for 5pm departure

For the tests of Simultaneous PFDynA in this section, PQI was added to the calculation of the indifference points ω^* as described in equation (25). As with the horizontally-differentiated attribute, we assume that the airline has knowledge of the underlying disutilities ψ faced by both business and leisure customers for each unit of PQI. We also test what happens when airlines overestimate or underestimate this value relative to the mean. The remaining parameters remain the same as in the tests of Simultaneous PFDynA with only horizontally-differentiated attributes: the value-of-time disutility parameter d is set to \$40 for business passengers and \$20 for leisure passengers, and the airline uses an input Q-multiplier of 2.0 for business passengers and 1.5 for leisure passengers. As in earlier tests, the underlying departure preference distribution Ω is known to the airline.

5.5.2 Simultaneous PFDynA with horizontally and vertically differentiated attributes

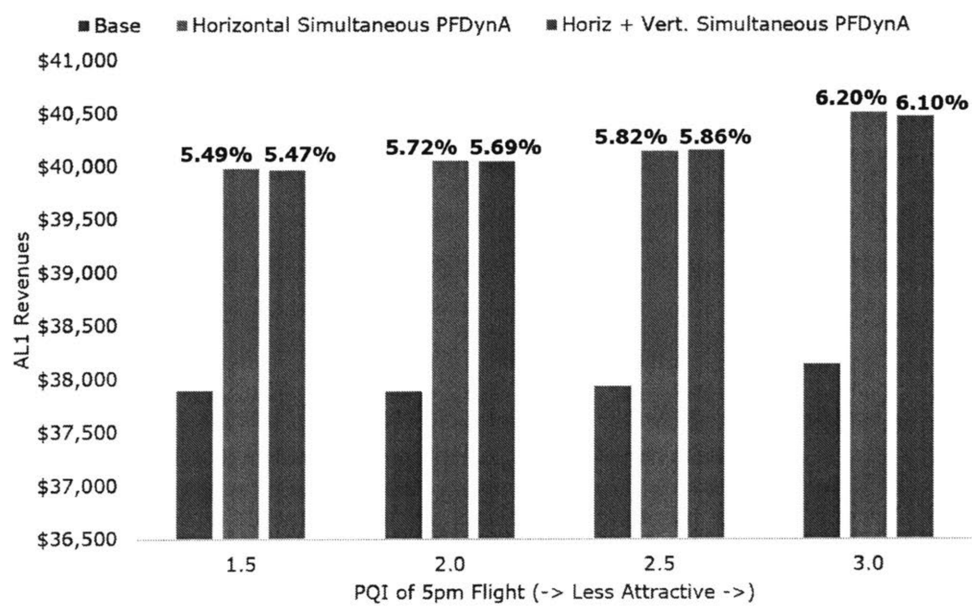

Figure 5.15: AL1 revenue when AL1 uses traditional RM (base), Horizontal, or Horizontal and Vertical Simultaneous PFDynA

In Figure 5.15, we show the performance of Simultaneous PFDynA with several different values of PQI for Flight 2. In the middle bar, we test Simultaneous PFDynA that considers only the horizontally-differentiated departure time attribute. The rightmost bar in each group uses Simultaneous PFDynA with both horizontal and vertically differentiated attributes, as described in this section. Note that regardless of the PQI of Flight 2, the performance of both heuristics is more or less identical. Both heuristics result in increases in revenue of between 5.5% and 6.2% over the base case, but there is not a significant difference

($\leq 0.1\%$) between incorporating or not incorporating the vertically differentiated attribute into the dynamic price calculation.

This is because both versions of Simultaneous PFDynA result in roughly the same patterns of discounting and incrementing, as shown in Figure 5.16. While Simultaneous PFDynA of either variety does lead to a slight change in discounting and incrementing patterns from Two-Way PFDynA, which prices both flights independently, accounting for the vertically-differentiated attribute does not significantly change the behavior of Simultaneous PFDynA with horizontal differentiation alone.

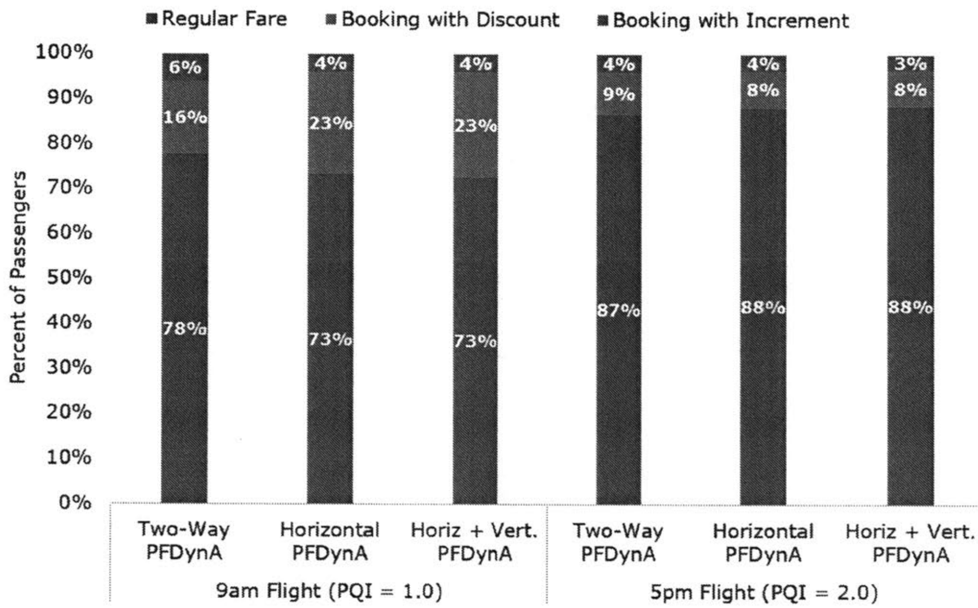

Figure 5.16: Percent of AL1 passengers booking with increments or discounts when AL1 uses Two-Way, Horizontal, or Horizontal and Vertical Simultaneous PFDynA

This may be surprising until we consider the effects of the vertical attribute on the Simultaneous PFDynA process. Recall that we accounted for the vertical attribute by incorporating its disutility into the customer's perceived price. We then updated the calculation of the indifference point ω^* (Equation (25)) to consider not only the horizontal attribute, but also the vertical attribute. We reproduce Equation (25) as Equation (26) below:

$$\omega^* = \frac{(D_1 + D_2)}{2} + \frac{((f_2 + \psi v_2) - (f_1 + \psi v_1))}{2d} \tag{26}$$

Consider our most "extreme" path quality difference, where Flight 1 has a PQI of 1.0 and Flight 2 has a PQI of 3.0. Compared to the original Simultaneous PFDynA model that considered only horizontal attributes (Equation (18)), adding the vertical attribute leads to a change in ω^* of $(3\psi - \psi)/2d = \psi/d$. For business passengers with $\psi = \$30$ and $d = 40$, this would move ω^* by $(30/40) = 0.75$ hours, or 45 minutes relative to the horizontal-only model. For leisure passengers with $\psi = \$10$ and $d = 20$, this would move ω^* by $(10/20) = 0.5$ hours, or 30 minutes relative to the horizontal-only model. With these values, ω^* does not change enough to result in a noticeable difference from the horizontal-only model.

Shifting the indifference point left or right by less than one hour does not result in significant changes to the price recommendations set by Simultaneous PFDynA. In other words, the presence of the horizontal attribute is more central to the calculation of the optimal price adjustments than the presence of the vertical attribute. This holds even if we assume much higher values for the PQI disutility ψ, as shown in Figure 5.17. Even with an input ψ as high as \$120 for business passengers and 100% segmentation accuracy, the revenue performance of Simultaneous PFDynA with horizontal and vertical differentiation does not change much (-0.13%) from the version of the heuristic that incorporates horizontal differentiation alone.

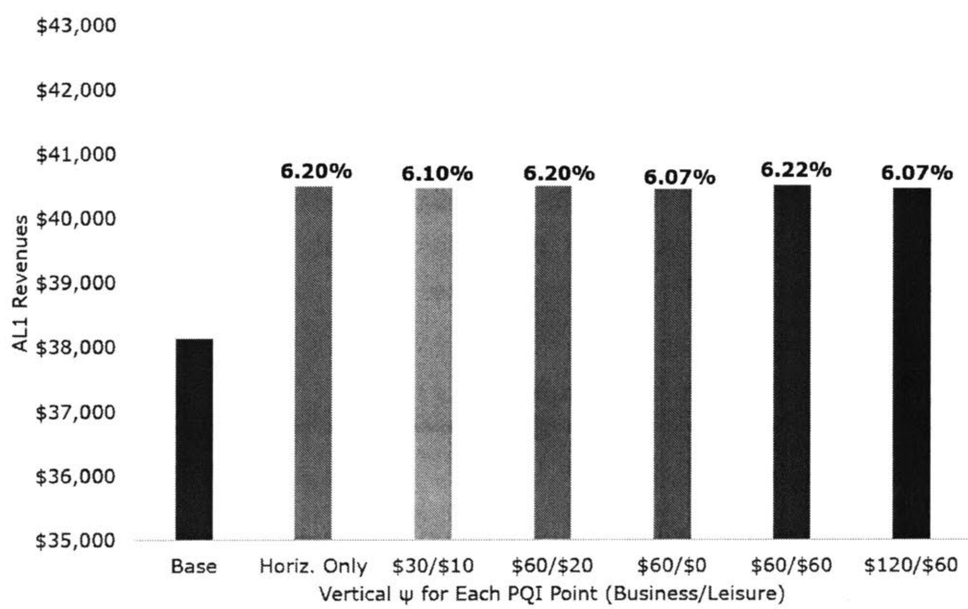

Figure 5.17: AL1 revenues when AL1 uses Horizontal and Vertical Simultaneous PFDynA with different vertical attribute disutilities, compared to traditional RM base and Horizontal Simultaneous PFDynA (Flight 1 PQI = 1.0, Flight 2 PQI = 3.0)

This result, combined with the results of horizontal differentiation in the previous section, suggest that the majority of the revenue gains of dynamic price adjustment mechanisms like PFDynA may be able to be achieved without complex choice models that attempt to describe customers' choices between diverse sets of offerings. Even with a significant amount of information—the actual disutilities that customers face for horizontal and vertical attributes, and the actual underlying departure time preference distribution—the heuristics incorporating itinerary attributes showed improvements of up to 1% over heuristics that priced each flight individually. Since the availability of data is likely to be lower in the real world than in the PODS environment, the marginal benefit of incorporating additional itinerary attributes in dynamic pricing choice functions may not outweigh the marginal costs of additional complexity from an optimization time and data requirement perspective.

5.6 Simultaneous dynamic pricing of three or more flights

The examples shown in the previous section focused on scenarios with *two* differentiated flights. In theory, any number of substitutable flights with differentiated attributes could be priced with this model. The extension to multiple flights is cumbersome mathematically, and since simultaneous dynamic pricing seemed to present only marginal gains over flight-by-flight pricing, the benefits of simultaneously pricing flights in more complex scenarios may be minimal. However, we briefly discuss in this section how the Simultaneous PFDynA model could be extended mathematically to three or more flights.

Our exposition will focus on a case with three horizontally-differentiated flights, as shown in Figure 5.18. The flights 1, 2, and 3 in this example depart at times D_1 = 9am, D_2 = 12pm, and D_3 = 8pm. We assume that the flights are only horizontally differentiated in this example, but vertically differentiated attributes could also be added as in the previous section.

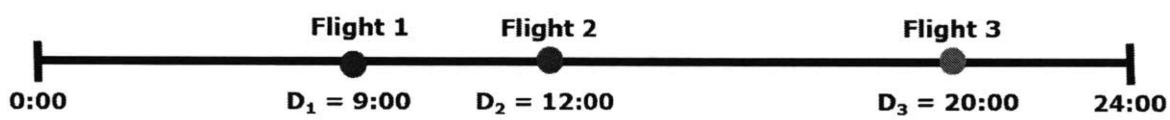

Figure 5.18: Hotelling line with three horizontally-differentiated flights

Our goal is to find the triplet of prices $[f_1^*, f_2^*, f_3^*]$ that maximizes the following expression:

$$\begin{aligned}[f_1^*, f_2^*, f_3^*] = \arg \max_{f_1, f_2, f_3} [(f_1 - BP_1) \cdot Prob(1|f_1, f_2, f_3) \\ + (f_2 - BP_2) \cdot Prob(2|f_1, f_2, f_3) \\ + (f_3 - BP_3) \cdot Prob(3|f_1, f_2, f_3)]\end{aligned} \quad (27)$$

Consider the probability $Prob(1|f_1, f_2, f_3)$ that Flight 1 will be purchased.[28] To compute

[28]The probabilities for Flights 2 and 3 can be constructed by symmetry in a similar manner.

this probability, we first need to construct the customer's choice set by considering which flights he will be able to afford.

First, the customer may be able to only afford Flight 1. This occurs if his maximum WTP θ exceeds the price of Flight 1 (f_1), but is less than the prices of Flights 2 and 3. If $f_2 < f_1$ or $f_3 < f_1$, it will never be the case that only Flight 1 is affordable. Mathematically, this is equivalent to:

$$\text{Prob(Only 1 is affordable)} = \begin{cases} \text{Prob}(f_1 < \theta < \min[f_2, f_3]) & \text{if } f_1 < \min[f_2, f_3] \\ 0 & \text{otherwise} \end{cases} \tag{28}$$

It is also possible that exactly two flights—for instance, Flights 1 and 2—will be affordable, but Flight 3 is unaffordable. This occurs with the following probability:

$$\text{Prob(Only 1 and 2 are affordable)} = \begin{cases} \text{Prob}(\max[f_1, f_2] < \theta < f_3) & \text{if } f_3 > \max[f_1, f_2] \\ 0 & \text{otherwise} \end{cases} \tag{29}$$

The customer could also potentially be able to afford Flights 1 and 3 only, with Flight 2 being unaffordable. This occurs with the probability:

$$\text{Prob(Only 1 and 3 are affordable)} = \begin{cases} \text{Prob}(\max[f_1, f_3] < \theta < f_2) & \text{if } f_2 > \max[f_1, f_3] \\ 0 & \text{otherwise} \end{cases} \tag{30}$$

Finally, it may the case that the customer can afford all three flights—that is, that his maximum WTP exceeds the prices of each of the flights. This happens with the following probability:

$$\text{Prob(All three flights are affordable)} = \text{Prob}(\theta > \max[f_1, f_2, f_3]) \tag{31}$$

Within each of these possible choice sets, we need to compute the probability that the customer will purchase Flight 1 out of all of his affordable options. First, if only Flight 1 is affordable, he will purchase that flight with probability 1. If the customer can afford more than one flight, we can compute an indifference point ω^* between each pair of flights, as shown in Figure 5.19.

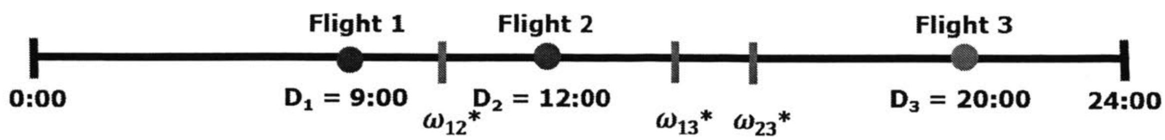

Figure 5.19: Indifference points for three horizontally-differentiated flights

For instance, suppose Flights 1 and 2 are affordable. The customer will choose to purchase Flight 1 in this scenario when his departure time preference ω is to the left of the indifference point ω_{12}^*, where:

$$\omega < \omega_{12}^* = \frac{D_1 + D_2}{2} + \frac{f_2 - f_1}{2d} \tag{32}$$

Similarly, if only Flights 1 and 3 are affordable, the customer will choose Flight 1 if $\omega < \omega_{13}^*$, where:

$$\omega < \omega_{13}^* = \frac{D_1 + D_3}{2} + \frac{f_3 - f_1}{2d} \tag{33}$$

Finally, if all three flights are affordable, the customer will:

- Choose Flight 1 if $\omega < \omega_{12}^*$.
- Choose Flight 2 if $\omega_{12}^* < \omega < \omega_{23}^*$, where $\omega_{23}^* = \frac{D_2+D_3}{2} + \frac{f_3-f_2}{2d}$.
- Choose Flight 3 if $\omega > \omega_{23}^*$.

Given all of these probabilities, we can construct the complete purchase probabilities for each flight for use in the dynamic pricing equation in Equation (27):

$$\begin{aligned} Prob(1|f_1, f_2, f_3) =& \text{Prob(Only 1)} + \text{Prob(1 and 2)} \cdot \text{Prob}(\omega < \omega_{12}^*) \\ &+ \text{Prob(1 and 3)} \cdot \text{Prob}(\omega < \omega_{13}^*) + \text{Prob(All)} \cdot \text{Prob}(\omega < \omega_{12}^*) \end{aligned} \tag{34}$$

$$\begin{aligned} Prob(2|f_1, f_2, f_3) =& \text{Prob(Only 2)} + \text{Prob(1 and 2)} \cdot \text{Prob}(\omega > \omega_{12}^*) \\ &+ \text{Prob(2 and 3)} \cdot \text{Prob}(\omega < \omega_{23}^*) + \text{Prob(All)} \cdot \text{Prob}(\omega_{12}^* < \omega < \omega_{23}^*) \end{aligned} \tag{35}$$

$$\begin{aligned} Prob(3|f_1, f_2, f_3) =& \text{Prob(Only 3)} + \text{Prob(1 and 3)} \cdot \text{Prob}(\omega > \omega_{13}^*) \\ &+ \text{Prob(2 and 3)} \cdot \text{Prob}(\omega > \omega_{23}^*) + \text{Prob(All)} \cdot \text{Prob}(\omega > \omega_{23}^*) \end{aligned} \tag{36}$$

The previous sets of equations provide a complete description of the probabilities that would need to be calculated to simultaneously price three horizontally differentiated flights. The probabilities do not require any additional information over the two-flight case, but the addition of the third flight significantly increases the complexity and dimensionality of the problem. Namely, searching for a triplet of optimal prices is likely to be more computationally intensive than in the two-flight case. Pricing four, five, or more flights simultaneously is theoretically possible, but will become even more computationally intensive as each additional flight exponentially increases the number of possible choice sets that must be considered when computing purchase probabilities and finding optimal prices.

5.7 Conclusions

In this chapter, we introduced and tested an extension to PFDynA that simultaneously priced multiple differentiated flights. We used a locational choice model to frame customers' purchasing decisions based on their willingness to pay and departure time preferences. We then integrated this choice model into the PFDynA dynamic pricing heuristic to directly incorporate the differentiated attributes of the flights into the calculation of price adjustments.

We found that depending on the demand scenario, Simultaneous PFDynA led to revenue gains of between 5% and 7% in PODS over the base case in a small network with one airline, one market, and two non-stop flights. This heuristic performance was at the upper range of other dynamic pricing heuristics reported in the literature, but also assumed that airlines could accurately segment booking requests into leisure and business categories, and possessed accurate information about departure time preferences and value-of-time disutilities.

Simultaneous PFDynA was able to produce revenue gains because it stimulated new bookings through targeted discounts to leisure customers, increased yields for business customers by incrementing some fares, and led to forecast spiral up as an increase in bookings in higher classes resulted in less availability in less-expensive classes. It also resulted in shifts of price-sensitive demand from more-attractive itineraries to less-attractive itineraries, increasing availability and yields for schedule-sensitive passengers who preferred the more attractive flights. In a simple competitive environment with two airlines each offering two flights, the use of Simultaneous PFDynA by both airlines produced revenue gains of about 6%.

However, Simultaneous PFDynA produced only marginal gains of less than one percent in the single-airline case over Two-Way PFDynA, which priced each flight independently without directly considering the presence or attributes of alternative flights. That was because Two-Way PFDynA was also able to stimulate demand, increase yields, and generate forecast spiral-up, even without directly incorporating the substitutability between the flights. Adding horizontal attributes into the PFDynA choice function produced only a marginal shift in demand from more-attractive to less-attractive itineraries.

We also tested a version of the heuristic that incorporated vertically-differentiated attributes, such as path quality, into the dynamic pricing calculation. Tests of Simultaneous PFDynA with vertically and horizontally differentiated attributes show very little change in revenue performance over the version of the heuristic with horizontal attributes alone. This is because differences in vertical attributes do not significantly change the location of the indifference point ω^*, and as a result do not significantly change the dynamic price adjustments suggested by the PFDynA mechanism.

Finally, we described how to modify the price calculations for Simultaneous PFDynA to dynamically price three or more flights. Extending the heuristic is a cumbersome mathematical exercise, as the number of possible choice sets increases exponentially with each additional flight. Given these challenges of estimating and implementing a simultaneous dynamic pricing engine in practice, and given the limited gains of simultaneous dynamic pricing versus flight-by-flight dynamic pricing, practitioners may be content to rely on the more simple flight-by-flight heuristic, at least for early adaptations of dynamic pricing.

6 Implications of Dynamic Pricing for the Airline Industry

In Chapters 4 and 5, we showed through a number of simulations that the dynamic pricing heuristics developed in Chapter 3 can lead to revenue gains, even in competitive environments. However, some airlines have raised concerns about the feasibility and prudence of dynamic pricing. These objections relate mostly to real-world implementation concerns about dynamic pricing that are not present in a simulation environment.

Specifically, some practitioners worry that dynamic pricing will not be implementable in practice due to legality, privacy issues, or the possibility of customer abuse. Others are concerned that implementing dynamic pricing in the airline industry would be irresponsible because it would lead to a revenue-negative situation called the *race to the bottom* that would leave airlines worse off than with traditional revenue management and pricing practices.

In this chapter, we consider each of these possible implications of dynamic pricing for the airline industry. We discuss the foundations of each concern, and review relevant economic literature in consideration of whether the concern would likely manifest itself should airlines continue to move down the road of dynamic pricing. Indeed, some airlines, vendors, and industry groups are already starting to develop systems for dynamic pricing and advocate for their use in the industry. This section explores the factors that could pose obstacles to widespread adoption of these systems in practice.

6.1 Will dynamic pricing cause a race to the bottom?

6.1.1 What is the race to the bottom?

Consider a simple example of a single market served by two airlines, each of which offers a single non-stop flight. Suppose that the flights offered by the two airlines are identical: they depart at the same time, offer the same level of service and amenities, and have the same elapsed time. In our choice model, customers are indifferent between the two flights if they are priced at the same fare, and otherwise will purchase the flight with the lower price, as long as that price is lower than their maximum WTP θ.

The "race to the bottom" refers to the potential for each airline to undercut the price of the other by an arbitrarily small amount ϵ in order to capture all of the market demand. Suppose both airlines initially price the flight at the monopoly price p:

$$p^* = \arg\max_p (p - c) \cdot Prob(\theta > p) \tag{37}$$

where c is the airline's marginal cost[29] and θ is the passenger's maximum WTP. In this case, it is in Airline 1's interest to lower its price by one unit of currency (for instance, to $p^* - \$0.01$) as long as this new price is greater than Airline 1's marginal cost. At $p^* - \$0.01$, customers will always choose to purchase Airline 1's flight, since the flights are otherwise identical. Airline 1 will capture all of the market demand, and Airline 2 will receive no customers.

But Airline 2 also has an incentive to respond in this situation. Specifically, Airline 2's best response is to undercut the new price of Airline 1 by one cent. When Airline 2 sets its price at $p^* - \$0.02$, it will now capture all of the market demand. To this move, Airline 1 is once again incentivized to respond by undercutting by a single cent, and so forth.

This undercutting will only stop when one airline reaches its marginal cost. At that point, it does not make sense for the airline to continue to undercut, because it will receive negative expected revenue from selling a seat below marginal cost. In other words, it is better to leave the seat empty than to take a passenger at a price less than marginal cost and incur a loss. When one airline reaches its marginal cost, the other airline will price at one cent below this amount, as long as this amount is still greater than the airline's own marginal cost. At this point, no further undercutting will take place.

In our simple example, if the airlines have equal marginal costs, both airlines will price at marginal cost and split the demand. The airlines will lose money overall if fixed costs are high (as is the case in the airline industry). If one airline has a lower marginal cost than the other, that airline will capture the market demand and the airline with the higher marginal cost will suffer losses equal to their fixed costs.

This undercutting pattern—the race to the bottom—is the classic result of so-called *Bertrand competition* (Tirole, 1988). In Bertrand competition, firms compete by setting prices for their goods.[30] If the products are homogeneous, as they are in the situation described above, the race to the bottom will result. The implication of Bertrand competition, therefore, is that in markets with as few as two firms, prices will be driven down to the marginal cost of the more costly firm.

Economic theory would predict that in competitive markets, airlines would eventually end up pricing at marginal cost. However, the structure of the current pricing and revenue management environment and the nature of competition in today's airline industry prevents this outcome from occurring in many cases. The aim of this section is to discuss why this

[29]In this chapter, the term "marginal cost" will refer to the sum of the airline's marginal cost of capacity (i.e., the bid price) given the network structure, fleet, time remaining to departure, and forecast of demand-to-come, plus any marginal costs of carriage (e.g., bag handling, check-in costs, catering, etc.) for a single passenger.

[30]The "opposite" of Bertrand competition is Cournot competition, in which firms are price takers according to a market demand curve but compete by selecting the quantity to produce (Tirole, 1988).

is the case and evaluate if elements of dynamic pricing could potentially change this status quo, leading to a downward spiral of both prices and airline revenues.

6.1.2 The race to the bottom in the airline industry

Airlines are businesses with high fixed costs and low marginal costs. Assuming that the airline's flight network is fixed in the short term, the marginal costs associated with selling an additional seat include a *marginal cost of capacity* for each flight leg (i.e., the bid price), as well as a small *marginal cost of carriage* (e.g., an onboard snack, bag handling costs for a checked bag, etc.). These marginal costs from selling an additional seat are small relative to the airlines' fixed costs of acquiring aircraft and operating a network of flights (Belobaba, 2016). If the race to the bottom were to occur in the airline industry and result in marginal cost pricing, we would expect to see significant revenue losses throughout the industry.

While airlines were relatively profitable globally at the end of 2017, it does not take a deep look into history to find a time when this was not the case. The airline industry has been plagued since its inception by periods of instability and bankruptcies. From 2000 – 2005, U.S. airlines alone lost $30 billion dollars (Goetz and Vowles, 2009). The reasons for airline bankruptcies are numerous: high and volatile costs of inputs like fuel or labor; depressed demand due to exogenous macroeconomic shocks like the Great Recession or the terrorist attacks of 9/11; mismanagement or poor business strategies; and periods of excessive capacity growth that drive down prices.

To survive, airlines have undergone a series of mergers, increasing market concentration. At the same time, facing competition from ultra-low-cost carriers with unbundled service offerings, large airlines have reduced service quality and discontinued complimentary amenities in the economy cabin. In such an environment, air transportation became more like a commodity, where airlines' offerings are not differentiated based on service quality. This "commoditization" of the flying experience leads to an environment that, based on the theory of Bertrand competition, could result in a race to the bottom.

The global airline industry's current profitability seems to run contrary to the predictions of Bertrand competition. Yet the absence of a race to the bottom is not uncommon in competitive markets in which it appears the rules of Bertrand competition should apply. That is, in many situations where economists would expect a race to the bottom and marginal cost pricing, firms are still able to make economic profits. This is often called the *Bertrand paradox*; despite Bertrand-like settings, the Bertrand outcome often does not manifest itself in real life (Tirole, 1988; Dufwenberg and Gneezy, 2000).

To explain the Bertrand paradox, economists have often noted that real-life competitive markets, like the airline industry, do not exactly fit all of the assumptions that undergird Bertrand competition. When one or more deviations occur from perfect Bertrand competition, the race to the bottom is not assured, and equilibrium outcomes in which prices are set above marginal cost (and even as high as the monopoly price) are possible. The current structure of the airline industry possesses at least five features that could help it avoid a race to the bottom:

- Airlines face *capacity constraints*, in terms of total aircraft in their fleets and seats available on each individual flight, in the short- and medium-term. This causes airlines to face increasing marginal costs of capacity (bid prices) as planes become more full.
- Airlines have *incomplete and imperfect information* regarding their competitors' bid prices and marginal costs of carriage.
- Through the use of branded fares and fare families, airlines are offering *increasingly differentiated products* that remove some incentives for marginal-cost pricing.
- Airlines interact with their competitors *repeatedly and indefinitely* in a wide variety of markets. As a result, the outcomes of the pricing and revenue management games are different than if the interactions occurred only once.
- Publicly filed fare structures allow airlines to *perfectly monitor* when their competitors make pricing changes and react accordingly, including the possibility of responding with fare wars or punishments when undercutting is detected.

In this section, we explain the economic theory behind each of these features of the airline industry and how they may help the industry avoid the race to the bottom. Following this, we discuss how dynamic pricing could potentially result in changes to these features, and whether these changes could result in a new race to the bottom in the industry.

Capacity constraints

In the airline industry, each airline may be limited in its ability to satisfy the entire market demand when price equals marginal cost. This is because each flight is capacity constrained. If one airline operating a single flight was to undercut its competitors to try to capture all of the demand for a given market, it would likely be unable to accommodate all of the passengers who want to travel on its single flight.

Capacity constraints have two effects on Bertrand competition. First, and particularly in the airline industry, capacity constraints mean that marginal costs are not constant throughout the entire selling period. Specifically, as a flight becomes more full and (more precisely) as remaining capacity decreases relative to forecast demand, the bid price for that flight will increase. This increases the marginal cost of capacity for the flight.

This also changes the behaviors of Bertrand competition; if marginal costs are increasing, then a firm that lowers its price to attempt to capture all of the market demand will soon see its marginal cost increase above that of its other competitors. As marginal costs rise, a single firm will not necessarily capture all the demand, and prices will not always be set at levels that eliminate economic profit in the long term.

Even if marginal costs do not increase over time, capacity constraints can change the outcome of a Bertrand game (Pepall et al., 2011). Consider two airlines in Bertrand competition selling a homogeneous flight, and suppose the airlines have capacities CAP_1 and CAP_2. Let $Q(p)$ represent the market demand curve at price p; that is, the number of people that would travel in that market at price p. Suppose for now that both airlines have equal marginal costs c that do not change as flights become more full.

If $CAP_1 > Q(c)/2$ and $CAP_2 > Q(c)/2$, then the Bertrand equilibrium $p_1 = p_2 = c$ would result. Each airline has enough seats to serve the entire market at the Bertrand equilibrium price. However, if $CAP_1 < Q(c)/2$, then the Bertrand outcome is not an equilibrium.

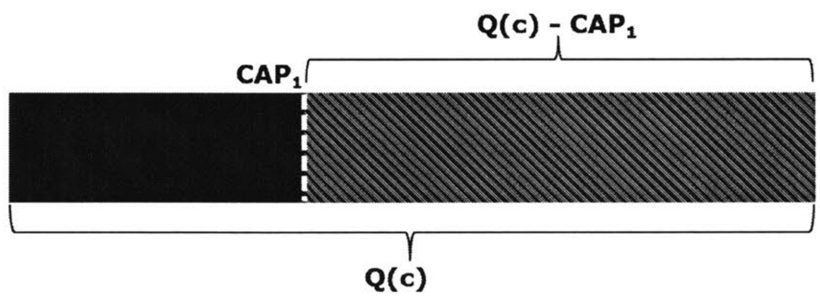

Figure 6.1: Residual demand when Airline 1 is capacity constrained

To see why, suppose Airline 1 chooses $p_1 = c$. Then Airline 1 will face demand $Q(c)/2$, but can only serve CAP_1 passengers. Airline 2 then faces residual demand of $Q(c) - CAP_1$, represented by the shaded box on the right side of Figure 6.1. This residual demand is entirely under the control of Airline 2; Airline 1 is sold out, so it can accept no more passengers. If Airline 2 sets a price $p_2 = c$, then it will make no profit from the residual demand.

But if Airline 2 sets a price $p_2 > c$, then it will make economic profit on each passenger it accepts. Airline 2 will not be able to capture the entire segment of residual demand at a price above c, but it will make profit on each passenger it does accept, because its prices are above its marginal cost. This strategy strictly dominates the strategy of setting price equal to marginal cost. Therefore, in this scenario, the Bertrand outcome where both firms price at marginal cost is not an equilibrium (Kreps and Scheinkman, 1983; Pepall et al., 2011).

If both airlines are capacity constrained ($CAP_1 < Q(c)/2$ and $CAP_2 < Q(c)/2$), then neither airline can serve the entire market (or its portion of the market if the firms have equal marginal cost) at the Bertrand equilibrium price. In this case, both airlines can take advantage of the residual demand remaining from the other firm and charge a price higher than marginal cost to make economic profit from these customers. Here, the Bertrand outcome of pricing at marginal cost is also not an equilibrium.

As a brief extension, consider the case in which firms can make decisions about quantity before they make decisions about price. This closely mirrors the airline environment in the medium-term: network planning teams decide on the capacity of each flight by assigning aircraft of different sizes, and the pricing and revenue management teams then take that capacity as given when computing the marginal cost of capacity and selecting the price to charge at any moment.

In a classic paper, Kreps and Scheinkman (1983) suggest that if capacity is costly to acquire, then firms in a duopoly will set $CAP_i < Q(c)/2$. This is because setting excess capacity $CAP_i > Q(c)/2$ will result in the Bertrand equilibrium and require pricing at marginal cost, at which point there is no economic profit. It is much better for firms to restrict their capacity such that they are unable to satisfy the entire market demand. In this scenario, as described above, no firm can satisfy the entire market demand by undercutting its competitors, and prices will be greater than marginal cost at equilibrium.

The Kreps and Scheinkman (1983) result suggests that capacity-setting is an important component of the airline price competition game. The outcome of the game suggests that keeping capacity relatively low can lead to higher prices and higher economic profits than flooding the market with capacity. As a result, airlines may be incentivized to encourage other competitors to keep capacity growth limited. In the period from 2010 - 2016, following a spate of airline bankruptcies, a reduction in demand, and a spike in fuel prices, the airline industry in the United States did exactly that in an era referred to as the "capacity discipline" period (Wittman, 2014).

With capacity discipline, airlines grew domestic capacity by less than would have been expected given the rate of domestic economic growth. This strategy resulted in both higher load factors and higher yields for U.S. airlines, and has been cited as a primary reason for the high level of stable profits in the U.S. industry during this period (Wittman, 2014).

However, with increasing rates of capacity growth (particularly in the Middle East and from the U.S. to Europe), airlines may worry that capacity constraints will no longer serve as a barrier to a race to the bottom in some markets. In an environment with excess capacity, airlines could more easily accommodate the extra demand resulting from lower prices, which increases the incentives to discount. Wall Street analysts often worry about the effects of unchecked capacity growth on industry profitability; in 2015, Southwest Airlines' stock fell by 9.1% in a single day after it increased its domestic capacity growth target by a single percentage point from 7% to 8% (Rich, 2015).

Unknown competitor costs and the winner's curse

Yet even in airline markets with excess capacity, other factors could still prevent a race to the bottom. For instance, in Bertrand competition, it is assumed that each airline's marginal costs of capacity and marginal costs of carriage can be observed by all airlines. In a one-shot game, there is no ambiguity about what prices to charge: the "undercutting" pattern described above is not dynamic but instead static. Firms can arrive at the Bertrand equilibrium without any period of learning in regards to their competitors' pricing decisions or costs.

What would happen if marginal costs of other firms were not known in a one-shot Bertrand game? This is the case in the airline industry: airlines have a good idea of their own marginal costs (both the costs of carriage and the bid prices associated with each leg of an itinerary) but have imperfect and incomplete information about other airlines' costs. Airlines typically do not know the bid prices for each flight and departure date of other airlines, and may have only rudimentary or aggregate information about other costs, perhaps through annual reports to shareholders or filings with the government.

Spulber (1995) explores such a scenario in a one-shot Bertrand game. He finds that when costs are unknown to other firms, the Bertrand outcome of pricing at marginal cost is not an equilibrium. Instead, firms competing *a la* Bertrand will price above marginal cost in equilibrium. In the Spulber (1995) model, prices rise above the marginal cost that would be predicted by Bertrand competition, but are not as high as the monopolistic or Cournot prices, contra Kreps and Scheinkman (1983).

The main intuition behind the Spulber (1995) model is that when other the firms' costs are unknown, each firm faces competing incentives when they set a price. On one hand, firms want to set a low price, to increase the chance that they will win the entire market.[31] On the other hand, firms want to maximize their payoffs if they do end up winning the market. If the firm sets a price equal to marginal cost, they make no economic profit even if they win the game and capture all the demand.

[31] Note that there are no capacity constraints in the Spulber (1995) model that would prevent a firm from satisfying the entire market demand at the Bertrand price.

If other firms' costs are uncertain, it may indeed be the case that the given firm's costs are lower than all of its competitors. In this case, setting a price above marginal cost (in the interval between the firm's marginal cost and the higher costs of its competitors) will still result in winning the market, but also generate positive economic profit. Setting a price above marginal cost may reduce the probability that a firm will win the market, but it will also lead to non-zero payoffs in the case that the firm does win. As a result, setting prices above marginal cost in a mixed strategy equilibrium always leads to higher expected payoffs than setting price equal to marginal cost, and the Bertrand equilibrium is nullified.

This scenario is similar to a *first-price sealed bid auction.* In a first-price sealed bid auction, firms submit hidden bids to a central authority (in this case, the buyer). The buyer then selects the bid with the lowest price among the sealed bids. It is well known that in this environment, a *winner's curse* emerges. If a firm submits the lowest bid (perhaps as a result of having the lowest cost), then it regrets not setting its bid closer to the next-higher bid (Easley and Kleinberg, 2010).

As a numerical example, suppose that Airline 1 has a marginal cost of \$100 and Airline 2 has a marginal cost of \$150, and that each airline cannot observe the costs of its competitor. Suppose that for a given passenger, both airlines "bid" their marginal costs. Airline 1 would "win" the auction and gain the passenger at \$100, but it would have been better off if it had bid \$149—a higher bid than its marginal cost. This uncertainty about the marginal costs of its competitors will incentivize Airline 1 to *shade* its bid higher than its marginal cost, to increase its potential reward if it wins the auction.

The incentive for bid shading is caused by the imperfect information available to firms. If firms' costs were known, it would not make sense for firms to shade their bids, because it would be clear who the winner of the auction would be at the start of the game. This represents a potentially counterintuitive scenario where *less* available information about other firms' behaviors can help escape a race to the bottom. If airlines knew each others' exact bid prices for each flight departure, it would be easier for an airline with a lower bid price to undercut a competitor with a higher bid price for an equivalent flight departure.

Product differentiation

The standard Bertrand competition model assumes that the products offered by both firms are identical commodities. At the same price, customers are indifferent between the products offered by both firms; otherwise, the customers will always prefer the product offered at the lower price. While air travel became increasingly commoditized in the early 2000s (Peterson, 2010), as we discussed in Chapter 5 there are often features or components of itineraries that will be distinct amongst alternatives, including departure time, elapsed time, fare restrictions, branded fare components, or service quality. Also, the New Distribution Capability will allow airlines to present more information to consumers than just a schedule and price—for instance, a picture of an extra-legroom seat or a description of premium cabin services. This could also increase differentiation. With differentiated products, the Bertrand equilibrium also may not hold.

The following example is adapted from Pepall et al. (2011), but takes a familiar form to the models presented in Chapter 5. Suppose that two airlines offer horizontally differentiated flights at two ends of a Hotelling line, with customers' preferences evenly and uniformly distributed across the line. In this case, the horizontal differentiation could refer to attributes like departure time, but also to customers' idiosyncratic tastes for various airlines, due perhaps to differences in service quality or membership in a frequent flyer program.

Let θ_i represent the maximum WTP for customer i, ω_i represent customer i's position on the Hotelling line, and p_1 and p_2 represent the prices of Airlines 1 and 2, respectively. As we found in Chapter 5 (Equation (18), the indifference point is:

$$\omega^* = \frac{D_1 + D_2}{2} + \frac{p_2 - p_1}{2d}$$

where D_1 and D_2 are the departure times of flights 1 and 2 and d is the disutility from moving one unit away from the customer's preference location on the line. Without loss of generality, assume unit market demand and normalize the ends of the Hotelling line to $D_1 = 0$ and $D_2 = 1$. The indifference point ω^* is then equal to equal to $\frac{p_2 - p_1 + d}{2d}$.

To compute the profit functions for each firm, first note that if both airlines face marginal cost c, airline i will receive revenue equal to $(p_i - c)$ from selling each seat. Therefore, the profit functions for Airlines 1 and 2 are:

$$\Pi_1(p_1, p_2) = (p_1 - c) \cdot \frac{p_2 - p_1 + d}{2d} \tag{38}$$

$$\Pi_2(p_1, p_2) = (p_2 - c) \cdot \frac{p_1 - p_2 + d}{2d} \tag{39}$$

To find the profit-maximizing price p_1^*, we can differentiate Π_1 with respect to p_1 and set the result equal to 0.

$$\frac{\partial \Pi_1}{\partial p_1} = \frac{p_2 - p_1 + d}{2d} - \frac{p_1 - c}{2d} = 0$$

$$p_2 - p_1 + d = p_1 - c$$

$$p_1^* = \frac{p_2 + c + d}{2} \tag{40}$$

It is also possible to show by symmetry that:

$$p_2^* = \frac{p_1 + c + d}{2} \tag{41}$$

Then, substitution yields:

$$\begin{aligned} p_1^* &= \frac{\frac{p_1^* + c + d}{2} + c + d}{2} \\ 4p_1^* &= p_1^* + 3c + 3d \\ p_1^* &= c + d \end{aligned} \tag{42}$$

and by symmetry:

$$p_2^* = c + d \tag{43}$$

The optimal equilibrium prices for both firms are $p_1^* = p_2^* = c + d$. These prices are greater than what Bertrand competition predicts ($p_1^* = p_2^* = c$). Specifically, the prices are marked up above marginal cost by an amount proportional to the disutility faced by customers.

In more complex environments with more competitors, or where firms can choose where they position themselves along the line, different equilibria may result. For instance, as more firms are added into the model, the pricing game becomes more competitive and prices begin to approach the Bertrand equilibrium (Pepall et al., 2011). In general, however, the presence of differentiated attributes (in this case, departure time) give firms some pricing power, avoiding the Bertrand outcome.

In the mid-2010s, airlines trended towards increasing differentiation through the introduction of branded fares. These products, such as "basic economy" and fare families, allow airlines to provide various combinations of amenities for the same seat in the same cabin. These developments also help to increase product differentiation which, as described above, can break down the Bertrand equilibrium.

Repeated interactions, indefinitely repeated games, and perfect monitoring

The final attribute of airline competition we will consider is perhaps the most important reason why airlines today appear to be avoiding the race to the bottom outcome. The Bertrand game, along with the games described by Hotelling (1929); Kreps and Scheinkman (1983); Spulber (1995); and Pepall et al. (2011), are all one-shot games in a single market. In reality, airlines compete with each other repeatedly in a wide variety of markets each day (Morrison and Winston, 1996; Ciliberto and Williams, 2014). These repeated multimarket interactions continue indefinitely into the future. As a result, the airline competitive price-setting game is in fact not a one-shot game, but an infinitely (or indefinitely) repeated game.

In indefinitely repeated games, the equilibrium outcomes can change from the static game. Generally, indefinitely repeated games have more equilibria than finitely-repeated or one-shot games. In some types of indefinitely repeated games, there are actually an infinite number of possible equilibria. This intuition is formalized in the Folk Theorem, of which various flavors have been attributed to and proved by many authors, including Friedman (1971) and Fudenberg and Maskin (1986). We present a general description below.

The Folk Theorem:
Suppose a game has a feasible set of strategies Γ for which all players' payoffs exceed the payoffs from the game's one-shot Nash equilibrium. Then for some discount factor $\delta < 1$, the payoffs from Γ are an equilibrium outcome for the indefinitely-repeated game.

The Folk Theorem says that if there is a feasible set of strategies that leave the players better off than the one-shot equilibrium strategies, the payoffs from this set of strategies can be an equilibrium result in the indefinitely-repeated game as long as the players are patient enough. One example with which to illustrate this point is the classic Prisoner's Dilemma.

	Cooperate	Defect
Cooperate	(10,10)	(-100, 100)
Defect	(100, -100)	(-10, -10)

Table 6.1: A simple Prisoner's Dilemma game

Consider the two-player Prisoner's Dilemma in Table 6.1. In the table, the cells represent the payoffs for Player 1 and Player 2, respectively. If the players both choose to Cooperate, they each receive a payoff of +10. Yet it is a best response for each player to choose to Defect, because if the other player still decides to Cooperate, the defecting player receives a payoff of +100. As a result, the Nash equilibrium of the one-shot or finitely-repeated Prisoner's Dilemma game is (Defect, Defect), which leads to a payoff of -10 for each player.

The Folk Theorem implies that in an indefinitely repeated Prisoner's Dilemma game, the payoffs (10, 10) from the (Cooperate, Cooperate) strategy can be sustained as an equilibrium for some discount factor $\delta < 1$.[32] One strategy that can lead to this outcome is the so-called "grim trigger." Players using the grim trigger strategy choose to play Cooperate indefinitely, unless they observe their opponent playing Defect. The first time the opponent plays Defect, the trigger is sprung, and the player chooses to play Defect for all future periods. In this way, the temptation of "cheating" is outweighed by the threat of punishment when the cheating is detected—namely, by the playing of Defect for all future periods, forcing the game into the inefficient (Defect, Defect) equilibrium.

[32] The discount factor describes how players value payoffs in future periods relative to payoffs in the current period. Higher discount factors imply that players care more about future payoffs.

As long as players care enough about future outcomes, the short term gain of 100 by defecting in a single period is outweighed by the -10 losses in all future periods as a result of the grim trigger. As a result, players will never choose to Defect. In the Bertrand game, an example of a grim trigger strategy is shown in Figure 6.2.

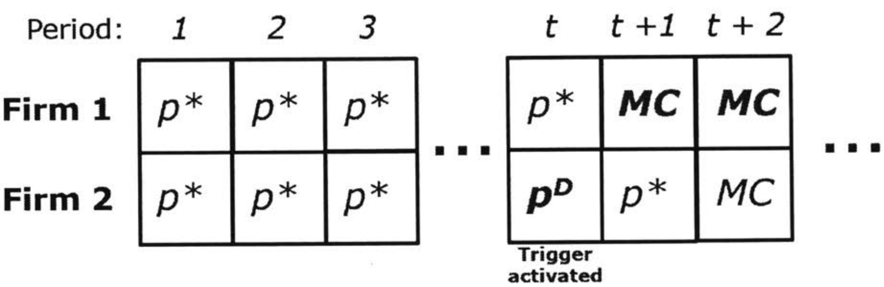

Figure 6.2: The grim trigger strategy for the indefinitely-repeated Bertrand pricing game

Firms begin by playing the monopoly price p^* indefinitely, enjoying high profits and capturing demand equal to $Q(p^*)/2$. However, it is each firm's incentive to deviate from the monopoly price and play a lower deviation price $p^D < p^*$. The deviating firm would capture all of the market demand $Q(p^D) > Q(p^*)/2$ and increase its profits for that period. But such an action will activate the grim trigger, causing the first firm to price at marginal cost $(MC < p^D < p^*)$ for the remainder of the game. If both players play the grim trigger strategy, the cooperative outcome is an equilibrium of this game, and the Bertrand equilibrium of pricing at marginal cost can be avoided (Pepall et al., 2011).[33]

The grim trigger strategy is not the only mechanism that can sustain payoffs equal to or close to the cooperative equilibrium. Other modified strategies, in which punishment for cheating is only applied for a certain number of periods (as opposed to for the remainder of the game) are also possible. One example of a modified triggering strategy is a price war.

In a price war, firms punish deviations from the competitive equilibrium by charging a low price for a fixed number of periods T, leading to low or negative profits for all firms. At the end of the price war, the punishing firm reverts back to the cooperative equilibrium until the next deviation is detected. Price wars can occur for many reasons, but in the industrial organization literature they typically serve as punishment mechanisms for firms that break the cooperative equilibrium in a indefinitely repeated game (Green and Porter, 1984; Rotemberg and Saloner, 1986; Pepall et al., 2011).

Price wars are a well-documented behavior in competitive airline markets (Brander and

[33]It is important to note that the Folk Theorem only proves the *existence* of equilibria, and it does not imply anything about which equilibrium will actually result. In this game, both the cooperative outcome and the Bertrand outcome can be equilibria of the game, but it is up to the players to decide which equilibrium will result when the game is played indefinitely.

Zhang, 1993; Morrison and Winston, 1996; Busse, 2002; Zhang and Round, 2011; Ciliberto and Williams, 2014). These behaviors can be very costly to airlines; Morrison and Winston (1996) found that price wars lowered U.S. airline profits by $8 billion between the years of 1979 and 1995.

Despite their costs, Morrison and Winston (1996) identify how price wars are used in the airline industry as a punishment mechanism to respond to deviations from the cooperative equilibrium in markets that are served by multiple airlines:

> "Multimarket contact could stimulate fare wars because carriers engage in 'price disciplining,' where they respond to price cuts by a rival in their most profitable markets by cutting prices in their rival's most profitable markets. This behavior could escalate into a fare war." (Morrison and Winston (1996), p. 98)

Additional work has provided evidence of other types of price wars. In a classic model, Rotemberg and Saloner (1986) find that price wars are more often to occur in "boom" parts of business cycles when demand is high. Zhang and Round (2011) found empirical evidence of this phenomenon in China, where price wars were more common during periods of the year with higher demand, although Morrison and Winston (1996) find that price wars also occur to fill excess capacity when airlines make errors in estimating supply versus demand.

According to some economists, fare wars are a natural outcome of competitive markets with repeated interactions. As we will discuss later, this is particularly the case when a firm's actions can not be perfectly monitored by other firms (Green and Porter, 1984; Matsushima, 2004). An additional hypothesis is that a fare war could occur due to a change in an airline's discount factor, which would change the strategies that can be sustained in equilibrium.

For instance, if an airline is facing financial difficulties, or if a new management team is pressured to improve financial performance, short-term results may become more important, reducing the discount factor and making it more difficult to sustain the cooperative equilibrium. Busse (2002) found evidence that airlines in financial distress are more likely to engage in fare wars, providing some evidence for this hypothesis.

Ultimately, the presence of the Folk Theorem in many indefinitely repeated Prisoner's Dilemma games means that the race to the bottom in competitive markets with repeated interactions is not assured. However, as mentioned above, the Folk Theorem only provides us with the *existence* of equilibria; it does not tell us which equilibrium will be present in any given instance of the game. Different market structures and competitive environments could also lead to the emergence of different equilibria. To address this point, we next discuss how dynamic pricing could change the competitive structure of the airline industry.

6.1.3 How could dynamic pricing change airline competition?

In the section above, we provided explanations of several features of airline markets—capacity constraints, product differentiation, imperfect information, and repeated multi-market interactions—that could give airlines incentives to price above marginal cost and avoid the race to the bottom. Given the airline industry's high levels of profitability in the mid-2010s, it is reasonable to assume that these factors are currently mitigating this risk in today's industry environment.

However, airlines have expressed concerns that dynamic pricing *specifically* will spur a race to the bottom. If this is correct, it must be the case that one or more features of dynamic pricing would change the ways that airlines interact competitively. In this section, we investigate several features of dynamic pricing that could result in changes to competitive environments and discuss how these changes may or may not change competitive outcomes.

Consider a dynamic pricing heuristic such as PFDynA. In PFDynA, the existing structure of pricing and revenue management remains largely intact. Airlines file fares publicly, and then use existing revenue management systems to determine which products are made available. Dynamic pricing adds a third stage: each airline then individually decides whether to adjust the price of the lowest available filed fare, either up or down, for each booking request. This new stage leads to three distinct differences from status quo pricing and RM:

- **Continuous price points:** With dynamic pricing, prices are no longer limited to a discrete, pre-defined set.
- **Segmentation and customization:** Customers of different types are offered different prices, which may or may not be customized.
- **Imperfect private monitoring of dynamic pricing actions:** Dynamically adjusted prices may not be filed, or be filed as private fares.[34] At the limit, dynamic prices may be visible only to a single customer and not observable by other airlines.

First, consider continuous price points. Suppose, at the limit, that there was a fare class for each unit of currency in a reasonable range, such as every dollar from [$50, $1000] in a domestic market. If all of those price points were filed by the airlines as separate fare classes, it would be more difficult to monitor changes in pricing actions by looking at changes in published fares. For instance, if an airline decides to change its lowest possible price point in a market from $150 to $149, it would simply open the $149 class in its inventory system.

[34]Private fares are fare products that are not distributed publicly via ATPCO. These fares are common for use with specific travel agencies or corporate clients that may receive special negotiated rates (Vinod, 2010).

In this case, no new fares would be filed, and pricers would not be alerted to the change. This action would still be reasonably public, insofar as airlines would be able to monitor availability to detect the presence of undercutting behavior. It is not clear that continuous price points alone, with or without other features of dynamic pricing, would lead to a significant change to the status quo. However, it would require frequent monitoring of competitor availability to ensure that changes to lowest-offered price points are quickly detected.

Next, consider the effects of segmentation and personalization. In Chapter 2, we discussed some of the findings from the literature regarding targeted and customized pricing. We review one of those results (Thisse and Vives, 1988) briefly here.

In Figure 6.3, we return to our familiar setup of horizontally differentiated flights on a normalized Hotelling line. Suppose again that the flights are located at the ends of a [0,1] line segment and that customers, who are normalized to unit demand for convenience, are uniformly distributed along the line. We found above in equations (42) and (43) that the equilibrium prices for this game with product differentiation are $p_1^* = p_2^* = c + d$, where d is the disutility from moving away one unit from a customer's preferred attribute location.

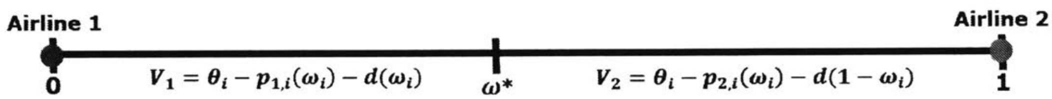

Figure 6.3: A Hotelling line with personalized pricing

Now suppose that the firms can exactly observe each customer's position on the line. That is, each firm can observe ω_i for each customer i and charge personalized prices $p_{i,1}(\omega_i)$ and $p_{i,2}(\omega_i)$. A customer i at location ω_i will receive utility $V_1 = \theta_i - p_{1,i}(\omega_i) - d(\omega_i)$ from purchasing Airline 1's flight at price $p_{1,i}$, and utility $V_2 = \theta_i - p_{2,i}(\omega_i) - d(1 - \omega_i)$ from purchasing Airline 2's flight at price $p_{2,i}$; he will choose the flight that gives him higher utility. At equal prices, the indifference point ω^* is equal to $\frac{1}{2}$.

Consider Airline 1's pricing game. If $\omega_1 > \frac{1}{2}$, then the customer is located in Airline 2's area of influence, and Airline 1 will not be able to attract her to book away from Airline 2 at a price of $c + d$. As a result, Airline 1 lowers the price for this customer to its marginal cost c.

If $\omega_i < \frac{1}{2}$, Thisse and Vives (1988) show that the optimal price is $p_1 = c + d(1 - 2\omega_i)$. In either case, the prices charged are less than $p_1^* = c + d$, which is the competitive equilibrium without personalized pricing. As a result, in this game, personalized pricing results in lower prices and lower revenues.

However, the Thisse and Vives (1988) game makes several important assumptions that may not be the case in actual airline markets. The first is that no additional customers will choose to enter the market as a result of the lower prices. That is, there is no demand stimulation, which is unlikely to be the case in the real world.

Thisse and Vives (1988) also assume perfect targetability of customers. Firms in their model can see exactly where each customer is located on the line, and the game devolves into Bertrand competition for each customer. This is not realistic in practice; customers will likely only be able to be segmented into general groups, and the exact preferences or behavior of individual customers is unlikely to be known. This can actually be a good thing for airlines; research by Chen et al. (2001), Chen and Iyer (2002), and Shaffer and Zhang (2002) have found that imperfect targetability can actually lead to higher revenues than perfect targetability when it comes to personalized pricing.

Furthermore, airlines are already practicing segmentation through their design of fare restrictions and itinerary products. The presence of fare restrictions like Saturday night minimum stays and nonrefundability are an attempt by airlines to practice third-degree price discrimination and segment customers into groups. This behavior can lead to higher fares for certain groups of customers, and provides a countervailing force to the race to the bottom. As a result, it is unclear that the imposition of segmented pricing directly in a dynamic pricing algorithm would result in a significant change from the status quo, since airlines are already practicing segmentation today.

Finally, we arrive at imperfect private monitoring of pricing actions. With dynamic pricing mechanisms, airlines may be able to offer fares that are not filed publicly and which may not be visible to other firms. For example, if a dynamic price adjustment is targeted to a specific customer or segment of customers, it may not be visible in an anonymous internet search or GDS availability display, and pricing tools may not be able to immediately identify the change in pricing behavior. The lack of public monitoring of dynamic pricing actions has often been cited by airlines as a key concern in reference to the race to the bottom.

In contrast to the other two elements of dynamic pricing, which may result in limited effects on the competitive landscape, a move away from publicly filed fares is the facet of dynamic pricing that has the potential to lead to the greatest changes to airline competition in regards to the race to the bottom.

6.1.4 Indefinitely repeated games with imperfect and/or private monitoring

The economics literature has a rich set of results regarding indefinitely repeated games with imperfect and/or private monitoring. These papers generally focus on proving Folk Theorems in situations when firms' actions are either imperfectly visible or invisible to other firms. In many cases, Folk Theorems are possible to prove even without perfect public monitoring, suggesting that non-Bertrand equilibria are still possible in these environments.

Consider again a generic indefinitely repeated Prisoner's Dilemma game. Strategies like the grim trigger, discussed in the previous section, assume that firms can observe each others' actions at each stage of the game and, if necessary, apply appropriate punishment for deviations. In the case of the grim trigger, a firm will choose to defect in perpetuity after the first time it observes its competitor choosing to defect. But what if the competitor's actions are not perfectly observable? If the firm cannot observe when its competitors are cheating, it may be more difficult to apply punitive actions, weakening cooperative equilibria.

A branch of economics literature has evolved to consider situations when competitor actions cannot be monitored. First, let us begin with a taxonomy of monitoring:

Perfect vs. Imperfect Monitoring

A game possesses *perfect monitoring* when during each period t of the game, all firms can observe the set of actions $A = (a_{1,t}, a_{2,t}, \dots, a_{n,t})$ played by all n players. A game possesses *imperfect monitoring* when firms can only observe signals $\Xi = (\xi_{1,t}, \xi_{2,t}, \dots, \xi_{n,t})$ that may be correlated with actions $a_{i,t}$.

Public vs. Private Monitoring

A game possesses *public monitoring* when all firms receive the same monitoring information (either actions A or signals Ξ) in each stage t of the game. A game possesses *private monitoring* when firms receive private or limited information about actions or signals that may be different from information received by other firms.

The typical Bertrand competition scenario can be seen as a game with perfect public monitoring. In each stage, each firm can directly observe the actions of all competitors, and can use that information to decide whether to continue to cooperate or to play a punishment on the next turn. In a variation, suppose instead that firms cannot perfectly observe opponents' actions, but instead receive a common signal that is correlated with the actions played by each firm. For instance, firms may not be able to directly observe its competitors' pricing actions, but could observe aggregate industry sales figures in each period.

Green and Porter (1984) present a canonical example of a game with imperfect public monitoring. Their game models Cournot competition, where firms set quantity production and then observe a public market price, but the insights can be easily translated into a Bertrand price-setting environment (Tirole, 1988; Ellison, 1993). For instance, suppose two firms set their (private) market prices for a homogeneous good in each time period. Firms cannot observe their competitors' pricing actions directly, but they can observe a common signal $Q(p)$

about the aggregate quantity sold in the market. Suppose also that there is a possibility of an exogenous upward demand shock in the market, which occurs with probability α.

When firms both play the cooperative strategy and there is no demand shock, the total quantity sold in the market will equal $Q(p^*)$. However, suppose that in a given period the firms observe a signal of $\xi = Q(p^?) > Q(p^*)$. Two things could have happened here: the market could have experienced an upward demand shock, or one of the firms could have deviated and played a lower price, stimulating demand. Since the players' actions are not public, they have only this imperfect public signal available to tell them whether a firm deviated in the previous period.

Green and Porter (1984) and Tirole (1988) show that in this environment, the perfectly cooperative equilibrium is not possible. However, a partially-cooperative equilibrium is. Suppose that firms play using this strategy: they begin by playing the cooperative price p^* until they observe a higher quantity sold in the market than expected. At this point, the firms enter a price war for T periods during which they set $p = c$ (marginal cost). At the end of the price war, they return to the cooperative equilibrium until they observe the high-quantity signal again.

In this model, firms will not be able to avoid price wars at least some of the time (Ellison, 1993). However, the presence of the public signal means that punishment can still be applied when a signal that does not meet the conditions of the cooperative equilibrium can be observed by all firms. In this game, the threat of punishment will still induce firms to play the cooperative action most of the time. Green and Porter (1984) show that price wars will happen less often and for shorter periods T when the chance of the demand shock α is low, because it will be more obvious when firms are deviating from the cooperative equilibrium. Therefore, in the Green and Porter (1984) model, cooperative outcomes can still be maintained most of the time even when perfect monitoring is not possible.

When firms receive only private information about the actions or signaled actions of competitors, cooperation becomes even more difficult. This is due to the lack of consistency in the information received by players (Kandori, 2002). When monitoring is public, players can be certain or nearly certain when to enter the punishment phase. When monitoring is private, only certain players may receive information of a defection. In this case, coordinating a punishment action is more difficult (Matsushima, 2004).

Consider a variant of the Green and Porter (1984) model where some firms do not receive any information about the total quantity sold in the previous period. Since they do not observe the public signal, these firms would not know when or whether to enter the punishment phase. The difficulty of punishment in this setting increases the incentives to defect.

Until recently, the existence of a Folk Theorem in indefinitely repeated games with private monitoring was uncertain (Kandori, 2002). However, some progress has been made in this area. Ely and Välimäki (2002) and later Matsushima (2004) in a general case have shown the Folk Theorem is still valid in a two-player game with imperfect private monitoring. In the words of Matsushima (2004), cooperative equilibria are "sustainable even if the price level for each firm is not observable to its rival and each firm has the option of making secret price cuts" (p.843). This is a good representation of the dynamic pricing scenario with privately filed fares—airlines would be unable to see each others' prices, and could choose to charge prices lower than the filed fares for some customers.

Matsushima (2004)'s analysis focuses on "review periods" during which firms consider a certain number of privately-observed signals and decide whether to respond with a punishment. He shows that both the general prisoner's dilemma and the price-setting game with secret price cuts allow for the cooperative equilibrium, as long as private signals received by firms are only correlated through the unobservable random shocks.

This stands in contrast to Green and Porter (1984), whose punishment phases result from firms being unsure whether a negative signal was a result of a macro shock or a deviation on the part of a competitor. In the Matsushima (2004) model, firms' signals do not depend on the random shocks, allowing punishment to be applied. Therefore, when firms identify instances of higher-than-expected demand during its review phases, it can recognize that these signals are the result of its opponent's deviation and apply punishment appropriately. This mechanism, like the grim trigger, allows the cooperative equilibrium to remain intact.

Since Matsushima (2004)'s groundbreaking paper, further work has proven Folk Theorems in more general environments (Awaya and Krishna, 2016). Yamamoto (2009) later extended the Matsushima (2004) result to N-player prisoner's dilemma games, and Sugaya (2011) later showed that the Folk Theorem holds in general N-player games, as long as players receive a large number of signals and players are patient enough. Recall, however, that the Folk Theorem does not tell us that the cooperative payoffs *will* result from an indefinitely repeated game with private monitoring; just that such an equilibrium *could* result. While the mathematical analysis behind the Folk Theorem shows that cooperative payoffs are possible in equilibrium, the strategies that lead to these payoffs are often extraordinarily complicated, and it is unclear that they would occur naturally in the real world.

Theoretical research has shown several factors that may make it easier for the cooperative equilibrium to result from an indefinitely repeated prisoner's dilemma with private monitoring. The first is a mechanism for communication between players. If players are able to communicate their actions to each other, this both improves the quality of the signals provided to the players and removes the coordination problem with private monitoring described above. Compte (1998), Kandori and Matsushima (1998) and Obara (2009) have proven Folk Theorems for repeated games with private monitoring and communication.

Direct communication between firms regarding pricing is generally illegal for reasons of antitrust. Yet this does not mean that monitoring is impossible. Technology companies currently exist that monitor the lowest available fares of airlines and sell that information to competitors. It is very likely that similar technology would emerge to monitor dynamic pricing behavior. Customers themselves could also serve as a monitoring device. Airline customers are remarkably vocal about prices, often discussing on online forums and blogs when new, lower prices are made available to even a single customer. Such communication could also serve as a signal to competitors if one airline is engaging in undercutting behavior.

Even if this monitoring is not conducted in real-time, it may still serve as a deterrent to short-term undercutting. Compte (1998) shows that delayed communication can actually improve outcomes for firms, because it maintains the threat of punishment without the short-term inefficient costs of enforcement. That is, even delayed information about pricing behavior can provide a mechanism to discourage defection. A market for information as described above could move the game closer to one of perfect public monitoring. More public information reduces the incentives to defect, because it would make it more obvious which players defected in each period and easier for punishments to be applied.

Summing up, the move from perfect public monitoring to imperfect private monitoring can change the set of equilibria from indefinitely repeated Prisoner's Dilemma games. Without perfect signals about which players defected in each period, it is more difficult to apply punishments to defecting players. This increases the incentives to defect.

However, economists have shown that Folk Theorems still hold in many indefinitely repeated games with imperfect private monitoring. That is, cooperative payoffs are still possible in equilibrium in these games. The strategies that bring about these equilibria may be very complicated, however, so it is likely that new markets of information would emerge to improve the monitoring quality. Economic theory has shown that these types of markets for information can make it less likely that the race to the bottom will occur.

6.1.5 The possibility of irrational agents: What if airlines do not behave as economic theory would suggest?

While the economic theory reviewed above suggests that airlines could potentially avoid the race to the bottom even in an environment with dynamic pricing, the possibility remains that dynamic pricing could spur revenue-damaging pricing responses from airlines. Particularly, if some airlines do not behave in the ways in which economic theory would predict, there may be short-term costs as airlines transition from legacy patterns of competition to a new world of dynamic pricing.

It is possible that some airlines will not understand how to react to new dynamic pricing technologies practiced by their competitors. Airlines that are used to monitoring changes in prices through public fare filings may be confused and surprised if a competitor begins to sell dynamically-adjusted fares that are not published through normal means. It is possible that airlines could respond bluntly to this type of action, broadly instituting lower fares across the marketplace as a means of discouraging asymmetries in pricing technology.

We have also assumed that even in Bertrand competition, airlines will never choose to price below their marginal cost (i.e., their bid prices for each flight departure). But this may not be a reasonable assumption in all cases. Airlines have been shown to practice "predatory pricing" below marginal costs to discourage competitive entry in certain markets (Robenalt, 2007; Sagers, 2009). If marginal costs do not serve as a lower bound to a race to the bottom, the Bertrand equilibrium may be even more problematic for industry profitability.

As airlines start to return to a capacity growth regime, they may have more incentive to provide discounts to fill empty seats. The results from Chapters 4 and 5 show that intelligent discounting can increase industry revenues even if practiced by all airlines—namely, giving discounts to certain leisure-type customers when prices are ordinarily high. But it is not guaranteed that airlines with a dynamic discounting capability will use it scientifically or rationally; if dynamic pricing were to lead to widespread discounts below the lowest filed fares, it could lead to revenue losses for airlines.

These risks are harder to model through economic theory, since they presuppose agents that are not acting "rationally" in an economic sense. Pricing below marginal costs to send a "message" to a competitor with a pricing technology that an airline does not understand may not be the economic "best response," but it may align with the human psychology of airline managers. As dynamic pricing mechanisms become more widespread in the coming years, airlines may need to reconsider how they monitor and respond to the pricing actions of their competitors.

6.1.6 Wrapping up the race to the bottom

In this section, we discussed various economic factors related to the race to the bottom: the incentive for firms selling homogeneous products to undercut each other on price until one firm reaches its marginal cost. While the race to the bottom has been observed in the past in the airline industry, high levels of profitability suggest that the race to the bottom may not currently be in effect. However, some airline managers are concerned that dynamic pricing could reintroduce the race to the bottom to the airline industry.

These concerns align with theoretical economic literature that has explored why the race to the bottom may not always occur in practice. We discussed five features of airline markets: capacity constraints, imperfect information, product differentiation, repeated interactions, and perfect public monitoring, that could explain why airlines are at times able to avoid price wars and marginal cost pricing. Most importantly, repeated multimarket interactions and the public nature of filed fare structures serve as monitoring and punishment mechanisms for airlines that defect from the cooperative equilibrium. Some studies that suggest that occasional price wars can be a natural part of indefinitely repeated games, even without dynamic pricing.

The feature of dynamic pricing that could most likely lead to a race to the bottom is if dynamic pricing requires airlines to partially move away from publicly-filed fares. In this case, it may be more difficult for airlines to observe the pricing actions of their competitors, and could allow airlines to practice secret price cuts. In the economic literature, this situation is called imperfect private monitoring.

In these environments, economists have proven some Folk Theorems, which state that any individually rational payoff from a one-shot game can be sustained in equilibrium in an indefinitely repeated game as long as players are sufficiently patient. This means that in indefinitely repeated games, the cooperative equilibrium is still possible, and the race to the bottom is not assured. To maintain cooperative equilibria, firms in indefinitely repeated games often practice "trigger strategies," in which an observed deviation from the competitive outcome is punished by a temporary or permanent shift to an inefficient outcome in future periods. In the airline industry, these punishments could be applied through price wars or through network planning actions like adding capacity to a competitor's market.

Economists have shown that even in games with imperfect private monitoring, the race to the bottom is avoidable as long as firms' private signals contain enough information about competitor deviations. But this work tells us only that such equilibria are possible, not likely, and imperfect private monitoring may make it more difficult to sustain these equilibria relative to a public monitoring case, such as the current publicly filed fare environment.

If airlines do begin practicing dynamic pricing, it is likely that technologies will evolve to monitor competitors' pricing actions and to enforce punishments on players that defect from cooperative equilibria (for instance, by setting arbitrarily low Q-multiplier inputs to generate more discounts than would be advisable in a revenue-positive scenario in competition). Economic theory suggests that increasing monitoring or communication technologies will make it more likely that a cooperative equilibrium will be sustained in an indefinitely repeated game. Yet there may be some short-term costs to this transition, particularly if airlines respond in "irrational" ways to new technologies or pricing strategies that run counter to the current status quo of airline competition.

6.2 Regulatory and legal implications of dynamic pricing

Historically, regulators have seemed to accept one of the basic tenants of airline pricing and revenue management: that different customers booking at different times can be charged different prices. However, will customized or personalized pricing—offering the exact same product to different customers at different prices at the same time—pass the same legal tests?

6.2.1 The DOT's Order to Show Cause for IATA Resolution 787

While neither the Department of Transportation nor the Department of Justice has launched a full legal review of dynamic pricing, the DOT did consider some arguments for and against the technology in the run-up to approval of IATA Resolution 787, which authorized airlines to begin development work on the New Distribution Capability. In response to the DOT's Notice of Proposed Rulemaking, several consumer advocates filed complaints during the comment period, alleging that NDC could harm competition and lead to higher prices for consumers.

In its Notice to Show Cause from 2014, the DOT put forth many important policy positions in reference to dynamic pricing. The DOT acknowledged that revenue management systems currently result in many variations in prices charged:

> "Prices paid by individual passengers already vary widely due to carriers' virtually universal use of yield management techniques, though this price variation has generally been based on variation in the nature of the trip (including when it was booked) rather than variation in the nature of the person taking the trip. We are tentatively not prepared to prohibit future innovations that may better match capacity with demand." (Department of Transportation (2014), p. 14)

However, the Department also acknowledged that "a number of aspects of NDC as presented in Resolution 787 created serious concerns in the related areas of privacy and anonymous shopping." Specifically, consumer advocacy groups worried that "requiring passengers to disclose information such as age, marital status, type of trip (e.g. leisure or business), frequent flyer status, or nationality as a condition for receiving a quote would enable carriers to engage in more perfect price discrimination, raise fare levels, and harm competition" (Department of Transportation (2014), p. 12).

In response to these claims, the U.S. DOT instituted a requirement that anonymous shopping (in terms of voluntary passenger disclosure of information) would always be required. That is, customers would never be required to submit any identifiable personal information in the course of a shopping request. Of course, this does not prevent airlines from using other information about the request—for instance, trip origin, point of sale, search parameters, or search history—to segment customers or personalize offers.

From past precedent, the Department of Transportation has been fairly permissive in allowing airlines to segment their demand and perform price discrimination. As mentioned above, the Department of Transportation "has not found fare differences associated with certain status indicators, such as family fares, companion fares, affinity travel, and corporate and government travel management arrangements, to be unreasonably discriminatory" (Department of Transportation (2014), p. 13), and third-degree price discrimination in terms of senior fares and point-of-sale pricing has been largely permitted. However, airlines are not permitted to discriminate against passengers based on several protected dimensions, such as age, sex, color, religion, or national origin. Furthermore:

> "...whether other potential bases for price discrimination, such as income level, marital status, and trip purpose, would be unreasonably discriminatory or constitute an unfair or deceptive trade practice we leave to future determination." (Department of Transportation (2014), p. 13).

Finally, the DOT closes the relevant section of its Order to Show Cause by dismissing a legal challenge that personalized pricing would not be allowed because it would allow carriers to charge fares that were not publicly filed. As the Department states, such a requirement has not existed in the U.S. air transportation industry since its deregulation in 1978, and that open skies agreements in other countries has also eliminated the requirements for publicly filed fares in international markets. That is, the move towards imperfect private monitoring of airline prices would appear to not be rejected on regulatory grounds, based on the language in the Show Cause order (Department of Transportation, 2014).

This Order to Show Cause, produced by a Department of Transportation in a Democratic administration, suggests an agency that may be relatively tolerant of dynamic pricing—even those mechanisms that may be used with more highly-targeted segmentation practices. Shifting political attitudes could change the willingness of regulators to allow targeted dynamic pricing, particularly as the U.S. airline industry becomes more consolidated. Moreover, in other parts of the world such as Europe, regulators are highly protective of customer rights and data. In these locations, personalized or even segmented pricing may face legal headwinds if these practices are viewed as discriminatory.

6.2.2 Legality of dynamic pricing and price discrimination in other industries

Particularly in online markets, price discrimination has faced legal challenges in other U.S. industries. In a case in 1996, a plaintiff filed suit against the lingerie retailer Victoria's Secret, alleging among other complaints that female customers received catalogues with lesser discounts than catalogues sent to men (Weiss and Malhotra, 2001). This lawsuit was rejected in what is seen as a benchmark ruling; as in the airline industry, courts have typically upheld the rights of businesses to practice third-degree price discrimination.

But despite the legal precedent, retailers have often faced consumer pressure when evidence of dynamic pricing is discovered. In a Wall Street Journal article in 2012 (Mattioli, 2012), the online travel agency Orbitz admitted that users of Macintosh computers were shown more expensive sets of hotels than users of Windows PCs, and a local television station investigation in Minneapolis in the same year (Collin, 2012) identified cases of Delta Air Lines presenting higher prices to customers that were logged in to their SkyMiles loyalty program accounts (Delta blamed the issue on "computer problems").

Retailers in other industries appear to have already begun targeted dynamic pricing based on customer information. In a separate Wall Street Journal investigation (Valentino-DeVries et al., 2012), journalists found that the office supply retailer Staples charged different prices for identical products on its website based on customer locations. Staples claimed the differences were due to "a variety of factors" including "the costs of doing business" in different locations.

Other retailers found to be practicing this type of dynamic pricing in the Wall Street Journal investigation included Discover, Office Depot, Rosetta Stone, and Home Depot (Valentino-DeVries et al., 2012). This type of dynamic pricing is not *per se* illegal, but media investigations can cause negative publicity for companies that often apologize publicly for the pricing tests, as Amazon did in 2000 after media reports uncovered that different customers were charged different prices for the same DVDs.

Some firms have been more upfront about their price discrimination strategies. For years, the ride-hailing Uber displayed "surge pricing" multipliers to inform customers when their ride had been marked up, ostensibly due to higher levels of demand (although the actual mechanism behind Uber's surge pricing has never been publicly disclosed). Uber faced consumer backlash when high multipliers were displayed in certain high-demand periods, such as New Year's Eve. As a result, in 2016, this overt notification was removed from the Uber app, which then displayed a price quote with a small icon indicating the presence of surge pricing instead of a specific multiplier (Hawkins, 2016).[35]

Platforms like Uber have served to normalize the idea that prices for transportation can change from second to second, and may vary for different customers making the same request. In a future world, airline shopping may work much in the same way: a request for itineraries may be met with customized offers with dynamic, offer-specific prices that could change for different customers, or if the request is resubmitted.

The existing legal guidance by the Department of Transportation and the U.S. judiciary has seemed to indicate that such actions would not be *per se* illegal; firms are allowed to practice price discrimination, and the Department of Transportation in a Democratic administration has seemed to permit a surprising degree of flexibility for airlines to segment customers to "better match capacity with demand."

However, it is likely that a widespread implementation of dynamic pricing in the U.S. airline industry would trigger demands for a legal review by consumer advocacy groups. While such actions may not be illegal, airlines may need to prepare themselves for a media backlash if dynamic pricing mechanisms are portrayed in a negative light.

Specifically, fare incrementing strategies (that do not include additional ancillary services in the offer) could prove to be problematic publicly, and could lead to customer gaming as discussed in the next section. Framing dynamic pricing in terms of discounting in relation to a filed reference price may be more palatable for consumers; many customers are used to receiving targeted coupons or discounts through email or in-store promotions.

Ultimately, just as Uber has been able to train consumers to accept dynamically changing prices, so too could airlines reframe the way shopping for air transportation is currently seen by consumers. The end result could be a more flexible pricing environment for airlines that could ultimately allow them to offer products that are targeted to the customers making the request.

[35]Incidentally, this change happened shortly after an expert in airline revenue management was hired as the Head of Dynamic Pricing Research at Uber.

6.3 Potential for customer abuse and the framing of dynamic offers

As we showed in Chapter 4, much of the success of the PFDynA dynamic pricing heuristics comes from the ability to segment customers into different types and to charge each segment a potentially unique price. Some airlines have brought up the question of customer "cheating" with dynamic pricing. That is, what if customers are able to trick or fool the dynamic pricing engine into thinking they belong to a different (leisure) segment, and are thus eligible for a lower price?

Gaming could occur if customers become aware of discounts provided to certain segments and are able to fool the dynamic pricing system into believing that they are eligible for these discounts. Frequent-flyer websites are constantly on the lookout for these types of deals or actions in today's environment, which could involve behaviors like accessing a foreign version of an airline's website to change the point-of-sale of the ticket. Some travel agents have also been known to practice this type of point-of-sale abuse in order to obtain lower prices.

To combat this, airlines often segment based on the characteristics of the *request*, not characteristics of the *customer*. For instance, a customer searching for a two-night trip from Monday to Wednesday from New York to Washington, DC would likely be a business traveler, whereas a customer requesting a price for a two-week long trip from a small community to a leisure destination is likely to be a leisure passenger. Airlines often practice this type of segmentation today through the use of fare rules. This type of segmentation is more difficult to game, since the characteristics of the request cannot be modified by using a different browser or an anonymous shopping option.

In our simulation results in Chapter 4, we also investigated this question through our experiments with identification accuracy. While an airline's identification accuracy could be seen as the chance that the airline will make a mistake in identifying a customer, the same tests could be reframed as the chance that a customer will be able to fool the system into thinking she belongs to a different segment. In our simulations in Sections 4.3.4, 4.4.5, and 4.5.5, we found that dynamic price adjustment was still revenue-positive in competition even with an identification accuracy rate of about 80% or lower.

If incrementing heuristics were used and segmentation was practiced based on the characteristics of the *customer*, gaming to avoid the increment could be very easy. Just as the case of the Delta frequent flyer in Minneapolis who found that he was being charged a higher price due to his log-in information (Collin, 2012), such customers could instantly access the "normal" unincremented prices either through the use of an anonymous or incognito web browser, or through practicing anonymous shopping as is guaranteed through the U.S. Department of Transportation's adoption of IATA Resolution 787 (Department of Transportation, 2014).

No consumer would want to self-identify any information that would lead to being charged a higher price. But this does not mean that incrementing strategies will be impossible to implement. Airlines could also combine incremented prices with additional ancillary services to create dynamic, customized offers. If the dynamic pricing engine suggests that a higher price for a certain customer, that higher price could be combined with additional features and services, such as early boarding, lounge access, or additional frequent flyer miles. Airlines would not charge higher prices for exactly the same offer, but would offer different combinations of products at different prices in response to different requests. These offers may contain products that are more relevant to customers with a higher willingness-to-pay.

Reasonable framing of dynamic pricing by airlines is also likely to be important to the success of dynamic pricing. The experience of Uber with surge pricing suggests that customer backlash may prevent airlines from explicitly incrementing prices for certain customers. However, some elements of the Uber pricing model could be relevant for airlines. Customers may benefit from certain limited amounts of information presented along with price quotes. For instance, icons that suggest that a certain itinerary is facing "higher demand than normal," or that a certain request is "popular on your dates," as is currently done on some hotel booking sites, could inform customers about why the prices they face are higher for certain itineraries or at certain times.

Some experimental studies have also shown that customers are more likely to consider price differentiation as fair if the prices are presented in the form of discounts from a higher price. For instance, participants in one study (Weisstein et al., 2013) were shown a webpage with a USB flash drive for sale. The price of the product was shown as either \$12 or \$15 with a \$3 discount. They were then told that a friend had paid \$10 for the same flash drive. The participants were asked if they felt the amount they had to pay for the flash drive (in either case, \$12) was fair.

Participants that were shown the website with the flash drive priced at \$15 with a \$3 discount were more likely to agree that their price was fair than customers who were shown the \$12 price. Framing the \$12 price as a discount from a higher reference price changed customers' perceptions of price fairness. As behavioral economics would suggest, changing the manner in which a price is presented to a customer can change not only how the customer makes decisions about a transaction, but also how she feels about the transaction after the fact.

6.4 Conclusions: Will practical concerns prevent the implementation of dynamic pricing in the airline industry?

Despite the revenue potential of dynamic pricing, concerns exist within the airline industry that dynamic pricing technology is either unrealistic or borderline irresponsible—that it could irreparably change the fundamentals of airline competition and ultimately leave airlines worse off than traditional revenue management and pricing practices. In this chapter, we considered several common concerns with dynamic pricing, specifically focusing on how dynamic pricing could result in a change from the status quo.

First, we discussed the prospects for a race to the bottom in an increasingly commoditized air transportation landscape. We reviewed some concepts in economic theory that hypothesize why a race to the bottom may not occur in practice, with a particular focus on indefinitely repeated pricing interactions. Economic research, both theoretical and empirical, has suggested that repeated multimarket contact can create incentives for airlines to avoid the race to the bottom. In these types of economic games, the threat of punishment for deviation from a cooperative strategy can avoid an inefficient equilibrium like the race to the bottom.

A shift to dynamic pricing could lead to an environment of imperfect private monitoring of pricing behaviors, in which firms cannot observe the pricing actions made by their competitors. In this environment, identifying deviations from the cooperative equilibrium and applying punishments for those deviations becomes more difficult. However, economists have shown that so-called Folk Theorems, which prove the existence of cooperative equilibria, are still valid in many types of indefinitely repeated games with imperfect private monitoring. It is likely that new public monitoring technologies or data exchanges, such as web scrapers or price monitoring services, would emerge if dynamic pricing technology were to become widespread. Such marketplaces for information could also serve to prevent race to the bottom outcomes, although there may be some short-term costs as airlines transition to a new competitive landscape.

We also briefly examined the potential legality of dynamic pricing and price discrimination in general. Typically, lawsuits brought against companies for practicing price discrimination have not been successful, and companies have been given rights to charge different prices to different customers, as long as the discrimination does not fall along the lines of protected features like race or gender. The U.S. Department of Transportation has also appeared willing to allow airlines to customize offers to "better match capacity with demand." Yet dynamic pricing may well trigger potential legal reviews, and airlines should be prepared for potential customer backlash as customers become adjusted to more dynamically changing airline prices along the lines of other transportation platforms like Uber or Lyft.

Finally, we discussed the possibility for gaming on the part of consumers tricking the system in terms of their eligibility for discounts or more attractive offers. This is not a new phenomenon with dynamic pricing, as airlines currently fight point-of-sale abuse, hidden-city ticketing, and other gaming behaviors in existing distribution systems. Our simulation results from Chapter 4 found that the dynamic pricing heuristics are fairly robust to identification errors or customers gaming the system. Clever segmentation techniques focused on the characteristics of the request or upon passenger characteristics like frequent flyer status that cannot be forged could serve to lessen the possibility of abuse of dynamic pricing systems.

Moving to new technologies is always risky, and airlines that choose to practice dynamic pricing will not immediately be able to do so perfectly or completely. As dynamic pricing systems are rolled out, airline pricing, revenue management, and distribution departments will need to adapt to new ways of doing business. Errors will inevitably be made, and legal and media reviews of dynamic pricing techniques will likely ensue. It will be up to the airlines to decide whether these risks are worth the potential gains of dynamic pricing, increased segmentation, and offer customization.

7 Conclusions, Extensions, and Current Industry Developments

As airlines have gained access to increasingly detailed information about their customers and have begun to sell even more varied and complex products, many carriers have found themselves limited by the rigid structure of legacy distribution technologies. These constraints have served as barriers to pricing and revenue management innovation, since any move towards more complex dynamic pricing mechanisms would need to be interoperable with existing systems in the short term.

In this dissertation, we have motivated, designed, and evaluated the first methods for integrating dynamic pricing mechanisms into existing airline pricing and revenue management processes. This work has helped to answer some pressing questions about the implications of using dynamic pricing to adjust pre-filed fares. Our development and simulation of these next-generation pricing methods indicates that they have the potential to increase airline revenues when used by a single airline and in competition. They also have the potential to significantly change both the ways in which consumers shop for airline tickets and the ways in which airlines interact in competitive markets.

In this chapter, we review the key contributions, findings, and implications of the dissertation, and provide some ideas for future research and extensions of the models proposed in this work.

7.1 Review of key contributions, findings, and implications

We began our discussion with an overview of existing pricing and revenue management practices in the airline industry. Currently, airlines are limited by pre-Internet distribution architecture to offer only a finite number of price points in each market. Airlines publicly file pre-priced fare products with central distribution agencies, which are then communicated to customers as well as to other airlines. Decades of development in revenue management research have provided airlines with algorithms to select which fare products to make available as a function of forecasted demand relative to remaining capacity.

Recent advancements in distribution technology could give airlines more flexibility in the ways in which they price and sell their products. For instance, the New Distribution Capability (NDC) is an XML-based distribution standard which would, among other features, allow airlines to distribute and sell product offers with unique prices that are not necessarily tied to filed fare products. These new developments have caused some practitioners to imagine a "world without booking classes" in which prices for each shopping session are generated on-the-fly based on characteristics of individual customers (Westermann, 2013).

Since so many airline commercial processes are tied to current distribution standards, switching to NDC is a complicated and risky endeavor for an airline. In the short term, airlines have begun to investigate ways to implement dynamic pricing of itineraries without a full implementation of NDC. One possible solution is the so-called "dynamic pricing engine" (DPE), which would allow airlines to modify the prices of filed fare products for certain transactions (Ratliff, 2017). Messaging standards for DPEs are currently under development in the industry, but the scientific methods that would drive the price modifications, as well as the performance of DPEs relative to existing revenue management techniques, have remained unexplored.

The primary contribution of this dissertation was to design new dynamic pricing mechanisms that are compatible with existing revenue management systems, and to compare the performance of these methods to existing RM practices. By applying targeted increments or discounts to certain customers in certain situations, the mechanisms aim to increase airline revenues by stimulating new bookings from price-sensitive passengers while increasing yield from more price-inelastic customers. In many cases, the mechanisms suggest no adjustment be made to the price of the product, and the customers book at the filed fare. In this way, traditional pricing, revenue management, and distribution systems can still be used to determine fare product availability, yet airlines are given more control over the prices that are offered in response to each shopping request.

To motivate the development of the dynamic pricing mechanisms, we first created a definitional framework that allowed us to compare pricing practices in various industries. As described earlier, airlines currently select prices by choosing among a limited set of pre-defined price points—a technique that we call *assortment optimization*. Assortment optimization is also common in other transportation-related businesses such as passenger rail and car rental agencies. On the other end of the spectrum is *transactional continuous pricing*, in which prices are selected from a continuous range of possible values for each individual transaction. This type of pricing is currently practiced in some markets by the on-demand transportation app Uber, which computes prices for each ride based on the current state of market demand and characteristics of the route.

An intermediate step between assortment optimization and transactional continuous pricing is *dynamic price adjustment*. With dynamic price adjustment, firms start with a pre-defined set of possible price points, but then can increment or decrement those prices for specific transactions. This is an attractive proposition for an airline implementation of dynamic pricing, since it allows the existing architecture of filed fare products to be used in conjunction with session-specific price adjustments. Both Uber and the online retailer Amazon have used dynamic price adjustment mechanisms to mark up or mark down their prices. For example, Uber uses "surge pricing" to mark up the normal price of a ride in periods of high demand or low driver supply.

We then turned our attention to previous work on dynamic pricing in the academic literature. Our review focused on two different disciplines' approaches to dynamic pricing: economics and operations research.

In the economics literature, dynamic pricing is framed through the lens of price discrimination, in which firms can charge different prices by segmenting customers into different categories. Airlines currently practice so-called *third-degree price discrimination* by designing products to separate leisure and business demand. As airlines gain more capabilities to customize offers, they will begin to move closer to *first-degree price discrimination*, in which each customer is given a unique price equal to his or her willingness to pay.

We also discussed the literature on behavior-based price discrimination, in which firms' pricing decisions are influenced by signals that customers send to the firms about their preferences or behavior. While this type of price discrimination can sometimes increase revenues, the literature also suggests that giving firms too much information about each customer can lead to a pattern called the *race to the bottom*, where firms undercut each other on price until they reach their marginal costs. A race to the bottom can be damaging to profitability, particularly in industries with low marginal costs and high fixed costs, such as the airline industry. While the economics literature provides good insight into microeconomic market principles, the models described in these papers are often highly stylized and have limited direct applicability to the airline revenue management setting.

The operations research literature also contains many papers on dynamic pricing. These papers typically focus on capacity-constrained dynamic pricing problems, in which there are only a limited number of resources available for sale. Models and heuristics are developed to determine what prices to charge based on remaining capacity, demand forecasts, or other information. However, competitive outcomes are rarely considered, and these approaches often have limited applicability for real-life airline revenue management. An optimal OR solution to an airline dynamic pricing problem would not be practical if it could not be implemented within the current context of fare class-centric revenue management and distribution.

As airline practitioners determine whether to begin practicing dynamic pricing, a key question is how the revenue performance of dynamic pricing would compare to existing RM approaches. Few papers in the literature have aimed to answer this question. Those that have tried (e.g., Zhang and Lu (2013); Fiig et al. (2016)) have claimed revenue gains of about 1 - 6% for dynamic pricing over traditional revenue management systems. Yet these papers do not conduct a thorough investigation of the revenue management outcomes of these mechanisms, including impacts on yields, load factors, and fare class mix. They also do not consider the implications of dynamic pricing on airline competition, consumers, and regulators. This dissertation aims to fill these gaps in the literature.

We next turned our attention to developing our dynamic price adjustment mechanism for airline revenue management. We began by creating a new model to describe how customers make selections amongst itineraries. In our model, customers have a maximum willingness-to-pay (WTP) for each itinerary, as well as valuations for various itinerary attributes, such as schedule quality or fare restrictions. The customers first evaluate affordable options with prices less than their maximum willingness-to-pay, then deterministically select the itinerary that maximizes their utility given their preferences and the characteristics of each itinerary.

A key finding from this exercise was the fact that a customer's WTP for an itinerary depends on the other affordable alternatives in his choice set. For instance, suppose a customer is willing to pay up to $300 for a given itinerary. If another airline introduces an identical itinerary for $200, the customer will no longer be willing to pay $300 for the first itinerary. In other words, the customer's WTP for a product *conditional on* the other options in his choice set is different from his maximum WTP. This was a key insight, since previous work has assumed that customers have only a maximum WTP for air travel that does not change contextually. Using maximum WTP instead of conditional WTP could cause a dynamic pricing algorithm to overestimate customer valuations, particularly in competitive scenarios.

By proving a straightforward theorem, we related choice behavior in our model to our concept of conditional WTP. Specifically, we found that the probability that a customer selects an option from a choice set is equal to the probability that the product's price is less than the customer's conditional WTP for that product. This means that we can use conditional WTP as a proxy for customer choice in our dynamic price adjustment calculations.

From this insight, we then described our method for dynamic price adjustment, which we call Probabilistic Fare-Based Dynamic Adjustment, or PFDynA. With PFDynA, a traditional airline revenue management system first determines the availability of each of an airline's filed fare products in a market. Then, PFDynA decides whether to make a price adjustment (either an increment or a discount) to the price of the lowest-available fare product for each itinerary. The decision of whether or not to make a price adjustment is determined in part by the bid prices from the revenue management system. For instance, if the flight is forecast to be very full, the bid price will be high, and the airline will be less likely to give a discount.

The price adjustment is also a function of the airline's estimate or guess of conditional WTP for passengers in that market. We assumed that airlines could segment incoming requests into categories, such as leisure and business, with some degree of accuracy. Airlines already practice some segmentation today based on characteristics of the booking requests; new technologies like NDC could also allow for segmentation on a customer level based on frequent flyer history.

In our model, the airline uses a parameter called a Q-multiplier to represent its estimate of conditional WTP for passengers of each type in each market. We assumed that conditional WTP followed a Gaussian distribution, which is a common assumption in airline revenue management literature. The mean of the conditional WTP distribution is computed by multiplying the airline's input Q-multiplier by the lowest filed fare in the airline's fare structure. In this way, the probability that a customer will purchase the product can be easily computed using the cumulative density function of the Gaussian distribution.

To operationalize the model, we defined several variants of PFDynA that could be used depending on the airline's commercial strategy and objectives. For instance, the airline could choose to provide only discounts in an effort to attract additional bookings from lower-WTP leisure-type customers. Alternatively, the airline could increment prices for certain business-type requests to try to increase yield. We referred to a model in which airlines both incremented and discounted prices in the ways described above as "Two-Way PFDynA." We also considered variants in which segmentation accuracy was not perfect, or when booking requests could not be segmented at all.

We tested the PFDynA dynamic price adjustment heuristic in PODS—a complex, agent-based simulation that models the interactions between passengers and airlines. Passengers in PODS make choices in a similar manner as in the choice model described earlier, and airlines use traditional RM methods to determine which of a set of pre-defined fare products to make available for purchase. PODS allows us to test how dynamic price adjustment mechanisms perform relative to traditional airline RM, and how these mechanisms result in changes to fare class mixes, yields, load factors, and competitive outcomes.

We tested variants of PFDynA in a variety of monopolistic and competitive networks. Our tests assumed that airlines can segment between different types of booking requests—business and leisure—with a certain degree of accuracy. In the simulation studies, airlines use Increments-Only PFDynA to compute increments for requests identified as the higher-WTP business type, and/or Discounts-Only PFDynA to compute discounts for requests identified as the lower-WTP leisure type. Most booking requests receive no price adjustments, meaning passengers book at normal filed fares.

Our tests of PFDynA in PODS showed that dynamic price adjustment mechanisms can lead to revenue gains over traditional airline RM methods. The results of our simulations generally agree with the 1 - 6 % revenue gains that were shown in past work on dynamic pricing in the airline industry. However, the detailed nature of PODS allowed us to draw some additional insights that had previously been absent in the literature.

Dynamic price adjustment leads to revenue increases for four main reasons. First, providing discounts to low-WTP requests in certain situations (namely, when fares are relatively high) leads to stimulation of demand from passengers who ordinarily would not have been able to travel. These additional bookings are made mostly in higher fare classes.

This led to attractive side effects in the airline's revenue management system. The increase in higher-class bookings increases the airline's demand forecasts for those higher classes, leading its RM optimizer to protect fewer seats for the least-expensive fare classes for later flight departures. This mechanism, which we call *forecast spiral-up*, can allow an airline practicing discounting strategies to see increases in both load factor and yield.

Incrementing strategies also increased revenue for airlines. Incrementing high-WTP requests (particularly when fares are relatively low) leads to an increase in yield, as well as an increase in higher-class bookings in restricted fare structures. This also led to forecast spiral-up, with fewer seats available for the least-expensive fare classes.

In competitive environments, price adjustments could cause passengers to change their booking decisions between airlines. For instance, if only a single airline in a competitive environment is discounting, passengers who would have ordinarily booked with other airlines may choose to book with the discounting carrier. This further increases that carrier's load factors, and leads to share shift from other airlines. Conversely, if only one carrier is incrementing, that carrier lost high-class bookings to other airlines that were not incrementing prices for high-WTP requests. When all of the airlines in a competitive scenario practiced PFDynA, the mechanisms still produced revenue gains due to demand stimulation, forecast spiral-up, and/or increases in yield.

Sensitivity analyses of PFDynA found that the mechanisms produced revenue gains when used with a variety of common RM optimization methods, as well as with advanced forecasting methodologies. Revenue gains are also possible when the airlines used less-restricted fare structures. We also found that some degree of segmentation was critical to the performance of PFDynA. However, this segmentation does not have to be perfect—depending on the method and competitive environment, revenue gains from PFDynA were still possible even if the airline made segmentation mistakes 20% of the time or more. Finally, we found that removing bounds on the price adjustments increased airline revenues, but made the heuristics more sensitive to inaccuracies in demand segmentation or WTP estimation.

We then extended the PFDynA price adjustment method to price multiple substitutable itineraries simultaneously. We used the economic concept of a Hotelling line to represent customers' trade-offs between different itinerary attributes. These included horizontally-differentiated attributes like departure time, for which preferences are directionally heterogeneous, and vertically-differentiated attributes like schedule quality, for which preferences are directionally homogeneous.

We tested the performance of this "Simultaneous" version of PFDynA in PODS against versions of the heuristics that priced each flight independently. We found that the simultaneous dynamic pricing model did increase the revenue performance of PFDynA, mostly by targeting discounts to encourage price-sensitive passengers to switch to unattractive itineraries while incrementing prices for schedule-sensitive travelers on attractive itineraries. However, these revenue gains were on the order of less than 1% over flight-by-flight PFDynA in a simple network. The relatively small marginal benefits of simultaneous pricing suggest that practitioners may wish to remain with the simpler flight-by-flight PFDynA methods, which are more computationally tractable and require fewer data inputs.

Finally, we considered some real-world implications of the introduction of dynamic pricing mechanisms in the airline industry. This discussion was motivated by airline concerns that dynamic pricing could lead to customer abuse, legal challenges, and a revenue-negative competitive outcome called the race to the bottom. This last implication is particularly troubling for the industry, since a race to the bottom would mean airlines would undercut each other to price closer to their marginal costs. Since airlines are businesses with high fixed costs and low marginal costs, a race to the bottom would likely mean significant harm to industry profitability.

While the economic theory of Bertrand competition suggests that the race to the bottom can occur in markets with as few as two firms selling homogeneous products, we rarely see evidence of these outcomes in real-world markets. This so-called Bertrand paradox can be explained by differences between real-world markets and the assumptions made in Bertrand competition. Particularly, the presence of capacity constraints, unknown competitor costs, differentiated products, and repeated multi-market interactions in the airline industry help prevent the risk of a race to the bottom. The high current level of industry profitability suggests that airlines are avoiding marginal cost pricing in practice.

Dynamic pricing would introduce many more price points into air transportation markets, and could allow airlines to better segment their demand. However, the facet of dynamic pricing that poses the greatest risk to the status quo of airline competition is the reduced ability of airlines to monitor the pricing actions of their competitors. Unlike traditional pricing and RM, in which filed fare structures can easily be observed by all airlines, dynamic price adjustments may be targeted to specific customers, making it more difficult for airlines to know when their competitors are discounting. This could lead to the possibility of price wars, as airlines react to real or perceived pricing actions by their competitors.

These risks could be mitigated through the creation of a market for information for dynamic price adjustment decisions. By giving airlines additional information about the pricing actions of their competitors, airlines could maintain pricing discipline without costly price wars. Academic research has suggested that even if monitoring of pricing actions is delayed, firms could still maintain cooperative equilibria in competitive markets that otherwise would race to the bottom (and, a delay in monitoring allows for these outcomes without the short-term costs of enforcement).

Since technologies already exist to monitor changes in airline prices in the current environment, it is quite likely that new monitoring technologies would be created for a world with dynamic pricing. This would help prevent the race to the bottom as a result of dynamic pricing in the long run, although there may still be short-term instability as airlines learn new ways to monitor their competitors' pricing behavior.

We also briefly reviewed the possible legal implications of dynamic pricing in various industries. In 2014, the U.S. Department of Transportation reviewed the possibility of customized pricing in the airline industry in their review of the resolution that authorized NDC, and issued a relatively permissive statement that the agency was unwilling to prohibit innovations to "better match capacity with demand" (Department of Transportation (2014), p. 14). Nevertheless, it is likely that dynamic pricing mechanisms could trigger legal reviews in the U.S. and in other countries, particularly if airlines generate prices using data that is not willingly supplied by customers.

In other industries, firms found to be practicing transactional dynamic pricing have been subject to negative media coverage and customer backlash. Yet reactions to dynamic pricing are highly dependent on to the ways in which these practices are framed. For instance, customers often react positively to targeted discount or coupon programs that provide lower prices to certain customers for certain products. Alternatively, higher prices for high-WTP requests could be bundled in combination with bundles of ancillary services.

Ultimately, each airline will need to perform its own evaluation of whether dynamic pricing mechanisms are worth the possible short-term commercial risks. These mechanisms would require the development of new core competencies among RM and pricing professionals, and would lead to changes in the current status quo. Since the implementation of dynamic pricing mechanisms is likely to be incremental across an airline's customer base or network, the revenue gains in the short-term may be lower relative to the short-term start-up costs. However, the results of this dissertation suggest that if practiced carefully, integrating dynamic pricing mechanisms into traditional airline RM systems has the potential to increase airline revenues, even in competitive environments.

7.2 Future research directions: Dynamic pricing in a new industry landscape

This dissertation's development of the PFDynA dynamic price adjustment method serves as a first step towards the integration of dynamic pricing into traditional pricing and airline revenue management practices. There are many ways in which researchers and industry stakeholders should consider extending this work to further explore the mechanics and implications of dynamic price adjustment.

- This dissertation introduced the concept of *conditional WTP*, recognizing that a customer's willingness-to-pay for an item changes depending on the other alternatives in his choice set. In PFDynA, conditional WTP was parameterized by an input Q-multiplier, which allowed us to create distributions of WTP for a variety of markets with a single parameter. We typically used very simple input Q-multiplier schemes that were either constant over time or increased linearly.

  Future work might consider testing PFDynA using input Q-multipliers that change more dynamically in response to different market conditions. For instance, lower input Q-multipliers could be used in highly-competitive markets in which there are many itinerary options, since passengers would be more likely to have lower conditional WTP in these markets. Conversely, in monopolistic markets, higher Q-multipliers could be used. Dynamically adjusting input Q-multipliers based on the alternatives available in the customer's choice set could potentially improve the performance of PFDynA, since it would more closely capture the underlying dynamics of conditional WTP.

- In Chapter 5, we tested a more complex customer choice model with Simultaneous PFDynA, which priced multiple differentiated itineraries simultaneously. While Simultaneous PFDynA did not result in much higher revenue performance than the more simple flight-by-flight approach tested in Chapter 4, it is possible that different customer choice models could produce higher revenue gains. It could be interesting to test Simultaneous PFDynA with other assumptions about customer choice, including models that may make different assumptions about how customers decide amongst alternatives.

- Since dynamic price adjustments are made in reference to an underlying fare structure, the design of the fare structure itself is critical to PFDynA performance. While we tested PFDynA with various restricted and less-restricted structures, it would be interesting to extend the work by also adjusting the prices of the fare structure itself. For instance, does PFDynA provide higher or lower gains when the initial fare structure has prices that are far apart versus close together? And, if there were more price points included in the underlying fare structure, what would be the effect on PFDynA performance?

- Machine learning has been an increasingly popular topic in recent years, and dynamic pricing could be a natural place to apply these techniques. Future work could focus on dynamically learning the Q-multipliers for a population of customers by testing different combinations of parameters. Machine learning could also be applied to the segmentation aspect of dynamic pricing. Methods like k-means clustering could be used to identify the number of customer types in a given population, and could be combined with other approaches to improve an airline's segmentation accuracy. As we have shown, better segmentation accuracy will likely lead to higher PFDynA revenues.

- Finally, dynamic price adjustment could serve as an intermediate step towards other more complex dynamic pricing mechanisms. One example would be *continuous pricing*, in which prices are freely selected from a continuous range instead of adjusted from a pre-priced itinerary product. Challenges of this approach include computing accurate bid prices for use in dynamic pricing, as well as the substantial challenge of distributing continuous prices without references to a fare class in current distribution architecture. Continuous pricing would also require enormous changes to existing pricing and RM systems and practices. The development of NDC would help airlines move closer to a "world without booking classes" (Westermann, 2013), but its full implementation is likely many years away.

- An even more complex (but potentially more lucrative) mechanism is *dynamic offer generation*. Dynamic offer generation combines the price selection and product creation processes into a single mechanism. Airlines could create bundles of itineraries and add-on services—such as extra legroom seats, checked bags, and lounge access—and then dynamically price these bundles. At the limit, dynamic offer generation could occur on a transaction-by-transaction basis, with each shopping request receiving a unique bundle with a dynamically calculated price.

  NDC would likely be necessary for a full implementation of dynamic offer generation. However, short-term solutions may be possible by combining some of the lessons learned from dynamic price adjustment with the economic literature regarding bundling. Some early-stage research has started to investigate the dynamic offer generation problem in simple scenarios (such as a single flight with a single ancillary) with promising preliminary results (Wittman and Bockelie, 2017). Much additional work is necessary to determine the science behind dynamic offer generation and how it could be integrated into legacy airline processes.

These new lines of scientific inquiry will become increasingly important as the airline industry starts to move away from the status quo of traditional pricing and revenue management towards next-generation dynamic pricing mechanisms. As of early 2018, there are signals that this shift has already begun. Some airlines are beginning to participate in pilot programs for new Dynamic Pricing Engine standards for dynamic price adjustment, and even more carriers have started to incorporate elements of the New Distribution Capability into their pricing, revenue management, and distribution systems. Given these recent trends, it would not be surprising to see at least one major carrier begin practicing a dynamic price adjustment method similar to PFDynA prior to the year 2020.

While there may be some short-term turbulence as airlines figure out how to adapt to and compete in a world with next-generation pricing, it is becoming increasingly apparent that this world will not be avoided entirely. Traditional airline commercial strategies may not be compatible with the fast-paced environment of dynamic pricing. Instead, new mindsets for airline competitive interaction will need to evolve among senior managers, pricing teams, and revenue management professionals in order to succeed in an environment where some carriers are practicing new and potentially unfamiliar pricing strategies. Given the first-mover advantage and potential revenue benefits of dynamic pricing demonstrated in this dissertation, airlines that stay ahead of the curve in terms of technological development will be best positioned to thrive in this new industry landscape.

References

Abeliuk, A., G. Berbeglia, M. Cebrian, and P. Van Hentenryck. 2016. Assortment optimization under a multinomial logit model with position bias and social influence. Journal of Operations Research 14: 57-75.

Akçay, Y., H.P. Natarajan, and S.H. Xu. 2010. Joint Dynamic Pricing of Multiple Perishable Products Under Consumer Choice. Management Science 56(8): 1345-1361.

Allingham, M. 1999. Rational Choice. St. Martin's Press, New York, NY.

Alptekinoğlu, A. and J.H. Semple. 2016. The exponomial choice model: A new alternative for assortment and price optimization. Operations Research 64(1): 79-93.

Anderson, S.P., A. de Palma, and J.F. Thisse. 1992. Discrete Choice Theory of Product Differentiation. MIT Press, Cambridge, MA.

Ariely, D. 2010. Predictably Irrational: The Hidden Forces that Shape Our Decisions. Harper Perennial, New York.

Awaya, Y. and V. Krishna. 2016. On Communication and Collusion. The American Economic Review 106(2): 285-315.

Aydın, G. and J.K. Ryan. 2000. Product line selection and pricing under the multinomial logit model. Proceedings of the 2000 MSOM Conference, Ann Arbor, MI. 1-49.

Aydın, G. and S. Ziya. 2009. Technical Note – Personalized Dynamic Pricing of Limited Inventories. Operations Research 57(6): 1523-1531.

Bala, P.K. 2014. Synergizing RM science and airline distribution technologies: A pricing example. Journal of Revenue and Pricing Management 13: 402-410.

Balcombe, K., I. Fraser, and L. Harris. 2009. Consumer willingness to pay for in-flight service and comfort levels: A choice experiment. Journal of Air Transport Management 15(5): 221-226.

Becker, G. 1976. The Economic Approach to Human Behavior. Chicago University Press, Chicago.

Belobaba, P. 1992. Optimal vs. heuristic methods for nested seat allocation. Proceedings of AGIFORS Reservations and Yield Management Study Group. Brussels, Belgium.

Belobaba, P. 2016. Optimization models in RM systems: Optimality versus revenue gains. Journal of Revenue and Pricing Management 15(3): 229-235.

Belobaba, P.P. and C. Hopperstad. 2004. Algorithms for revenue management in unrestricted fare markets. Meeting of the INFORMS Section on Revenue Management. Massachusetts Institute of Technology, Cambridge, MA.

Belobaba, P., A. Odoni, and C. Barnhart. 2016. The Global Airline Industry. John Wiley and Sons, London.

Ben-Akiva, M., and S.R. Lerman. 1985. Discrete Choice Analysis: Theory and Application to Travel Demand. MIT Press, Cambridge, MA.

Bernstein, F., A.G. Kök, and L. Xie. 2015. Dynamic Assortment Customization with Limited Inventories. Manufacturing and Service Operations Management 17(4): 538-553.

Bertsimas, D. and S. de Boer. 2005. Simulation-based Booking Limits for Airline Revenue Management. Operations Research 53(1): 90-106.

Besbes, O. and D. Sauré. 2016. Product Assortment and Price Competition under Multinomial Logit Demand. Production and Operations Management 25(1): 114-127.

Bitran, G.R. and S.V. Mondschein. 1997. Periodic Pricing of Seasonal Products in Retailing. Management Science 43(1): 64-79.

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

Brander, J.A. and A. Zhang. 1993. Dynamic oligopoly behavior in the airline industry. International Journal of Industrial Organization 11: 407-435.

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

Brunger, W.G. 2010. The impact of the Internet on airline fares: the 'Internet Price Effect'. Journal of Revenue and Pricing Management 9(1): 66-93.

Busse, M. 2002. Firm financial condition and airline price wars. RAND Journal of Economics 33(2): 298-318.

Carrier, E. 2003. Modeling airline passenger choice: Passenger preference for schedule in the passenger origin-destination simulator (PODS). Unpublished Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Chen, M. and Z.-L. Chen. 2015. Recent Developments in Dynamic Pricing Research: Multiple Products, Competition, and Limited Demand Information. Production and Operations Management 24(5): 704-731.

Chen, Q., S. Jasin, and I. Duenyas. 2015a. Real-Time Dynamic Pricing with Minimal and Flexible Price Adjustment. Management Science 62(8): 2437-2455.

Chen, S., G. Gallego, M.Z.F. Li, and B. Lin. 2010. Optimal seat allocation for two-flight problems with a flexible demand segment. European Journal of Operations Research 201: 897-908.

Chen, X., Z. Owen, C. Pixton, and D. Simchi-Levi. 2015b. A Statistical Learning Approach to Personalization in Revenue Management. Unpublished Working Paper.

Chen, Y., C. Narasimhan, and Z.J. Zhang. 2001. Individual Marketing with Imperfect Targetability. Marketing Science 20(1): 23-41.

Chen, Y. and G. Iyer. 2002. Consumer Addressability and Customized Pricing. Management Science 21(2): 197-208.

Choudhary, V., A. Ghose, T. Mukhopadhyay, and U. Rajan. 2005. Personalized Pricing and Quality Differentiation. Management Science 51(7): 1120-1130.

Ciliberto, F. and J.W. Williams. 2014. Does multimarket contact facilitate tacit collusion? Inference on conduct parameters in the airline industry. RAND Journal of Economics 45(4): 764-791.

Cléaz-Savoyen, R. 2005. Airline revenue management methods for less-restricted fare structures. Unpublished Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Collin, L. 2012. 2 Same Flights, 2 Different Prices: Frequent Flyer Discrepancies. WCCO CBS Minnesota. 15 May 2012.

Collinson, P. 2010. Beware the cookies: They can cost you money. The Guardian. 6 August 2010.

Compte, O. 1998. Communication in Repeated Games with Imperfect Private Monitoring. Econometrica 66(3): 597-626.

Cooper, W.L., T. Homem-de-Mello, and A.J. Kleywegt. 2006. Models of the Spiral-Down Effect in Revenue Management. Operations Research 54(5): 968-987.

Deneckere, R.J. and R.P. McAfee. 1996. Damaged Goods. Journal of Economics and Management Strategy 5(2): 149-174.

Department of Transportation. 2014. Agreement among Member Carriers of the International Air Transport Association concerning an agreement (Resolution 787) of the Passenger Services Conference. Order to Show Cause. Docket OST-2013-0048. 21 May 2014.

Dong, L., P. Kouvelis, and Z. Tian. 2009. Dynamic Pricing and Inventory Control of Substitute Products. Manufacturing and Service Operations Management 11(2): 317-339.

Dufwenberg, M. and U. Gneezy. 2000. Price competition and market concentration: an experimental study. International Journal of Industrial Organization 18(1): 7-22.

Easley, D. and J. Kleinberg. 2010. Networks, Crowds, and Markets: Reasoning about a Highly Connected World. Cambridge University Press, Cambridge, UK.

Ellison, G. 1993. Theories of cartel stability and the Joint Executive Committee. RAND Journal of Economics 25(1): 37-57.

Elmaghraby, W. and P. Keskinocak. 2003. Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices, and Future Directions. Management Science 49(10): 1287-1309.

Elster, J. 1986. Rational Choice. New York University Press, New York.

Ely, J.C. and J. Välimäki. 2002. A robust folk theorem for the prisoner's dilemma. Journal of Economic Theory 102: 84-105.

Espino, R., J.C. Martin, and C. Roman. 2008. Analyzing the effect of preference heterogeneity on willingness to pay for improving service quality in an airline choice context. Transportation Research Part E 44: 593-606.

Esteves, R.-B. 2009. Price discrimination with partial information: Does it pay off? Economics Letters 105: 28-31.

Esteves, R.-B. 2014. Price Discrimination with Private and Imperfect Information. Scandinavian Journal of Economics 116(3): 766-796.

Federgruen, A., and M. Hu. 2015. Multi-Product Price and Assortment Competition. Operations Research 63(3): 572-584.

Fiig, T., U. Cholak, M. Gauchet, and B. Cany. 2015. What is the role of distribution in revenue management? – Past and future. Journal of Revenue and Pricing Management 14(2): 127-133.

Fiig, T., O. Goyons, R. Adelving, and B. Smith. 2016. Dynamic pricing — The next revolution in RM? Journal of Revenue and Pricing Management 15(5): 360-379.
Fiig, T., K. Isler, C. Hopperstad, and P. Belobaba. 2010. Optimization of mixed fare structures: Theory and applications. Journal of Revenue and Pricing Management 9(1): 152-170.
Friedman, J.W. 1971. A Non-cooperative Equilibrium for Supergames. The Review of Economic Studies 38(1): 1-12.
Fudenberg, D. and E. Maskin. 1986. The Folk Theorem in Repeated Games with Discounting or Incomplete Information. Econometrica 54(3): 533-554.
Fudenberg, D. and J. M. Villas-Boas. 2006. Behavior-Based Price Discrimination and Customer Recognition. In: Handbooks in Information Systems, Volume 1. T. Hendershott (ed.) Elsevier B.V. 377-436.
Gallego, G. and M. Hu. 2014. Dynamic Pricing of Perishable Assets Under Competition. Management Science 60(5): 1241-1259.
Gallego, G., G. Iyengar, R. Phillips, and A. Dubey. 2004. Managing flexible products on a network. Working paper.
Gallego, G., A. Li, V.-A. Truong, and X. Wang. 2016. Online Personalized Resource Allocation with Customer Choice. Working Paper. http://arxiv.org/abs/1511.01837v1
Gallego, G. and G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. Management Science 40(8): 999-1020.
Gallego, G. and G. van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. Operations Research 45(1): 24-41.
Gallego, G. and R. Wang. 2014. Multiproduct Price Optimization and Competition Under the Nested Logit Model with Product-Differentiated Price Sensitivities. Operations Research 62(2): 450-461.
Gaur, V. and D. Honhon. 2006. Assortment Planning and Inventory Decisions Under a Locational Choice Model. Management Science 52(10): 1528-1543.
Gillen, D. W., W. G. Morrison, and C. Stewart. 2003. Air travel demand elasticities: Concepts, issues, and ameasurement. Department of Finance Report, Government of Canada.
Goetz, A.R. and T.M. Vowles. 2009. The good, the bad, and the ugly: 30 years of US airline deregulation. Journal of Transport Geography 17(4): 251-263.
Golrezaei, N., H. Nazerzadeh, and P. Rusmevichientong. 2014. Real-Time Optimization of Personalized Assortments. Management Science 60(6): 1532-1551.
Gorin, T.O. 2000. Airline Revenue Management: Sell-Up and Forecasting Algorithms. Unpublished Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.
Grauberger, W. and A. Kimms. 2016. Airline revenue management games with simultaneous price and quantity competition. Computers and Operations Research 75: 64-75.
Green, E.J. and R.H. Porter. 1984. Noncooperative Collusion under Imperfect Price Information. Econometrica 52(1): 87-100.
Gunther, D., R. Ratliff, and A. Sylla. 2010. Airline Distribution. In: Quantitative Problem Solving Methods in the Airline Industry. C. Barnhart and B.C. Smith (Eds.) Springer Science, New York.
Hawkins, A.J. 2016. Uber is trying to make you forget that surge pricing exists. The Verge. 23 Jul 2016.
Hotelling, H. 1929. Stability in Competition. Economics Journal 39: 41-57.
Hoyles, Y. 2015. New Distribution Capability (NDC): Strategy Paper (Version 1.2). IATA. http://www.iata.org/whatwedo/airline-distribution/ndc/Documents/ndc-strategy-paper.pdf.
International Air Transport Association. 2016. NDC Change Readiness Guide for Airlines. Edition 2. `http://www.iata.org/whatwedo/airline-distribution/ndc/Documents/ndc-change-readiness-guide-airlinesV2.pdf`.
Isler, K. and E. D'Souza. 2009. GDS capabilities, OD control and dynamic pricing. Journal of Revenue and Pricing Management 8: 255-266.
Isler, K. and H. Imhof. 2008. A game theoretic model for airline revenue management and competitive pricing. Journal of Revenue and Pricing Management 7(4): 384-396.
Jagabathula, S. and P. Rusmevichientong. 2016. A Nonparametric Joint Assortment and Price Choice Model. Management Science 63(9): 3128-3145.
Jansen, B. 2017. Airlines gauge success of Basic Economy by how many passengers avoid it. USA Today. 26 October 2017.

Kambour, E. 2014. A next generation approach to estimating customer demand via willingness to pay. Journal of Revenue and Pricing Management 13(5): 354-365.
Kandori, M. 2002. Introduction to Repeated Games with Private Monitoring. Journal of Economic Theory 102: 1-15.
Kandori, M. and H. Matsushima. 1998. Private observation, communication and collusion. Econometrica 66(3): 627-652.
Kincaid, W.M. and D.A. Darling. 1963. An Inventory Pricing Problem. Journal of Mathematical Analysis and Applications 7: 183-208.
Kothari, A., M. Manireddy, and R. Sundarajan. 2016. Discovering patterns in traveler behavior using segmentation. Journal of Revenue and Pricing Management 15(5): 334-51.
Kreps, D.M. and J.A. Scheinkman. 1983. Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes. The Bell Journal of Economics 14(2): 326-337.
Kumar, R., A. Li, and W. Wang. 2017. Learning and Optimizing through Dynamic Pricing. Journal of Revenue and Pricing Management. Available online 2 October 2017.
Kumar, R. and D. Walczak. 2016. Dynamic Pricing and Classless Revenue Management. Presentation at INFORMS Revenue Management and Pricing Section Conference, New York, NY. 17 June 2016.
Kumar, V., B.M. Horio, A.H. DeCicco, S. Hasan, and V.L. Stouffer. 2015. Understanding Air Transportation Market Dynamics Using a Search Algorithm for Calibrating Travel Demand and Price. AIAA Modeling and Simulation Technologies Conference, Dallas, TX. June 2015.
Lacroix, D. 1996. Web Services Ubiquity in the Travel & Tourism Industry. Presentation at W3C Seminar, Paris, France.
Lautenbacher, C.J. and S. Stidham. 1999. The Underlying Markov Decision Process in the Single-Leg Airline Yield- Management Problem. Transportation Science 33(2): 136-146.
Levin, J. and P. Milgrom. 2004. Introduction to Choice Theory. Lecture Notes. Stanford University.
Levina, T., Y. Levin, J. McGill, and M. Nediak. 2009. Dynamic Pricing with Online Learning and Strategic Consumers: An Application of the Aggregating Algorithm. Operations Research 57(2): 327-341.
Lewis, M. 2016. The Undoing Project: A Friendship That Changed Our Minds. W.W. Norton & Company, New York.
Li, G., P. Rusmevichientong, and H. Topaloglu. 2015. The d-Level Nested Logit Model: Assortment and Price Optimization Problems. Operations Research 63(2): 325-341.
Lin, K.Y. and S.Y. Sibdari. 2009. Dynamic price competition with discrete customer choices. European Journal of Operational Research 197: 969-980.
Liu, Q. and G. van Ryzin. 2008. On the Choice-Based Linear Programming Model for Network Revenue Management. Manufacturing & Service Operations Management 10(2): 288-310.
Lua, W.F. 2006. Matching of Lowest Fare Seat Availability in Airline Revenue Management Systems. Unpublished Master's Thesis. Massachusetts Institute of Technology.
Mahajan, S. and G. van Ryzin. 2001. Stocking Retail Assortments Under Dynamic Consumer Substitution. Operations Research 49(3): 334-351.
March, J.G. 1978. Bounded Rationality, Ambiguity, and the Engineering of Choice. Bell Journal of Economics 9: 587-608.
Martínez-de-Albéniz, V. and K. Talluri. 2011. Dynamic Price Competition with Fixed Capacities. Management Science 57(6): 1078-1093.
Matsushima, H. 2004. Repeated Games with Private Monitoring: Two Players. Econometrica 72(3): 823-852.
Mattioli, D. 2012. On Orbitz, Mac Users Steered to Pricier Hotels. The Wall Street Journal. 23 Aug. 2012.
McAfee, R.P. and V. te Velde. 2007. Dynamic Pricing in the Airline Industry. In Handbook on Economics and Information Systems. Ed. T.J. Hendershott. Elsevier.
McGee, B. 2013. Do travel deals change based on your browsing history? USA Today, 3 April 2013.
McGill, J.I. and G.J. van Ryzin. 1999. Revenue Management: Research Overview and Prospects. Transportation Science 33(2): 233-256.
Morrison, S.A. and C. Winston. 1996. Brookings Papers on Economic Activity: Microeconomics. 85-131.
Mussa, M. and S. Rosen. 1978. Monopoly and Product Quality. Journal of Economic Theory 18: 301-317.
Netessine, S. and R.A. Shumsky. 2005. Revenue management games: horizontal and vertical competition. Management Science 51(5): 813-831.
Newcomer, E. 2017. Uber Starts Charging What It Thinks You're Willing to Pay. Bloomberg. 19 May 2017.

Nicolae, M.L., M.E. Ferguson, and L.A. Garrow. 2016. Measuring the Benefit of Offering Auxiliary Services: Do Bag-Checkers Differ in their Sensitivities to Airline Itinerary Attributes? Production and Operations Management 25(10): 1689-1708.

Niketic, N. and G.R.J. Mules. 1993. How Qantas Airways Uses EDI for Travel Reservation Processing. Journal of Systems Management 44: 9-13.

Njegovan, N. 2006. Elasticities of demand for leisure air travel: A system modelling approach. Journal of Air Transport Management 12(1): 33-39.

Obara, I. 2009. Folk Theorem with Communication. Journal of Economic Theory 144(1): 120-134.

Pepall, L., D. Richards, and G. Norman. 2011. Contemporary Industrial Organization: A Quantitative Approach. John Wiley & Sons, Hoboken, N.J.

Peterson, S. 2010. Airlines 2020: Substitution and commoditization. IBM Institute for Business Value.

Ratliff, R. 2017. Industry-standard Specifications for Air Dynamic Pricing Engines: Progress Update. Proceedings of the 2017 AGIFORS Revenue Management Study Group Meeting, San Francisco, CA

Ratliff, R., B.V. Rao, C.P. Narayan, and K. Yellepeddi. A multi-flight recapture heuristic for estimating unconstrained demand from airline bookings. Journal of Revenue and Pricing Management 7(2): 153-171.

Rice, K. 2015. Senators express concerns over airlines' use of customer data. Travel Weekly. 4 May 2015.

Rich, G. 2015. Airline Stocks Plunge as Southwest Boosts Capacity. Investor's Business Daily. 20 May 2015.

Robenalt, J.L. 2007. Predatory Pricing in the Low-Fare Airline Market: Targeted, Discriminatory, and Achieved with Impunity. Ohio State Law Journal 68: 641-677.

Rotemberg, J.J. and G. Saloner. 1986. A Supergame-Theoretic Model of Price Wars during Booms. The American Economic Review 76(3): 390-407.

Rusmevichientong, P., Z.-J. M. Shen, and D.B. Shmoys. 2010. Dynamic Assortment Optimization with a Multinomial Logit Model and Capacity Constraint. Operations Research 58(6): 1666-1680.

Sagers, C. 2009. "Rarely Tried, and...Rarely Successful": Theoretically Impossible Price Predation Among the Airlines. Journal of Air Law and Commerce 74: 919-958.

Samuelson, P. 1938. A Note on the Pure Theory of Consumers' Behaviour. Economica 5(17): 61-71.

Shaffer, G. and Z.J. Zhang. 2002. Competitive One-to-One Promotions. Management Science 48(9): 1143-1160.

Shy, O. and R. Stenbacka. 2015. Customer Recognition and Competition. Journal of Economics Management and Strategy 25(3): 539-562.

Smith, B.C. and C.W. Penn. 1988. Analysis of alternative origin-destination control strategies. AGIFORS Symposium Proceedings 28: New Seabury, MA.

Spulber, D.F. 1995. Bertrand Competition when Rivals' Costs are Unknown. The Journal of Industrial Economics 43(1): 1-11.

Su, X. 2007. Intertemporary Pricing with Strategic Customer Behavior. Management Science 53(5): 726-741.

Sugaya, T. 2011. Folk Theorem in Repeated Games with Private Monitoring. Economic Theory Center Working Paper No. 011-2011.

Suh, M. and G. Aydın. 2011. Dynamic pricing of substitutable products with limited inventories under logit demand. IIE Transactions 43: 323-331.

Sumers, B. 2017. United and American Just Started Selling Super Restrictive Discount Fares. Skift, 21 February 2017.

Talluri, K.T. and G.J. van Ryzin. 1998. An Analysis of Bid-Price Controls for Network Revenue Management. Management Science 44(11): 1577-1593.

Talluri, K.T. and G.J. van Ryzin. 2004. Revenue Management Under a General Discrete Choice Model of Consumer Behavior. Management Science 50(1): 15-33.

Talluri, K.T. and G.J. van Ryzin. 2005. The Theory and Practice of Revenue Management. Springer, New York, NY.

Taubmann, H. 2016. One year since Lufthansa's Distribution Cost Charge: what we've learned. Amadeus Blog. Amadeus IT Group, 29 June 2016.

Teichert, T., E. Shehu, and I. von Wartburg. 2008. Customer segmentation revisited: The case of the airline industry. Transportation Research Part A: Policy and Practice 42(1): 227-242.

Thisse, J.-F. and X. Vives. 1988. On the Strategic Choice of Spatial Price Policy. The American Economic Review 78(1): 122-137.

Tirole, J. 1988. The Theory of Industrial Organization. MIT Press, Cambridge, MA.

Tversky, A. and D. Kahneman. 1981. The Framing of Decisions and the Psychology of Choice. Science 211: 453-458.

Ulu, C., D. Honhon, and A. Alptekinoğlu. 2012. Learning Customer Tastes Through Dynamic Assortments. Operations Research 60(4): 833-849.

U.S. Department of Transportation. 2016. The Value of Travel Time Savings: Departmental Guidance for Conducting Economic Evaluations Revision 2 (2016 Update). 27 September 2016.

Valentino-DeVries, J., J. Singer-Vine, and A. Soltani. 2012. Websites Vary Prices, Deals Based on Users' Information. The Wall Street Journal. 24 Dec. 2012.

Villas-Boas, J.-M. 1999. Dynamic competition with customer recognition. RAND Journal of Economics 30(4): 604-631.

Vinod, B. 2010. The complexities and challenges of the airline fare management process and alignment with revenue management. Journal of Revenue and Pricing Management 9(1-2): 137-151.

Vinod, B. 2015. The expanding role of revenue management in the airline industry. Journal of Revenue and Pricing Management 14(6): 391-399.

Walczak, D. and S. Brumelle. 2003. Semi-Markov information model for revenue management and dynamic pricing. Operations Research Spectrum 29: 61-83.

Weiss, R.M. and A.K. Malhotra. 2001. Online Dynamic Pricing: Efficiency, Equity and the Future of E-commerce. Virginia Journal of Law and Technology 11.

Weisstein, F.L., K.B. Monroe, and M. Kukar-Kinney. 2013. Effects of price framing on consumers' perceptions of online dynamic pricing practices. Journal of the Academy of Marketing Science 41: 501-514.

Westermann, D. 2006. (Realtime) dynamic pricing in an integrated revenue management and pricing environment: An approach to handling undifferentiated fare structures in low-fare markets. Journal of Revenue and Pricing Management 4(4): 389-405.

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

Williamson, E.L. 1992. Airline Network Seat Inventory Control: Methodologies and Revenue Impacts. Unpublished Ph.D. dissertation. Massachusetts Institute of Technology.

Wittman, M.D. 2014. The effects of capacity discipline on smaller US airports: trends in service, connectivity, and fares. Massachusetts Institute of Technology. Unpublished Master's Thesis.

Wittman, M.D. and P.P. Belobaba. 2017a. Dynamic availability of fare products with knowledge of customer characteristics. Journal of Revenue and Pricing Management 16(2): 201-217.

Wittman, M.D. and P.P. Belobaba. 2017b. Personalization in airline revenue management Heuristics for real-time adjustment of availability and fares. Journal of Revenue and Pricing Management 16(4): 376-396.

Wittman, M.D. and P.P. Belobaba. 2017c. Customized Dynamic Pricing of Airline Fare Products. Journal of Revenue and Pricing Management. Available online 9 October 2017.

Wittman, M.D. and A.S. Bockelie. 2017. Dynamic Offer Generation Initial Results: One Airline, One Ancillary. MIT PODS Research Conference, Cambridge, MA. November 2017.

Yamamoto, Y. 2009. A limit characterization of belief-free equilibrium payoffs in repeated games. Journal of Economic Theory 144(2): 802-824.

Zhang, D. and W.L. Cooper. 2009. Pricing substitutable flights in airline revenue management. European Journal of Operational Research 197: 848-861.

Zhang, D. and Z. Lu. 2013. Assessing the Value of Dynamic Pricing in Network Revenue Management. INFORMS Journal on Computing 25(1): 102-115.

Zhang, Y. and D.K. Round. 2011. Price wars and price collusion in China's airline markets. International Journal of Industrial Organization 29(4): 361-372.

Zhao, W. and Y.-S. Zheng. 2000. Optimal Dynamic Pricing for Perishable Assets with Nonhomogeneous Demand. Management Science 46(3): 375-388.