Efficient Formulations for Next-generation Choice-based Network Revenue Management for Airline Implementation

by

Michael C. Clough

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved April 2016 by the
Graduate Supervisory Committee:

Esma Gel, Co-Chair
Timothy Jacobs, Co-Chair
Ronald Askin
Douglas Montgomery

ARIZONA STATE UNIVERSITY

May 2016

ABSTRACT

Revenue management is at the core of airline operations today; proprietary algorithms and heuristics are used to determine prices and availability of tickets on an almost-continuous basis. While initial developments in revenue management were motivated by industry practice, later developments overcoming fundamental omissions from earlier models show significant improvement, despite their focus on relatively esoteric aspects of the problem, and have limited potential for practical use due to computational requirements. This dissertation attempts to address various modeling and computational issues, introducing realistic choice-based demand revenue management models. In particular, this work introduces two optimization formulations alongside a choice-based demand modeling framework, improving on the methods that choice-based revenue management literature has created to date, by providing sensible models for airline implementation.

The first model offers an alternative formulation to the traditional choice-based revenue management problem presented in the literature, and provides substantial gains in expected revenue while limiting the problems computational complexity. Making assumptions on passenger demand, the Choice-based Mixed Integer Program (CMIP) provides a significantly more compact formulation when compared to other choice-based revenue management models, and consistently outperforms previous models.

Despite the prevalence of choice-based revenue management models in literature, the assumptions made on purchasing behavior inhibit researchers to create models that properly reflect passenger sensitivities to various ticket attributes, such as price, number of stops, and flexibility options. This dissertation introduces a general framework for airline choice-based demand modeling that takes into account various ticket attributes in addition to price, providing a framework for revenue management mod-

els to relate airline companies product design strategies to the practice of revenue management through decisions on ticket availability and price.

Finally, this dissertation introduces a mixed integer non-linear programming formulation for airline revenue management that accommodates the possibility of simultaneously setting prices and availabilities on a network. Traditional revenue management models primarily focus on availability, only, forcing secondary models to optimize prices. The Price-dynamic Choice-based Mixed Integer Program (PCMIP) eliminates this two-step process, aligning passenger purchase behavior with revenue management policies, and is shown to outperform previously developed models, providing a new frontier of research in airline revenue management.

DEDICATION

I dedicate this work to my parents, Deborah and Harry Sullivan. Your unwavering support for my decision to pursue my PhD and confidence in my success made all of this possible.

ACKNOWLEDGMENTS

Over the past few years I have had great support and guidance allowing me to complete my dissertation and PhD. To start, I would like to thank my co-chairs, Dr. Esma Gel and Dr. Timothy Jacobs. Dr. Gel has been a mentor and a colleague, and has provided immeasurable assistance in directing my career, personal life, and research. Dr. Jacobs provided me with professional experience and insight into my research, and many days were spent bouncing ideas off of one another leading to the models in this dissertation. I would also like to thank Dr. Ronald Askin and Dr. Douglas Montgomery for taking part in my dissertation committee, providing direction while developing the scope of my research.

I would like to personally thank Austin Bren for joining me for the many lunches where we discussed ideas. His willingness to listen to my difficulties and provide his perspective allowed me to move forward with my research and gave me a different view on what I was trying to accomplish. Additionally, this work could not have been completed without the support of my family. They provided the impetus to begin my PhD, and reassured me I was able to complete this process. I could not have completed my PhD and dissertation without their support and dedication to my research and career.

TABLE OF CONTENTS

LIST OF TABLES

Table Page

Table Page

Table Page

LIST OF FIGURES

Chapter 1

INTRODUCTION

Since the deregulation of the airline industry in 1978 airline carriers have had the opportunity to control more aspects of their business. These decisions have become more difficult as networks connections grow and technology expands with a goal of providing a fully interconnected network of flights allowing passengers to fly virtually anywhere in the world. The introduction of more advanced decision making models and supportive technology has assisted the airline industry with their expansion, generating a necessity to optimally solve the decisions related to their operations. Some of these decisions include airport selection, fleet requirements, pricing structures, collective agreements with other carriers, and allocation of space on a plane. Although carriers have been solving these problems for decades with heuristics and assumption-filled models, advances in technology allow for modeling approaches that can improve the current methods and better represent the behavior airline demand typically follows.

The two most important decisions an airline makes day-to-day are centered around the pricing structure and ticket availability for customer purchase. The ticket price and consumption of space when tickets are purchased have some of the largest impacts on revenue gains when compared to other revenue earning practices, such as up-sale opportunities and in-flight options. With the progression of network development, a need for complex pricing models arose, leading the way to yield management. Yield management, formally, was the process of estimating and anticipating customer demand, utilizing these estimates to price tickets appropriately, assuming certain fixed costs were in place. Over the years, yield management evolved into the more common term used today: *revenue management*. Revenue management is at the core of airline operations today; expensive software is utilized and proprietary algorithms alongside heuristics are implemented to solve these difficult problems. In more advanced carriers, revenue management is incorporated into many decisions including fleet and crew assignment, emphasizing the importance of having a well-rounded revenue management solution methodology.

1.1 Revenue Management

1.1.1 Revenue Management Framework

There is a fundamental framework that exists among the literature concerning revenue management in the airline industry. We assume a network exists, consisting of *origins* and *destinations*, with no variable costs of flying in or out of an origin or destination. This assumption lets us optimize expected revenue rather than profit; the costs of including an origin or destination into the network have been handled, as well as the operational costs associated with flying. Between every origin and destination is a *leg* in which a plane has already been assigned, implying the capacity of the plane is known prior to solving a problem. Under these assumptions, we have a capacitated network containing multiple legs with no cost implications. There may be multiple *paths* between each origin and destination which consume space on multiple legs, while each path typically has a value assigned to it known as a *fare* or *price*. The combination of paths, fares, and flight schedule are represented as *itineraries*, and are purchased as *tickets*. As the number of paths and fare options (often referred to as fare classes) increase, the number of itineraries grows exceptionally fast. This is a fundamental problem in revenue management as many papers direct their focus on minimizing the number of itineraries to consider to make it possible for the industry to solve this combinatorial problem.

The objective of revenue management is to determine the number of seats (allocations) to sell at a particular fare to maximize the network revenue. This can be done in many ways including protection levels, seat assignments, policy implementation, ticket availability, or pricing thresholds (often called *bid prices*). A protection level is the number of seats which you "reserve" for your higher paying passengers. For instance, if the expected demand for a higher paying passenger was D_1 and the fare for that passenger was R_1, we can continue to accept passengers at a lower fare, R_2, only until the expected value of the higher paying passengers exceeds the lower fare. Thus, when $R_2 \geq R_1 P(D_1 > x)$ is no longer true for a capacity x, it is advantageous to reject the lower fare passenger in favor of the probability a higher fare passenger shows up. Rearranging this equation we can calculate protection levels for the higher paying passengers as $y_1 = P^{-1}(R_2/R_1)$, which yields the number of seats we should reserve for only the higher paying passengers. Once protection limits are reached, lower fare classes are closed and all arrivals requesting the lower fare tickets are rejected. This series of equations leads to one of the original models developed for revenue management called Littlewood's Rule (Littlewood, 2005), and laid the groundwork for future revenue management models.

Seat assignments are the number of seats the airline is willing to reserve for each fare class and can be determined from the protection levels, or vice versa. Once seats in a particular fare class (sometimes called a bucket) are sold out, that fare class is considered closed. Ideally, once a fare class is closed it is generally not re-opened as it would net a lower revenue than waiting for a higher paying passenger, but due to customer behavior there are instances where opening previously closed fare classes could be advantageous. Seat assignments are nested in fare class, such that allocation for the higher fare classes contain the sum of the allocations for the lower fare classes plus their protected seats. As an example, if I had two fares, high and low, and 10 seats on a plane, I would allow the entire plane to be sold at the high fare class. However, I would lose money if I sold all the tickets to the lower fare class, so I would put a cap on what they could buy, let's say 5. This would equate to seat allocations of 10 for the high fare class and 5 for the low fare class. Converting this to protection limits, I would have a protection level of 5 for the high fare class since I'm only willing to sell 5 seats to the lower fare class.

A policy, on the other hand, doesn't explicitly state how many seats are to be sold in any given fare class. It states for each itinerary in the system whether it is available or not. This, in effect, will govern what fare classes are available for purchase at any point in time. A bid price, on the other hand, is a hurdle rate associated with a given capacity in the network. This price is generalized to be the marginal value of a single seat for a leg. If a passenger requests a ticket below the bid price, we stand to make more money by rejecting him and awaiting a new arrival. If the passenger requests a ticket at or above the bid price, we will sell the ticket and consume the capacity along the path selected. Effectively, bid prices determine which fare classes are available by closing those whose marginal value is below the stated bid price. Implementing a bid price is similar to that of a policy, as bid prices doesn't explicitly state which seats are available for each fare class. Bid prices are often represented as a vector of prices, one for each leg, where flights with multiple legs are priced according to the sum of their leg's bid prices.

The final component of revenue management is the demand aspect. Revenue management has become more complex over the years, but can be separated into two categories: Independent Demand and Choice-Based Demand (Dependent Demand). Independent demand refers to a system where the observed demand does not change as a function of the options available. For instance, if a passenger was looking to book a specific flight at a fare of $250, this passenger, under the independent demand assumption, would not care that an earlier flight might be available. This assumption segregates the population into buckets with their requested path, as well as other ticket attributes, and the fare associated with the path. This was the assumption used for the majority of revenue management up until recently, when choice-based modeling of demand became more prevalent. Choice-based demand refers to a system where the observed demand is dependent on the itineraries and options available. Given a set of choices, a passenger's behavior is modeled by a set of probabilities corresponding to selecting each one of these choices, assuming these choices coincide with their preferred origin and destination and other ticket attributes. Aside from the origin and destination, common ticket attributes like time of day, price, number of stops, and refundability, play an important role in modeling passenger purchasing behavior. The large number of origin-destination combinations with these ticket attributes complicate the problem of revenue management greatly, as we now have to consider what ticket attributes are included in a purchase, and how these attributes effect our demand estimates. The independent demand assumption ignores the fact that multiple choices exist and the role these attributes have on purchasing behavior, whereas choice-based demand can take all of these effects into account.

1.1.2 Independent Demand Models

Revenue management optimization has been around since the early 1970's. In one of the first models, Littlewood formulated a single product problem (one leg), in which he looked to determine the optimal number of seats to protect based on two fares. He reduced the problem to a classic news vendor problem, in which we choose to sell the lower fare class ticket only if the expected revenue of the higher fare class ticket was lower. Applying the news vendor problem results, Littlewood developed an equation for generating protection levels, which was later named Littlewood's Rule (Littlewood, 2005):

$$y_1^* = F^{-1}\Big(1 - \frac{R_2}{R_1}\Big). \tag{1.1}$$

As previously discussed regarding protection limits, Littlewood's rule calculated the protection levels, y_1^*, for the higher fare class based on the fares of the two fare classes, namely, R_1 and R_2. This is often referred to as the first revenue management model, and became the foundation for some of the more popular models. One of these more popular models was developed by Peter Belobaba (Belobaba, 1989). His model, the so-called Expected Marginal Seat Revenue (EMSR) model expanded on Littlewood's formulation to take into account more than two fare classes. In his original model, EMSR-a, Belobaba makes comparisons between fare classes that are adjacent to one another, utilizing Littlewood's Rule. The model then aggregates protection levels as you move up the fare classes, to create a set of protection levels for all available fare classes (Belobaba, 1989). In his secondary model, EMSR-b, the demand for higher fare classes is included in the calculations for lower fare classes, as oppose to aggregating protection limits, and an average fare is considered when adjacent comparisons are being made through Littlewood's rule (Belobaba, 1992). The final results of both EMSR-a and EMSR-b were a set of protection levels, or booking limits, for each fare class on a given leg. Another model that expanded on Littlewood's original paper was that of Brumelle *et al.* (1990). In their model, they took Littlewood's formulation and introduced stochastic dependence for the demand between fare classes. They then examined the full fare spillage (demand that is not met due to seat limitations) and vertical shifts (the process by which demand shifts from one fare class to another when their original fare class is not available) from the lower valued, discount fare classes (Brumelle *et al.*, 1990).

Since the EMSR methodology was simplistic and easy to solve, researchers began expanding on Belobaba's original formulation and adapting it to other slightly more complicated settings. In Williamson's PhD dissertation, a method which prorates revenue across legs in an itinerary containing at least one connection was introduced (Williamson, 1992). This is an important concept, as revenue is determined by the itinerary, not the value across each leg of that itinerary. This method allowed seat allocations and protection limits to be calculated using the EMSR methodologies while still having a more complex network of flights (Williamson, 1992). Another model that stemmed from Littlewood's rule and is very similar to the EMSR methods was developed by Wollmer (1992). Instead of calculating protection levels, Wollmer determines the critical value of seats for each fare class. This critical value is analogous to the booking limitations determined by the EMSR methods.

Around the same time Belobaba was developing EMSR, Glover *et al.* (1982) developed one of the first network formulations to solve the RM problem. They modeled the problem as a seat allocation network flow problem, and proceeded to solve it to maximize profitability. The system bounds were determined by the demand for a given leg and the capacity of the plane on that leg (Glover *et al.*, 1982). This model then led to other more complex models such as multicommodity flow problems (Dror *et al.*, 1988) and alternate network flow formulations which focused on shadow prices (Simpson, 1989). Shortly after these network formulations began developing, associated linear programs were being developed to solve the network revenue management problem. Curry (1990) developed a linear program which utilizes a piece-wise approximation of the marginal seat revenue as an objective function. The model solved for seat allocations directly and included origins and destinations which could be nested within each other (Curry, 1990). A linear programming formulation originally investigated in Smith and Penn (1988) was later analyzed and formally documented by Talluri and van Ryzin (1999), which became a popular method called the Randomized Linear Program (RLP). In the RLP, a basic seat allocation deterministic linear program was solved repetitively with randomized demand estimates from a known distribution of demand. These solutions were then averaged to generate bid prices for each leg within the system. This method of randomizing the demand generates more realistic results and more reliable bid prices (Smith and Penn, 1988; Talluri and van Ryzin, 1999).

Another large movement in revenue management came from American Airlines. Smith *et al.* (1992) describes the optimization process they created for American Airlines. This process, called Dynamic Inventory and Maintenance Optimizer (DINAMO), was a segregated approach to solving many problems in their airline. DINAMO split up the decisions to be made into three distinct sets: overbooking, discount allocations, and traffic management. Using the optimal solutions of these individual problems, DINAMO then found the optimal seat allocations for the network (Smith *et al.*, 1992). Other models took a different approach and focused on optimal pricing policies to solve the RM problem. These pricing models varied, such as models that determine optimal prices to offer products for (Gallego and van Ryzin, 1997), models that select the duration of offering a known pricing point (Feng and Xiao, 2000b,a), and models that dictate when the price of a product should change (Feng and Gallego, 2000). Jacobs *et al.* (2010) considered the relationship between pricing, revenue management controls, and the scheduled capacity to create a statistic for evaluating the quality of an airline's strategy called the "price balance statistic", as well as an algorithm to optimize the relationship between these decisions. Overall, these models were different in how they solved the problem since they didn't determine actual allocations, only pricing policies. In addition to these pricing models, stochastic formulations (Moller *et al.*, 2007; Topaloglu, 2008; Erdelyi and Topaloglu, 2008; Chen and de Mello, 2010b), relaxation methods (Kunnumkal and Topaloglu, 2010a;

Topaloglu, 2009), and simulation models (Klein, 2007; Gosavi *et al.*, 2007; van Ryzin and Vulcano, 2008b) were also developed to handle the network revenue management problem. In addition to these optimization models, earlier revenue management literature began focusing on dynamic programming (DP) formulations (Lee and Hersh, 1993; Gallego and van Ryzin, 1994; Subramanian *et al.*, 1999; Liang, 1999; Lautenbacher and S. Stidham, 1999). These DP formulations would be the building blocks of more complicated methods, leading up to some of the key revenue management models we see today.

1.1.3 Choice-based Demand Models

As network formulations became more advanced, there was a growing need to better estimate the demand over the network. Independent demand assumptions failed to take into account the true purchasing behavior of passengers, resulting in sub-optimal decisions from revenue management models. Eventually, choice-based models made their way into revenue management, expanding on the ideas from the last two decades to create more advanced models that yield better pricing and ticket availability policies. These choice-based demand models, though, come at a cost of computational complexity and can be difficult to implement into current systems. Incorporation of passenger purchasing behavior based on ticket attributes and availability created complex demand models, requiring new revenue management models for utilizing this type of demand modeling.

One of the first papers incorporating choice-based demand models into revenue management was that of Gallego *et al.* (2004). In their paper, they developed a linear program that solved the network revenue management problem with a general discrete choice model. Their model took into account the probability that purchases were made over a set of disjoint options, and then optimized how many passengers would be allowed on each of the legs within the network (Gallego *et al.*, 2004). This model would later be called the Choice-Based Deterministic Linear Program (CDLP), and has become a benchmark model for revenue management solution methodologies. In the same year, another important model was developed by Talluri and van Ryzin (2004a). In their paper, they introduced a dynamic program that takes into account a general choice-based demand model, as well as a solution methodology to combat the curse of dimensionality common in dynamic programs. Talluri and van Ryzin (2004a) were the first to develop the concept of efficient sets, which limited the number of possible solutions to be evaluated. Other formulations utilizing dynamic programs were developed as well, including Markov decision processes (Zhang and Cooper, 2005, 2006; Secomandi, 2008; Zhang and Adelman, 2009) and solution methodologies using both dynamic programming formulations and linear programming formulations in conjunction with one another (Farias and van Roy, 2007; Adelman, 2007).

With these more complex and realistic models being developed, researchers began evaluating solution methods and altering the formulations themselves to create more robust and computationally efficient models. Kunnumkal and Topaloglu (2008) made an alternate version of the CDLP providing better bounds on the problem when policy decisions were being used. Liu and van Ryzin (2008) redefined the original CDLP model and developed an iterative approach to applying the bid prices from the CDLP to solve a leg-level dynamic program. The results of their paper yielded functional policy decisions that were both time and capacity dependent. Additionally, they expanded on the notion of efficient sets, applying them to their formulation of the CDLP. Following the development of the new CDLP, Bront *et al.* (2009) developed a column generation algorithm, which efficiently solved the CDLP with non-disjoint market segments. In their paper, they considered situations where market demand can overlap, a situataion that can complicate choice probabilities. Their paper outlines the details on applying the column generation algorithm, alternate ways to solve the sub problem, as well as mathematical tractability of their formulation. Talluri (2011) creates a method for solving the CDLP as well, called the segment-based deterministic concave-program (SDCP). His method is a relaxation of the CDLP, and provides looser upper bounds to the original problem. Shortly later, the SDCP was improved on by Meissner *et al.* (2013), in which they included constraints on the product selections to create an extended-SDCP (ESDCP).

Many researchers focused on solving the dynamic programming formulations developed over the years. Kunnumkal and Topaloglu (2010b) develop a dynamic programming decomposition method that solves a single leg DP with revenue estimates for each leg in an itinerary. These estimates are determined through an optimization model that takes into account the probabilistic choice-based demand. In Huang and Liang (2011), the authors develop their own dynamic programming formulation that solves for seat control policies. Their solution method for their DP estimates the value function of revenue for the problem and then solves the DP with a parametrized function and a sampling methodology. Zhang (2011) developed a new method to solve Talluri and van Ryzin's dynamic programming formulation, in which his model yielded tighter bounds on revenue than the decomposition and CDLP methods originally explored. Another method developed was that of Kunnumkal (2011), who took a two step approach where the first step relaxed the flight leg capacity constraints via Lagrangian relaxation, while the second step solves the problem with perfect information, yielding a final solution that determines capacity dependent policies. Meissner and Strauss (2012b) consider inventory sensitive bid prices, and developed a dynamic programming approach of their own. Their model estimates the value function of the Markov decision process and then solves for the bid prices appropriately.

Other models that have been developed under the choice-based assumptions range from heuristics to mixed integer programs. van Ryzin and Vulcano (2008a) develop an optimization model that solves for nested protection levels. Their model assumes a general choice-based demand model, which separates the choice model from the optimization problem itself. Their results are computationally efficient, which implies practical use in the airline industry. Chaneton and Vulcano (2011) convert the choice model into a continuous demand estimate, in which they develop a sub gradient algorithm to find the stationary point. Their model allows for partially accepted itineraries, similar to that of Topaloglu (2009), and yields mixed results over the CDLP. Chen and de Mello (2010a) develop an optimization model which allows the customers to work their way up the fare classes. Each customer has a finite probability of buying up the fare class buckets, creating a demand stream for a set of optimization problems to solve. Gallego *et al.* (2011) introduce a generalized attraction model, and show the relationship between their generalized attraction model and the more specific independent demand and basic attraction models. Their model was developed to overcome the complexity of the CDLP resulting in a sales-based linear program (SBLP) which utilizes the previously introduced general attraction model. Since a large majority of the research results in optimal bid prices, Meissner and Strauss (2012a) develop a heuristic that improves on the initial bid prices from any model. Their method covers general choice models, and shows revenue gains over available alternatives with a low computational burden.

Meissner and Strauss (2010) also develop a mixed integer program where policy decisions on restricted fare classes are determined simultaneously with pricing decisions on unrestricted fare classes. This problem formulation looks at networks where some fares are determined in advance (restricted), and others have a set of available options (unrestricted). Kunnumkal (2011) develop a two-step method for the choice-based revenue management problem. The first step of their method solves an MIP that selects the best policies, similar to that of the CDLP. The second step then determines the bid prices based on the policies selected through an LP. This LP can be randomized, similar to that of the RLP, and provides good solutions compared to that of the CDLP. Meissner and Strauss (2011) also develop a mixed integer program, under the assumption that there is weak market segmentation. Their model proves to be computationally intractable, yet the authors provide solution methods that trade off run time for computational accuracy.

1.1.4 Other Demand Models

Although independent demand and choice-based demand covers the vast majority of revenue management literature available, there are other demand models that have been used to produce solutions. Topaloglu (2009) uses Lagrangian relaxation to determine bid prices that are dependent on both capacity and time. His model allows for a single leg to be accepted out of a multi-leg itinerary, and then decomposes the problem into a leg-level heuristic. He argues this makes it computationally tractable, since the size of the problem is not constrained by the complexity of the network. Dynamic programming decomposition methods also exist for other demand models, including ones that solve for both overbooking and seat allocations (Erdelyi and Topaloglu, 2010), as well as pricing decisions where demand is dependent on the prices being offered (Erdelyi and Topaloglu, 2011). In Song *et al.* (2010), the authors build a mixed integer linear program to solve the network revenue management problem. Their model uses a stochastic estimation of demand through a linear approximation. A step function is used to estimate the demand, which then allows them to evaluate revenue as a uniform distribution. They found their MILP generated upper and lower bounds on the original randomized linear program (Song *et al.*, 2010). Another model, created by Perakis and Roels (2010), uses the decision criterion of maximin and minimax. They consider multiple control sets such as partitioned booking limits, nested booking limits, and fixed bid prices, and then generate an optimal solution based on the maximin and minimax criteria.

1.2 Contributions of this Dissertation

As airline revenue management has advanced from independent demand models into more realistic dependent demand models, a gap has developed between academic research and industry implementation. Models found in the literature focus on minimizing assumptions and properly modeling demand, while industry practice is relegated to sub-optimal models due to airline network complexity, lack of ability to change the methods by which they estimate demand, or lack of pragmatic solutions. The goal of this dissertation is to reduce this gap, and introduce alternative formulations that are both implementable for airline use and can progress airline RM research into the next frontier.

My first contribution in this dissertation is a mathematical formulation for network revenue management utilizing a MNL demand model. The formulation is substantially less complex than traditional dependent demand RM models in the literature, and consistently outperforms current industry practice. With numerous examples, I show the flexibility of the formulation while highlighting the considerable gains in computational complexity, eventually solving a large network example that is virtually unsolvable by one of academic literature's best performing models in a reasonable amount of time. The reduction of complexity paired with the performance of the formulation provides an applicable model for airline implementation as well as a foundation to build upon for future revenue management research.

My next contribution is centered around the MNL demand model framework and proper implementation for airline revenue management. Traditional dependent demand revenue management models utilize static assumptions on pricing and ticket preferences, creating demand estimates that are independent of revenue management controls, despite changes in price and ticket attributes directly influenced by these controls. I introduce an MNL framework for network RM that incorporates passenger sensitivities to price and other important ticket attributes, creating a demand model that properly responds to revenue management controls while addressing passenger purchasing behavior. The framework is easily implemented as there are tools for fitting these models based on sales data, resulting in demand estimates that reflect true passenger behavior without assuming static prices or ticket attributes.

The final contribution in this dissertation incorporates the previous framework into a choice-based mixed integer non-linear programming formulation for airline revenue management. Building upon the complexity gains from the first formulation and the dynamic nature of demand from the framework, I introduce a model that maximizes expected revenue by adjusting ticket availability and prices, simultaneously. Different from other revenue management models, my second formulation can adjust prices while accommodating changes in demand, and set ticket availability based on different ticket attributes commonly seen in the airline industry. The flexibility of this formulation leads to gains in expected revenue when compared to the first model, as well as a post-RM pricing method. Despite the non-linear nature of this formulation, the complexity gains and ease of solving this formulation make it possible for industry implemented, as shown by large examples based on the Southwest Airlines ticket model.

1.3 Dissertation Organization

The remainder of this dissertation is organized as follows. In Chapter 2, I introduce a mixed integer programming model that incorporates choice-based demand,

and compare it against the popular models in revenue management literature and practice. The contents of this chapter was published in October 2014 in the Journal of Revenue and Pricing Management (doi:10.1057/rpm.2014.17), and has been reproduced for my dissertation. Chapter 3 builds upon observations from Chapter 2, focusing on formally defining the multinomial logit choice demand framework for airline revenue management, integrating ticket attributes and price for passenger preference. The framework lays the groundwork for airline specific demand models, taking into account each airline's ticket definitions and display structure to generate unique passenger preference utilities. Chapter 4 utilizes the previous revenue management framework and builds upon Chapter 2's mixed integer program to introduce a price-dynamic choice-based mixed integer non-linear programming model for airline revenue management. The price-dynamic model is able to simultaneously solve for ticket availability and price, showing considerable gains versus other revenue management models. Chapter 5 closes the dissertation, highlighting the results from previous chapters and indicating a direction the field of revenue management can progress with these contributions.

Chapter 2

A CHOICE-BASED MIXED INTEGER PROGRAMMING FORMULATION FOR THE NETWORK REVENUE PROBLEM

2.1 Introduction

From its inception, airlines have used Revenue Management (RM) techniques to improve their revenue performance or yield by optimizing the passenger mix through fare class seat availability or bid price hurdle rates. Both leg-based and Origin-Destination (O&D)-based approaches have used a common assumption that the passenger demand associated with a given flight or O&D path and fare class are known and forecast independent of other options within the market. For example, demand forecasts for full fare passengers on BOS-PHX-LAX are based on historical traffic observations on that specific path and do not explicitly account for the passenger demand associated with other paths in the market like BOS-ORD-LAX. In addition, most RM approaches used in practice today assume that fare classes are mutually exclusive of one another when optimizing seat allocations or bid prices.

These assumptions preclude the demand interactions between different routes, fare classes and competition from other carriers in the same markets, and limit the quality of the optimization results and controls. To remove the limitations of these assumptions, the demand forecasts and optimization must consider the interactions between the different fare classes and routes available to potential passengers at the point of sale.

In this chapter, we propose a mixed integer programming formulation that explicitly incorporates the fare class and routing interactions using a MultiNomial Logit

(MNL) choice model. This formulation, which we refer to as the Choice-based Mixed Integer Program (CMIP), represents an alternative formulation to the Choice-based Deterministic Linear Program (CDLP), as proposed by Liu and van Ryzin (2008). The fundamental difference between the two formulations is that the one proposed here considers individual market strategies as variable options rather than network level strategies. CMIP shows similar revenue performance to the CDLP while enjoying significant reductions in the number of decision variables, which is shown to result in significant computational advantages in the problem instances that we have tested. CMIP is also shown to yield improvements over popular leg-based EMSR models and O&D network-based models covered by the literature and used in practice.

The remainder of this chapter is organized as follows. Section 2.2 reviews the recent literature and developments associated with incorporating passenger choice into the RM process, Section 2.3 presents the CMIP formulation, Section 2.4 solves an illustrative example and compares performance to other models, Section 2.5 solves larger examples and compares performance, and Section 2.6 highlights the conclusions and potential future research directions.

2.2 Literature Review

Although many RM models have been developed over the past 30 years, the following literature review focuses on the choice-based demand approaches that aim to model consumer behavior more accurately for the network RM problem.

To provide a framework for choice-based modeling, we first present an overview of some of the leg-and O&D-based methods available. Two leg-based independent demand methods worth noting, however, are that of Littlewood's 1972 paper (which was later republished in 2005), and Belobaba (1989). Often, Littlewood (2005) is cited as being one of the first models to solve the RM problem. His model determined the necessary protection limits by comparing two products' expected demand and fares. Littlewood (2005) proposed a rule, termed Littlewood's Rule, which determines protection levels for the higher fare classes. Belobaba (1989) expanded on Littlewood's research, and created the Expected Marginal Seat Revenue (EMSR) model. His first model, EMSR-a, executed pair-wise comparisons to determine how many seats to reserve for higher fare classes. The EMSR-a model could compare any number of pairs, and would aggregate protection levels as it moved up the fare class buckets (Belobaba, 1989). Later, Belobaba expanded on his own model, creating the EMSR-b methodology. EMSR-b, instead of aggregating over protection levels, aggregates the demand for higher paying passengers, and calculates a weighted fare for them (Belobaba, 1992). This weighted fare is then used for a comparison, and protection levels are calculated. For a more complete history of independent demand RM models, we refer the reader to Weatherford and Ratliff (2010). These independent demand models, like that of Littlewood (2005) and Belobaba (1989), were computationally efficient, but they lack the network interactions and competitive effects present in today's complex airline markets. To this end, research moved towards choice-based modeling techniques for solving the network RM problem.

Two of the first, and possibly most influential, papers in choice-based modeling for network RM were Gallego *et al.* (2004), and Talluri and van Ryzin (2004a). In Gallego *et al.* (2004), the authors propose a linear program that solves the network RM problem with a general discrete choice model. Using the probability of a purchase as a parameter, the model determines the amount of a time that each set of policies, defined by the itinerary and fare, is to be offered. This method maximizes the revenue across the entire network, by selecting a subset of available policies, constrained by the available space consumed on a leg (Gallego *et al.*, 2004). This model is later developed into the Choice-based Deterministic Linear Program (CDLP), and has become a benchmark for the testing of newer models in the field of network RM. In particular, the CDLP determines the optimal amount of time to offer a set of policies, S. This set S is comprised of open and closed policies for each O&D fare class combination within the network.

Talluri and van Ryzin (2004a) formulated the problem as a dynamic program, which modeled the probabilities of different purchases using a general discrete choice model, and determined which policy sets to offer based on the available capacity, similar to the model in Gallego *et al.* (2004). Talluri and van Ryzin introduced the concept of efficient sets, which allowed for search techniques to manage the complex nature of the solution space. From this point on, research in choice-based RM has gone in one of three directions: solution methodologies for the CDLP, approaches to solve dynamic programming formulations, or formulations that are new altogether.

2.2.1 CDLP Solution Methodologies

In Kunnumkal and Topaloglu (2008), the authors created an alternative form of the CDLP, and obtained better results through the solution of the primal. Liu and van Ryzin (2008) expanded on the original CDLP, and developed an iterative approach by applying the bid prices generated from the CDLP to a leg-level decomposition approach to Talluri and van Ryzin (2004a)'s dynamic program. The results of their method provided capacity and time-dependent bid prices, which are useful for industry application. The authors also expanded on the notion of efficient sets, and applied them to the CDLP, generating methods for solving this complex problem (Liu and van Ryzin, 2008). Shortly thereafter, Bront *et al.* (2009) developed a column generation algorithm to solve the CDLP, in the special case of having non-disjoint markets. Their model considered situations where market demand can overlap, and competition can arise between O&D's as well as pricing options. The authors provide details on how to solve the column generation algorithm for the CDLP, as well as provide two methods for solving the subproblem of determining which set to introduce into the reduced primal problem (Bront *et al.*, 2009). Talluri (2011) relaxed the CDLP, and solved a Segment-based Deterministic Concave-Program (SDCP), which provided looser upper bounds to the original problem. Following this relaxation, Meissner *et al.* (2013) expanded on the model to include constraints on the product selections, creating the extended-SDCP.

2.2.2 Approaches to Solve Dynamic Programming Formulations

Although Talluri and van Ryzin (2004a) provided one of the first models utilizing dynamic programming formulations for network RM, others also explored this approach. Zhang and Cooper (2005) offered a different perspective, and created a Markov Decision Process (MDP) formulation for cases where multiple flights are being offered between O&D's in short time spans. Later, Zhang and Cooper (2006) developed an MDP model that allowed for substitution to take place between flights. Although both of these MDP formulations could be solved via dynamic programming, more efficient methods were found in the form of inventory-pooling (Zhang and Cooper (2005)) and heuristics (Zhang and Cooper (2006)). Some models were developed in conjunction with dynamic programs, like in Farias and van Roy (2007) and Adelman (2007). In Farias and van Roy (2007), the authors model the network RM problem as a dynamic program, and then solve it with a linear programming approximation. Their model is unique, as it solves for the bid prices directly, rather than producing policy-based decisions. Adelman (2007) utilized an affine approximation for the value function of his dynamic program. Similar to Farias and van Roy (2007), his model determines the bid prices for the network, and generates a dynamic set of bid prices. Kunnumkal and Topaloglu (2010b) developed their own dynamic programming decomposition method, which solves the single-leg decomposition with revenue estimates for each leg in an itinerary. The revenue estimates were generated ahead of time through an optimization model utilizing the choice-based modeling schema.

Some models shift the focus from policy decisions and generate solutions that determine seat allocation policies. Huang and Liang (2011) developed a dynamic programming formulation, which they solve by estimating the value function of the dynamic program (DP) with a sampling technique. Their model solves for the seat control policies, rather than open or closed fare class decisions. In Zhang (2011), the authors proposed an alternative way to solve the dynamic programming formulation of Talluri and van Ryzin (2004a), and provided better bounds on the optimal solution for the original problem. Kunnumkal (2011) took a different approach to solving the dynamic program, and offered two approximation models for solving the choice-based network RM. Lagrangian relaxations were done for both methods, one based on relaxing the flight leg capacities and the other based on perfect demand information. His model generates capacity-dependent policies, similar to that of the original dynamic programming formulation (Kunnumkal, 2011). Another unique formulation is found in Meissner and Strauss (2012b), in which they develop a dynamic programming formulation that takes into account inventory sensitive bid prices. Their model estimates the value function of an MDP to determine capacity-dependent bid prices.

2.2.3 Alternative Formulations

Other models different from the typical dynamic programming formulations and CDLP were also developed. van Ryzin and Vulcano (2008a) developed an optimization model that solves the choice model independently from the optimization model itself, creating an easier and quicker solution methodology to the problem. Their model solves for nested protection levels, rather than policy or bid price optimization. Chaneton and Vulcano (2011) sought to simplify the problem by changing the formulation of the choice-based demand models. They estimate the choice-based demand by applying a linear approximation to the demand, creating a continuous function with a stationary point found using a sub-gradient algorithm. Their model allows for partially accepted itineraries in which the passenger requests can be accepted on legs within the itinerary, but not the entire itinerary itself (Chaneton and Vulcano, 2011). This approach is similar to that of Topaloglu (2009), in which a bid price solution methodology is developed by applying a leg-level decomposition approach. Chen and de Mello (2010a) developed a formulation that modeled the buy-up behavior directly. Their model allows for passengers to step up in fare classes if their desired fare class is unavailable. From this buy up pattern, the authors were able to determine a demand stream, which then was used to solve a set of optimization problems.

Gallego *et al.* (2011) introduced the generalized attraction model, which can be applied to any demand input. The independent demand, as well as basic attraction models, were found to be special cases of this generalized attraction model. They develop their model to combat the complexity of the CDLP, resulting in a new formulation known as the Sales Based Linear Program (SBLP). Another mathematical model, in the form of a mixed integer program, was developed by Meissner and Strauss (2010). Their model solved for both policy decisions on restricted fare classes (i.e., fare classes in which discrete fares are determined in advance), as well as pricing decisions on unrestricted fare classes (i.e., fare classes in which a continuous range of available prices exist). Kunnumkal (2011) developed a two-step method for solving the network RM problem. His method first determines which policies are optimal through a choice-based mixed integer program, followed by a linear program that determines the marginal value of seats. He argues that the linear program can be randomized, and provides good solutions compared to those obtained by the CDLP (Kunnumkal, 2011). Meissner and Strauss (2011) also developed a mixed integer program, under the assumption that market segmentation is weak. This creates ambiguity in the demand stream, and their model proved to be computationally intractable. The authors provided alternative solution methods for solving their mixed integer program, citing cases where shorter run times were more advantageous than computational accuracy. In the following section, we present a new formulation that improves on the research reviewed here.

2.3 Mathematical Model

We consider a network with legs $l \in L$, and containing multiple markets defined by set J. A market j represents an O&D pair. There exists a set of policies, I_j, defined for each market $j \in J$, where a policy $i \in I_j$ is defined as a pair of itinerary and fare class assignments. Since each market can have multiple itineraries (i.e., paths) and fare classes, we define the set K_{ji} that contains all defined fare class-itinerary pairs for market $j \in J$ and policy $i \in I_j$. The planning horizon associated with this model can be viewed as the time to departure. We discretize time into periods and denote the index set of time periods by T. Having defined the network parameters, the parameters and decision variables for our mathematical formulation are formally defined in Table 2.1.

As an example, consider the network given in Figure 1, which is the same example used in Bront *et al.* (2009) and Liu and van Ryzin (2008). The network contains three nodes and leg capacities of 10, 5 and 5 seats for legs 1, 2, and 3, respectively. Each leg represents a single flight, thus there are no parallel flights for this network. Table 2.2 includes further data on this example; eight products were defined by what the authors refer to as "O&D path" and fare class combinations.

λ_t	Number of customer requests for flights to the network in period t, $t \in T$
P_j	Probability of an arrival for market j, $j \in J$
S_{ji}	Probability that a purchase is made for market j under policy $i \in I_j$
$P_{k\|j,i}$	Probability of a purchase on fare class-itinerary $k \in K_{ji}$, given purchase is made for market j under policy $i \in I_j$
R_k	Revenue for a purchase on fare class-itin. k given policy i is used
A_{kl}	Binary parameter representing consumption of leg l for fare class-itin. $k \in K_{ji}$
c_l	Capacity of leg $l \in L$ available at the beginning of the planning horizon
Z_{jit}	Fraction of period t demand for market j served under policy $i \in I_j$
X_{jit}	Binary decision variable to use policy $i \in I_j$ for market j in period $t \in T$

Table 2.1: Table of Notations Used in CMIP

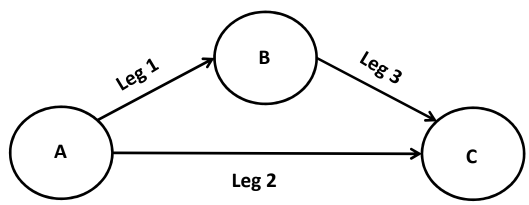

Figure 2.1: Illustrative Example: Three-leg Network

Table 2.3 includes data on customer preferences and utilities of the different products for each of the five segments. The preference vectors represent the utility that the products in the consideration set provides for the segment. For example, segment 1 has a consideration set of $\{1, 5\}$, and a preference vector of $(5, 8)$. This means that the first segment has a utility of 5 for product 1 (i.e., the A-C itinerary with a cost of \$1200) and a utility of 8 for product 5 (i.e., the A-C itinerary with a cost of \$800). The larger utility value implies that customer segment 1 prefers the A-C itinerary with a cost of \$800 over the A-C itinerary with a cost of \$1200. The no-purchase utility corresponds to the option of not purchasing either product.

For our formulation, the network illustrated in Figure 1 would result in three markets, i.e., $J = \{\text{AB}, \text{BC}, \text{AC}\}$. Note that market AC contains two different itineraries,

Product	Origin-Dest. Path	Class	Fare
1	A - C	High	$1200
2	A - B - C	High	$800
3	A - B	High	$500
4	B - C	High	$500
5	A - C	Low	$800
6	A - B - C	Low	$500
7	A - B	Low	$300
8	B - C	Low	$300

Table 2.2: O&D Paths and Fare Classes for the Illustrative Network Example (Liu and van Ryzin, 2008)

Segment	Arrival Rate	Consideration Set	Preference Vector	Utility of No Purchase
1	0.15	$\{1, 5\}$	(5, 8)	2
2	0.15	$\{1, 2\}$	(10, 6)	5
3	0.20	$\{5, 6\}$	(8, 5)	2
4	0.25	$\{3, 7\}$	(4, 8)	2
5	0.25	$\{4, 8\}$	(6, 8)	2

Table 2.3: Data on Demand and Customer Preferences for the Illustrative Network Example (Liu and van Ryzin, 2008)

the direct path from A to C as well as the path containing the connection A-B-C. For each market, $j \in J$, there is a set of available policies, I_j. Each policy lists the available options, each defined by an itinerary and the fare classes. For instance, the AC market has two competing paths, AC and ABC, thus the available policies for the AC market would be all combinations of high and low fare class options, as well as the possibility of closed itineraries.

When implementing policies in real life airline RM systems, certain fare classes are nested within their lower fare class counterparts. For instance, the policy containing AC high and AC low open simultaneously would be equivalent to opening only AC low, since policy implementation is generally based on bid prices. That is not to say the higher fare class is closed, just that there is no situation where an airline would refuse a higher paying passenger just because the policy only defines AC low as being open. This natural nesting among the fare classes eliminates the need to separately define policies in which AC high and AC low are open simultaneously. The elimination of these simultaneous policies, however, removes any buy up potential, thus the model makes a conservative assumption that buy up is negligible.

Finally, the set K_{ji} includes all fare class-itinerary pairs defined for market j under policy $i \in I_j$. The first component, market, is defined by the available market set J. The policy component is defined by the set of available policies I_j. The itinerary path is determined by the structure of the network itself. The combination of appropriate market-policy-itinerary path groupings generates the set K, which is referred to as the fare class-itinerary. For the illustrative example, the values of these sets can be seen in Table 2.4, in the fourth column.

We can now determine the values of our parameters λ_t, P_j, A_{kl}, R_k, $P_{k|j,i}$ and S_{ji} for the illustrative network example. For this example, we assume that the arrival rate, or the number of unit demand arrivals per period, stays constant at the values listed under the column labeled "Arrival Rate" in Table 2.3 for each customer segment. We actually use these values for $\lambda_t P_j$ for each market. From Table 2.3 the arrival rate for market AC, for example, will be equal to the sum of the arrival rate values listed for customer segments 1, 2, and 3, i.e., $\lambda_t P_{\mathrm{AC}} = 0.50$ customers for each time period. Similarly, $\lambda_t P_{\mathrm{AB}} = 0.25$ and $\lambda_t P_{\mathrm{BC}} = 0.25$ for all $t \in T$. Note that in the example, the time is scaled so that the total arrival rate, $\lambda_t = 1$.

The values of A_{kl} can be determined by examining the network and fare class-itineraries. If fare class-itinerary $k \in K_{ji}$ consumes space (i.e., 1 unit of capacity) on leg l, then A_{kl} is assigned a value of 1. Otherwise, A_{kl} is assigned a value of zero. Since R_k represents the revenue earned for fare class-itinerary k being purchased, those values can be read directly from the pricing table.

The values for $P_{k|j,i}$ were determined through conditioning. For instance, if the policy available was AC High/ABC High, then a fraction of the purchases would purchase the AC itinerary while others would purchase the ABC itinerary. Since we are conditioning on the fact that a purchase was made, we merely need to determine what fraction of passengers purchased the AC High option (or, "fare class-itinerary") and what fraction of passengers purchased the ABC High option. To do this, we must first determine which customer segments, as defined by Table 2.3, prefer each of the options. For the AC High option (defined as product 1), we can see that customer segments 1 and 2 have utility values for this product (as defined by their consideration sets). Likewise, for the ABC High option (defined as product 2), we can see that only customer segment 2 has preference for this product.

We first calculate S_{ji}, which denotes the probability of a purchase by an arriving market j customer under policy $i \in I_j$. For the above example of policy AC High/ABC High for the AC market, the probability of purchase by a market AC customer under policy $i \in I_j$ can be calculated as a weighted average of the purchase probabilities that can be calculated from the given utilities of the products available under the AC High/ABC High policy, and that of no purchase. That is,

$$S_{\text{AC,AC High/ABC High}} = \frac{5}{5+2}\left(\frac{0.15}{0.50}\right) + \frac{10+6}{10+5+6}\left(\frac{0.15}{0.50}\right) + 0\left(\frac{0.20}{0.50}\right) = 0.443\ ,$$

by using the given utilities and conditioning on the event that the AC customer is from segment 1, 2 or 3, respectively.

Note that the utility values used above represent the respective e^{u_l} terms used in the MNL model. Essentially, the ratio $5/(5+2)$ can be rewritten in traditional MNL format as $e^{1.609}/(e^{1.609}+e^{0.693})$, where the values of 1.609 and 0.639 would represent the MNL utilities for the AC market for AC High and the no purchase option, respectively. We use this simplified notation for ease of representation, following the tradition set by previous papers in this area.

Then, the probability that the customer bought fare class-itinerary k, such that $k \in K_{\text{AC,AC High/ABC High}}$ given that a purchase was made by an AC customer under policy AC High/ABC High can be calculated by

$$P_{\text{AC High|purchase under AC High/ABC High}} = \frac{\frac{5}{5+2}\left(\frac{0.15}{0.50}\right) + \frac{10}{10+5+6}\left(\frac{0.15}{0.50}\right)}{0.443} = 0.806\ .$$

Similarly, we can calculate

$$P_{\text{ABC High|purchase under AC High/ABC High}} = \frac{\frac{6}{10+5+6}\left(\frac{0.15}{0.50}\right)}{0.443} = 0.194\ ,$$

or, simply by observing that for this policy with two fare class-itinerary options,

$$P_{\text{ABC High|purchase under AC High/ABC High}} = 1 - P_{\text{AC High|purchase under AC High/ABC High}}\ .$$

Continuing similarly, we obtain the values in Table 2.4 for this illustrative example. Note that in columns one through five, the table provides the markets (i.e., $j \in J$), the policies defined for each market, (i.e., $i \in I_j$), the purchase probability for market

j under each policy $i \in I_j$, (i.e., S_{ji} values), the defined fare class-itineraries (i.e., set K_{ji}) for each market j under policy $i \in I_j$, and finally, the conditional probability that the option given by a particular fare class-itinerary $j \in K_{ji}$ will be selected, given that a purchase for market j was made under policy $i \in I_j$.

Finally, our formulation uses the following decision variables. X_{jit} represents the binary variable that takes on a value of 1 if the decision is to use policy i for market j in period t, and Z_{jit} represents the fraction of market j demand served under policy i in period t. To expand on the variable Z_{jit}, consider an example where, for a given market policy i and five time periods, Z_{jit} takes on values of $(0, 0, 1, 1, 0.172)$. This vector would represent the following set of decisions. For time periods 1 and 2, policy i is not available and no demand for market j would be served under this policy. During time periods 3 and 4, policy i is available, and any arriving demand for market j would be served. Finally, during time period 5, policy i is available, but only 17.2% of the potential demand should be served. Note that the term "served" here does not necessarily mean that they will be purchasing a ticket; it basically means that they get to consider the various options available to them for market j, under policy $i \in I_j$. As a result of this consideration, they may or may not purchase a ticket on market j. Using these two decisions variables, and the parameters previously defined, we formulate the *Choice-based Mixed Integer Program (CMIP)* as follows.

Market (Set J)	Policies (Set I_j)	S_{ji}	Fare Class-Itineraries (Set K_{ji})	$P_{k\|ji}$
AB	AB High	0.667	AB High	1
	AB Low	0.857	AB Low	1
BC	BC High	0.750	BC High	1
	BC Low	0.875	BC Low	1
AC	AC High/ABC High	0.443	AC High	0.806
			ABC High	0.194
	AC High/ABC Low	0.761	AC High	0.556
			ABC Low	0.444
	AC Low/ABC High	0.809	AC Low	0.681
			ABC High	0.319
	AC Low/ABC Low	0.607	AC Low	0.773
			ABC Low	0.227
	AC High/ABC Closed	0.414	AC High	1
			ABC Closed	0
	AC Low/ABC Closed	0.580	AC Low	1
			ABC Closed	0
	AC Closed/ABC High	0.279	AC Closed	0
			ABC High	1
	AC Closed/ABC Low	0.347	AC Closed	0
			ABC Low	1

Table 2.4: Set Definitions and Calculated Parameters for the Illustrative Example Adapted from Liu and van Ryzin (2008)

$$\text{Maximize} \qquad \sum_{t \in T} \sum_{j \in J} \lambda_t P_j \sum_{i \in I_j} Z_{jit} \, S_{ji} \sum_{k \in K_{ji}} R_k P_{k|j,i} \tag{2.1}$$

Subject to:

$$\sum_{t \in T} \sum_{j \in J} \sum_{i \in I_j} \sum_{k \in K_{ji}} \lambda_t P_j Z_{jit} S_{ji} P_{k|j,i} A_{kl} \leq c_l, \qquad \text{for all } l \in L, \tag{2.2}$$

$$\sum_{i \in I_j} X_{jit} \leq 1 \ , \qquad \text{for all } j \in J, t \in T, \tag{2.3}$$

$$Z_{jit} \leq X_{jit} \ , \qquad \text{for all } j \in J, i \in I_j, t \in T, \tag{2.4}$$

$$Z_{jit} \in R^+, X_{jit} \in \{0, 1\}, \qquad \text{for all } j \in J, i \in I_j, t \in T. \tag{2.5}$$

The objective function (2.1) represents the total expected revenue across all time periods for the decision variable Z_{jit}. The objective function can be broken into two main components: the arrival rate of demand per market and the expected revenue for a given policy. The first component, the arrival rate of demand per market, is the product of the expected number of customer requests in period t, (i.e., λ_t) and the probability that an arrival demands a ticket for market j, (i.e., P_j). This product, $\lambda_t P_j$, represents the expected number of arrivals in time period $t \in T$ for market j. The second component, the expected revenue obtained from market j under policy $i \in I_j$, $E[R_j(i)]$, assuming that customer preferences remain unchanged throughout the planning horizon, can be calculated as follows.

$$\begin{aligned} E[R_j(i)] &= (1 - S_{ji})\, 0 + S_{ji} \, E[\text{Revenue} \mid \text{purchase in market } j \text{ under policy } i \in I_j] \\ &= S_{ji} \sum_{k \in K_{ji}} R_k P_{k|j,i} \ , \end{aligned} \tag{2.6}$$

where R_k denotes the revenue from a sale on fare class-itinerary $k \in K_{ji}$, $P_{k|j,i}$ denotes the conditional probability of a purchase on fare class-itinerary $k \in K_{ji}$, given a purchase for market j is made under policy $i \in I_j$, and finally, S_{ji} is the probability that a purchase for market j is made under policy $i \in I_j$.

Constraint set (2.2) ensures that the capacity constraints on the legs are not violated. The continuous decision variable, Z_{jit}, allows for the partial accommodation of demand, and facilitates the determination of bid price values for the flight leg, $l \in L$. The purpose of a bid price is to represent the marginal value of an extra seat on a given leg. In the event constraint set (2.2) is binding, we can increase the capacity of a leg to determine what impact this increase would have on the objective function. The value of variable Z_{jit} could be increased a marginal amount, no greater than one, if the current value is less than one. In the event Z_{jit} was already at a value of one, then the model could select a different Z_{jit} to improve the objective function. This would force the constraint to be binding, again, and the objective function, which also contains the Z_{jit} variable, would increase appropriately. This increase would be analogous to the shadow price of a linear program, thus it can be used as the marginal value of a seat on a given leg. The marginal value of a seat is then translated into the bid price for the leg, and could be used for bid price based control policies.

One difference between our formulation and other formulations stems from the fact that the other models account for all market combinations in the form of sets, whereas CMIP combines policies for O&D markets to determine the overall policy for the network. For instance, the CDLP selects which sets are optimal, while the CMIP model selects, individually, which O&D fare class combinations optimize our revenue. Since the CMIP focuses on a market-by-market level, the total number of variables for the problem is greatly reduced, which results in reasonable solution times for larger networks, as we show for a large network instance. As mentioned above, the complexity of the network greatly impedes the quality of the solutions that one can obtain from the CDLP formulation within a reasonable run time. Hence, having fewer variables in the CMIP formulation allows for the modeling of larger networks with solution times that are implementable for industry use.

To illustrate the magnitude of the difference in variables, consider the small network that we considered earlier, depicted in Figure 1. In this network, for a single time period, the CMIP has a total of 24 variables and 18 constraints. The CDLP, for the same scenario, has a total of 255 variables and 4 constraints. As we increase the number of time periods, the CMIP increases in both variable and constraint totals, while the CDLP does not. The advantage of the CMIP, however, is when the complexity of the network is increased. Adding just one more node with respective high and low fare classes and connections (assuming this node is independent of the markets currently in the network), would only increase the CMIP to 28 variables and 22 constraints. This same network for the CDLP formulation would have 16,383 variables and 5 constraints. The CDLP has a smaller constraint set, yet the variable space is exponentially increasing as the complexity of the network is increased. The CMIP has a much smaller variable space, and a reasonably sized constraint space. As the complexity of the network is increased, the variable and constraint space does not increase in an exponential fashion; the total number of variables for the CMIP, however, would increase multiplicatively for each additional time period.

2.4 Solution of the Illustrative Example

We solved the CDLP and CMIP formulations for the example presented above (see Figure 1), and implemented the obtained policy decisions in a simulation to compare the performance of the two approaches. We used AMPL and Gurobi 5.0.1 to build and solve the formulations. We assumed that buy-up did not happen, as there only is a very small probability that a passenger will purchase a higher priced ticket if a lower priced ticket is available.

We programmed the simulation in MATLAB following a traditional Monte Carlo simulation approach. First, the simulation takes bid prices as the control, and generates the available fare classes in each time period. We assume a stationary arrival rate of customers (I.e., $\lambda_t = \lambda$ for every time period $t \in T$) and generate exponential customer interarrival times with this rate. For each customer, the simulation model generates the identity of the market that the customer is interested in purchasing, as well as what, if anything, the customer purchases using P_j, S_{ji} and $P_{k|ji}$ in a relatively standard random number generation scheme. In case of a purchase, the capacity of the legs for the requested itinerary is reduced, and the total revenue is updated. We ran the simulation for 2000 iterations for each of the network instances tested. This simulation is used for all of the results following the illustrative example.

The solutions from the two formulations generated similar bid prices across the majority of the tests. The average total revenue values obtained with the two approaches were also comparable. As seen from Table 2.5, the CDLP and CMIP reach identical bid prices in every case except for $T = 5$ and $\lambda = 5$. Due to the nature of this illustrative example, the leg AC only has pricing options of \$800 and \$1200. Hence, having a bid price of \$750 implies both pricing options are to be open. Note that obtaining a bid price of \$0 (which can be observed in the case of the CMIP for $T = 5$ and $\lambda = 5$), would have the same effect as having a bid price of \$750. This implies the CDLP and the CMIP generate identical bid price control strategies across all combinations for this example.

For the results presented in the last two columns of Table 2.5, the CMIP increased expected revenue by an average of 2.72% when compared to the CDLP. Although these models are similar, the way they handle expected traffic is different. The CDLP uses the direct probability of purchases generated from the MNL choice model. The CMIP uses the probability of purchases for a given fare class-itinerary, conditioned on a purchase being made. These differences are subtle, yet impact the CDLP and CMIP objective functions, so comparison on these values alone is insufficient. To this end, we simulated the policy for the instance of $\lambda = 5$ and $T = 5$, to see whether the differences between revenues would continue to hold. This problem instance was

		CMIP Bid Prices			CDLP Bid Prices			Obj. Fn.	Obj. Fn.
λ	T	AB	AC	BC	AB	AC	BC	CMIP	CDLP
	1	\$0	\$0	\$0	\$0	\$0	\$0	\$515	\$497
1	5	\$0	\$0	\$0	\$0	\$0	\$0	\$2,577	\$2,485
	10	\$0	\$0	\$0	\$0	\$0	\$0	\$5,155	\$4,971
	1	\$0	\$0	\$0	\$0	\$0	\$0	\$2,577	\$2,485
5	5	\$0	\$0	\$500	\$0	\$750	\$500	\$10,664	\$10,064
	10	\$300	\$1200	\$500	\$300	\$1,200	\$500	\$13,168	\$13,167
	1	\$0	\$0	\$0	\$0	\$0	\$0	\$5,155	\$4,971
10	5	\$300	\$1,200	\$500	\$300	\$1,200	\$500	\$13,168	\$13,167
	10	\$500	\$1,200	\$500	\$500	\$1,200	\$500	\$13,500	\$13,500

Table 2.5: Results from CMIP and CDLP Formulations

chosen since this was the only case where the bid prices differed between the two models. We determined the 95% confidence intervals around the expected revenue for both simulations. The CMIP resulted in an interval of $(\$6,959, \$10,071)$, while the CDLP resulted in an interval of $(\$6,950, \$10,088)$. As expected, the results we observed for the CMIP and CDLP were very close to one another. Based on the results of the simulation, we conclude that, in this set of problem instances, the CDLP and CMIP provide similar solutions, and can be used interchangeably.

2.4.1 Implementing the Solution

One advantage of the CMIP formulation is the fact that both the policy and bid price controls are useful for industry application. The solution to the CMIP indicates which policy should be offered for each market during a particular time period. In the three-leg network problem instance, the solution would instruct, for instance, to open the high fare class for itinerary ABC during time periods 1, 2 and 3. Additionally, it would indicate to open the low fare class for itinerary AB during time periods 1 and 2, while opening the high fare class for time period 3. A reservation system could directly interpret this decision to open and close these particular fare classes, following the guidance of the CMIP solution. The reservation system could then generate cut offs for certain fare classes and calculate protection levels based on the airline's current system, if necessary.

An effective bid price control can be derived from the solution to this formulation as well. Effectively, the lowest open fare class-itinerary on a leg, at optimality, reflects the marginal value of that leg. In terms of right-hand side sensitivity, adding an additional seat to this leg would increase the overall network revenue by the value of that itinerary as long as the optimal basis of the integer solution does not change. This marginal value is analogous to a bid price that can be used for inventory control. Since both the policy and bid price controls from the CMIP are implementable, this model could be utilized for either reservation system, as well as a reservation system that utilizes both solutions for pricing and capacity controls.

2.4.2 Comparison to Other Network RM Methods

Two common models currently being used in the industry include a stochastic network flow formulation and the EMSR-b model, discussed in the literature review. The network flow formulation represents one approach commonly used and is a popular O&D RM strategy. The network formulation, as seen in Appendix A of Jacobs *et al.* (2008), represents a stochastic passenger flow model, solved using a Lagrangian relaxation approach with a sub-gradient algorithm. The network flow formulation solves for the bid prices associated with each leg, and calculates the protection limit for each fare class on each leg using Littlewood's Rule.

The EMSR-b model represents a leg-based control strategy which estimates the bid prices using protection limits based on Littlewood's Rule. To account for the connecting traffic between O&D pairs, the revenue of the connecting fare was prorated and allocated to each leg in the O&D. The industry uses various versions of the EMSR-b and network formulation, making these industry models a good technique to compare against.

The EMSR-b model and the network formulation were solved for multiple passenger demand scenarios. The results show, using simulations of nine combinations of λ and T, that the CMIP performs better than both models. The CMIP outperforms the EMSR-b by 11.82% in mean revenue over all of the nine cases given in Table 2.6. A standard z-test on the difference of each set of means for the results shown in Table 2.6 illustrated that the differences in the expected revenue of all nine combinations are statistically significant, with p-values less than 0.001.

In some situations the CMIP simulation resulted in larger confidence intervals, but this is due to the highly segmented nature of fares. The difference between a sale and no sale is at least \$300 (in the case of the lowest fare for legs AB and BC), creating a large gap between revenues among simulations, yielding large standard deviations in relation to overall revenue. It is important to note, however, that as a larger amount of demand enters the network, the width of the confidence intervals for the CMIP reduces drastically; this is the exact opposite of the EMSR-b simulation, where the confidence intervals become wider as more demand enters the network.

For the network formulation, the gains were slightly less since the network formulation tends to perform better than the EMSR-b. The expected revenue showed an average increase of 9.60% over the nine cases presented in Table 2.7. Similar to the statistical tests of the EMSR-b, these nine cases show that the expected revenues are significantly different between the CMIP and network formulation, with p-values less

λ	T	CMIP 95% Confidence Interval	EMSR-b 95% Confidence Interval	% Increase Over Mean
	1	($0, $1,414)	($0, $903)	10.8
1	5	($228, $4,468)	($1,175, $3,191)	7.0
	10	($1,891, $7,213)	($2,692, $5,430)	10.8
	1	($261, $4,465)	($1,175, $3,191)	7.7
5	5	($6,959, $10,071)	($3,915, $9,663)	20.3
	10	($10,229, $12,199)	($7,834, $11,488)	13.8
	1	($1,891, $7,221)	($2,692, $5,430)	10.9
10	5	($10,186, $12,220)	($5,164, $8,818)	13.8
	10	($13,174, $13,772)	($10,069, $13,799)	11.4

Table 2.6: Expected Revenue Confidence Intervals by Bid Pricing Controls Simulation (CMIP vs. EMSR-b)

than 0.001. All three model simulations were run together, yet we chose to display the ESMR-b and network formulation results separately for easier comparison. Similar to the EMSR-b, the network formulation uses a segmented demand model for predicting the expected number of passengers. This causes the network formulation to open up the lower fare classes earlier than the CMIP does. Since the CMIP keeps the lower fare classes closed for a longer period of time, a higher overall revenue is earned.

The results show a significant performance difference between the four tested network RM models. Figure 2.2 illustrates the gains by each of the models as one increases the total number of passengers introduced into the system over the entire time horizon. At the lowest level of passenger demand, all four models behave similarly: they sell to any passenger that shows up. At the highest level of passenger demand, all the models again behave similarly: only sell to the highest paying pas-

λ	T	CMIP 95% Confidence Interval	Network Formulation 95% Confidence Interval	% Increase Over Mean
	1	($0, $1,414)	($0, $903)	10.8
1	5	($228, $4,468)	($1,169, $3,203)	6.9
	10	($1,891, $7,213)	($2,659, $5,521)	10.2
	1	($261, $4,465)	($1,169, $3,203)	7.5
5	5	($6,959, $10,071)	($3,949, $9,971)	19.3
	10	($10,229, $12,199)	($8,771, $12,517)	5.1
	1	($1,891, $7,221)	($2,659, $5,521)	10.2
10	5	($10,186, $12,220)	($8,791, $12,537)	5.0
	10	($13,174, $13,772)	($10,069, $13,799)	11.4

Table 2.7: Expected Revenue Confidence Intervals by Bid Pricing Controls Simulation (CMIP vs. Network Formulation)

sengers. However, as one moves from zero demand to a higher demand, the models begin to deviate from one another. The two dominating curves, the CMIP and CDLP, produces higher revenues compared to the EMSR-b and the network formulation. In fact, as seen in Figure 2.2, the CMIP and CDLP perform quite similarly.

2.5 Additional Examples

2.5.1 Small Network Instance

In addition to running the model on the three leg network seen above, we also tested it on another network given in Liu and van Ryzin (2008). This network, depicted in Figure 2.3, is a small 22 product network, consisting of 7 legs.

The network contains a direct flight from A to B, with competition from A to B through the hub, H. There are two flights from each of the direct legs between A, H,

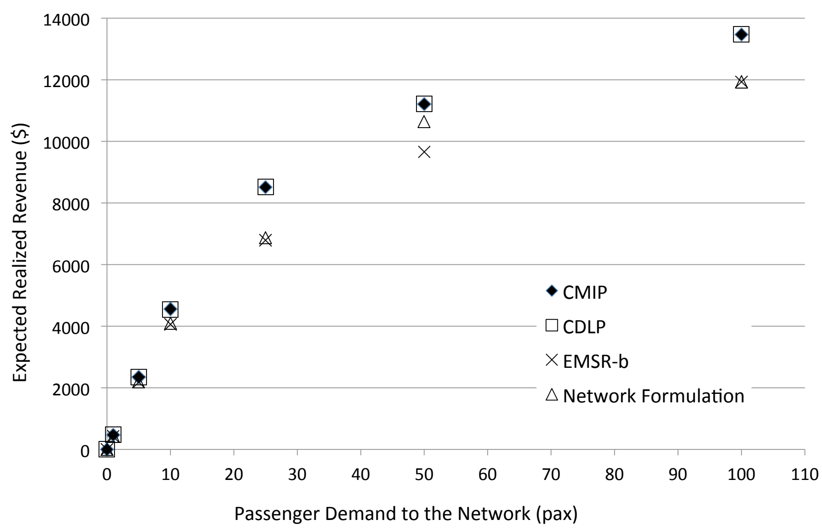

Figure 2.2: Passenger Demand to the Network Versus Expected Revenue Under Different RM Methods

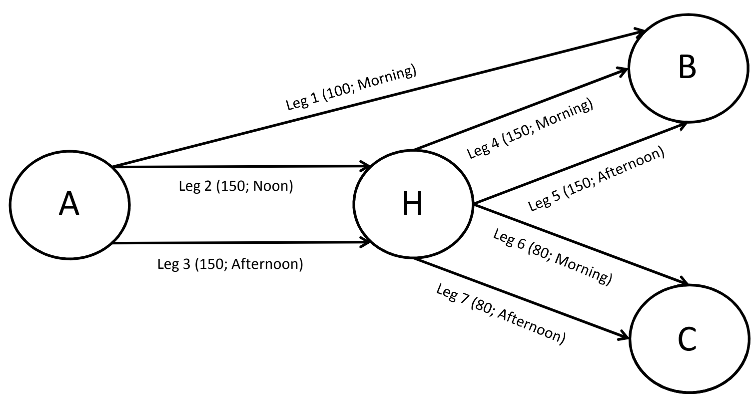

Figure 2.3: Small Network Instance - Adapted from Liu and van Ryzin (2008)

B, and C: an early flight and a later flight. The network instance data, including all of the MNL choice parameters, can be found in Tables A.1 and A.2 in the Appendix for the reader's convenience. The small network instance was ran for 1,000 time periods, with λ_t equal to 0.91 for each time period. This would represent a total of 910 customers introduced into the network.

We ran our model, and compared the results to those given in Bront *et al.* (2009) for the example in Liu and van Ryzin (2008), and saw that our model performed similarly to the CDLP. Across the five tests, each for a different fraction of total network capacity, our model performed 9.0% better, on average, than the CDLP, while still maintaining similar levels of network Load Factor (LF). The LF, defined as the average across all legs of the ratio of seats taken to total capacity, represents how many seats, on average, are consumed across the entire network. In addition to comparing it to the CDLP, we also compared it to the model solved using the independent demand assumption (referred to as the INDEP model) found in Bront *et al.* (2009), in which a deterministic linear program is solved with demand values generated under the assumption that all of the products are simultaneously open.

The first column of Table 2.8 indicates the percentage of the base capacity used for both the model solution as well as the simulation. This value represents an increase or decrease in the amount of capacity available, while maintaining the demand over the time horizon. The table includes the expected revenues obtained by the CDLP and CMIP solutions, as well as the percent increase they offer over the INDEP model. The table entries for CDLP and INDEP come from the simulation results reported in Bront *et al.* (2009). As the capacity in the network increases, the network load factor should decrease, as the demand introduced into the network does not increase, although the available space does. As seen in Table 2.8, the load factors decrease as the amount of capacity increases for all of the models, as expected.

Percent of	CMIP			CDLP			INDEP	
Base Cap.	Rev.	Inc. (%)	LF (%)	Rev.	Inc. (%)	LF (%)	Rev.	LF (%)
60	$224,114	30.0	98.5	$207,890	20.6	91.3	$172,362	97.7
80	$278,241	36.0	92.1	$261,264	27.7	85.6	$204,572	94.6
100	$297,752	31.7	83.6	$277,738	22.9	80.8	$226,002	87.7
120	$315,832	29.5	77.0	$282,842	16.0	71.6	$243,930	82.5
140	$318,153	22.8	70.2	$285,417	10.2	62.0	$259,039	77.0

Table 2.8: Expected Revenues and Percent Increase Over INDEP (Small Network Instance)

2.5.2 Large Network Instance

We finally use a large network instance with realistic aspects to further test the performance of the CMIP. The network structure, as well as the revenue values associated with each itinerary can be found in Jacobs *et al.* (2008). This network contains 48 legs, each with an initial capacity of 200 seats, joining 10 cities (see Figure 2.4). Each leg represents a single flight. There are no parallel flights available in this network. There are 178 itineraries, with three fare classes each, denoted as Y, M, and Q, for the given itinerary. There are a total of 90 O&D markets, with markets containing either one, two, or three possible paths between O&D. For a single time period, the CMIP model has a total of 3,598 variables. Note that for this example, the CDLP would result in a total of $2^{534} - 1$ (or, about 5.6E160) variables. The total number of arrivals to the system was set to 9,750 for a single time period. A single time period was used to see a representation of a solution for an entire booking horizon. The MNL values associated with each choice were arbitrarily generated, as well as the individual arrival probabilities per market. These values can be found in three tables (Tables A.3, A.4, and A.5) given in the Appendix, and are separated by how many competing itineraries existed between O&D, for easier classification. The preference vectors in Table A.3 represent the utility of the three fare classes, Y, M, and Q. The preference vectors in Table A.4 represent the utility of the three fare classes for each of the two itineraries available for that market. The first three utility values refer to the Y, M, and Q fare classes of the first itinerary, while the last three utility values refer to the Y, M, and Q fare classes of the second itinerary. Similarly, Table A.5 displays the preference vectors for the three itineraries with respect to the Y, M, and Q fare classes.

As indicated by the results given in Table 2.9, the time required to solve this

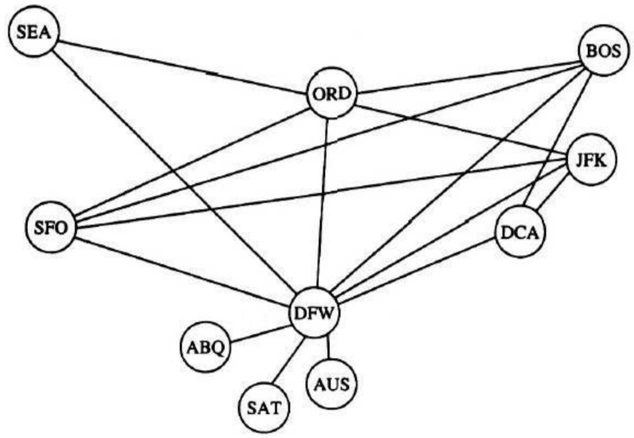

Figure 2.4: Large Network Instance, Adapted from Jacobs *et al.* (2008)

model is quite reasonable. As the capacity becomes more constraining, the model needs more time to find an optimal solution. However, it can still be solved in a reasonable amount of time. Since the CDLP is expected to have a total of 5.6E120 variables, we did not program the decomposition approach presented in Bront *et al.* (2009), as it would have proved to be computationally prohibitive.

In addition to the revenue and network LF, the available seat mile (ASM) and revenue per available seat mile (RASM) are also reported. The ASM is calculated by the number of seats available on a leg, multiplied by the distance traveled by the flight on that leg, then summed for all flight legs. The RASM is the total expected revenue divided by the ASM. This value is often used in the industry, and although the values seen in the table come from a fabricated instance, the RASMs are in-line with what is seen in industry practice.

2.5.3 Computational Complexity

Many of the existing RM approaches, such as the dynamic program proposed in Talluri and van Ryzin (2004a) generate an exponential number of solutions by

Percent of	CMIP				
Base Capacity	Revenue	LF (%)	ASM	RASM	Elapsed Time (sec.)
60	$1,018,043	82.2	6,380,160	$0.16	142.87
80	$1,146,655	78.5	8,506,880	$0.13	66.04
100	$1,231,055	70.6	10,633,600	$0.12	3.07
120	$1,283,682	62.6	12,760,320	$0.10	3.65
140	$1,322,589	55.2	14,887,040	$0.09	1.59

Table 2.9: Expected Values for the Large Network Instance

explicitly combining market policy strategies together across the network. This is problematic for industry use, as networks and fare class buckets have grown to create many itineraries. One example of this complexity issue is apparent in the CDLP's solution set, S. Since set S contains all the fare class and O&D controls for the network, S is dependent on the network definition, including prices and control strategies. For example, a small three node, four itinerary, two fare class network yields 255 decision variables for the CDLP formulation. However, if we increase the complexity of the network to four nodes, seven itineraries, and keep the two fare classes, the model has a total of 16,383 decision variables. This number continues to grow exponentially when any parameter of the network is increased, which can be seen in Table 2.10. This table illustrates the growth in complexity of the CDLP versus that of the CMIP for the previous examples. As one can see, the CMIP doesn't increase in size as fast, which would allow for consideration of being tractable for industry use.

2.6 Conclusions and Future Work

The proposed CMIP formulation uses an MNL model to explicitly model the impact of network-wide offerings on the probability of purchase to better reflect cus-

Network Instance	Number of Products	Number of Variables CMIP	CDLP
Three Leg Example	8	24	255
Small Network Example	22	116	4.2 x 10^6
Large Network Example	534	3598	5.6 x 10^{160}

Table 2.10: Variable Complexity of CMIP and CDLP for a Single Time Period

tomer behavior. Using problem instances of varying size, we have shown that CMIP outperforms both the EMSR-b and a basic network formulation. Another model used as a benchmark for performance was the CDLP, which utilizes the same MNL model to develop its probability of purchase but yields a solution that is somewhat difficult to decipher and implement. The CMIP in comparison offers an easy interpretation of its solution. The CMIP and CDLP performed similarly in both model solutions and simulation. The advantage of the CMIP, however, is based on the model's complexity. The CMIP is much smaller in size and easier to solve in most cases. As the network becomes more complex, the CMIP does not exponentially increase in size as the CDLP does.

From a pragmatic perspective, the CMIP approach builds on the advantages of previous models by addressing passenger choice in a computationally more efficient manner. Future work includes full scale tests of the approach and calibration of the passenger choice model needed to drive the optimization. Another aspect for future research includes the consideration of non-stationary demand to incorporate variation in the market demand over time. Other areas include expanding the model to handle bookings of multiple passengers at once, and time dependent demand utilities and pricing.

With the results from the examples and the possibility of many industry specific extensions, the CMIP looks promising for future research. The CMIP could improve on the models currently being used by leading airline companies today as well as be the groundwork for further development in the area of RM. Utilizing choice modeling and mathematical programming, the choice-based mixed integer program successfully optimizes the network RM problem for the airline industry. Further development of this model could assist in changing the way the industry solves their network RM problems.

Chapter 3

A MULTINOMIAL LOGIT FRAMEWORK FOR AIRLINE TICKET ATTRIBUTE AND PRICE SENSITIVITIES

3.1 Introduction

Airline revenue management (RM) research is typically divided into two categories: independent and dependent demand models (Talluri and van Ryzin, 2004b). Independent demand models, originally formulated in the mid 70's, assume demand for tickets can be segregated by price ranges, called fare classes, and that the probability of purchase is independent of which tickets are available at the time of purchase. Dependent demand models for airline RM, on the other hand, assume demand for tickets is dependent on what is offered at the time of purchase, and account for fare class competition within similar origin-destination combinations through utility-based choice models. Outclassed by fewer assumptions and a better representation of passenger purchasing behavior, independent demand models have been virtually replaced by dependent demand models in academic research. Despite the shift towards dependent demand in academic research, many industry applications still utilize models based on independent demand assumptions (Talluri and van Ryzin, 2004b).

Most dependent demand RM models utilize static assumptions on purchasing utilities and probabilities of purchase, failing to account for the diverse nature of purchasing behavior present in today's competitive airline industry. These RM models, mostly centered around linear and dynamic programming techniques, fail to address the details of the demand models and their impact on the solutions set forth by the RM methodology. These RM models make assumptions on purchasing behavior, such as static pricing options and known utilities, without considering the truly dynamic nature of airline ticket purchases. In this chapter, we introduce a framework for airline revenue management demand modeling, incorporating multinomial logit (MNL) choice integrated with price and ticket attribute sensitivities. This framework removes the static assumptions traditional dependent demand research makes on airline networks, and directly incorporates important attributes passengers focus on when making purchasing decisions.

The remainder of this chapter is organized as follows. In Section 3.2 we formally introduce the multinomial logit choice model and discuss how RM addresses ticket utility. In Section 3.3, we highlight the diversity of ticket attributes present in today's airline networks and discuss the importance of pricing decisions in RM. Section 3.4 introduces the MNL framework incorporating price and ticket attributes, providing examples of potential application on a real network. Section 3.5 reiterates the importance of including price and ticket attributes in RM decision making, highlighting the direction our framework can move RM research.

3.2 Multinomial Logit Choice Model

Multinomial logit choice models utilize a ratio of predicted values, often called utilities, to determine the likelihood of purchases given a set of available options. More specifically, multinomial logit models determine the probability of selecting a particular option out of a set of alternative options by a ratio of exponential terms in a logistic equation. The set of alternative options, θ, can be indexed 1 to M, such that $m \in \{1, 2, ..., M\}$ represents a particular option. For all $m \in \theta$, we can then model the probability an arbitrarily chosen person will select a particular option m

as

$$P(y = m|\theta) = \frac{e^{\boldsymbol{x_m}'\boldsymbol{\beta}}}{\sum_{m=1}^{M} e^{\boldsymbol{x_m}'\boldsymbol{\beta}}}, \quad \forall\, m \in \theta, \tag{3.1}$$

where $\boldsymbol{x_m}$ is a vector of attributes for option m and $\boldsymbol{\beta}$ is a vector of sensitivities for each attribute present in any option. We are interested in modeling N attributes, thus $\boldsymbol{\beta} = (\beta_1, \beta_2, \ldots, \beta_N)$, where each element of $\boldsymbol{\beta}$ refers to the sensitivity of an arbitrarily chosen person for that attribute. Although not explicitly modeled here, a no- purchase option is often included to account for the probability that no purchase is made. This option can be incorporated by introducing a utility value for not purchasing any option, which is called v_0. In this case, the probability that option m is selected becomes

$$P(y = m|\theta) = \frac{e^{\boldsymbol{x_m}'\boldsymbol{\beta}}}{\sum\limits_{i=1}^{M} e^{\boldsymbol{x_i}'\boldsymbol{\beta}} + v_0}, \quad \forall\, m \in \theta. \tag{3.2}$$

This method of incorporating the no-purchase option provides freedom pertaining to the probability no option is selected, since v_0 could be set to any value that fits the data.

Airlines have created their own unique product classifications to market tickets to the general public. These different products are made up of distinct ticket attributes, and play a role in how passengers choose to purchase tickets. For our purposes, the model expressed in Equation 3.1 is characterized by these ticket attributes including, price, time of day, path from origin to destination, and number of connections. Due to the wide array of product classifications, a useful MNL model would have to incorporate all relevant ticket attributes into the equation, creating a detailed demand forecast based on each airline's product classification.

Fitting these models can be computationally difficult, but methods exist for determining the values of $\boldsymbol{\beta}$ based on sales data. Currently, the best method for fitting these models is a two-step parameter estimation technique developed by Newman *et al.* (2014), which utilizes a log-likelihood expectation maximization technique with linear separability. Their method efficiently finds the parameters for the utility model, in addition to a no-purchase utility, resulting in a practical application of choice-based demand modeling for the revenue management. Traditional methods, like the expectation maximization and log-likelihood maximization techniques found in Talluri and van Ryzin (2004a) also work, but can be computationally prohibitive, making it difficult for airline's to implement for practical use.

When it comes to MNL choice models, revenue management literature has taken a simpler approach to the problem by reducing the complexity of the MNL equation. The exponential values, $e^{\boldsymbol{x_m}'\boldsymbol{\beta}}$, are interpreted as product utilities, where these utilities become the parameters for nearly all dependent demand revenue management models. These revenue management models reduce the complex nature of the MNL equation, assuming the form of the equation is known, by reducing the problem into a ratio of utilities of the form

$$P_m = \frac{u_m}{\sum_{i=1}^{M} u_i + v_0}, \quad \forall\, m \in \theta, \tag{3.3}$$

where u_m represents the utility for product m, v_0 represents the no-purchase utility, and P_m represents the probability of purchase for product m (Talluri and van Ryzin, 2004a; Gallego *et al.*, 2004; Bront *et al.*, 2009). The assumption that allows for this simplification of the MNL model deals with static prices and fare class offerings: each fare class has a price that cannot be changed and each origin-destination path can have multiple fare classes. With these assumptions, the static-utility models predicting passenger behavior are accurate, resulting in implied prices based on ticket offering decisions.

Due to this simplification, though, most choice-based revenue management models in existence are unable to solve for price and ticket availability, despite the impact these decisions have on passenger choice, especially pertaining to those attributes with the biggest effect (such as price). Under the static environment, revenue management models make no attempt at balancing price alongside ticket attributes, leaving a gap between true purchasing behavior and modeled purchasing behavior. Consequently, the resulting bid prices established by revenue management models based on these static-utility assumptions fail to take into account the impact of price on passenger preference, despite being the primary mechanism for pricing in the airline industry.

3.3 Current Airline Implementation of Ticket Attributes and Pricing

As previously mentioned, every airline has established its own product classification, as seen by the wide array of ticket displays. Based on these ticket displays, a ticket is made up of the product classification and specific flight information. While mostly a marketing decision, these product classifications have a large impact on the way passengers purchase tickets, implying the attributes included in each product should play a role in how the MNL model is characterized. Thus, identifying the attributes associated with the products, and subsequently including these attributes into our MNL model, would lead to a fine-tuned choice-based demand model, eventually leading to better ticket availability and pricing decisions under a RM methodology.

We can compare major airlines and their ticket displays side-by-side, and see the wide array of products passengers have to consider when purchasing a ticket. Despite providing the same fundamental service, a flight from an origin to a destination, American Airlines (AA) (Table 3.1) appears to have more products than its competitor, Delta Airlines (DA) (Table 3.2), with American Airlines offering seven different products compared to that of Delta Airline's three. These products, made up of different attributes, segments their demand in a way each airline has decided upon. American Airlines spreads out the attributes by offering products with flexibility versus non-flexible attributes, and offers more product options differentiated by refundability. Delta Airlines, on the other hand, categorizes their products by cabin, separating out the basic economy from the Main Cabin, eventually leading to first class. The fundamental service, a flight from an origin to a destination, is the same, despite the varying products offered by each airline.

The differences in attribute selection for products are even more pronounced if we compare AA to Southwest Airlines (Table 3.3) or Jet Blue (Table 3.4). Southwest Airlines and Jet Blue have fewer product classifications, while still providing the same flight options. Southwest Airlines focuses on product classifications based on attributes as opposed they have deemed important for each demand segment, while Jet Blue has taken a simplistic approach by offering only two products. All of the products offered by these airlines are made up of different ticket attributes, like refundability and number of free bags, and impact the purchasing behavior of potential passengers. Despite this impact, dependent demand RM models fail to address the different attributes for each product, and adhere to their static utility models, creating a gap between literature and implementation.

Ticket Attribute	Lowest Fare		Refundable			Business/First	
	Choice	First	Choice	Fully Flexible	First Flexible	First	First Flexible
Refundable	No	No	Yes	Yes	Yes	No	Yes
Transferable	No	Same Day	Yes	Yes	Yes	Same Day	Yes
Priority Check-in	No	Yes	No	Yes	Yes	Yes	Yes
Priority Boarding	No	Yes	No	Yes	Yes	Yes	Yes
Free Bags	0	3	0	2	3	3	3
Bonus Miles	No	Yes	No	No	Yes	Yes	Yes

Table 3.1: American Airlines Product Attributes (American Airlines, 2016)

Ticket Attribute	Basic Economy	Main Cabin	First/ Business
Refundable	No	With Fee	With Fee
Transferable	No	With Fee	With Fee
Priority Boarding	No	Yes	Yes
Seat Assignments	At Check-in	At Purchase	At Purchase
Loyalty Program Benefits	No	Yes	Yes
Upgraded Food Options	No	No	Yes

Table 3.2: Delta Airlines Product Attributes (Delta Airlines, 2016)

Ticket Attribute	Wanna Get Away	Anytime	Business Select
Refundable	No	Yes	Yes
Transferable	With Fees	Yes	Yes
Priority Boarding	No	No	Yes
Priority Security Lane	No	No	Yes
Loyalty Miles	Regular	Bonus	Bonus
Free Bags	2	2	2
Complementary Drinks	No	No	Yes

Table 3.3: Southwest Airlines Product Attributes (Southwest Airlines, 2016)

Ticket	Lowest	Refundable
Refundable	No	Yes
Transferable	With Fees	Yes
Priority Boarding	With Fees	With Fees
Free Bags	1	1

Table 3.4: Jet Blue Product Attributes (Jet Blue Airlines, 2016)

Since demand for airline tickets comes from the general population, it is apparent that demand for a specific ticket comes from a common customer pool, indifferent of airline. Although some airlines have frequent fliers, passengers typically focus on the price of a ticket before other attributes, including the airline of choice. Thus, each airline should be generating a unique choice-based demand model based-off of their own ticket display, as these ticket displays define the attributes assigned to each product. To do this, a general choice-modeling framework incorporating ticket attributes and price must be developed.

3.4 MNL for Ticket Attributes and Pricing

To incorporate ticket attributes and price into the multinomial logit choice model, we consider a network of airports connected with legs $l \in L$, and multiple demand markets defined by the set J, such that a market $j \in J$ is defined by an origin and destination combination. Each customer is assigned to an origin-destination market, j, determined by a probability λ_j, where each customer arrives from a common demand pool with an overall arrival rate per time period of γ_t. For each market j, we have a super-set of ticket availability matrices which we refer to as I_j. This super-set is made up of multiple ticket availability matrices, $i \in I_j$, where i represents a single ticket matrix availability for market j. For each ticket availability matrix i, we have an offer set of available tickets, K_{ji}, such that $k \in K_{ji}$ represents a ticket in ticket availability matrix i for origin-destination market j.

To better explain this set notation, consider the example for the American Airlines flights from PHX to JFK, seen in Figure 3.1. We would define the set of legs, L, as {PHX-JFK,PHX-DFW,DFW-JFK,PHX-CLT,CLT-JFK}, as these are the only legs present in this example. The market, j, defined by the origin-destination combination, would be PHX-JFK. For this example, there would only be one market, so J would be made up of one element. Figure 3.1 depicts one ticket availability matrix; the available tickets are offered at a price, and the unavailable tickets are indicated by

Phoenix, AZ → New York, NY (JFK Airport)

Date: Monday, March 02, 2015

Flight # / Carrier	Depart	Arrive	Stops	Travel time	Coach Non-refundable	Coach Flexible	Business Non-refundable	Business Flexible
425	9:58 AM PHX	4:54 PM JFK	(0)	4h 56m Seats	$374	$565	Not Available	Not Available
630	1:50 PM PHX	8:30 PM JFK	(0)	4h 40m Seats	$374	$565	$623	$814
622	4:30 PM PHX	11:09 PM JFK	(0)	4h 39m Seats	$374	$565	$623	$814
1336 Operated by American Airlines 64 Operated by American Airlines	6:10 AM PHX 10:30 AM DFW	9:23 AM DFW 2:53 PM JFK	(1)	6h 43m	$382	$573	$658	$2,113
1808 639	6:15 AM PHX 1:25 PM CLT	12:02 PM CLT 3:16 PM JFK	(1)	7h 1m Seats	$381	$572	$656	$847
1808 2017	6:15 AM PHX 2:20 PM CLT	12:02 PM CLT 4:05 PM JFK	(1)	7h 50m Seats	$381	$572	$656	$847

Figure 3.1: American Airlines Example Offer Set for PHX to JFK (American Airlines, 2015)

"Not Available". Note that in this example one can construct $2^{24} - 1$ different ticket availability matrices, that is, $i \in I_j$, where $I_j = \{1, 2, \ldots, 2^{24} - 1\}$, but most of these are invalid for revenue management applications. Nesting of ticket products would have to be preserved, so ticket availability matrices that have products like "Coach Non-Refundable" available while making products like "Coach Flexible" unavailable would be dismissed. Nesting of ticket products greatly reduces the number of ticket availability matrices, while preventing bizarre combinations airlines refuse to offer. Each flight would have a total of seven distinct nested product offerings, resulting in a total of $6^7 = 279,936$ ticket availability matrices, compared to the $2^{24} - 1 = 16,777,216$ ticket availability matrices without nesting.

To see an alternative ticket availability matrix, note that copying the set of ticket availability displayed in Figure 3.1 and replacing Flight 425's Coach Non-refundable ticket to "Not Available" would produce a unique alternative ticket availability matrix that preserves the nesting of ticket options. Traditional revenue management models focus on optimizing the availability of these tickets, thus we denote the tickets listed in i as the set K_{ji}, where $k \in K_{ji}$ indicates a specific ticket, such as Flight 622, Coach Flexible. From the ticket display in Figure 3.1, we can see that a ticket could be defined by the intersection of a flight and product, and that there are a total of 24 tickets, some of which are not available. Thus, each ticket is made up of a combination of product attributes, as well as flight attributes.

With these definitions, we can incorporate price and other ticket attributes into the MNL model previously discussed. To do this, we must break down the MNL model by characterizing the regression component of our MNL, namely, $\boldsymbol{x_n'}\boldsymbol{\beta}$, into components more fitting for airline revenue management. As previously discussed, price is an important component to revenue management, thus separating price and its sensitivity from the remainder of the regression equation would create a model that can be used to solve for price directly, while considering the dynamic nature of purchasing behavior.

We can define a predicted response for ticket k as α_k, that represents the overall desire to purchase ticket k based on the various attributes. Under this definition, α_k can be characterized as

$$\alpha_k = Y_k\beta_p + \boldsymbol{x_k'}\boldsymbol{\beta}, \quad \forall\, k \in K_{ji}, \tag{3.4}$$

where β_p is the sensitivity to price for a ticket and Y_k is the price for ticket k. Similar to the MNL definition in Equation 3.1, $\boldsymbol{x_k}$ is a vector of attributes associated with ticket k, but without price, and $\boldsymbol{\beta}$ is a vector of sensitivities for these attributes, but without price. Based on Figure 3.1, if the two attributes we were concerned with were "Refundable" and "Number of Stops", $\boldsymbol{x_k} = (1, 0)$ for the "Coach Refundable", Flight 425 ticket since that ticket is refundable and has no stops. Conversely, the "Coach Non-refundable" ticket for Flights 1336 and 64 would yield an $\boldsymbol{x_k} = (0, 1)$

since that ticket is non-refundable and has a stop. Although this is a small example, any ticket attributes could be incorporated into $\boldsymbol{x_k}$ in this manner. These ticket attributes could be departure time, number of stops, or travel time, in addition to the ticket attributes characterized by American Airline's Coach Non-Refundable, Coach Refundable, Business Non-Refundable, and Business Refundable products.

To compare this to the static-utility MNL format, we can define u_k as the utility of ticket k, evaluated as

$$u_k = e^{\alpha_k}, \quad \forall\, k \in K_{ji}. \tag{3.5}$$

As opposed to the static utility model, u_k has dependency on ticket price and attributes, thus any revenue management model utilizing these utilities must dynamically determine the utilities while solving for ticket price and availability.

Implementation of this model is rather straight forward, and we can utilize it to determine probabilities of purchase for any demand component. We can calculate the probability a customer from market j will purchase a ticket in the ticket availability matrix i, P_{ji}, as

$$P_{ji} = \frac{\sum_{k \in K_{ji}} u_k}{\sum_{k \in K_{ji}} u_k + v_j}, \quad \forall j \in J, i \in I_j, \tag{3.6}$$

where v_j is the no-purchase utility for origin-destination market j. Since the MNL model deals with availability when determining purchase probabilities, the probability a ticket is purchased is directly related to what tickets are available at that time. Based on a given ticket availability matrix i, we can also calculate the probability a customer purchased ticket k as

$$Q_{k|ji} = \frac{u_k}{\sum_{k \in K_{ji}} u_k + v_j} \left(\frac{1}{P_{ji}} \right), \quad \forall\, j \in J, i \in I_j, k \in K_{ji}. \tag{3.7}$$

This value, $Q_{k|ji}$, could then be used for a revenue management model that takes price and ticket attributes into account when determining ticket availability based solely on ticket purchase probabilities, rather than ticket utilities. Since $Q_{k|ji}$ is dependent on price, as well as other attributes, the relationship between ticket price and purchase probability is accounted for, as opposed to the static utilities of previous revenue management techniques.

3.4.1 *Implementation of MNL Framework*

We can implement this framework and display its usefulness with an example of two competing airlines. Suppose we have two flights from PHX to JFK for each airline, American Airlines and Delta Airlines, with product definitions, ticket attributes, and pricing seen in Table 3.5. Under static utility models, AA and DA would expect similar utilities for the purchase of each ticket, since they are offering similar products: a ticket from PHX to JFK at similar times. These similar utilities would result in similar probabilities of purchase, under the static utility models, which should result in similar revenue management policies.

Airline	Departure Time	Number of Stops	Ticket Purchase Options & Price	
AA	9:00 AM	0	Lowest Fare - Choice $250	Refundable - Choice $325
AA	4:30 PM	1	Lowest Fare - Choice $265	Refundable - Choice $340
DA	9:15 AM	0	Basic Economy $235	Main Cabin $340
DA	5:00 PM	1	Basic Economy $270	Main Cabin $335

Table 3.5: Ticket Options for PHX to JFK

In reality, though, the ticket offerings for each airline are different, due to the different product classifications each airline has taken. The Lowest Fare - Choice product for AA contains no special ticket attributes, similar to the Basic Economy product of DA, whereas the Refundable - Choice product contains a mixture of ticket attributes like Refundable and Transferable, similar to the Main Cabin product of DA. There are differences, though, between the products of each airline, and these differences could impact the probability of purchase. We can use the MNL framework described above to generate utilities dependent on ticket attributes and price based on a given set of sensitivities, taking into account these different product classification schemes. For instance, consider the attribute sensitivities seen in Table 3.6, which were entirely fabricated.

Using these sensitivities for both airlines and the market population, we can generate utilities for each ticket for both AA and DA based on our MNL framework, seen in Table 3.7, and compare each airline's probabilities of purchase assuming all four tickets are being offered for each airline (Table 3.8). In this example, we assume the utility for no-purchase, that is a passenger choosing not to purchase a ticket, is zero. This implies that when a passenger shows up to purchase a ticket, a purchase will always be made. Since each airline would develop their demand models independently of one another and the probability of no-purchase is zero, the probabilities of purchase within each airline should sum to one.

As Table 3.8 suggests, the probabilities of purchase for each airline are substantially different despite offering similar ticket options. AA's demand is somewhat evenly distributed between the four options, with the later flight Lowest Fare - Choice product yielding the smallest probability of purchase. DA's demand heavily favors the Main Cabin product, primarily due to the priority boarding and seat assignments at purchase, with more than 10 times the probability of purchasing these Main Cabin products versus their Basic Economy counterparts.

Since the majority of revenue management models utilize these probabilities of purchase, implementation of the MNL framework suggested here would generate

Ticket Attribute	Attribute Sensitivity
Departure Times	
00:00 - 06:00	+1.5
06:01 - 12:00	+2.0
12:01 - 18:00	+1.7
18:01 - 24:00	-0.2
Ticket Changes	
Transferable (Free)	+0.60
Transferable (With Fee)	+0.25
Refundable	+2.1
Day of Flight	
Priority Check-in	+0.50
Priority Boarding	+0.75
Number of Free Bags	+0.50 per bag
Loyalty Program Perks	+0.25
Other Attributes	
Seat Assignment (At check-in)	+0.0
Seat Assignment (At purchase)	+2.1
Number of Stops	-0.5 per stop
Ticket Price	-0.01 per $

Table 3.6: Ticket Attributes and Passenger Sensitivities

Airline	Departure Time	Number of Stops	Ticket Options & Utility	
AA	9:00 AM	0	Lowest Fare - Choice 4.95	Refundable - Choice 4.71
AA	4:30 PM	1	Lowest Fare - Choice 2.59	Refundable - Choice 3.49
DA	9:15 AM	0	Basic Economy 0.70	Main Cabin 14.15
DA	5:00 PM	1	Basic Economy 0.45	Main Cabin 9.03

Table 3.7: Calculated Ticket Utilities for PHX to JFK

Airline	Departure Time	Number of Stops	Ticket Options & Purchase Probability	
AA	9:00 AM	0	Lowest Fare - Choice 0.31	Refundable - Choice 0.30
AA	4:30 PM	1	Lowest Fare - Choice 0.16	Refundable - Choice 0.22
DA	9:15 AM	0	Basic Economy 0.03	Main Cabin 0.58
DA	5:00 PM	1	Basic Economy 0.02	Main Cabin 0.37

Table 3.8: The Probabilities of Purchase Implied by the Utilities Given in Table 3.7

unique ticket availability and pricing solutions for AA and DA, whereas the static utility models currently implemented would result in similar ticket availability and pricing for each airline. The differing attributes for each product and inherent attributes within each ticket create unique purchase probabilities that our framework can determine, without over-complicating the MNL demand model and tailoring it to an individual airline's use. These purchase probabilities can then be fed into a revenue management model, leading to unique policies for each airline.

3.4.2 *Solving for RM Policies Considering Ticket Attributes*

In this section, we demonstrate the impact of using a carefully constructed MNL model that more accurately reflects customer preferences. To do this, we use the utilities developed in the previous section to solve the Choice-based Mixed Integer Program (CMIP) from Chapter 2. Solving the CMIP for both AA and DA, separately, will highlight the impact these differing utilities and ticket attributes have on the revenue management controls, which ultimately impact expected revenue.

We assume each itinerary (four tickets for each airline) only has a single seat available, and a total of five potential passengers are arriving ($\gamma_t = 5$) in a single time unit ($T = 1$). Since each airline has two flight options, this means there is a total of two seats available per airline when a passenger shows up to purchase a ticket. The earlier flight for each airline has no stops, as indicated by Table 3.9, while the later flights have a single stop. We only consider a single seat on each of these flights to represent the immediate decision an airline would have to make for a set of potential passengers. To this end, we solve the CMIP and simulate the results in the same fashion as Chapter 2.

Since we are only concerned with the PHX-JFK market, the connecting flight implied by the single stop for the later flights and its competitive demand is ignored. Although this example is simple, it will allow us to show the impact different revenue management controls can have on identical demand streams considering we have different utilities for each airline, as opposed to traditional static-utility based revenue management models. The results from the CMIP model are summarized below in Table 3.9, in a manner the airline industry would typically use.

Table 3.7 makes it readily apparent that the price and attributes included in a ticket can play a large role in the distribution of utilities, whereas Table 3.9 reinforces

Airline	Departure Time	Number of Stops	Ticket Options & Purchase Probability	
AA	9:00 AM	0	Lowest Fare - Choice Not Available	Refundable - Choice $325
AA	4:30 PM	1	Lowest Fare - Choice Not Available	Refundable - Choice $340
DA	9:15 AM	0	Basic Economy Not Available	Main Cabin $340
DA	5:00 PM	1	Basic Economy $270	Main Cabin $335

Table 3.9: Ticket Availability for PHX to JFK Tickets Based on CMIP Solution

this impact by demonstrating the differing revenue management policies employed by each airline. Under the utilities in Table 3.7, AA earns an expected revenue of $577 for their two seats, and chooses to close the lower priced tickets for each of their flights. DA earns an expected revenue of $562 for their two seats, slightly less than their AA counterpart, while choosing to close only one ticket option, leaving the other three options open for purchase. Although there is a difference in expected revenue, a total of 2.5%, the important piece of this result is the different policies. An identical demand stream was used for each airline, yet, due to the attributes included in each product and ticket, their policies are substantially different. These differences would play a huge role in revenue generation, and would only be compounded with a larger network.

3.5 Conclusion

The use of dependent demand models is clearly the direction revenue management should continue to move, but the assumptions of static utility models prevent RM practitioners from properly modeling purchase behavior, producing sub-optimal ticket availability and pricing decisions. We have introduced an MNL framework that incorporates ticket attributes, including price and other ticket options, better representing the true purchase behavior of customers and opening the door for future revenue management models to incorporate better estimates of ticket utilities.

By directly incorporating price and common ticket attributes into the MNL framework, our framework allows for RM models based off of responsive demand models to be developed, hopefully leading to RM models that can manage both ticket pricing as well as ticket availability. These two decisions, pricing and ticket availability, are the crux of an airline's revenue management system, but are often solved separately due to the complex nature of purchasing behavior. This framework gives researchers and practitioners of RM an opportunity to formulate RM models with more accurate demand predictions while taking into account all ticket attributes designated by the airline, as well as the products classified by the airline's ticket display. The flexibility of regression equations in the MNL framework allows each airline to generate their own utility models based on their products, generating a unique demand stream for optimization.

Additional applications of this framework, outside of RM, could be focused around product definitions. Each airline has chosen attributes to generate their product classifications, but don't necessarily consider how these decisions impact revenue. Utilizing this framework, a model could be developed that determines which attributes to incorporate into a product definition, while considering purchase behavior based on these attributes. A model of this type would bridge the gap between marketing decisions and revenue management, fueled by the MNL framework suggested here.

The next step for RM is implementing this framework into a model that can effectively handle both ticket availability and pricing decisions, while dynamically determining demand based on these decisions. Incorporating the nature of passenger preference based on price and other ticket attributes directly into the RM model allows for better demand estimates and the resulting controls airlines use for ticket display and pricing. Additionally, this MNL framework could be used as a marketing tool, generating demand estimates for different assignments of ticket attributes to products. One could generate probabilities of purchase to determine the distribution of demand based on different attributes, selecting the product classification that maximizes either revenue or overall market share.

Chapter 4

ADDRESSING TICKET ATTRIBUTE AND PRICE SENSITIVITIES IN CHOICE-BASED REVENUE MANAGEMENT

4.1 Introduction

In general, airline Revenue Management (RM) is centered around the control of inventory by focusing on ticket availability or seat protection levels. In traditional RM, important factors such as price are incorporated by implementing "fare class buckets", representing subsets of the market willing to purchase at a range of prices. To address the complex nature of purchasing behavior, practitioners of RM have increased the number of fare classes for each ticket type in an attempt to capture the widest range of potential purchases. This proliferation of fare class buckets has created an overly-complex environment, forcing the industry to take sub-optimal approaches in managing their inventory levels.

In this chapter, we introduce a new model that eliminates the need for fare classes within product classification, and addresses the nature by which customers purchase tickets. By manipulating prices while simultaneously assigning availability, our model is able to incorporate customer sensitivities to common ticket attributes, such as price and travel perks, while maximizing revenue in a choice-based demand setting. Contrary to other choice-based demand RM models, our dynamic-pricing model doesn't rely on complex fare class structures, and can set prices explicitly without the added complexity of the current fare class system. Our results indicate that price is an important factor governing purchasing behavior, and that traditional RM models aren't satisfactory in modeling this sensitivity with fare classes, and thus generate sub-optimal solutions.

The remainder of this chapter is structured as follows. In Section 4.2 we provide a brief literature review centered around choice-based models to create the framework of traditional RM, with some emphasis given to the process of pricing products. In Section 4.3 we introduce the Price-Dynamic Choice-Based Mixed Integer Program (PCMIP) alongside the demand model definition and problem framework. In Section 4.4, we solve the PCMIP on problem instances of varying size and complexity, and compare them against the Choice-based Mixed Integer Program found in Chapter 2. Section 4.5 contains a short summary of our work, conclusions drawn from the results, and some insight into the future of pricing and its importance within RM.

4.2 Literature Review

Revenue management is not a new field of research by any stretch of the imagination. Older models, focused around independent demand assumptions, have been developed and used for decades and built the framework we now use to further advance the field. Our research, although developed on the backbone of original RM research, is primarily focused on choice-based demand and the subsequent pricing of tickets, thus the literature review will be focused around these principles.

4.2.1 Revenue Management Models

Some of the groundbreaking work completed in choice-based revenue management is found in Gallego *et al.* (2004) and Talluri and van Ryzin (2004a). In Gallego *et al.* (2004), the authors consider a network of flexible products, under both independent and dependent demand models, and introduce a choice-based deterministic linear program (CDLP) that closely approximates the stochastic optimization problem. Their linear program is easily solvable under independent demand assumptions, and they provide a framework for column generation under the more complicated dependent demand formulation (Gallego *et al.*, 2004). Talluri and van Ryzin (2004a) take a different approach and introduce a dynamic programming formulation that incorporates a general discrete choice model while determining optimal offer sets. The authors present an analysis of the efficient frontier, a subset of ordered policies for which all other policies are sub-optimal, in a single leg environment. Due to the difficult nature of parameterizing choice-models, Talluri and van Ryzin (2004a) also develop an expectation maximization technique for fitting the demand model, given information from a single firm.

Following the work of Gallego *et al.* (2004), Liu and van Ryzin (2008) expand on the linear programming formulation and investigate the effects of increased network complexity and the scaling of demand and capacity. The authors extend the previous notion of efficient sets into a network framework, and offer a heuristic to convert the static CDLP solution into a dynamic policy, solidifying its presence in the literature. Bront *et al.* (2009) also introduce a means to solve the CDLP by constructing a column generation algorithm under disjoint demand assumptions. Their column generation algorithm provides a strong approximation to the dynamic program found in Talluri and van Ryzin (2004a), while efficiently providing a means of solving the complex dependent demand version of the CDLP found in Gallego *et al.* (2004). Due to the size of the CDLP, Gallego *et al.* (2011) introduce a sales-based linear program that reduces the number of variables under a general attraction model, while still converging to the same expected revenue as that of the CDLP. Similarly, Clough *et al.* (2014) offer an alternative formulation in the form of a mixed integer program. Clough *et al.* (2014) consider a multinomial logit choice model in an airline network setting, and assume market demand in which origins and destinations are separate can be solved independently. Their model, the Choice-based Mixed Integer Program (CMIP), offers a smaller solution space under these assumptions, and maintains revenue performance compared to the CDLP and traditional airline control mechanisms.

4.2.2 Dynamic Pricing Models

Dynamic pricing in a choice-based environment is a relatively new area of research within RM. Prior to choice-based demand models being in the spotlight, some research had focused on the importance of pricing, as seen in Jacobs *et al.* (2010), where the authors investigate the relationship between pricing and revenue management controls. Jacobs *et al.* (2010) consider the impact capacity has on the dynamic pricing problem, creating a "price balance statistic" used for evaluating the quality of a strategy and finding the optimal mix of pricing, scheduled capacity, and RM controls.

Choice-based dynamic pricing research has taken a different approach. Aydin and Ryan (2000) consider a retail setting where consumers choose products based on a pre-defined selection and pricing setup. The authors examine multiple environments including the introduction of new products with associated pricing and the selection of a pre-selected set of product and price options. Zhang and Cooper (2009) introduce a Markov decision process to address dynamic pricing in an MNL environment. They consider multiple substitutable flights in a single O&D market, and show their model is intractable for realistic settings. In a similar approach, Dong *et al.* (2009) develop a dynamic programming formulation in a retail environment. They consider an environment with a long lead time and short selling environment, where the retailer must determine both inventory and pricing. More recently, Zhang and Lu (2013) introduce a dynamic programming formulation for dynamic pricing and offer a nonlinear programming approximation approach. They compare their methods against both static pricing models and other choice-based demand models, concluding that dynamic pricing could have substantial gains versus their static counterparts.

We believe these dynamic pricing models have only touched the surface of what we can achieve with pricing and revenue optimization. In the following section, we introduce a model that builds upon the previous literature, incorporating choice-based demand and price sensitivities directly into a price-dynamic revenue management model.

4.3 Model Formulation

Similar to the previous chapter, we consider a network of airports connected with legs $l \in L$, and multiple demand markets defined by set J, such that a market $j \in J$ is defined by an origin and destination combination. Customers arrive to the system in multiple time periods, indexed 1 to T, with a rate of γ_t, $t = (1, 2, \ldots, T)$. These customers are subsequently assigned to a market j, where λ_j refers to the probability a passenger is assigned to market j. For each market j, there is a super-set of policies defining ticket availability, I_j, where $i \in I_j$ represents a specific matrix of ticket availability for market j, similar to Chapter 3. For each ticket availability matrix i, there is an offer set of available tickets, K_{ji}, such that $k \in K_{ji}$ represents a ticket in ticket availability matrix i for origin-destination market j.

Since price is of concern, we must consider the possibility that prices can change with time. In Equation (3.6), time has no bearing on the probabilities of purchase. We can incorporate the time component into our MNL probabilities by including the possibility for prices to change in different times units. All other ticket attributes would remain the same, so the only adjustment that needs to be made is to index price by both time and ticket k, resulting in overall ticket sensitivities dependent on ticket attributes, price, and time period, characterized as

$$\alpha_{kt} = Y_{kt}\beta_p + \boldsymbol{x_k'}\boldsymbol{\beta}, \quad \forall\, k \in K_{ji}, t \in \{1, 2, \ldots, T\}, \tag{4.1}$$

where Y_{kt} represents the price for ticket k in time period t, and β_p, $\boldsymbol{x_k}$, and $\boldsymbol{\beta}$ have the same meaning as in Chapter 3. Since α_{kt} now has a time component, the utilities for each product would also have a time component, as seen in Equation (4.2).

$$u_{kt} = e^{\alpha_{kt}}, \quad \forall\, k \in K_{ji},\; t \in \{1, 2, \ldots, T\}. \tag{4.2}$$

Under this setup, we can modify the probabilities defined in Chapter 3 to incorporate time periods, resulting in the probability a purchase is made in market j in ticket availability matrix i in time period t, as

$$P_{jit} = \frac{\sum_{k \in K_{ji}} u_{kt}}{\sum_{k \in K_{ji}} u_{kt} + v_j}, \quad \forall\, j \in J, i \in I_j,\; t \in \{1, 2, \ldots, T\}, \tag{4.3}$$

and the probability a passenger purchases ticket k in time period t, as

$$Q_{kt|ji} = \frac{u_{kt}}{\sum_{k \in K_{ji}} u_{kt} + v_j} \left(\frac{1}{P_{jit}} \right), \quad \forall\, j \in J, i \in I_j, k \in K_{ji},\; t \in \{1, 2, \ldots, T\}. \tag{4.4}$$

These probabilities will allow us to incorporate price and ticket attribute sensitivities directly into a revenue management model. Building upon the complexity gains of the CMIP, we introduce the *Price-dynamic Choice-based Mixed Integer Program*

(PCMIP):

$$\text{Maximize} \quad \sum_{t \in T} \gamma_t \sum_{j \in J} \lambda_j \sum_{i \in I_j} P_{jit} \sum_{k \in K_{ji}} Y_{kt}\; Q_{kt|ji} \tag{4.5}$$

Subject to:

$$\sum_{t \in T} \gamma_t \sum_{j \in J} \lambda_j \sum_{i \in I_j} P_{jit} \sum_{k \in K_{ji}} Q_{kt|ji} A_{kl} \leq c_l, \quad \text{for all } l \in L, \tag{4.6}$$

$$\sum_{i \in I_j} X_{jit} \leq 1\ , \quad \text{for all } j \in J, t \in \{1, 2, \ldots, T\}, \tag{4.7}$$

$$P_{jit} = \frac{\sum_{k \in K_{ji}} u_{kt}}{\sum_{k \in K_{ji}} u_{kt} + v_j}, \quad \text{for all } j \in J, i \in I_j, t \in \{1, 2, \ldots, T\}, \tag{4.8}$$

$$Q_{kt|ji} = \frac{u_{kt}}{\sum_{k \in K_{ji}} u_{kt} + v_j} \left(\frac{1}{P_{jit}} \right),$$

$$\text{for all } j \in J, i \in I_j, k \in K_{ji}, t \in \{1, 2, \ldots, T\}, \tag{4.9}$$

$$\text{LB}_k \leq Y_{kt} \leq \text{UB}_k, \quad \text{for all } k \in K_{ji}, t \in \{1, 2, \ldots, T\}, \tag{4.10}$$

$$X_{jit} \in \{0, 1\}, \quad \text{for all } j \in J, i \in I_j, t \in \{1, 2, \ldots, T\}. \tag{4.11}$$

There are two decision variables for the PCMIP. X_{jit} represents the binary decision to select ticket availability matrix i for origin-destination market j in time period t, which allows any ticket $k \in K_{ji}$ to be purchased. Y_{kt} represents the price for ticket k in time period t, thus the model can manipulate Y_{kt} directly to reach capacity as opposed to opening and closing a set of fare classes, as often seen in traditional RM models. With these decision variables, the PCMIP will select a ticket availability matrix as well as assign the prices to each available ticket in that matrix.

The objective function, given in Equation (4.5), represents the expected revenue under the decision variable Y_{kt}, taking into account the probability a ticket purchase is made, determined by P_{jit} and $Q_{kt|ji}$. Constraint set (4.6) bounds the expected demand to that of the capacity on a leg, c_l. Constraint set (4.7) ensures the model can only select one ticket availability matrix of the super-set I_j. Constraint sets

(4.8) and (4.9) define P_{jit} and $Q_{kt|ji}$, based on Equations (4.3) and (4.4), respectively. Constraint set (4.10) establishes lower and upper bounds on the pricing decision for each ticket, while constraint set (4.11) defines the binary restriction for X_{jit}.

The PCMIP is unique when compared to other choice-based revenue management models previously developed. First, it determines the ticket availability matrix to be used for any time period, while simultaneously setting ticket prices. Compared to other revenue management models, such as the CMIP, the PCMIP has flexibility in pricing instead of being limited to a static price. The PCMIP eliminates the need for a two-step optimization process by combining the decisions of ticket availability and pricing into a single, concise model. Secondly, similar to the CMIP, the PCMIP reduces the number of ticket availability matrices to consider when compared to other choice-based revenue management models by separating origin-destination markets that are independent of one another. Although airline networks are large due to a highly connected network, the PCMIP is able to manage these complex networks under this independent market assumption.

Computational complexity, though, is a hurdle that needs to be addressed. As modeled, the PCMIP is a mixed integer non-linear program (MINLP), which can be extremely difficult to solve. Fortunately, the structure of this problem leads to a relatively easy solution with standard branch-and-cut algorithms. The structure of the PCMIP, similar to the of the CMIP, only needs to consider a single ticket availability set in each market in an optimal solution. The selection between a mutually exclusive set, like that of the super set of policies I_j, is quick to solve on its own, resulting in a search for optimal prices. In a worst case scenario, an algorithm could consider all ticket availability matrices $i \in I_j$, and then determine the optimal prices given that availability set. Once these values have been determined, the algorithm would simply pick the availability set for each market j that maximizes the total expected revenue. Realistically, though, branch-and-cut algorithms would be able to set the ticket availability and prices in a traditional manner, as the selection of ticket availability matrices are mutually exclusive. Thus, implementation of this model is still feasible, with additional work needed on the design of an algorithm.

4.4 Computational Results

To test the efficacy and quality of the PCMIP formulation, we have constructed four examples with varying parameters. The first example is based on the parallel flight network seen in Liu and van Ryzin (2008) and Bront *et al.* (2009), and will focus on how the PCMIP behaves compared to the CMIP when the only behavioral factors are price and time of day preference. The second example, common among many RM papers, is the familiar three leg network seen in Liu and van Ryzin (2008). We compare the PCMIP to the CMIP under many fare class options, illustrating the importance of price optimization over fare class selection. The third example is also based on networks seen in Liu and van Ryzin (2008) and Bront *et al.* (2009). We compare the PCMIP to the CMIP in a hub-and-spoke network under preferences associated with the Southwest Airlines business model. We consider many fare class options when evaluating the CMIP, and show the continuous price options of the PCMIP are superior to that of the CMIP under multiple fare class options. The final example is an expanded version of the large network example seen in Liu and van Ryzin (2008). We consider a slightly larger hub-and-spoke network, with added direct flights for competition in addition to time of day preference. In all cases, the policies of the PCMIP and CMIP are simulated in a traditional Monte Carlo simulation of 1000 runs, where the passengers arrive according to a Poisson arrival process, and each passenger considers a ticket out of the ticket availability matrix.

4.4.1 Parallel Flights Example: Time of Day Effects

In this example (Figure 4.1), we have a single origin and destination with three flights representing an early (Leg 1), mid-day (Leg 2), and late (Leg 3) set of flights. Each flight leg has its own capacity, with the early flight having 30 seats available, the mid-day flight having 50 seats available, and the late flight having 40 seats available. We tested two preference scenarios in this network; one where price is the only factor and one where price and time of day (ToD) are factors. In both cases, the sensitivity to price was set to $\beta_p = -0.002$, representing a negative utility to price. When ToD is included as a factor, we put preference on the early flight, and used the mid-day flight as a reference. Thus, in the case where ToD and price are the factors that influence purchasing behavior, we reach a regression model of the form

$$\alpha_{kt} = -0.002Y_{kt} + \mathbb{I}_{\text{Leg 1}} + 0.5\mathbb{I}_{\text{Leg 3}} \;,\forall k \in K_{ji}, t \in \{1, 2, \ldots, T\}, \tag{4.12}$$

where $\mathbb{I}_{\text{Leg i}}$ represents a 0/1 selection of Leg i. In the case where ToD is not included, implying price is the only concern for purchase behavior, we have a regression model of the form

$$\alpha_{kt} = -0.002Y_{kt}, \forall k \in K_{ji}, t \in \{1, 2, \ldots, T\}. \tag{4.13}$$

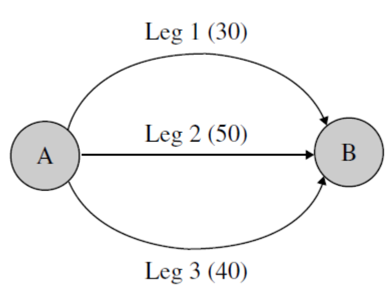

Figure 4.1: Parallel Flight Network - Adapted from Liu and van Ryzin (2008)

In each scenario, we tested the network for five time units ($T = 5$), with $\gamma = (6, 12, 24, 48, 96)$, for a total of 186 customers introduced to the network. The no-purchase utility, v_j, was set to be 0.5 for all markets. Similar to the example in Liu and van Ryzin (2008), each flight leg has two fare classes, High and Low, and the bounds on the fares can be seen in Table 4.1. These bounds were generated such that the median of the upper and lower bound match the original prices found in Liu and van Ryzin's example. Under these conditions, we solved the PCMIP for the two preference scenarios, and compared the results to that of the CMIP. As the results from Chapter 2 suggest, the CMIP does as well or better than other models, thus if the PCMIP is able to consistently do better than the CMIP, it follows that the PCMIP will consistently do better than other models. Additionally, to see the efficacy of optimal pricing while selecting ticket availability, we needed a similar model to highlight the increase. Since the PCMIP is an extension of the CMIP and the ticket availability sets are coded in a similar fashion, comparison of these two models would be sufficient. On its own, the CMIP doesn't have the opportunity to set prices directly, so we set the prices for each fare class to the original prices given by Liu and van Ryzin (2008), and used these values to determine ticket availability from the perspective of the CMIP.

The first scenario, with price being the only factor, converged to the solution shown in Tables 4.2 and 4.3, for each of the models. The prices of each ticket on each leg are given, as these represent both the availability of a ticket and the price at which it is sold. The PCMIP is able to select its own pricing policy, whereas the CMIP must select the prices given to it.

The simulation of the solutions from Tables 4.2 and 4.3 resulted in the PCMIP consistently outperforming the CMIP, as displayed in Table 4.4 and Figure 4.2. The PCMIP yielded a 6.4% increase in revenue overall, while maintaining a similar network

Leg & Fare Class	Lower Bound	Original Price	Upper Bound
Early High	$600	$800	$1000
Early Low	$200	$400	$599
Mid-day High	$800	$1000	$1200
Mid-day Low	$300	$500	$700
Late High	$400	$600	$800
Late Low	$100	$250	$399

Table 4.1: Bounds and Original Prices for Parallel Network

Table 4.2: PCMIP Solution - No ToD Pref.

Time Period	Early	Mid-day	Late
1	$945	$945	$800
2	$945	$945	$800
3	$945	$945	$800
4	$945	$945	$800
5	$945	$945	$800

Table 4.3: CMIP Solution - No ToD Pref.

Time Period	Early	Mid-day	Late
1	Sold Out	$1000	$600
2	Sold Out	$1000	Sold Out
3	$800	$500	Sold Out
4	$800	$1000	$600
5	$800	$1000	$600

load factor (PCMIP 78% versus CMIP 82%) and traffic. Since the PCMIP is free to select price, and thus less constrained than the CMIP, the higher revenue gains were expected, and the PCMIP dominated the CMIP in cumulative revenue over the five time periods.

The second scenario for this network, where price and time of day preferences were considered, yielded a considerably more dynamic solution, as seen in Tables 4.5 and 4.6. Now that price isn't the only factor to consider, the PCMIP must balance the sensitivity to price along with the time of day preferences, resulting in a more

Time Period	PCMIP		CMIP	
	Rev.	Traffic	Rev.	Traffic
1	$2,627	2.96	$1,997	2.76
2	$4,896	5.52	$2,346	2.35
3	$10,236	11.53	$7,459	12.30
4	$20,883	23.55	$19,623	26.24
5	$41,375	46.54	$38,902	51.05
Total	$80,016	90.10	$70,327	94.70

Table 4.4: Expected Revenue and Traffic per Time Period of Simulated Policies (No ToD Preference)

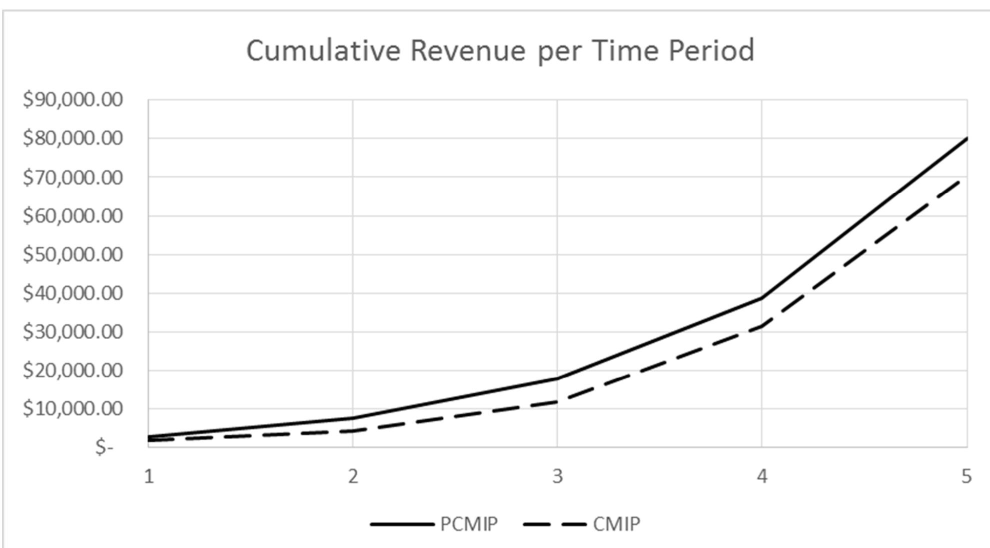

Figure 4.2: Cumulative Revenue per Time Period of Simulated Policies (No ToD Preference)

complex pricing structure. The obtained prices were simulated, and the PCMIP outperformed the CMIP by 13.8%, while managing to achieve a higher load factor of 94% versus that of the CMIP's 86% (see Table 4.7 and Figure 4.3). It is worth noting that the "bump" in the cumulative revenue (Figure 4.3) at time period 3 stems from the difference in solutions. The PCMIP favors closing the early and late flights in anticipation of future demand, whereas the CMIP only closes the early flight. Ultimately, though, the PCMIP's solution yields higher revenue and dominates the CMIP solution in cumulative revenue for all time periods, similar to when price was the only consideration.

Table 4.5: PCMIP Solution - With ToD Pref.

Time Period	Leg 1	Leg 2	Leg 3
1	$1000	$608	Sold Out
2	$1000	$589	$800
3	Sold Out	$700	Sold Out
4	Sold Out	$666	$800
5	$1000	$800	$800

Table 4.6: CMIP Solution - With ToD Pref.

Time Period	Leg 1	Leg 2	Leg 3
1	$800	$500	Sold Out
2	Sold Out	$500	$600
3	Sold Out	$1000	$600
4	Sold Out	$1000	Sold Out
5	$800	$500	$600

This example showed that, in the simplest of cases, the PCMIP's ability to select price in addition to availability outperforms the static-price version of the problem, as solved by the CMIP. Under the conditions of the two scenarios, the PCMIP's flexibility dominated the expected revenue of the CMIP while maintaining comparable load factors. Even in situations where the policies were greatly different, as seen in time period 3 of the second scenario, the cumulative revenue of the PCMIP never fell below that of the CMIP, suggesting the PCMIP is better equipped to handle the necessary decisions in a parallel network of this type.

Time Period	PCMIP		CMIP	
	Rev.	Traffic	Rev.	Traffic
1	$2,782	3.36	$2,601	3.83
2	$6,491	8.08	$3,884	6.97
3	$6,816	9.74	$8,809	12.85
4	$22,001	29.70	$10,059	10.06
5	$53,266	60.65	$43,895	68.09
Total	$91,356	111.52	$69,248	101.79

Table 4.7: Expected Revenue and Traffic per Time Period of Simulated Policies (With ToD Preference)

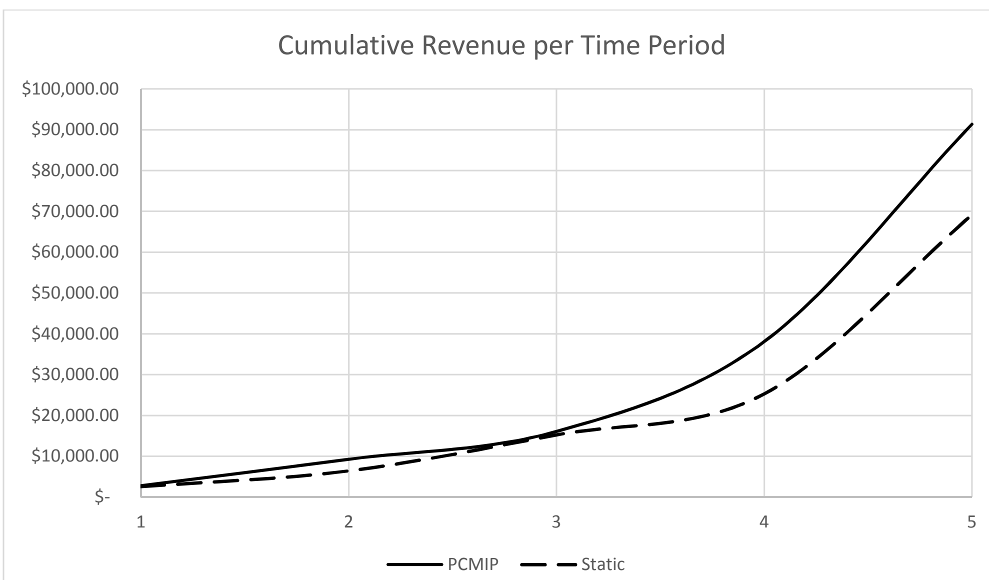

Figure 4.3: Cumulative Revenue per Time Period of Simulated Policies (With ToD Preference)

4.4.2 Small Network Example: Fare Class Impact

This example was modeled off of the three-leg network example seen in Liu and van Ryzin (2008) and Bront *et al.* (2009), seen in Figure 4.4. We considered multiple factors for purchasing behavior, including price, itinerary, and direct versus indirect. The sensitivity to price was set to $\beta_p = -0.003$, again, representing a negative sensitivity to price. The differences between markets were set to (AB,BC,AC) $= (2, 1, 3)$, while the effect of direct versus indirect was set to $\beta_{\text{Indirect}} = -3.5$, implying a negative utility associated with having intermediate stops. Under these values, the regression equation would read

$$\alpha_{kt} = -0.003Y_{kt} + 2\mathbb{I}_{\text{AB}} + \mathbb{I}_{\text{BC}} + 3\mathbb{I}_{\text{AC}} - 3.5\mathbb{I}_{\text{Indirect}} \ , \quad \forall k \in K_{ji}, t \in \{1, 2, \ldots, T\}, \tag{4.14}$$

where $\mathbb{I}_j$ represents the 0/1 selection for each market j, and $\mathbb{I}_{\text{Indirect}}$ represents the 0/1 selection of indirect versus direct for a given ticket k. The purpose of $\mathbb{I}_j$ is to account for different sensitivities across different markets. Simply put, $\mathbb{I}_j$ adjusts the magnitude of α_{kt} for each market so that every market isn't constrained to have the exact same utility under similar preferences. Note, though, that these values were created for this example and do not represent the true effects of a system, but are merely designed to show the use and quality of the PCMIP formulation.

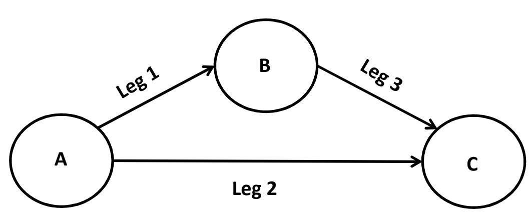

Figure 4.4: Small Network Example - Adapted from Liu and van Ryzin (2008)

The small network example was solved for five time units (i.e., $T = 5$), with $\gamma \in (1, 2, 3, 10, 25)$, representing a total of 41 customers entering the system. The market arrival rates and no purchase utilities are defined in Table 4.8, while the bounds and static prices can be found in Table A.6 in the Appendix. This network was solved under multiple fare class assumptions, ranging from two fare classes up to six fare classes. As more fare classes were added, the static prices of the additional fare classes were evenly distributed around the static price of the two fare class example. For instance, origin-destination path A-C has two prices under the two fare class example of \$1000 and \$800. Expanding this to a three fare class example, path A-C would now have three prices of \$1000, \$900, and \$800. This same process was repeated for all origin-destination paths and fare classes so that the CMIP had the widest array of options for optimizing revenue. The goal of using multiple fare class examples is to show that the gains of having more options for the CMIP are outclassed by the flexibility of the PCMIP.

O&D Market	Market Probability (λ_j)	No-purchase Utility (v_j)
AC	0.50	1.25
AB	0.25	2
BC	0.25	1

Table 4.8: Market Parameters for Small Network Example with Two Fare Classes

As seen from Table 4.9, the PCMIP outperforms the CMIP regardless of the number of fare classes the CMIP considers. In the simplest case, with two fare classes, the PCMIP outperforms the CMIP by 3.4%, while maintaining a similar network load factor of 63% versus the CMIP's 64%. These results highlight an important feature of the PCMIP: additional fare classes don't provide an opportunity for increased revenue when both price and availability are being optimized. As expected, the PCMIP

Time Period	PCMIP	CMIP				
		Number of Fare Classes				
		Two	Three	Four	Five	Six
1	$398	$167	$176	$169	$173	$171
2	$293	$118	$123	$124	$123	$128
3	$554	$244	$250	$251	$250	$250
4	$1924	$2031	$1998	$2025	$2027	$2007
5	$5196	$5535	$5539	$5523	$5522	$5526
Total	$8367	$8095	$8086	$8093	$8094	$8081

Table 4.9: Simulated Expected Revenue for a Selected Number of Fare Classes

represents an infinite number of fare classes since it is able to select an infinite number of prices. This feature essentially eliminates the need for fare classes in development of the data set, which is a common issue in traditional RM implementation where fare classes for each ticket type have grown significantly over the years. The CMIP, on the other hand, could show gains when adding fare classes, but in this example, under the preference conditions stated, there isn't an advantage to increasing the number of fare classes.

4.4.3 Large Network Example: Implementation of Southwest Airlines Ticket Attributes

In this example, we chose a large hub-and-spoke network and incorporated the Southwest Airlines product classification for passenger preferences seen in Table 4.10, reproduced from Chapter 3. The network, seen in Figure 4.5, contains eight legs and 20 origin-destination markets, and is based on the network found in Liu and van Ryzin (2008). Each leg has an identical capacity of 200, and the arrival rates for each market can be found in Table A.7 in the Appendix. We ran this example for a single time unit ($T = 1$) and a $\gamma = 2000$, representing a total of 2000 customers introduced to the system.

Ticket Attribute	Wanna Get Away	Anytime	Business Select
Refundable	No	Yes	Yes
Transferable	With Fees	Yes	Yes
Priority Boarding	No	No	Yes
Priority Security Lane	No	No	Yes
Loyalty Miles	Regular	Bonus	Bonus
Free Bags	2	2	2
Complementary Drinks	No	No	Yes

Table 4.10: Southwest Airlines Product Attributes (Southwest Airlines, 2016)

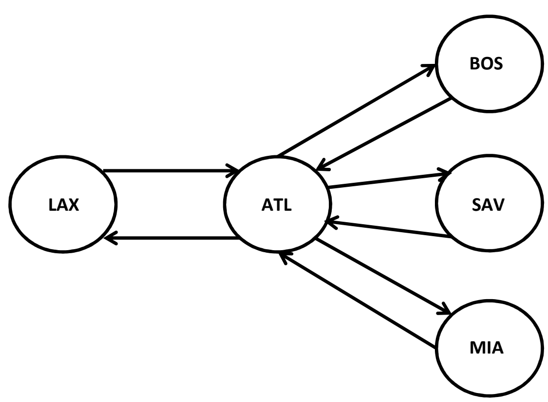

Figure 4.5: Large Network Example - Adapted from Liu and van Ryzin (2008)

Passenger choice was modeled according to the Southwest Airlines product classification found on their website. We considered eight ticket attributes along with three products, described in Table 4.11. As the table suggests, "Business Select" represents all ticket attributes, while "Anytime" and "Wanna Get Away" are made up of a subset of these ticket attributes. The regression coefficients used to develop the probabilities of purchase are also indicated in Table 4.11, and represent a passengers sensitivity to each ticket attribute. In addition to these attributes, we also considered price sensitivity and set $\beta_p = -0.015$, while the no-purchase utilities can be found in Table A.7 in the Appendix.

Ticket Attributes	Business Select	Anytime	Wanna Get Away	Regression Coefficient
Refundable	✓	✓		1
Reusable Funds	✓	✓	✓	2
Same Day Changes	✓	✓		1.5
Priority Boarding	✓			1
Priority Security Lane	✓			0.5
Two Free Checked Bags	✓	✓	✓	2
Complimentary Premium Drink	✓			0.5

Table 4.11: Southwest Ticket Definitions and Regression Coefficients for Large Network Example

With these regression coefficients and the parameters of the network, we solved the PCMIP and CMIP with a single fare class option, based on the price bounds seen in Table A.8 in the Appendix. Again, the static price for the CMIP comes from the midpoint between these price bounds. Both the PCMIP and CMIP were solved, and in the form of an airline representation, the solutions for each model are displayed in Table 4.12. As you can see, the solutions are quite different, including situations where the PCMIP chooses to close a ticket type (ATL-MIA and MIA-ATL), whereas the CMIP chooses to leave them open. Additionally, the pricing options are vastly different, with the PCMIP dynamically setting prices to account for their negative utility.

Itinerary	PCMIP			CMIP		
	Business Select	Anytime	Wanna Get Away	Business Select	Anytime	Wanna Get Away
ATL-BOS	$379	$370	Sold Out	$690	$315	Sold Out
ATL-LAX	$445	$366	$340	$723	$405	$183
ATL-MIA	$366	Sold Out	Sold Out	$632	$227	Sold Out
ALT-SAV	$369	$346	Sold Out	$685	$268	Sold Out
BOS-ATL	$379	$370	Sold Out	$690	$315	Sold Out
BOS-ATL-LAX	$616	$500	$86	$1,058	$558	$250
BOS-ATL-MIA	$471	$338	$217	$736	$404	$169
BOS-ATL-SAV	$416	$271	$265	$708	$341	$133
LAX-ATL	$445	$366	$340	$723	$405	$183
LAX-ATL-BOS	$616	$500	$103	$1,058	$558	$250
LAX-ATL-MIA	$836	$411	$100	$1,168	$623	$205
LAX-ATL-SAV	$719	$475	$84	$1,110	$597	$237
MIA-ATL	$366	Sold Out	Sold Out	$632	$227	Sold Out
MIA-ATL-BOS	$471	$338	$200	$736	$404	$169
MIA-ATL-LAX	$836	$411	$100	$1,168	$623	$205
MIA-ATL-SAV	$591	$336	$128	$796	$463	$168
SAV-ATL	$371	$368	Sold Out	$685	$268	Sold Out
SAV-ATL-BOS	$416	$266	$211	$708	$341	$133
SAV-ATL-LAX	$719	$475	$91	$1,110	$597	$237
SAV-ATL-MIA	$591	$336	$146	$796	$463	$168

Table 4.12: Solutions for the Large Network Example

In simulation, the PCMIP resulted in an expected revenue of $525,420, while the

CMIP resulted in an expected revenue of $352,885. The difference between the two models, a 48.9% increase in revenue, is primarily due to the flexibility of the PCMIP and the use of any pricing structure bounded by the upper and lower bounds for each ticket, whereas the CMIP must select ticket availability and is automatically forced to pick a price based on that selection. In reality, though, airlines don't use a single fare class for pricing of their products. Generally, each product has multiple fare classes within it, each with their own respective price ranges. To mimic this sort of system and give the CMIP an opportunity to do better, we solved the CMIP under the same network and preference parameters, but gave it the option of two, three, or four fare classes. The prices for each of these fare classes were evenly distributed around the original bounds to give the CMIP the best spread of options. The solutions to each of these fare class options were simulated, and the results are tabulated in Table 4.13.

Metric	PCMIP	CMIP			
		Number of Fare Classes			
		One	Two	Three	Four
Expected Revenue	$525,420	$352,885	$464,610	$498,865	$485,378
Standard Deviation	$9890	$6401	$13,881	$12,833	$11,987
Coefficient of Variation	0.02	0.02	0.03	0.03	0.02
Expected Traffic	1127	1162	1031	1113	1081
Network Load Factor	98%	97%	89%	95%	94%

Table 4.13: Large Network Example Simulation Results for Varying Number of Fare Classes

As the table suggests, the PCMIP outperforms the CMIP in all situations, even when adding more fare class options. Even when the CMIP generates its highest expected revenue, with three fare class options per product, the PCMIP still outperforms the CMIP by 5.3%. Despite this significant increase in revenue, the PCMIP maintains a reasonable coefficient of variation and a higher load factor. Ultimately, the PCMIP is better able to balance the preferences of each ticket as well as price, resulting in a better mix of passengers while maintaining a high load factor and maximizing revenue.

This example was constructed to show two things: the PCMIP can maintain its edge in more complex networks and that we can model the business classification of tickets easily. Although this example doesn't portray a full network, it does take into account competition for seats across different itineraries and it manages to include different ticket attributes while solving the hub-and-spoke network. Under the Southwest Airlines product classification, the PCMIP consistently outperformed the CMIP under multiple fare class options, and maintained similar coefficient of variances and network load factors. In this example, the PCMIP dominates the CMIP in all aspects, again suggesting the decision of price is more important than the number of fare classes to consider.

4.4.4 Expanded Large Network Example: Impact of Additional Competition

Although the large network example previously discussed provided results that solidify the PCMIP's ability to select price in a complex network, it lacks some key factors in choice-based revenue management. For one, the competition between origin-destination combinations is missing, considering there are no direct routes between the spoke cities. Additionally, the flights within the network are arbitrarily timed, implying time of day preference isn't important when passengers select their tickets. To expand on the original large network example, we have added four direct flights (seen in Figure 4.6), with each additional leg having a capacity of 100 seats.

We also incorporated time of day preference into the expanded large network ex-

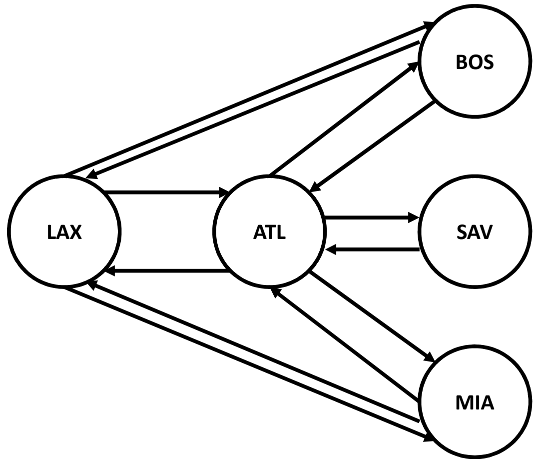

Figure 4.6: Expanded Large Network Example

ample. Any flight that has a later connection is considered to be an "early flight", as well as any direct flight between the spokes. For instance, all flights leaving LAX, BOS, SAV, and MIA are early flights. Conversely, all flights leaving ATL are considered "late flights", so that any issues with timings on connections are assumed to be satisfied. Keeping all of the regression parameters from the previous example the same, we set the sensitivity to early flights to 0.75, implying earlier flights have a positive effect on purchasing behavior. The additional flights were given bounds based on ticket type, governed by the Southwest Airlines product classification, and are displayed in Table 4.14.

Itinerary	Business Select	Anytime	Wanna Get Away
LAX-BOS	\$776-\$2000	\$651-\$775	\$0-\$650
LAX-MIA	\$991-\$1875	\$601-\$990	\$0-\$600
BOS-LAX	\$776-\$2000	\$651-\$775	\$0-\$650
MIA-LAX	\$991-\$1875	\$601-\$990	\$0-\$600

Table 4.14: Price Bounds on Additional Tickets

In solving this example, we considered three methodologies. First, we solved this example with the given parameters utilizing the PCMIP. Then, we solved this example with the CMIP, under a single fare class offering. Since we expect the PCMIP to consistently outperform the CMIP as it can select price, we also allowed a post-CMIP price optimization method to improve on the CMIP's original solution. The post-CMIP Price Optimal solution forces the PCMIP to use the CMIP solution, but allows the PCMIP to adjust prices. Essentially, we solved the PCMIP given the CMIP's solution for ticket availability, creating a two-step ticket availability and price optimization approach. After solving this example with these methods, we simulated the results of the PCMIP (Table A.9 in the Appendix), CMIP, and Post-CMIP Price Optimal methods (Table 4.16), and found that the PCMIP outperforms both the CMIP and Post-CMIP Price Optimal method, as seen in Table 4.15.

Metric	PCMIP	CMIP	Post-CMIP Price Optimal
Expected Revenue	\$637,190	\$494,489	\$620,321
Standard Deviation	\$19,918	\$23,843	\$22,290
Coefficient of Variation	0.03	0.05	0.04
Expected Traffic	1288	1247	1322
Network Load Factor	81%	78%	81%

Table 4.15: Extended Large Network Example Simulation Results

As the table suggests, the gains in revenue are drastic when comparing the CMIP solution to that of the PCMIP. The PCMIP outperforms the CMIP by 28.9%, showing an increase in overall traffic and network load factor. In spite of the poor performance of the CMIP compared to the PCMIP, utilizing the CMIP solution and selecting optimal prices shows a substantial gain in expected revenue. The Post-CMIP Price Optimal solution improves the CMIP expected revenue by 25.4%, while falling slightly below the PCMIP solution by 2.6%. Thus, under more competition, more flight options, and incorporating a time of day preference, the PCMIP outperforms the CMIP even when the CMIP's solution is optimized for price. Although the Post-CMIP Price Optimal solution does considerably better, this suggests the two-step optimization approach often incorporated in the airline industry is sub-optimal, considering the decisions of ticket availability and pricing can be solved simultaneously without adding fare class complexity.

It is clear that optimizing pricing decisions while simultaneously setting ticket availability produces better results compared to static models, but the results of the Post-CMIP Price Optimal method compared to that of the CMIP highlight the impact of pricing sensitivities on network revenue management. We can see the importance of pricing by paying attention to the differences in pricing decisions between the two methods. Table 4.16 contains the ticket availability and prices for the CMIP and Post-CMIP Price Optimal method that were simulated to achieve the results in Table 4.15. On average, the CMIP charges $122 more than the Post-CMIP Price Optimal method, deterring customers from making purchases in certain markets. Due to this reduced average price, the Post-CMIP Price Optimal method is able to accept more passengers but not at the expense of revenue.

These results show the added difficulty in optimizing revenue with more complex networks. As more competition and ticket preferences are added, the passengers have more choices to choose from, thus increasing the variability of their choice process. Nonetheless, the PCMIP was able to handle the added complexity better than the static model, yielding a sizable increase in expected revenue. As more complexity is added to the network in the form of competitive flight options, the PCMIP should be able to handle the complex network and maintain its edge compared to the static model, as seen in the diverse set of examples given.

Itinerary	CMIP			Post-CMIP Price Optimal		
	Business Select	Anytime	Wanna Get Away	Business Select	Anytime	Wanna Get Away
ATL-BOS	$690	$314	Sold Out	$493	$378	Sold Out
ATL-LAX	$723	$405	Sold Out	$457	$444	Sold Out
ATL-MIA	$632	$227	Sold Out	$388	$263	Sold Out
ATL-SAV	$685	$268	Sold Out	$461	$368	Sold Out
BOS-ATL	$690	$314	Sold Out	$517	$378	Sold Out
BOS-LAX	$1,388	$713	$325	$776	$651	$233
BOS-ATL-LAX	$1,058	$558	Sold Out	$739	$615	Sold Out
BOS-ATL-MIA	$736	$404	Sold Out	$651	$470	Sold Out
BOS-ATL-SAV	$708	$341	Sold Out	$645	$415	Sold Out
LAX-ATL	$723	$405	Sold Out	$497	$444	Sold Out
LAX-BOS	$1,388	$713	$325	$776	$651	$285
LAX-MIA	$1,433	$796	$300	$991	$601	$298
LAX-ATL-BOS	$1,058	$558	Sold Out	$623	$615	Sold Out
LAX-ATL-MIA	$1,168	$623	Sold Out	$836	$412	Sold Out
LAX-ATL-SAV	$1,110	$597	$237	$719	$475	$371
MIA-ATL	$632	Sold Out	Sold Out	$508	Sold Out	Sold Out
MIA-LAX	$1,168	$796	$300	$836	$601	$309
MIA-ATL-BOS	$736	$404	Sold Out	$553	$470	Sold Out
MIA-ATL-LAX	$1,168	$623	$205	$836	$411	$410
MIA-ATL-SAV	$796	$463	Sold Out	$591	$475	Sold Out
SAV-ATL	$685	$268	Sold Out	$488	$368	Sold Out
SAV-ATL-BOS	$708	$341	Sold Out	$611	$415	Sold Out
SAV-ATL-LAX	$1,110	$597	$237	$719	$475	$414
SAV-ATL-MIA	$796	$463	Sold Out	$591	$520	Sold Out

Table 4.16: CMIP and Post-CMIP Price Optimal Solutions for the Expanded Large Network Example

4.5 Conclusion

The extensive use of fare classes in the airline industry has created a computationally difficult problem when managing inventory to maximize revenue. In this chapter, we introduce a model that eliminates the need for fare classes, which optimizes ticket prices and availability for maximum revenue, creating a manageable non-linear mixed integer program. The model considers sensitivities to many ticket attributes in a multinomial logit model, including price and path, aligning it with the popular discrete choice modeling found in the current revenue management literature.

Compared against a static model, the Price-dynamic Choice-based Mixed Integer Program shows significant gains in revenue while maintaining computational efficiency. In the first example, a parallel network, the PCMIP consistently outperformed the CMIP when price and time of day were the only considerations of passenger choice. Even under situations where the two models converged to vastly different solutions, the PCMIP's decisions dominated those of the CMIP, yielding a 13% increase in expected revenue. With added complexity to the network, the PCMIP continued to outperform the CMIP, as seen in the second example. In this example we considered a situation where the CMIP was given a better opportunity to perform with added fare classes, highlighting that the PCMIP doesn't require this added complexity to yield better revenue results.

A large network example was also considered, in two different situations. First, the product classification of Southwest Airlines was applied to the PCMIP and CMIP, and solved for a hub-and-spoke network. The results showed that, even under a complex network with realistic ticket attributes, the PCMIP maintains its edge against a traditional static RM model. This edge continued to hold, despite added fare classes for the static model, a common practice in today's industry. Expanding on the large network, competition was added into the network by introducing direct connections between some of the spoke cities. Additionally, a time of day preference was added to the example to introduce a more realistic setting of passenger choice. Under these new additions, the PCMIP continued to dominate the CMIP as well as a post-CMIP price optimal solution, yielding an average increase of 28.9% compared to the CMIP and 2.7% compared to the post-CMIP price optimal solution. These results highlight not only the quality of the PCMIP, but also the importance of pricing in revenue management.

The PCMIP has the freedom to select price bounded by passenger preference, while determining the availability of each ticket type and is not bounded by static price assumptions found in traditional fare class-based models. Developing an algorithm that integrates the structure into a search method could lead to a tractable formulation for airline use, while the current formulation is easily solvable for moderately sized networks. As the progression of revenue management has slowed, incorporating price into controlling seat inventory provides an avenue for future development, potentially leading to airline implementation of choice-based models that incorporate pricing decisions alongside ticket availability.

Chapter 5

CONCLUSION

Despite revenue management being researched for over 35 years, this dissertation provides impetus for continued research in the field of airline revenue management. The need for easily implemented models that link important decisions such as ticket availability and price has created a dichotomous system between academic research and industry practice. Centered around computationally complex models and assumption-latent demand models, academic research in airline revenue management has become highly theoretical with little practical use for the airlines. Industry models, on the other hand, provide sub-optimal solutions, thanks to over-simplified demand assumptions and the use of quick-to-compute heuristics. This dissertation addresses this gap, providing computationally efficient mathematical formulations and an easily implementable dependent-demand framework that provides better solutions than traditional RM models.

The first model, introduced in Chapter 2, is a choice-based mixed integer program (CMIP) that incorporates dependent demand through itinerary utilities. The CMIP consistently outperformed industry standards such as the EMSR-b and network optimization techniques, as well as one of the more favored models, the Choice-based Deterministic Linear Program (CDLP). Although the gains were smaller when comparing the CMIP to the CDLP, the reduction in problem size and overall complexity allows for airlines to implement the CMIP, which is often a criticism of the CDLP. The combination of revenue performance and the size of the formulation led to a useful model for implementation, but turned the attention to the construction of the demand model, itself.

In Chapter 3, I formally define a multinomial logit (MNL) choice demand model for airline use, taking into account everyday ticket attributes and price. Previous revenue management models, including CMIP, failed to take into account the importance of price and other ticket attributes when generating the ticket utilities. The equations defined in Chapter 3 remove static assumptions, such as price-independent purchase utilities, allowing researchers to develop revenue management models that can account for purchase behavior while optimizing ticket availability and price. This MNL framework simplifies the process of modeling dependent demand, while maintaining the complex nature of passenger purchase behavior. Once the framework had been established, extending the CMIP from Chapter 2 into a price-sensitive model followed naturally.

The final model, the price-dynamic choice-based mixed integer program (PCMIP), integrated the demand model established in Chapter 3 into the CMIP. The PCMIP proved to be quite effective, optimizing both ticket availability and price, eliminating the need for complex fare class systems. The PCMIP outperformed all previous models, including the CMIP, since it had the advantage of dynamically setting prices in the deterministic environment. With more flexibility, the PCMIP was easily solved on multiple examples, including real-world examples on networks larger than most revenue management research, utilizing tickets characterized by Southwest Airlines purchasing options. Despite being a mixed integer non-linear program, the compact nature of the CMIP was leveraged in the PCMIP, leaving a compact formulation with minimal integer solutions while optimizing prices. The union of price and ticket availability optimization with a compact formulation resulted in a network revenue management model with practical application.

The two models presented in this dissertation, along with the revenue management dependent-demand framework, has the opportunity to drive revenue management research and industry application towards an agreement on proper methodologies. Implementing basic assumptions and eliminating the use of independent demand models and fare classes, this dissertation provides the groundwork in industry-applicable revenue management models, without having to generate computationally complex network formulations or make restrictive assumptions. Advancing the field of revenue management into one where academics and practitioners can agree on methodologies and computational complexity should provide many years of advancement in the field of network revenue management and airline operations optimization.

REFERENCES

Adelman, D., "Dynamic bid prices in revenue management", Operations Research **55**, 4, 647–661 (2007).

American Airlines, "American airlines trip planner", `http://www.aa.com`, accessed: 02-20-2015 (2015).

American Airlines, "American airlines fare comparison tool", `http://www.aa.com`, accessed: 01-14-2016 (2016).

Aydin, G. and J. K. Ryan, "Product line selection and pricing under the multinomial logit choice model", Working paper, Purdue University, Indianapolis (2000).

Belobaba, P., *Optimal versus heuristic methods for nested seat allocation* (AGIFORS Yield Management Study Group, 1992).

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

Bront, J. J. M., I. Mendez-Diaz and G. Vulcano, "A column generation algorithm for choice-based network revenue management", Operations Research **57**, 3, 769–784 (2009).

Brumelle, S. L., J. I. McGill, T. H. Oum, K. Sawaki and M. W. Tretheway, "Allocation of airline seats between stochastically dependent demands", Transportation Science **24**, 3, 183–192 (1990).

Chaneton, J. M. and G. Vulcano, "Computing bid-prices for revenue management under customer choice behavior", Manufacturing and Service Operations Management **13**, 4, 452–470 (2011).

Chen, L. and T. H. de Mello, "Mathematical programming models for revenue management under customer choice", European Journal of Operational Research **203**, 294–305 (2010a).

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

Clough, M. C., T. L. Jacobs and E. S. Gel, "A choice-based mixed integer programming formulation for network revenue management", Journal of Revenue and Pricing Management **13**, 366–387 (2014).

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

Delta Airlines, "Delta airlines onboard experience", `https://www.delta.com`, accessed: 01-14-2016 (2016).

Dong, L., P. Kouvelis and Z. Tian, “Dynamic pricing and inventory control of substitute products”, Manufacturing & Service Operations Management **11**, 2, 317–339 (2009).

Dror, M., P. Trudeau and S. P. Ladany, “Network models for seat allocation on flights”, Transportation Research - Part B **22B**, 4, 239–250 (1988).

Erdelyi, A. and H. Topaloglu, “Computing protection level policies for dynamic capacity allocation problems by using stochastic approximation methods”, Tech. rep., School of Operations Research and Information Engineering - Cornell University (2008).

Erdelyi, A. and H. Topaloglu, “A dynamic programming decomposition method for making overbooking decisions over an airline network”, INFORMS Journal on Computing **22**, 3, 443–456 (2010).

Erdelyi, A. and H. Topaloglu, “Using decomposition methods to solve pricing problems in network revenue management”, Journal of Revenue and Pricing Management **10**, 325–343 (2011).

Farias, V. F. and B. van Roy, “An approximate dynamic programming approach to network revenue management”, Working paper (available at http://www.standford.edu/ bvr/psfiles/adp-rm.pdf), Stanford University (2007).

Feng, Y. and G. Gallego, “Perishable asset revenue management with markovian time dependent demand intensities”, Management Science **46**, 7, 941–956 (2000).

Feng, Y. and B. Xiao, “A continuous-time yield management model with multiple prices and reversible price changes”, Management Science **46**, 5, 644–657 (2000a).

Feng, Y. and B. Xiao, “Optimal policies of yield management with multiple predetermined prices”, Operations Research **48**, 2, 332–343 (2000b).

Gallego, G., G. Iyengar, R. Phillips and A. Dubey, “Managing flexible products on a network”, CORC technical report tr-2004-01, Department of Industrial Engineering and Operations Research, Columbia University, http://www.corc.ieor.columbia.edu/reports/techreports/ (2004).

Gallego, G., R. Ratliff and S. Shebalov, “A general attraction model and an efficient formulation for the network revenue management problem”, Working Paper (2011).

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

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

Glover, F., R. Glover, J. Lorenzo and C. McMillan, “The passenger-mix problem in scheduled airlines”, Interfaces **12**, 3, 73–80 (1982).

Gosavi, A., E. Ozkaya and A. F. Kahraman, "Simulation optimization for revenue management of airlines with cancellations and overbooking", OR Spectrum **29**, 21–38 (2007).

Huang, K. and Y.-T. Liang, "A dynamic programming algorithm based on expected revenue approximation for the network revenue management problem", Transportation Research Part E **47**, 333–341 (2011).

Jacobs, T. L., R. Ratliff and B. C. Smith, "Understanding the relationship between price, revenue management controls and scheduled capacity - a price balance statistic for optimizing pricing strategies", Journal of Revenue and Pricing Management **9**, 356–373 (2010).

Jacobs, T. L., B. C. Smith and E. L. Johnson, "Incorporating network flow effects into the airline fleet assignment process", Transportation Science **42**, 4, 514–519 (2008).

Jet Blue Airlines, "Jet blue flight benefits", `https://book.jetblue.com`, accessed: 01-14-2016 (2016).

Klein, R., "Network capacity control using self-adjusting bid-prices", OR Spectrum **29**, 39–60 (2007).

Kunnumkal, S., "Randomization approaches for network revenue management with customer choice behavior", Www.isb.edu/faculty/Sumit_Kunnumkal/papers/choice_random.pdf (2011).

Kunnumkal, S. and H. Topaloglu, "A refined deterministic linear program for the network revenue management problem with customer choice behavior", Wiley InterScience Www.interscience.wiley.com (2008).

Kunnumkal, S. and H. Topaloglu, "Computing time-dependent bid prices in network revenue management problems", Transportation Science **44**, 1, 38–62 (2010a).

Kunnumkal, S. and H. Topaloglu, "A new dynamic programming decomposition method for the network revenue management problem with customer choice behavior", Prodcution and Operations Management **19**, 5, 575–590 (2010b).

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

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

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

Littlewood, K., "Forecasting and control of passenger bookings", Journal of Revenue and Pricing Management **4**, 2, 111–123 (2005).

Liu, Q. and G. van Ryzin, "On the choice-based linear programming model for network revenue management", Manufacturing & Service Operations Management **10**, 2, 288–310 (2008).

Meissner, J. and A. Strauss, "Improved bid prices for choice-based network revenue management", European Journal of Operational Research **217**, 417–427 (2012a).

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

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

Meissner, J. and A. K. Strauss, "Pricing structure optimization in mixed restricted/unrestricted fare environments", Journal of Revenue and Pricing Management **9**, 399–418 (2010).

Meissner, J. and A. K. Strauss, "Choice-based network revenue management under weak market segmentation", Tech. rep., Kuehne Logistics University and Lancaster University Management School (2011).

Moller, A., W. Romisch and K. Weber, "Airline network revenue management by multistage stochastic programming", Computer Management Science **5**, 355–377 (2007).

Newman, J. P., M. E. Ferguson, L. A. Garrow and T. L. Jacobs, "Estimation of choice-based models using sales data from a single firm", Manufacturing and Service Operations Management **16**, 2, 184–197 (2014).

Perakis, G. and G. Roels, "Robust controls for network revenue management", Manufacturing and Service Operations Management **12**, 1, 56–76 (2010).

Secomandi, N., "An analysis of the control-algorithm re-solving issue in inventory and revenue management", Manufacturing and Science Operations Management **10**, 3, 468–483 (2008).

Simpson, R. W., *Using network flow techniques to find shadow prices for market demands and seat inventory control* (MIT, Department of Aeronautics and Astronautics, Flight Transportation Library, 1989).

Smith, B. and C. Penn, "Analysis of alternative origin-destination control strategies", in "In Proceedings of the 28th Annual AGIFORS Symposium", pp. 123–144 (1988).

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

Song, Y. S., S. T. Hong, M. S. Hwang and M. G. Yoon, "Milp model for network revenue management in airlines", Journal of Business & Economics Research **8**, 2, 55–62 (2010).

Southwest Airlines, "Soutwhest airlines flight fare benefits", https://www.southwest.com, accessed: 01-14-2016 (2016).

Subramanian, J., J. Shaler Stidham and C. J. Lautenbacher, "Airline yield management with overbooking, cancellations, and no-shows", Transportation Science **33**, 2, 147–167 (1999).

Talluri, K., "A randomized concave programming method for choice network revenue management", Urlwww.econ.upf.edu/docs/papers/downloads/1215.pdf (2011).

Talluri, K. and G. van Ryzin, "A randomized linear programming method for computing network bid prices", Transportation Science **33**, 2, 207–216 (1999).

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

Talluri, K. T. and G. J. van Ryzin, *The theory and practice of revenue management* (International Series in Operations Research and Management Science, 2004b).

Topaloglu, H., "A stochastic approximation method to computer bid prices in network revenue management problems", INFORMS Journal on Computing **20**, 4, 596–610 (2008).

Topaloglu, H., "Using lagrangian relaxation to computer capacity-dependent bid prices in network revenue management", Operations Research **57**, 3, 637–649 (2009).

van Ryzin, G. and G. Vulcano, "Computing virtual nesting controls for network revenue management under customer choice behavior", Manufacturing and Service Operations Research **10**, 3, 448–467 (2008a).

van Ryzin, G. and G. Vulcano, "Simulation-based optimization of virtual nesting controls for network revenue management", Operations Research **56**, 4, 865–880 (2008b).

Weatherford, L. R. and R. M. Ratliff, "Review of revenue management methods with dependent demands", Journal of Pricing and Revenue Management **9**, 326–340 (2010).

Williamson, E. L., *Airline network seat inventory control: Methodologies and revenue impacts*, Ph.D. thesis, Massachusetts Institution of Technology (1992).

Wollmer, R. D., "An airline seat management model for a single leg route when lower fare classes book first", Operations Research **40**, 1, 26–37 (1992).

Zhang, D., "An improved dynamic programming decomposition approach for network revenue management", Manufacturing and Service Operations Management **13**, 1, 35–52 (2011).

Zhang, D. and D. Adelman, "An approximate dynamic programming approach to network revenue management with customer choice", Transportation Science **43**, 3, 381–394 (2009).

Zhang, D. and W. L. Cooper, "Revenue management for parallel flights with customer-choice behavior", Operations Research **53**, 3, 415–573 (2005).

Zhang, D. and W. L. Cooper, "Pricing substitutable flights in airline revenue management", Tech. rep., PROS Revenue Management (2006).

Zhang, D. and W. L. Cooper, "Pricing substitutable flights in airline revenue management", European Journal of Operational Research **197**, 848–861 (2009).

Zhang, D. and Z. Lu, "Assessing the value of dynamic pricing in network revenue management", INFORMS Journal on Computing **25**, 1, 102–115 (2013).

APPENDIX A

RAW DATA

Product	Legs	Class	Fare	Product	Legs	Class	Fare
1	1	H	1000	12	1	L	500
2	2	H	400	13	2	L	200
3	3	H	400	14	3	L	200
4	4	H	300	15	4	L	150
5	5	H	300	16	5	L	150
6	6	H	500	17	6	L	250
7	7	H	500	18	7	L	250
8	{2,4}	H	600	19	{2,4}	L	300
9	{3,5}	H	600	20	{3,5}	L	300
10	{2,6}	H	700	21	{2,6}	L	350
11	{3,7}	H	700	22	{3,7}	L	350

Table A.1: Product Definitions for the Small Network Instance - Adapted from Liu and van Ryzin (2008)

Segment	O-D	Consideration Set	Preference Vector	Utility of No Purchase	λ_t
1	A-B	{1,8,9,12,19,20}	(10,8,8,6,4,4)	1	0.08
2	A-B	{1,8,9,12,19,20}	(1,2,2,8,10,10)	5	0.20
3	A-H	{2,3,13,14}	(10,10,5,5)	1	0.05
4	A-H	{2,3,13,14}	(2,2,10,10)	5	0.20
5	H-B	{4,5,15,16}	(10,10,5,5)	1	0.10
6	H-B	{4,5,15,16}	(2,2,10,8)	5	0.15
7	H-C	{6,7,17,18}	(10,8,5,5)	1	0.02
8	H-C	{6,7,17,18}	(2,2,10,8)	5	0.05
9	A-C	{10,11,21,22}	(10,8,5,5)	1	0.02
10	A-C	{10,11,21,22}	(2,2,10,10)	5	0.04

Table A.2: Segment Definitions for the Small Network Instance - Adapted from Liu and van Ryzin (2008)

Single Itinerary Markets				
Market	Itinerary	Arrival Rate	Preference Vector	No Purchase Utility
SATBOS	SATDFW1DFWBOS1	0.002	(1, 2, 3)	4
SATSEA	SATDFW1DFWSEA1	0.005	(2, 2, 4)	5
SEAABQ	SEADFW1DFWABQ1	0.002	(3, 3, 3)	8
SEAAUS	SEADFW1DFWAUS1	0.002	(4, 4, 5)	10
SEADCA	SEADFW1DFWDCA1	0.005	(3, 5, 7)	11
SEAJFK	SEADFW1DFWJFK1	0.007	(1, 2, 3)	8
SEASAT	SEADFW1DFWSAT1	0.002	(2, 3, 8)	10
SEASFO	SEADFW1DFWSFO1	0.002	(1, 1, 5)	9
SFOABQ	SFODFW1DFWABQ1	0.002	(3, 4, 8)	9
SFOAUS	SFODFW1DFWAUS1	0.002	(2, 2, 7)	9
SFODCA	SFODFW1DFWDCA1	0.002	(1, 1, 4)	11
SFOORD	SFODFW1DFWORD1	0.005	(1, 2, 3)	5
SFOSAT	SFODFW1DFWSAT1	0.005	(1, 1, 2)	4

Table A.3: Consideration Sets and Utility Values for the Single Itineraries in the Large Network Example

Double Itinerary Markets						
Market	Itinerary		Arrival Rate		Preference Vector	No Purchase Utility
ABQAUS	ABQDFW1DFWAUS1	ABQDFW2DFWAUS2	0.037	0.012	(1, 2, 3, 4, 5, 6)	7
ABQBOS	ABQDFW1DFWBOS1	ABQDFW2DFWBOS2	0.002	0.049	(2, 3, 4, 5, 6, 7)	8
ABQDCA	ABQDFW1DFWDCA1	ABQDFW2DFWDCA2	0.002	0.007	(1, 3, 4, 5, 5, 6)	9
ABQDFW	ABQDFW1	ABQDFW2	0.005	0.005	(2, 2, 3, 3, 4, 5)	7
ABQJFK	ABQDFW1DFWJFK1	ABQDFW2DFWJFK2	0.01	0.01	(1, 1, 2, 3, 4, 7)	9
ABQORD	ABQDFW1DFWORD1	ABQDFW2DFWORD2	0.012	0.002	(2, 2, 3, 5, 10, 11)	12
ABQSAT	ABQDFW1DFWSAT1	ABQDFW2DFWSAT2	0.002	0.002	(1, 1, 2, 2, 3, 3)	4
ABQSEA	ABQDFW1DFWSEA1	ABQDFW2DFWSEA2	0.005	0.007	(1, 1, 1, 2, 4, 5)	9
ABQSFO	ABQDFW1DFWSFO1	ABQDFW2DFWSFO2	0.002	0.005	(1, 3, 3, 4, 5, 6)	6
AUSABQ	AUSDFW1DFWABQ1	AUSDFW2DFWABQ2	0.007	0.002	(1, 1, 2, 3, 4, 5)	5
AUSBOS	AUSDFW1DFWBOS1	AUSDFW2DFWBOS2	0.002	0.002	(1, 2, 3, 4, 5, 6)	7
AUSDCA	AUSDFW1DFWDCA1	AUSDFW2DFWDCA2	0.002	0.005	(2, 3, 4, 5, 6, 7)	8
AUSDFW	AUSDFW1	AUSDFW2	0.005	0.01	(1, 3, 4, 5, 5, 6)	9
AUSJFK	AUSDFW1DFWJFK1	AUSDFW2DFWJFK2	0.005	0.005	(2, 2, 3, 3, 4, 5)	7
AUSORD	AUSDFW1DFWORD1	AUSDFW2DFWORD2	0.002	0.002	(1, 1, 2, 3, 4, 7)	9
AUSSAT	AUSDFW1DFWSAT1	AUSDFW2DFWSAT2	0.01	0.005	(2, 2, 3, 5, 10, 11)	12
AUSSEA	AUSDFW1DFWSEA1	AUSDFW2DFWSEA2	0.002	0.005	(1, 1, 2, 2, 3, 3)	4
AUSSFO	AUSDFW1DFWSFO1	AUSDFW2DFWSFO2	0.007	0.002	(1, 1, 1, 2, 4, 5)	9
BOSABQ	BOSDFW1DFWABQ1	BOSDFW2DFWABQ2	0.005	0.002	(1, 3, 3, 4, 5, 6)	6
BOSAUS	BOSDFW1DFWAUS1	BOSDFW2DFWAUS2	0.002	0.002	(1, 1, 2, 3, 4, 5)	5
BOSDFW	BOSDFW1	BOSDFW2	0.002	0.005	(1, 2, 3, 4, 5, 6)	6
BOSJFK	BOSDFW1DFWJFK1	BOSDFW2DFWJFK2	0.01	0.002	(1, 3, 4, 5, 5, 6)	9
BOSSAT	BOSDFW1DFWSAT1	BOSDFW2DFWSAT2	0.002	0.002	(1, 2, 3, 4, 4, 5)	5
BOSSEA	BOSDFW1DFWSEA1	BOSDFW2DFWSEA2	0.005	0.002	(1, 1, 2, 2, 3, 4)	4
DCAABQ	DCADFW1DFWABQ1	DCADFW2DFWABQ2	0.002	0.005	(1, 3, 3, 4, 5, 6)	7
DCAAUS	DCADFW1DFWAUS1	DCADFW2DFWAUS2	0.002	0.01	(1, 1, 2, 2, 3, 3)	5
DCAORD	DCADFW1DFWORD1	DCADFW2DFWORD2	0.002	0.002	(1, 1, 4, 4, 6, 7)	8
DCASAT	DCADFW1DFWSAT1	DCADFW2DFWSAT2	0.005	0.002	(1, 2, 3, 4, 4, 5)	6
DCASEA	DCADFW1DFWSEA1	DCADFW2DFWSEA2	0.002	0.002	(1, 1, 3, 3, 4, 4)	5
DCASFO	DCADFW1DFWSFO1	DCADFW2DFWSFO2	0.005	0.002	(1, 1, 1, 2, 4, 5)	7
DFWABQ	DFWABQ1	DFWABQ2	0.002	0.005	(1, 2, 3, 3, 4, 5)	8
DFWAUS	DFWAUS1	DFWAUS2	0.002	0.002	(1, 4, 4, 5, 8, 9)	10
DFWBOS	DFWBOS1	DFWBOS2	0.005	0.005	(2, 2, 4, 5, 6, 6)	9
DFWJFK	DFWJFK1	DFWJFK2	0.002	0.012	(1, 1, 3, 4, 4, 9)	9
DFWORD	DFWORD1	DFWORD2	0.005	0.002	(1, 1, 1, 2, 3, 4)	5
DFWSAT	DFWSAT1	DFWSAT2	0.002	0.007	(1, 2, 2, 3, 3, 4)	5
DFWSEA	DFWSEA1	DFWSEA2	0.005	0.007	(1, 2, 2, 3, 6, 6)	10
DFWSFO	DFWSFO1	DFWSFO2	0.01	0.005	(1, 3, 3, 5, 6, 6)	9
JFKABQ	JFKDFW1DFWABQ1	JFKDFW2DFWABQ2	0.002	0.005	(1, 1, 3, 4, 5, 5)	9
JFKAUS	JFKDFW1DFWAUS1	JFKDFW2DFWAUS2	0.005	0.01	(1, 3, 4, 7, 8, 9)	10
JFKBOS	JFKDFW1DFWBOS1	JFKDFW2DFWBOS2	0.002	0.005	(2, 2, 5, 6, 6, 7)	9
JFKORD	JFKDFW1DFWORD1	JFKDFW2DFWORD2	0.005	0.005	(1, 2, 3, 3, 4, 5)	7
JFKSAT	JFKDFW1DFWSAT1	JFKDFW2DFWSAT2	0.002	0.005	(1, 2, 3, 4, 5, 5)	6
JFKSEA	JFKDFW1DFWSEA1	JFKDFW2DFWSEA2	0.005	0.01	(1, 2, 2, 3, 4, 5)	5
ORDABQ	ORDDFW1DFWABQ1	ORDDFW2DFWABQ2	0.002	0.005	(1, 2, 5, 5, 8, 9)	10
ORDAUS	ORDDFW1DFWAUS1	ORDDFW2DFWAUS2	0.005	0.01	(1, 1, 2, 2, 3, 3)	4
ORDBOS	ORDBOS1	ORDDFW1DFWBOS1	0.01	0.012	(1, 1, 5, 5, 8, 9)	10
ORDDCA	ORDDFW1DFWDCA1	ORDDFW2DFWDCA2	0.002	0.005	(2, 2, 4, 4, 6, 6)	8
ORDDFW	ORDDFW1	ORDDFW2	0.002	0.01	(1, 1, 3, 4, 5, 7)	8
ORDJFK	ORDDFW1DFWJFK1	ORDDFW2DFWJFK2	0.007	0.005	(1, 4, 5, 5, 6, 6)	10
ORDSAT	ORDDFW1DFWSAT1	ORDDFW2DFWSAT2	0.002	0.005	(1, 4, 4, 6, 10, 11)	12
ORDSEA	ORDDFW1DFWSEA1	ORDSEA1	0.01	0.005	(1, 4, 3, 8, 5, 10)	11
ORDSFO	ORDDFW1DFWSFO1	ORDDFW2DFWSFO2	0.01	0.002	(1, 2, 2, 3, 4, 4)	5
SATABQ	SATDFW1DFWABQ1	SATDFW2DFWABQ2	0.005	0.002	(2, 2, 4, 5, 5, 6)	9
SATAUS	SATDFW1DFWAUS1	SATDFW2DFWAUS2	0.002	0.01	(1, 4, 4, 9, 10, 11)	15
SATDCA	SATDFW1DFWDCA1	SATDFW2DFWDCA2	0.005	0.007	(2, 3, 3, 4, 4, 9)	9
SATDFW	SATDFW1	SATDFW2	0.002	0.005	(1, 1, 2, 2, 3, 3)	4
SATJFK	SATDFW1DFWJFK1	SATDFW2DFWJFK2	0.007	0.002	(1, 1, 4, 4, 6, 7)	8
SATORD	SATDFW1DFWORD1	SATDFW2DFWORD2	0.002	0.002	(3, 4, 8, 8, 10, 10)	15
SATSFO	SATDFW1DFWSFO1	SATDFW2DFWSFO2	0.005	0.01	(1, 1, 4, 4, 9, 9)	10
SEABOS	SEADFW1DFWBOS1	SEADFW2DFWBOS2	0.002	0.005	(1, 2, 2, 5, 5, 9)	10
SEADFW	SEADFW1	SEADFW2	0.002	0.01	(1, 4, 4, 5, 9, 9)	4
SEAORD	SEAORD1	SEADFW1DFWORD1	0.002	0.002	(1, 1, 2, 2, 4, 4)	3
SFODFW	SFODFW1	SFODFW2	0.002	0.005	(1, 1, 3, 3, 4, 5)	2
SFOJFK	SFODFW1DFWJFK1	SFOJFK1	0.005	0.002	(1, 1, 4, 5, 6, 7)	10
SFOSEA	SFODFW1DFWSEA1	SFODFW2DFWSEA2	0.002	0.002	(1, 2, 3, 4, 5, 6)	7

Table A.4: Consideration Sets and Utility Values for the Double Itineraries in the Large Network Example

Triple Itinerary Markets								
Market	Itinerary			Arrival Rate			Preference Vector	No Purchase Utility
BOSDCA	BOSDCA1	BOSDFW1DFWDCA1	BOSDFW2DFWDCA2	0.002	0.005	0.002	(1, 1, 2, 2, 3, 3, 4, 4, 5)	7
BOSORD	BOSORD1	BOSDFW1DFWORD1	BOSDFW2DFWORD2	0.005	0.002	0.012	(1, 2, 2, 3, 4, 5, 6, 6, 7)	7
BOSSFO	BOSDFW1DFWSFO1	BOSDFW2DFWSFO2	BOSSFO1	0.002	0.002	0.005	(1, 2, 4, 4, 7, 8, 9, 9, 10)	11
DCABOS	DCADFW1DFWBOS1	DCADFW2DFWBOS2	DCABOS1	0.024	0.027	0.049	(1, 1, 2, 4, 4, 5, 6, 6, 9)	9
DCADFW	DCADFW1	DCADFW2	DCABOS1BOSDFW2	0.002	0.005	0.01	(2, 2, 2, 3, 4, 6, 7, 8, 9)	9
DCAJFK	DCADFW1DFWJFK1	DCADFW2DFWJFK2	DCAJFK1	0.005	0.005	0.002	(1, 2, 3, 4, 4, 5, 6, 7, 8)	8
DFWDCA	DFWDCA1	DFWDCA2	DFWBOS1BOSDCA1	0.002	0.01	0.017	(1, 3, 3, 6, 6, 7, 8, 9, 10)	12
JFKDCA	JFKDFW1DFWDCA1	JFKDFW2DFWDCA2	JFKDCA1	0.005	0.005	0.002	(1, 2, 3, 4, 4, 5, 6, 7, 8)	8
JFKDFW	JFKDFW1	JFKDFW2	JFKDCA1DCADFW2	0.002	0.005	0.002	(1, 2, 3, 4, 4, 5, 7, 7, 8)	8
JFKSFO	JFKDFW1DFWSFO1	JFKDFW2DFWSFO2	JFKSFO1	0.002	0.005	0.002	(1, 3, 3, 4, 4, 6, 7, 8, 9)	9
SFOBOS	SFOBOS1	SFODFW1DFWBOS1	SFODFW2DFWBOS2	0.002	0.005	0.007	(1, 1, 1, 5, 5, 5, 8, 8, 8)	3

Table A.5: Consideration Sets and Utility Values for the Triple Itineraries in the Large Network Example

Product	Origin-Dest. Path	Class	Lower Bound	Static Price	Upper Bound
1	A - C	A	\$1000	\$1200	∞
2	A - B - C	A	\$650	\$800	∞
3	A - B	A	\$400	\$500	∞
4	B - C	A	\$400	\$500	∞
5	A - C	B	\$0	\$800	\$999
6	A - B - C	B	\$0	\$500	\$649
7	A - B	B	\$0	\$300	\$399
8	B - C	B	\$0	\$300	\$399

Table A.6: Ticket Prices and Bounds for Small Network Example with Two Fare Classes

Market	Arrival Rate (λ_j)	No-purchase utility (v_j)
ATLBOS	0.058	1
ATLLAX	0.058	1
ATLMIA	0.054	1
ATLSAV	0.046	1.5
BOSATL	0.058	1
BOSLAX	0.048	2
BOSMIA	0.044	1
BOSSAV	0.052	1
LAXATL	0.058	1
LAXBOS	0.048	1
LAXMIA	0.046	1.4
LAXSAV	0.053	1
MIAATL	0.054	1.2
MIABOS	0.044	1
MIALAX	0.046	1.5
MIASAV	0.040	1
SAVATL	0.046	2
SAVBOS	0.052	1
SAVLAX	0.053	1.1

Table A.7: Arrival Rates and No-purchase Utilities for Large Network Example

Itinerary	Business Select	Anytime	Wanna Get Away
ATL-BOS	$379-$1000	$250-$378	$0-$249
ATL-LAX	$445-$1000	$366-$444	$0-$365
ATL-MIA	$264-$1000	$190-$263	$0-$262
ATL-SAV	$369-$1000	$168-$368	$0-$167
BOS-ATL	$379-$1000	$250-$378	$0-$249
BOS-ATL-LAX	$616-$1500	$500-$615	$0-$499
BOS-ATL-MIA	$471-$1000	$338-$470	$0-$337
BOS-ATL-SAV	$416-$1000	$266-$415	$0-$265
LAX-ATL	$445-$1000	$366-$444	$0-$365
LAX-ATL-BOS	$616-$1500	$500-$615	$0-$499
LAX-ATL-MIA	$836-$1500	$411-$835	$0-$410
LAX-ATL-SAV	$719-$1500	$475-$718	$0-$474
MIA-ATL	$264-$1000	$190-$263	$0-$262
MIA-ATL-BOS	$471-$1000	$338-$470	$0-$337
MIA-ATL-LAX	$836-$1500	$411-$835	$0-$410
MIA-ATL-SAV	$591-$1000	$336-$590	$0-$335
SAV-ATL	$369-$1000	$168-$368	$0-$167
SAV-ATL-BOS	$416-$1000	$266-$415	$0-$265
SAV-ATL-LAX	$719-$1500	$475-$718	$0-$474
SAV-ATL-MIA	$591-$1000	$336-$590	$0-$335

Table A.8: Bounds on Prices for Large Network Example

Itinerary	PCMIP		
	Business Select	Anytime	Wanna Get Away
ATL-BOS	$503	Sold Out	Sold Out
ATL-LAX	$457	$444	Sold Out
ATL-MIA	$472	Sold Out	Sold Out
ATL-SAV	$454	Sold Out	Sold Out
BOS-ATL	$540	Sold Out	Sold Out
BOS-LAX	$776	$651	$247
BOS-ATL-LAX	$616	$510	Sold Out
BOS-ATL-MIA	$538	Sold Out	Sold Out
BOS-ATL-SAV	$547	Sold Out	Sold Out
LAX-ATL	$497	$444	Sold Out
LAX-BOS	$776	$651	$296
LAX-MIA	$991	$601	$295
LAX-ATL-BOS	$616	$558	Sold Out
LAX-ATL-MIA	$836	$442	Sold Out
LAX-ATL-SAV	$719	$475	$322
MIA-ATL	$513	Sold Out	Sold Out
MIA-LAX	$991	$601	$294
MIA-ATL-BOS	$553	$470	Sold Out
MIA-ATL-LAX	$836	$425	Sold Out
MIA-ATL-SAV	$591	$448	Sold Out
SAV-ATL	$512	Sold Out	Sold Out
SAV-ATL-BOS	$513	Sold Out	Sold Out
SAV-ATL-LAX	$719	$475	Sold Out
SAV-ATL-MIA	$591	$454	Sold Out

Table A.9: PCMIP Solution for the Expanded Large Network Example