Airline Revenue Management for Continuous Pricing: Class-Based and Classless Methods

by

Nicholas James Liotta

B.S., Aeronautical Engineering, Clarkson University, 2017

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

Master of Science in Transportation

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2019

Author . . . .

Department of Civil and Environmental Engineering
May 20, 2019

Certified by. . . . .

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

Accepted by . . . .

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

Airline Revenue Management for Continuous Pricing: Class-Based and Classless Methods

by

Nicholas James Liotta

Submitted to the Department of Civil and Environmental Engineering
on May 20, 2019 in Partial Fulfillment of the Requirements for the
Degree of Master of Science in Transportation

Abstract

The development of the New Distribution Capability for airlines has raised interest within the airline industry in "continuous pricing", where fares offered to customers are not limited to a set of pre-determined price points.

This thesis provides an overview of experiments on four revenue management (RM) methods proposed for the practical implementation of continuous pricing. Two of these methods, termed class-based RM for continuous pricing, utilize existing forecasting and seat protection optimization methods to determine what fares to offer. The other two methods, termed classless RM, calculate optimal fares based on the maximization of expected revenue contribution at a given point in time during the booking process. This thesis examines the performance of probabilistic bidprice and unbucketed dynamic programming methods for both the class-based and the classless methods for continuous pricing.

The continuous pricing methods are compared with traditional class-based methods in unrestricted fare structures using the Passenger Origin Destination Simulator. Compared to a baseline with six fare classes, when two competing airlines both implement class-based continuous pricing, revenues can increase by up to 1%, and, when both airlines implement classless pricing, they can gain up to 2% in revenue. When only one airline implements continuous pricing in a competitive setting, revenue gains of 10–13% are possible over the six-fare class baseline. These larger gains mostly come at the expense of the competitor, which loses revenue and bookings. For all cases, as the number of fare classes in the baseline increases, the revenue gains of continuous pricing are diminished and may even become revenue losses under certain conditions.

The positive results of the continuous pricing methods are a result of the increased price granularity offered by continuous pricing. It is this price granularity that causes most of the revenue gains when a competitor airline does not switch to continuous pricing. The price granularity effect also explains why increasing the number of fare classes with the traditional class-based RM methods can generate as much and sometimes more revenue than the continuous pricing methods.

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

Acknowledgements

Firstly, I would like to thank Dr. Peter Belobaba, my thesis supervisor and academic advisor. His introduction to and guidance through revenue management greatly helped me better understand the field and produce the research contained in this thesis. His knowledge revenue management and the airline industry has been invaluable, and his funding made my attendance of MIT possible.

I would also like to thank and remember Craig Hopperstad, who programmed the versions of the Passenger Origin Destination Simulator (PODS) used in this thesis, before passing away in January of 2019. His helpfulness and responsiveness to requests and questions was greatly appreciated, and he will be missed.

Additionally I would like to thank Dr. Michael Wittman and Dr. Adam Bockelie for helping two other first-year graduate students and me through the learning curve of PODS research, despite having to work on their own doctoral dissertations at the same time (especially Adam, who mentored the three of us on his own for a year after Mike graduated).

I would like to also thank the rest of the PODS student research crew, Akash Raigangar, Alexander Papen, Kevin Wang, and Tim Lu, as well as William Lee who also worked in the research group. Thanks for making the lab an interesting and entertaining place to be.

I would like to thank my parents for supporting me throughout my education process. It is safe to say that I would not be at this point without their help, love, and support.

Lastly, I would like to dedicate this thesis to my grandfather, who passed away in October of 2018. He would have loved to see me graduate from MIT, and I will miss him greatly.

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

The most recent seismic shift in commercial aviation was arguably the deregulation of the United States aviation market in 1978. Once separated from the relative stability of government control over routes and prices, airlines increasingly needed to find ways to increase revenue and cut costs. This led to many of the characteristics of air travel today, such as hub-and-spoke networks and revenue management systems. It is likely that the next great shift is upon the airline industry at this time. This shift is in the form of the arrival of the New Distribution Capability (NDC). With NDC, airlines anticipate being able to break free of many of the constraints that they have had in distributing their tickets. Among these constraints is the ability to increase the number of the fare classes (price points) that they sell, or, at the greatest extreme, eliminate fare classes altogether. There is hope that complete elimination of fare classes will allow for "continuous pricing" and increase revenue, as the theory is that allowing an airline to pick any price to offer will allow it to choose the exact one that maximizes revenue. However, there has not yet been a practical continuous pricing method proposed and tested.

1.1 Revenue Management

The inventory of seats on a flight is both constrained and perishable. Put in more accessible terms, the product which an airline sells, its seats, are both limited and must be sold by a certain point in time, departure, or else they go unsold and unused. The effect of this is that airlines have incentives to both refuse passengers who are not willing to pay higher fares, in order to save inventory for higher value passengers, as well as to accept those lower-valued passengers, in order not to have inventory be unused. It is in an attempt to perform this balancing act that revenue management becomes relevant.

The natural solution to the problem of trying to separate customers with a high willingness-to-pay from those with a low willingness-to-pay is to sell the same seats at different values. Traditionally, these values, also known as "fare classes" have been defined by the lower-priced classes have more rules and restrictions on them. For example, it may make sense to charge a lower fare paying passenger a fee to change what flight he or she is flying on, while allowing a higher fare passenger the ability to change flights with no extra charge. Other traditional restrictions have acted to separate business and leisure travelers, as business travelers tend to be more sensitive to time restrictions, such as when they fly or the duration of their trip, while being relatively insensitive to price since they are generally not personally responsible for paying for their trip. The converse is true for leisure travelers, who tend to be more open to adapting the exact dates and times they travel, but will notice price more. Simply having differentiated fare classes would be sufficient if air travel inventory were not constrained by capacity, but this is clearly not the case. Revenue management methods, therefore, try to conserve the higher-value seats for those passengers willing to pay the higher fares, while also not letting the inventory perish.

The backbone of most revenue management methods has been the Expected Marginal Seat Revenue (EMSR) method first developed in Belobaba (1987). This method protects a seat for a higher-valued fare class if the expected marginal revenue of protecting the seat is higher than the value of the same seat in the next fare class. The expected marginal revenue value of a seat is the probability of the seat being sold according to a demand forecast multiplied by the fare for that class. Importantly, EMSR nests protections to match airline inventory structures such that every seat protected for a given class is also protected and available for all the classes more valuable than it.

Since the introduction of EMSR, there have been many additions to expand and improve upon it. EMSR is designed to optimize fare class protection levels on a single leg, and cannot be transferred simply to network-scale problems. Various methods have been devised to allow EMSR to work at the leg level with network fares adjusted to account for how they impact other legs. These mechanisms, such as displacement-adjusted virtual nesting (DAVN) and probabilistic bidprice proration (ProBP) have been designed to adjust multi-leg itinerary fares in the optimization process to account for the fact that multi-leg itinerary fares displace local traffic.

Recently, there has also been more interest in dynamic programming that optimizes protection levels taking into account arrival order of passengers, particularly in academic theory but also to some extent among airlines and commercial RM systems. The core fundamentals and purpose, however, have always remained in all of these revenue management methods: to sell seats to passengers at prices they will pay by simultaneously protecting enough seats for high-valued passengers while not allowing too many seats to go unsold.

While the development of more advanced revenue management methods has been occurring, there has also been another force that has pressured airlines to remove restrictions from their fares. This has occurred as a result of the rise of low-cost carriers, which typically have greatly simplified fare structures. In the absence of restrictions and advance purchase requirements, revenue management methods must instead account for the fact that business passengers, who are generally willing to pay more, tend to book later than leisure passengers in order to separate the two passenger types. While the forecasting methods required changes to account for the lack of restrictions, the revenue management optimization methods themselves could remain.

A result of removing all restrictions and advance purchase requirements is that passengers will, without fail, always buy the lowest fare available. This effectively makes it so that there is one fare being offered at any given point in time. This has, in turn, raised the question if, rather than being bound to finite fare classes, an airline selling unrestricted fares could generate more revenue by offering fares from a continuous range of prices.

1.2 New Distribution Capability and Non-Static Pricing

For the last 30 years, the airline industry has relied on the same set of systems to sell its seats to passengers (Westermann, 2013). These systems, known as the Global Distribution Systems (GDSs), place on how airlines may sell their products. For example, the GDSs that an airline publish its fares in advance in fare classes, each of which must be represented by a letter (the result of this being that each airline only has 26 possible fare classes on a flight to split across all cabins and travel types), and each of which is associated with a fixed, published value for a given itinerary on a given day. This requirement has prevented any use of a continuous range of prices instead of fixed price points, and essentially requires airlines to use pricing methods that fall under the category of assortment optimization, where fares are selected from pre-defined price points (Wittman & Belobaba, 2018).

These limitations of GDSs have caused a push within the industry in recent years to develop a new system for distributing fares to passengers. The new standard that has been developed is known as New Distribution Capability (NDC). Information communicated by NDC is written in XML, which gives airlines far more flexibility in their offerings, and NDC also eliminates the need for advanced publishing of fares (Westermann, 2013). One major new ability of NDC, which is already being tested by some airlines, is the idea of dynamic price adjustment (Wittman & Belobaba, 2018). While dynamic price adjustment still relies on fare classes, it allows fares for these fare classes to be adjusted from their fixed points (Wittman, 2018). Beyond this, NDC also introduces the technical possibility of continuous pricing, where, as mentioned above, predefined fare classes are eliminated entirely and replaced with a continuous range of possible fares. There are many technical challenges to the continuous method, such as needing different revenue management, reservation, and distribution systems (Westermann, 2013), but using it has remained a goal of some airlines.

1.3 Motivation for Research

As mentioned previously, one of the many challenges of implementing continuous pricing is the lack of revenue management models capable of correctly controlling the price of seats within the continuous range. Nearly every revenue management method developed to this point has been designed in such a way that it is reliant on the existence of fare classes. While there have been algorithms developed in academic research for continuous pricing, these methods often require critical assumptions that make them less than practical for use in the industry.

In order to use continuous pricing in a more practical setting, a revenue management method capable of taking a forecast and determining an optimal fare from that forecast is required. One way of doing this is to take pre-existing forecasting and optimization methods that use fare classes and adapt them so that they determine an optimal continuous fare instead of picking from a pre-defined price point. However, since moving away from predefined price points is the objective of continuous pricing, it would theoretically be more appealing to develop forecasting and optimization methods that are classless, i.e. they do not use fare classes at all. It is worth noting, though, that whatever revenue management methods are developed for use in continuous pricing structures could also be used in assortment optimization pricing structures to determine which fixed-price points to offer in unrestricted fare structures. This results in four different classifications from the two types of pricing discussed being combined with class-based and classless forecasting and optimization. A visual interpretation of this is shown in Figure 1-1.

		Forecasting and Optimization	
		Class-Based	**Classless**
Pricing/Distribution	**Fixed Price Points**	**Traditional Class-Based RM**	**Classless RM for Fixed Price Points**
	Continuous	**Class-Based RM for Continuous Pricing**	**Classless RM for Continuous Pricing**

Figure 1-1: Combinations of Pricing and Forecasting and Optimization Methods

This thesis will discuss the testing of practical revenue management methods for both continuous pricing and classless forecasting and optimization. To determine their feasibility in practical situations, the methods were tested in the Passenger Origin Destination Simulator (PODS). Compared to many other stochastic simulation methods used in academic literature, PODS has several advantages. It allows a user to see the effects of competition on a revenue management method, as passengers generated by the simulator decide whether to buy at all, on which airline and at what price. Another advantage is that, rather than allowing the revenue management method to know with complete accuracy the distribution parameters of the forecast, the simulated RM method is dependent on a forecaster that, similarly to the real world, will not determine the demand distribution with perfect accuracy. By incorporating these elements, it is possible to determine whether proposed revenue management methods for continuous pricing have robustness against competition and uncertain demand.

1.4 Thesis Outline

The remainder of this thesis will be divided into five chapters. Chapter 2 will cover the literature existing on topics relevant to this thesis in revenue management. Topics covered will include forecasting, seat protection optimization methods, and fare quotation methods that are relevant to testing continuous pricing methods. Chapter 3 detail the forecasting and optimization methods that will be examined in this thesis. Chapter 4 will discuss the PODS software methodology to be used for testing the proposed methods for continuous pricing as well as discussing the network primarily being used for testing in this thesis. Chapter 5 will then present the results of the tests as well as perform an analysis of those results. This chapter will compare traditional class-based, continuously priced class-based, and classless optimization methods. Lastly, Chapter 6 will summarize the thesis as well as discuss potential future work.

Chapter 2: Literature Review

This chapter will review the literature available concerning continuous pricing pertaining to this thesis. This chapter is split into three sections, each one pertaining to the major part of a revenue management system: forecasting, seat protection optimization, and fare quotation. Each section will discuss past research pertaining to topics discussed in this thesis in order to provide a framework to work within. The forecasting and seat protection optimization sections will discuss the literature that forms the framework for the forecasting and seat protection optimization methods for continuous pricing being discussed in this thesis. The fare quotation section will discuss the basis of the framework for dynamic pricing, of which continuous pricing is a component, as well as discussing literature on previous attempts at continuous pricing.

2.1 Forecasting Methods

Demand forecasting is a crucial component of nearly any revenue management system. While a great deal of work has been put into studying optimal seat protection methods, these methods depend heavily in practical application on an accurate demand forecasting method. The forecasting methods used for testing the continuous pricing methods in this thesis will be discussed in the following sub-sections.

2.1.1 Pick-up Forecasting

Pick-up forecasting is a simple but useful forecasting method. At each time frame during the booking process, pick-up forecasting combines two elements to generate an expected total forecast for each origin-destination itinerary fare class. The first element is the number of bookings currently made for each fare class at the time frame in question. The second element involves determining how much demand is still expected to come based on historical booking data. By summing these two values, a total forecast for a flight is generated. For a thorough explanation of pick-up forecasting, see Gorin (2000).

It is not possible to estimate demand still to come based strictly on the average number of bookings that are made from the current time frame until departure, since the number of such bookings made is often constrained by availability of fare classes. As a result, the historical booking data must be detruncated. The method used for the experiments discussed in this thesis is known as expectation maximization (EM).

EM is an iterative, statistical approach to detruncation. In each iteration, EM produces a new estimate of demand in each closed observation by calculating the mean of a truncated normal distribution with a normal mean and variance from the recorded open observation bookings and the previous iteration's estimated closed observation demands. The lower tail of the distribution is truncated for each closed observation by the number of recorded bookings in said closed observation. After each iteration, the new estimates of demand for closed observations replace the previous iteration's estimates in the normal distribution, and this process is repeated until the estimates converge. An in-depth analysis of EM and other detruncation processes can be found in Weatherford & Pölt (2002).

2.1.2 Q-Forecasting

Pick-up forecasting is an effective forecasting tool when demand is independent for each individual fare class. While, in a fully restricted fare structure, this assumption can be a good approximation for how demand actually behaves, once restrictions are removed, any semblance of independent fare class demand being a reality disappears. In fact, in a fully unrestricted fare structure, fare class demands are dependent, as, in the absence of restrictions, passengers will always buy the lowest fare. When pick-up forecasting interacts with an unrestricted fare structure, an effect known as spiral-down occurs.

Since, for pick-up forecasting, expected bookings for each individual fare class are based on historical bookings in that fare class, the fact that bookings in unrestricted fare structures will always trend to lower fare classes, even if passengers might be willing to pay more, causes pick-up forecasting to over-report forecasts in lower classes and under report in higher classes. This then leads to the seat protection method to allow greater availability of lower value seats, which causes more seats to be sold in the lower value fare classes. This drives the forecasts generated by the pick-up method even lower, which leads to even greater availability of lower value seats. This cycle of forecaster and seat protection method driving protection levels lower and lower is known as spiral-down, and is illustrated in Figure 2-1. For a more detailed look at spiral-down its effects, see Cooper et al. (2006) and Cléaz-Savoyen (2005).

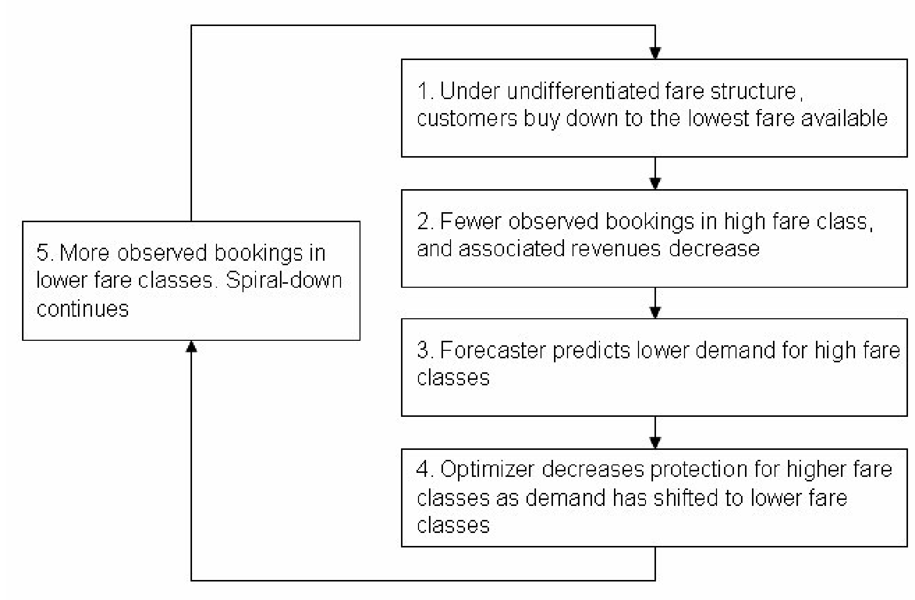

Figure 2-1: Spiral-down (Tam, 2008)

In order to overcome spiral-down, forecasting methods for use in unrestricted fare structures need to consider not what passengers have paid in the past, but rather what passengers will be willing to pay. Q-Forecasting was developed by Belobaba & Hopperstad (2004) to fulfill this task. Q-Forecasting works by taking historical bookings, and converting all of the bookings in the forecast into an equivalent number of Q-equivalent, or lowest-class equivalent, bookings. This conversion takes into account the complete lack of independence of fare class demand in an unrestricted fare structure.

In order to convert a fare class booking into an equivalent number of Q-Bookings, the sell-up rate from the lowest class to a given class must be estimated. Essentially, the question being asked by the conversion is how many passengers willing to pay at least the lowest fare are there for every passenger willing to pay the higher fare. Once the sell-up rate is estimated, it is a simple matter of dividing the historical bookings forecast by the sell-up rate from the lowest fare class to determine the equivalent number of Q-Bookings. It is worth noting that this sell-up rate and willingness-to-pay generally increase as departure approaches as a result of business passengers, who generally have a higher willingness-to-pay, arriving later in the booking process. While there are several methods for modeling sell-up (Cléaz-Savoyen, 2005), the one used in this thesis's Q-Forecasting will be a negative exponential distribution. Once the equivalent number of Q-Bookings has been determined, the forecasts are then detruncated and are then partitioned back into the original fare classes, once again using the sell-up rate, to determine a forecast that resists spiral-down. The algorithm for Q-Forecasting will be further discussed in Chapter 3 of this paper.

2.2 Fare Adjustment

Q-Forecasting is an effective counter against spiral-down, but it does nothing to counteract another issue that arises from the fact passengers will still only buy the lowest fare class available in an unrestricted fare structure. In addition to causing spiral-down, the non-independence of demand in unrestricted fare structures causes another issue that does not occur in fully restricted fare structures with independent demand.

When demand for each fare class is independent, there is sometimes reason to make some lower fare classes unavailable regardless of if they would typically be available according to the seat protection optimizer. On an aircraft with infinite capacity, it would only ever make sense to have all fare classes available in a restricted fare structure, as opening the lower fare classes would only accommodate more demand without impacting the demand with higher willingness-to-pay. Without the independence of demand afforded by a restricted fare structure, however, opening up lower fare classes, even on aircraft with infinite capacity, may actually lead to a decrease in revenue as a result of passengers who were willing to buy the higher class fares now buying down to the lowest available fare class.

As most seat protection models assume independent demand, it is necessary to have some way of accounting for this effect of non-independence of demand in an unrestricted fare structure, and Fiig et al. (2010) developed a method of "transforming the fare and the demand of a general discrete choice model to an equivalent independent choice model". Fare adjustment, as developed by Fiig et al. (2010) counteracts the threat of buy-down by reducing the fares as input into the seat protection optimizer to account for the opportunity cost of selling the seat. This opportunity cost is modeled as the marginal revenue of said seat. Further details on the fare adjustment algorithm can be found in Chapter 3.

2.3 Seat Protection Models

After the forecasting component, next step in most revenue management systems is seat protection. The seat protection step is the step where a revenue management system decides what seats to make available at what fares.

Seat protection models can be divided into two types: booking limit and bidprice controls. Booking limit methods decide how many seats to assign as available for each fare class, while bidprice controls set a threshold price to decide whether or not a fare class is available. Bidprices, or the expected value of the last seat on an aircraft, are often used as the minimum fare in the latter method, in which case the method is often referred to as a bidprice control method.

Another, and perhaps more useful, way to divide seat protection methods is into leg-based and origin-destination models. Leg-based models are conceptually and algorithmically simpler than origin-destination models, since, as their name implies, they aim to determine the optimal seat protections for each leg. However, only focusing on individual legs causes issues when deciding on protection levels for passenger itineraries that traverse multiple legs, as leg-based models have no way of determining whether allowing a passenger to travel in one seat on each of two legs is more valuable than allowing two single-leg passenger to travel in those seats on each leg. More information about the potential downsides of leg-based seat protection methods in a connecting flight network can be found in Belobaba et al. (2016).

2.3.1 Leg-Based Protection Methods

Leg-based protection models are the framework for most seat protection methods, as network-based models, which will be discussed in the next section, are often adaptations of leg-based models to work in a multiple leg network. The first practical model for seat protection came from Littlewood (1972), who showed that the optimal level of availability could be determined for two fare classes by using the expected marginal revenue, which would be determined from the forecast demand for each class, of each seat strictly reserved for the higher-priced fare class. The model developed by Littlewood (1972), which is a booking limit model, is provably optimal for a two fare class structure where the lowest fare paying passengers arrive first, but has no provisions to handle fare structures with more than two classes.

In order to extend Littlewood (1972)'s solution to a general solution, Belobaba (1987) introduced the expected marginal seat revenue method (EMSR), a heuristic that can determine protection levels for any number of fare classes. EMSR was later improved by Belobaba (1992) to a variant known as EMSRb, also a heuristic, which has since become widely used for leg-based seat protection determination. The key feature of EMSRb is that seat allocations for a fare class n are determined based on nested demands and fares for class n and all classes with higher fares, and then any seats protected for class n are also available to any fare class with higher fares. The purpose of this joint protection is so that, in the event that a higher fare paying passenger arrives when a lower paying passenger was expected, the higher fare passenger is not turned away. A visualization of joint protection/allocation is provided in Figure 2-2.

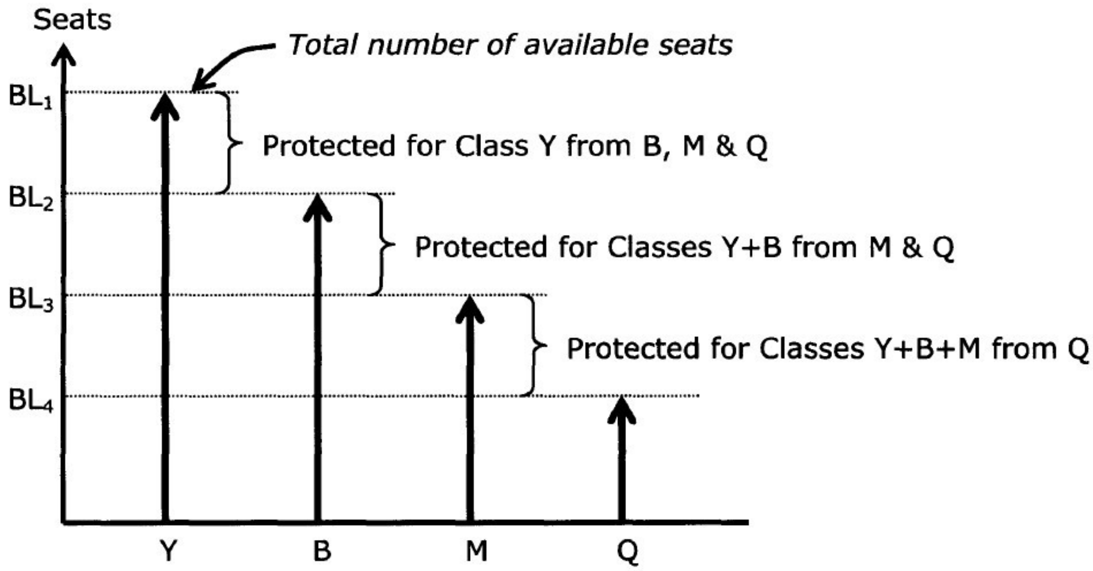

Figure 2-2: Joint Protection (Cléaz-Savoyen, 2005)

While EMSRb is a popular leg-based seat protection method, it is also what might be considered a static method. The solution provided is only a heuristic for optimality, and assumes passengers arrive in the assumed order of lowest fare paying to highest fare paying. While this is in practice addressed through frequent re-optimization, rerunning the EMSRb algorithm throughout the booking process, dynamic programming (DP) methods, which could provide seat optimization solutions for any combination of time remaining until departure and passenger arrival order, might also be considered for the optimization step. Lautenbacher & Stidham (1999) provides just such a DP.

Lautenbacher's DP, which can be applied as a bidprice control or seat allocation method, divides the booking window into a finite number of time slices, each of which is assumed to have no more than one arriving passenger. In each time slice, acceptance or rejection of a booking request is determined based on remaining time until departure and present capacity. By solving the DP recursively backwards from departure, seat protection allocations and bidprices for every capacity and time slice combination are determined. By having these allocations and bidprices, the

DP still provides a solution without reoptimization even if a high-paying passenger arrives earlier in the booking process than expected.

2.3.2 Origin-Destination RM Optimization Methods

As previously mentioned, the issue with attempting to strictly use leg-based protection methods is how to value an itinerary traversing two legs compared to two separate local itineraries on each leg. Most origin-destination methods developed, therefore, have tried to in some way account for this by reducing the optimization value of multi-leg itinerary fares to represent the fact that they take up space on multiple flights rather than just one. Williamson (1992) discusses many different methods of decrementing multi-leg itinerary fares, but a frequently used option is to solve a network linear program maximizing revenue for an entire network as if demand were deterministic.

While seat allocations obtained from doing so are unusable as controls in practice, the resulting shadow prices from the network LP are a decent estimation of how much one seat is worth on a capacity constrained flight. These shadow prices can then be used to decrement multi-leg itinerary fares to account for their extra displacement compared to leg-based fares. With the decrements in place, leg-based seat protection methods, such as the aforementioned EMSRb or Lautenbacher DP become usable. Typically, however, demands for individual origin-destination itinerary fares are too small for EMSRb to generate useful seat protection levels. This can be handled by implementing a fare bucketing system such as the one developed by Smith & Penn in 1988, where fares are of similar network values are grouped together for the process of determining fare class availability.

As an alternative to decrementing multi-leg fares for optimization purposes, it is also possible to prorate them across the legs they traverse. Bratu (1998) developed a method of doing so, known as ProBP, where network fares are iteratively prorated across multiple legs using bidprices generated from the EMSRb method on each leg. Theoretically, one major advantage ProBP would have over the previously mentioned decrementing methods would be that the proration method accounts for the variability of demand in calculating its bidprices. The bidprices generated from this proration method can be used to decrement fares, similar to the shadow prices from the deterministic linear program, or, alternatively, these bidprices could be used directly in a bidprice control method. It is the latter of the options, the bidprice control method, which will be used in this thesis.

2.3.3 Theoretical Optimality versus Practicality

It is worth noting that, for all of the theory about what theoretical advantages one method of seat protection has over another, these advantages often do not materialize when used in practical scenarios.

As one example of theory versus practice, while EMSRb is a heuristic solution, an optimal extension of Littlewood (1972) beyond two fare classes was determined around the same time as EMSR and EMSRb (Brumelle and McGill, 1993). However, in addition to being computationally difficult for its day when EMSRb was relatively simple, simulations showed that even the original EMSR method may be better than this method, known as OBL, in practical scenarios (Belobaba, 1992). Even though the OBL is theoretically better than EMSRb, it still makes some limiting assumptions. Like EMSRb, OBL assumes that the lowest fare class passengers arrive first and that willingness-to-pay then increases until the most expensive passengers arrive right before departure. This assumption, while based on general trends, is not strictly true in reality as some price-sensitive passengers will not look for a flight until very close to departure and some less price-sensitive passengers will finalize their plans long before departure, and, as a result, the advantage OBL had of being an optimal solution rather than a heuristic is lost. While situations such as higher paying passengers arriving earlier than expected during the booking process can and are generally addressed by rerunning the seat protection method, OBL required far too many re-optimizations to be considered usable. EMSRb was the more practical solution, even if it was not theoretically optimal.

A similar issue of theoretical optimality not transferring to practical performance has also been reported with DP-based revenue management. Since the Lautenbacher DP mentioned previously provides a solution for any arrival order of passengers, one might expect it to constantly outperform EMSRb for revenue. However, this has been shown to not necessarily be the case by Belobaba & Diwan (2010). In this case, the issue is two-fold. Firstly, the nature of the Lautenbacher DP is that it assumes a Poisson variance in its demand, which research has shown is not generally the case for demand variance in the real-world (Belobaba & Diwan, 2010). Secondly, and perhaps more importantly, the Lautenbacher DP is only provably optimal when the demand profile is known. As will be shown later, in more realistic simulations, the forecast demand is generally not completely reflective of the actual true demand. This particular issue with assumed knowledge of demand in theoretical work is frequently a major barrier between theoretical revenue management and practical implementations.

2.4 Fare Quote Generation

The last component of a revenue management system is fare quote generation. This is the step where a revenue management system picks what price to offer to prospective passengers. This can be as simple as simply quoting the price from the filed fare to the passenger, or it can be a more complicated process.

"Dynamic pricing" has become a popular phrase in the airline industry, although there is no consensus on what the term actually means. For the purpose of this thesis, the following definition will be used: "Firms practice dynamic pricing when they charge different customers different prices for the same product, as a function of an observable state of nature" (Wittman & Belobaba, 2018). This definition extends the dynamic pricing term to encompass nearly any revenue management method that controls fare class availability. Even an incredibly basic method, like arbitrarily deciding that 50% of the seats on the aircraft should be reserved for one fare class and the other 50% to another would fall under dynamic pricing by this definition since the price being is different for the two sets of customers, the product is the same, and the differences are based on an observable state of nature, which is how full the aircraft is. In addition to defining dynamic pricing, Wittman & Belobaba (2018) also define three mechanisms to implement dynamic pricing: assortment optimization, dynamic price adjustment, and continuous pricing.

Assortment optimization, defined as firms selecting prices from a finite set of prices (Wittman & Belobaba, 2018), is the category under which the previously discussed revenue management mechanisms described fall. The nature of Global Distribution Systems (GDSs), with their limited capabilities, has largely constrained practical revenue management to using assortment optimization mechanisms (Westermann, 2013). As a result, most of the research into revenue management have been into methods with this type of fare quotation. The previously discussed EMSR and Lautenbacher DP are part of this category. The introduction of New Distribution Capability, however, has made possible the widespread use of revenue management methods reliant on other types of dynamic pricing.

Another mechanism of dynamic pricing is to take the set of price points from assortment optimization and then to adjust those prices based on the observable state of nature. This is the definition of dynamic price adjustment (Wittman & Belobaba, 2018). The concept behind dynamic price adjustment is to give airlines greater flexibility in the prices they can offer, while still working within the existing revenue management systems to some extent. The effects of this form of dynamic pricing have been researched, although much of the research has not considered practical scenarios with competition and unknown demand. Research into more practical scenarios has indicated that, by using conditional willingness-to-pay models, revenue can be increased through the implementation of dynamic price adjustment (Wittman, 2018).

The third type of dynamic pricing, continuous pricing, eliminates the finite set of possible prices altogether and instead selects fares from a continuous price range (Wittman & Belobaba, 2018). Conceptually, this is a "cleaner" mechanism than dynamic price adjustment. Allowing a continuous range of fares allows price to be adjusted freely with the observable state of nature, and it is in this category that the new revenue management methods being discussed in this thesis fall. The fare optimization solution that maximizes revenue for a single flight with no competition and undifferentiated products has been determined when demand is Poisson (Gallego & van Ryzin, 1994). This single leg solution was later extended into a network solution (Gallego & van Ryzin, 1997), although the network solution is impossible to solve for any cases other than deterministic ones owing to the rate at which the complexity of the solution would grow. The solution provided by Gallego and van Ryzin is a DP similar to the Lautenbacher DP.

While there has been a good deal of research into continuous pricing methods, there are several issues with transferring the theoretical research into a practical method. Similar to the aforementioned issues of seat protection algorithm optimality, the Gallego and van Ryzin papers make two critical assumptions. Firstly, they assume that the demand profile is known by the airline, and, secondly, they focus on an airline operating without competition in the market. Neither of those assumptions is very realistic, and, as previously discussed, theoretical models often do not perform as well as heuristics when key assumptions are removed In assortment optimization cases, dynamic-programs, which are theoretically better, often perform worse than a frequently re-optimized static solution. For continuous pricing, however, there is no research concerning optimal models and heuristics in more realistic scenarios, and there are very few sources with any comparisons between a continuous pricing model and other fare quotation methods, even in the more idealized scenario.

One source that does compare continuous pricing with other methods is Zhang & Lu (2013). In addition to comparing multiple fare quotation strategies, Zhang & Lu (2013) also raises the issue of a network DP being too hard to solve by using a DP decomposition to approximate the DP. Zhang & Lu (2013)'s results do show that compared to several other methods, the continuous pricing model can increase revenue. However, like Gallego & van Ryzin (1994 & 1997) before it, Zhang & Lu (2013) still makes the two critical assumptions of known demand distribution parameters and non-competitive market scenarios.

2.5 Summary

This chapter documented the existing literature on three major parts of a revenue management system. It began by discussing forecasting, explaining pick-up forecasting and then discussing how Q-Forecasting and fare adjustment can be applied to pick-up forecasting to prevent spiral-down and buy-down. Next, seat protection methods were discussed. Both leg-based and origin-destination protection methods were examined, and it was discussed how theoretical optimality does not always yield a practical solution as result of assumptions made by the protection methods. Lastly, in a discussion of generating fare quotes, the idea of how nearly all revenue management involves some form of dynamic pricing, as defined by Wittman & Belobaba (2018) was established and previous attempts at continuous pricing were discussed in greater detail. In the next chapter, the forecasting and seat protection optimization algorithms that are used for the continuous pricing methods being examined in this thesis will be discussed in greater detail.

Chapter 3: Revenue Management Models for Continuous Pricing

In order to fulfil the purpose of this thesis, the analysis of several different new methods of revenue management (RM) for continuous pricing, it is necessary to review pertinent components of said RM methods. First, the forecasting methods relevant to the RM methods for continuous pricing will be described and explained. Second, the traditional class-based seat protection optimization models on which the RM methods for continuous pricing were based will be discussed, before then introducing the new methods of class-based and classless RM for continuous pricing.

3.1 Forecasting

A key component of RM is demand forecasting. Forecasts are used by the optimizer to determine fare availability. Typically, forecasts are generated from historical booking data for each origin-destination itinerary fare class (ODIF) according to different parameters such as day of the week and season of the year. The challenge for forecasting is to accurately determine how passenger purchasing behavior will be distributed, something that is made more challenging in the restriction-free fare structures that the continuous pricing methods considered in this thesis rely on.

3.1.1 Fare Restrictions

Restrictions, which include rules such as requiring passengers to purchase a round-trip ticket or stay a minimum number of nights, are a fundamental part of most airline fare structures. By applying more restrictions to the lowest priced fare classes, airlines are able to encourage passengers for whom the restrictions have a higher disutility to purchase more expensive fare classes.

While restrictions are generally useful in RM, the methods for continuous pricing considered in this thesis are not capable of employing them. All of the methods to be discussed are designed to offer one fare calculated to be revenue optimizing at any given time. While it would, in theory, be possible to develop continuous pricing methods with multiple offered fares, the complexity of the algorithms needed to do so would be very high compared to the single-fare methods. A single-fare continuous pricing method only needs to optimize its price with respect to the demand and capacity at a given time, while a multiple-fare method would need to optimize its prices with respect to not only the demand and capacity but with respect to each other. This necessary lack of restrictions for continuous pricing methods in turn influences the forecasting method used.

3.1.2 Q-Forecasting

A standard method for RM forecasting is to base forecasts on historical booking data. This is what the pick-up method of forecasting, used in the research for this thesis, uses. Put most simply, pick-up forecasting estimates how many bookings are expected from a given time frame to departure (L'Heureux, 1986). Pick-up forecasting does not take into account passengers who have already booked when projecting the number of expected remaining bookings-to-come.

In the absence of restrictions, the standard forecasting method does not function as effectively on its own. The standard forecasting model treats demand for each class as independent, and, while this is not entirely true in a restricted fare structure, the restrictions do separate the demand to some extent. Without the restrictions, any arriving demand will only ever purchase the lowest available fare, as there is no incentive for passengers to buy higher fares. Since forecasts are based on historical booking data, using standard forecasting will constantly see higher demand in lower fare classes and lower demand in higher fare classes. This causes the RM optimizer to allow greater lower class availability, which then in turn induces an even more skewed forecast. This feedback loop repeats until availability for lower fares is far higher than it should be and revenue is severely diminished. This effect is known as "spiral-down" (Cooper et al., 2006), and it requires a modification to the forecasting method to counter it.

One counter to spiral-down in unrestricted fare structures is known as Q-Forecasting (Belobaba & Hopperstad, 2004). On the broadest level, Q-Forecasting works by accounting for the fact that demands for each fare class are not independent (and, for completely unrestricted fare structures, there is no independence of demand for each fare class). Q-Forecasting converts historical bookings in each fare class into an equivalent number of bookings in the lowest fare class, which is denoted as "Q-Class" in this thesis. Pick-up forecasting can then be used on the Q-Class bookings and the forecast detruncated to determine the forecast the lowest fare class. This forecast can then be partitioned into forecast bookings in higher fare classes.

The fare class conversions for Q-Forecasting are performed using sell-up estimations. The method of estimating sell-up that will be used in this thesis is an exponential distribution constructed as follows. At a given time frame in the booking process, an estimate of the fare ratio between the lowest fare class fare and the fare which 50% of passengers will sell-up from the lowest fare to is input into the forecaster. This fare ratio estimate, known as the FRAT5 value, is typically user defined, although attempts have been made to estimate FRAT5 curves from data about passenger behavior. This FRAT5 value is then used to generate an estimate of the probability of sell-up from the lowest fare class (Q) to each fare class j, P_{Sup_j}, in the following exponential distribution:

$$P_{\mathrm{Sup}_j} = e^{\frac{\ln(0.5)(\mathrm{fare}_j - \mathrm{fare}_Q)}{(\mathrm{FRAT5}-1)(\mathrm{fare}_Q)}}$$

Q-Forecasting can then use this probability distribution to convert historical bookings for all fare classes into an equivalent number of Q-Bookings ($\mathrm{Book}_{Q_{eq}}$):

$$\mathrm{Book}_{Q_{eq}} = \sum_{j=1}^{Q} \frac{\mathrm{Book}_j}{P_{\mathrm{Sup}_j}}$$

It is sometimes desirable to, in practice, limit the amount any one booking may be scaled to eliminate outliers that may disproportionately affect the booking data. When this maximum scaling limit, known here as XSCL, is applied, the bookings conversion equation becomes:

$$\mathrm{Book}_{Q_{eq}} = \sum_{j=1}^{Q} \min\left(\frac{\mathrm{Book}_j}{P_{\mathrm{Sup}_j}}, \mathrm{Book}_j * \mathrm{XSCL}\right)$$

After applying detruncation and pick-up forecasting to the Q-bookings, a repartitioning step can then convert the volume of forecasted Q-Bookings for a specific future departure date back into the original fare classes. However, unlike the initial historical booking data, which was based on how passengers booked in the past, the repartitioning distributes the forecast based the willingness-to-pay (WTP) estimate.

$$\mathrm{Forecast}_j = \mathrm{Book}_{Q_{eq}}\left(P_{Sup_j} - P_{Sup_{j-1}}\right), \qquad \forall f$$

An example of Q-Forecasting is shown in Table 3-1 (with a FRAT5 of 1.5 and without any maximum scaling limit):

Fare	Avg. Bookings	Sell-up Prob.	Q-Bookings	Partition	Forecast
$600	0.000	0.001	0.000	0.001	0.047
$500	0.000	0.004	0.000	0.003	0.141
$400	0.000	0.016	0.000	0.012	0.563
$325	0.000	0.044	0.000	0.029	1.371
$250	1.000	0.125	8.000	0.081	3.879
$200	3.000	0.250	12.000	0.125	6.000
$150	8.000	0.500	16.000	0.250	12.000
$100	12.000	1.000	12.000	0.500	24.000
Total	**24.000**		**48.000**		

Table 3-1: Q-Forecasting Example (Belobaba and Liotta, 2017)

This repartitioning step is the critical element that makes Q-Forecasting resistant to spiral-down. As forecasts are based on how the bookings should be distributed based on an expected WTP estimate rather than how they are distributed, it prevents a skewing towards lower-value fares from occurring. It is possible to overestimate the WTP of passengers and instead make the higher-fare forecasts too high, but a good FRAT5 value will control the forecast scaling well.

3.1.2.1 FRAT5 Curves

A good general assumption, based on empirical observation, is that leisure passengers, who have lower WTPs, arrive earlier in the booking process and that business passengers, who have higher WTPs, arrive later. As a result, the overall WTP of the arriving passengers will increase as departure nears. In order to account for this, FRAT5 values should generally increase with each time frame as departure nears. This is achieved through the creation of FRAT5 curves, which specify the FRAT5 value for every time frame. In the simulation software used to test the continuous pricing methods, WTP varies with assumed network and passenger characteristics, and, as a result, FRAT5 curves are selected for the simulations through trial-and-error until which curve results in the most revenue is determined. The FRAT5 curves particular to the networks used to simulate the RM methods for continuous pricing will be discussed in Chapter 4.

3.1.3 Summary

Since the RM methods for continuous pricing to be presented require that only one price be offered at a time, all compatible fare structures must be unrestricted. As a result, forecasting for continuous pricing requires Q-Forecasting to prevent spiral-down. With the proper WTP estimation parameters, Q-Forecasting allows for better performance of the seat protection optimizer.

3.2 Seat Protection Optimization

After forecasting, the next step in RM is to use the forecasts to determine seat protection levels. Traditionally, this has been done by using pre-defined fare classes to determine which of those classes, each with pre-defined prices, will be offered. Recently, however, methods that use continuous pricing, where the prices that may be offered are taken from a continuous range rather than prescribed discrete points, and methods that are classless, where the optimization step is not based on individual fare classes, have been suggested.

Two examples of this traditional class-based optimization for network RM, probabilistic bidprice (ProBP) control and unbucketed dynamic programming (UDP), will be described, followed by descriptions of using Class-Based ProBP and UDP for continuous pricing, and then followed by classless versions of ProBP and UDP, which also use continuous pricing. However, before discussing the methods of optimization, it is necessary to discuss fare adjustment, a component used in the optimization process to account for the unrestricted fare structure necessary for the methods of continuous pricing being tested.

3.2.1 Fare Adjustment

While Q-Forecasting is a good technique to prevent spiral-down from occurring, there is one further feature of unrestricted fare structures that it does not resolve. Namely, Q-Forecasting does not completely take into account that passengers will still only buy the lowest available fare at a given time. The Q-Forecast might project greater demand for higher fare classes, but if the existing RM optimizer still allows empty seats to fall to the lowest fare class then there will be nothing to stop passengers from buying seats in the lowest fare class. As a result, it is possible to add a further component that accounts for the fact that, when passengers choose one fare, they are not buying another higher-value fare that they might have been willing to buy. The way of addressing this used in this thesis is known as fare adjustment (Fiig et al., 2010).

Marginal revenue fare adjustment modifies the fares as specified in the fare structure by subtracting the price elasticity cost of a seat from the filed price of the seat (i.e. the marginal revenue of selling a seat). The formula for this is as follows (Fiig et al., 2010):

$$\mathrm{fare}_j' = \mathrm{fare}_j - \frac{\mathrm{fare}_Q(\mathrm{FRAT5} - 1)}{-\ln(0.5)}, \qquad \forall j \in \{2, \dots, Q\}$$

The adjusted fares (fare_j') are then what are used by class-based seat protection optimizers in the RM process (classless methods do not need fare adjustment as there are no classes to protect demand for). The optimizer then performs its optimizations on the Q-Forecasts with the adjusted fares. Availability of fare classes at their filed values is based on their fare-adjusted values.

3.2.2 Traditional Class-Based Probabilistic Bidprice Method (ProBP)

To discuss ProBP (Bratu, 1998), it is first necessary to discuss EMSRb (Belobaba, 1992). As previously mentioned in Chapter 2, EMSRb is a common method of calculating protection levels for a given fare structure and forecast on a single leg. EMSRb is a heuristic approach which optimizes expected revenue by determining the proper joint protection levels for fare classes. In a less abstract sense, any seat protected for a fare class is jointly protected for any higher fare classes. The protections, therefore, are only against the lower-value fare classes. The number of seats jointly protected for each class and the classes above it is determined by calculating the point at which the expected marginal value of protecting an additional seat for said fare classes is equal to the value of the next lowest fare class. In a fare structure with μ_i mean demand, σ_i demand standard deviation, and fare f_i for each fare class i, the expected marginal revenue from protecting the π^{th} seat for classes 1 to n is calculated as follows:

$$\mu_{1,n} = \sum_{i=1}^{n} \mu_i$$

$$\sigma_{1,n} = \sqrt{\sum_{i=1}^{n} \sigma_i^2}$$

$$P_{1,n}(\pi_n) = \text{NormCDF}(\pi_n, \mu_{1,n}, \sigma_{1,n})$$

$$\overline{f}_{1,n} = \frac{\sum_{i=1}^{n}(f_i * \mu_i)}{\mu_{1,i}}$$

$$\text{EMSR}_{1,n}(\pi_n) = \overline{f}_{1,n} * P_{1,n}(\pi_n)$$

To determine the joint protection level for classes 1 through n, the expected marginal seat revenue from the incremental seat being protected should be equal to the next class's fare:

$$\text{EMSR}_{1,n}(\pi_n) = \overline{f}_{1,n} * \text{NormCDF}(\pi_n, \mu_{1,n}, \sigma_{1,n}) = f_{n+1}$$

This equation can then be rearranged to find the protection level π_i:

$$\pi_n = \text{invNorm}\left(\frac{f_{n+1}}{\overline{f}_{1,n}}, \mu_{1,n}, \sigma_{1,n}\right)$$

In order to translate the protection levels into something usable for an RM and seat inventory control system, the protection levels are subtracted from the capacity to determine the booking limits for each class:

$$\text{Limit}_n = \text{Capacity} - \pi_{n-1}$$

An example of EMSRb is shown in Table 3-2.

CAPACITY	135							
		FORECAST DEMAND				JOINT DEMAND		
BOOKING CLASS	FARE	MEAN	SIGMA	JOINT PROTECT	BOOKING LIMIT	MEAN	SIGMA	AVG FARE
1	$ 670	12	7	6	135	12	7	$670.00
2	$ 550	17	8	23	129	29	10.6	$599.66
3	$ 420	10	6	37	112	39	12.2	$553.59
4	$ 310	22	9	62	98	61	15.2	$465.74
5	$ 220	27	10	95	73	88	18.2	$390.34
6	$ 140	47	14		40	135	22.9	$303.19

Table 3-2: EMSRb Example (Belobaba et al., 2016, p. 111)

In addition to determining protection levels, there is one additional value that can be found from EMSRb which is of particular use to ProBP. This is the "critical EMSR" value, EMSRc, which is the expected marginal seat revenue of the last available seat on the aircraft. For n_{Capacity} being the last class where $\text{Limit}_n > 0$:

$$\text{EMSRc} = \text{EMSR}_{1,n_{\text{Capacity}}}(\text{Capacity})$$

In the example from Table 3-2, the EMSRc value can be found as:

$$\text{EMSRc} = \big(1 - \text{NormCDF}(135, 135, 22.9)\big) * \$303.19 = \$151.60$$

EMSRb leg/class booking limits are a very robust method when it comes to optimizing protection levels on a single leg. However, in a larger network with connecting passengers, EMSRb does not account for passenger paths that involve more than one leg. Any network optimization method needs to account for the effect a path has on all of the legs that it traverses. ProBP is one of several methods that can be used to address this problem (Bratu, 1998).

ProBP seeks to solve the above mentioned issue of valuing origin-destination itinerary fares (ODIFs) that traverse multiple legs. For ProBP, this is done by iteratively prorating ODIFs over the legs they traverse. In other words, the ProBP algorithm allocates a fraction of the whole ODIF to each leg it traverses, and it does this iteratively. This proration is done using the aforementioned EMSRc values, which are generated from the previous iteration in the ProBP process, except for in the first iteration of ProBP, where an ODIF can be arbitrarily divided evenly among its legs. For an ODIF f_j using path j, the proration for each leg $m \in L_j$ is performed by using the following equation:

$$\mathrm{PRF}_{j,m} = \frac{\mathrm{EMSRc}_m}{\sum_{m \in L_j} \mathrm{EMSRc}_m} * f_j$$

After the above formula has been used for all ODIFs, the next step in each iteration is to determine EMSRc on each leg using the newly prorated fares (or, in the first iteration, using the evenly distributed fares). Once an EMSRc value is determined for each leg, a convergence criterion that specifies a user defined minimum for change in EMSRc values between iterations is checked. If EMSRc values for all of the flight legs in the airline's network change by less than the convergence criterion when compared to the previous iteration, the ProBP algorithm is complete and returns a converged EMSRc bidprice for each leg. If the criterion is not met, the proration process iterates again with the updated EMSRc bidprices.

After the completion of the algorithm, the bidprices generated by ProBP can be used in a bidprice control method. Unlike a seat protection method, such as EMSRb, a bidprice control method does not attempt to determine how many seats of each fare class to protect, but instead establishes a minimum value above which fare classes remain open and below which fare classes close. EMSRc is effectively a bidprice, so the final EMSRc values are used to determine availability. In the case of ProBP, an ODIF is allowed to be available if its value is greater than the sum of the EMSRc values on the legs it traverses.

As an example of ProBP and the following bidprice control, consider an airline using an unrestricted fare structure with two flights, one from AAA–BBB with 20 seats available and one from BBB–CCC with 15 seats available. This airline serves three markets: the two local markets and AAA–CCC. The fares offered by this airline are shown in Table 3-3.

Fare Class	AAA–BBB	BBB–CCC	AAA–CCC
1	$325	$510	$700
2	$215	$375	$500
3	$150	$225	$300
4	$95	$160	$200

Table 3-3: Class-Based ProBP Fares

Assume that the Q-Forecast (with a FRAT5 of 1.5) for each ODIF is as shown in Table 3-4 and Table 3-5:

Fare Class	AAA–BBB	BBB–CCC	AAA–CCC
1	3.80	3.84	4.20
2	2.79	2.12	1.05
3	10.43	15.90	15.75
4	20.96	16.53	21.00

Table 3-4: Class-Based ProBP Q-Forecast Mean

Fare Class	AAA–BBB	BBB–CCC	AAA–CCC
1	1.38	1.39	1.45
2	1.18	1.03	0.72
3	2.28	2.82	2.81
4	3.24	2.88	3.24

Table 3-5: Class-Based ProBP Q-Forecast Standard Deviation

As the fare structure is unrestricted, fare adjustment should be used to reflect this. The adjusted fares can be found by using the marginal revenue fare adjustment equation shown in Section 3.2.1 on fare classes 2–4 in each market. In each market, fare_Q is equal to the offered fare for class 4. Using a FRAT5 of 1.5, this yields the adjusted fares shown in Table 3-6:

Fare Class	AAA–BBB	BBB–CCC	AAA–CCC
1	$325	$510	$700
2	$146	$260	$356
3	$81	$110	$156
4	$26	$45	$56
MR	$69	$115	$144

Table 3-6: Class-Based ProBP Adjusted Fares

The purpose of ProBP is to determine an EMSRc bidprice for each leg. Local traffic can be assigned to its corresponding leg for this purpose, but the AAA–CCC fares must be prorated across the two legs. For most iterations in ProBP, this proration is done using the last iterations

EMSRc bidprices, but, since there is no iteration before the first iteration, it is simplest to divide the AAA–CCC fares evenly between the two legs. Table 3-7 shows the fares as allocated to each leg by this division.

OD Fare Class	Prorated Fare
AAA–CCC 1	$350
AAA–BBB 1	$325
AAA–CCC 2	$178
AAA–BBB 2	$146
AAA–BBB 3	$81
AAA–CCC 3	$78
AAA–CCC 4	$28
AAA–BBB 4	$26

OD Fare Class	Prorated Fare
BBB–CCC 1	$510
AAA–CCC 1	$350
BBB–CCC 2	$260
AAA–CCC 2	$178
BBB–CCC 3	$110
AAA–CCC 3	$78
BBB–CCC 4	$45
AAA–CCC 4	$28

Table 3-7: Class-Based ProBP Iteration 1 Prorated Fares

These EMSRc bidprices obviously have no previous iteration's bidprices to be compared to, so no convergence criterion, in this case $5, can be met. As a result, ProBP moves to the next iteration. For iteration 2, iteration 1's EMSRc values can then be used to prorate the AAA–CCC fares to the two legs. The proration rates are calculated as follows:

$$\text{AAA–BBB rate} = \frac{\$81.40}{\$81.40 + \$105.47} = 0.436$$

$$\text{BBB–CCC rate} = \frac{\$105.47}{\$81.40 + \$105.47} = 0.564$$

Using these proration rates, new prorated fares, which are shown in Table 3-8, can be determined.

OD Fare Class	Prorated Fare
AAA–BBB 1	$325
AAA–CCC 1	$305
AAA–CCC 2	$155
AAA–BBB 2	$146
AAA–BBB 3	$81
AAA–CCC 3	$68
AAA–BBB 4	$26
AAA–CCC 4	$24

OD Fare Class	Prorated Fare
BBB–CCC 1	$510
AAA–CCC 1	$395
BBB–CCC 2	$260
AAA–CCC 2	$201
BBB–CCC 3	$110
AAA–CCC 3	$88
BBB–CCC 4	$45
AAA–CCC 4	$32

Table 3-8: Class-Based ProBP Iteration 2 Prorated Fares

The EMSRc values from this iteration can be found to be $75.73 and $111.06 for AAA–BBB and BBB–CCC respectively. These values are not within $5 of the previous iteration's bidprices, so the process iterates once again.

The above example terminates after the next iteration with EMSRc bidprices of $73.13 and $113.65. This concludes the ProBP algorithm, and starts the bidprice control step. For local traffic on each of the legs, fare classes are only offered if they are valued above the legs' respective bidprices, while AAA–CCC fares are available only if they are greater than the sum of the two bidprices. Table 3-9 and Table 3-10 show which fares are made available (in green) and which are closed (in red) by the bidprice control.

Fare Class	AAA–BBB	BBB–CCC	AAA–CCC
1	$325	$510	$700
2	$146	$260	$356
3	$81	$110	$156
4	$26	$45	$56
Bidprice	$73.13	$113.65	$186.78

Table 3-9: Class-Based ProBP Adjusted Fares Open and Closed

Fare Class	AAA–BBB	BBB–CCC	AAA–CCC
1	$325	$510	$700
2	$215	$375	$500
3	$150	$225	$300
4	$95	$160	$200

Table 3-10: Class-Based ProBP Original Fares Open and Closed

3.2.3 Traditional Class-Based Unbucketed Dynamic Programming (UDP)

EMSRb and ProBP are both "static" optimization methods. Once they generate protection levels or bidprice cutoffs, those solutions apply until the method is used again at a later time in the booking process to re-optimize and produce another, new solution. Additionally, both EMSRb and ProBP do not account for the different possible arrival orders for passengers into their solutions (i.e. they both assume that the bookings for least expensive fares will arrive first and will monotonically increase to the most expensive fares at the end).

As an alternative to EMSRb, the Lautenbacher dynamic program (LDP) offers an optimization approach that is dynamic and explicitly accounts for arrival order of passengers with different WTPs. LDP uses a DP under an assumption of a Markov Process for arrivals to decide what fares to make available (Lautenbacher, 1999). The time until departure is divided into n time slices (not to be confused with time frames) that are small enough that it can be assumed that there will not be more than one passenger request in a time slice. In each time slice, the decision is made whether to accept or reject a passenger based on what fare type they are requesting using the following recursive method.

In the dynamic program, p_{i_n} is defined as the probability that a request for class i (from 1 to m, with class 1 being the highest value fare class and class m being the lowest) will occur in time slice n. The probability that there is no request in n is defined as p_{0_n}. These probabilities are used to weight the expected revenue in the dynamic programming formulation (Lautenbacher, 1999):

$$U_n(x) = \sum_{i=1}^{m} \left(p_{i_n} \cdot \max\{f_i + U_{n-1}(x+1), U_{n-1}(x)\} \right) + p_{0_n} U_{n-1}(x)$$

Boundary conditions:

$$U_0(x) = \begin{cases} 0, & x \leq C \\ -f_1(x - C), & x > C \end{cases}$$

In the above formulation, f_i is the fare from class i, C is capacity, and $U_n(x)$ is the maximum revenue at time slice n (which counts down to zero as it approaches departure) with x bookings made.

The point of a dynamic program is to only accept a booking if the sum of the revenue gained from the booking and the expected revenue from the resulting diminished remaining capacity in the next time slice is greater than the expected revenue with the same capacity at the next time slice. This is solved recursively backwards from departure, or time slice 0, which is where the boundary conditions are implemented. The first boundary condition is straightforward, as it merely states that there is no further revenue to be gained. The second boundary condition exists to ensure the number of bookings does not exceed the declared capacity. The dynamic program, once solved, lends itself nicely to a bidprice control method, where the bidprice for time slice n is equal to $U_{n-1}(x) - U_{n-1}(x-1)$. As a result of its structure, LDP has a solution for every booking level and time slice that could occur, allowing it flexibility in the event of receiving no bookings or a higher value booking than expected.

Like EMSRb, LDP is only usable on a leg level. Unlike EMSRb, though, this is not because of any theoretical impossibility. Instead, the issue for LDP is that its computational complexity is such that, at the scale of a full network, applying it becomes a computational impracticality. However, a technique, known as displacement adjustment, can be used to allow LDP to be used on each individual leg in a network while taking into account the opportunity costs of having a passenger occupy a seat on multiple flights. Displacement adjustment, (Williamson, 1992) involves calculating a network bidprice for each leg, and then using those bidprices to adjust ODIFs that traverse multiple legs. While ProBP could be used to do this, a linear program is often used instead for dynamic programs in network O-D systems. The formulation for the linear program is as follows (where i is OD path, k is fare class, $R_{i,k}$ is the ODIF, $X_{i,k}$ is the number of passengers carried per ODI, $d_{i,k}$ is the demand for ODI, l is each leg, and C_l is the capacity of each leg):

$$\max \sum_{i,k} R_{i,k} X_{i,k}$$

s. t.

$$X_{i,k} \le d_{i,k}, \qquad \forall i, k$$

$$\sum_{i,k} X_{i,k} \le C_l, \qquad i \ni l, \qquad \forall i$$

On each leg, all ODIFs traversing the leg have their fares adjusted by subtracting the bidprices of the second set of constraints from the linear program of all legs other than the one in question. Thus, on leg l, any local fares remain unadjusted, but an ODIF that traverses both leg l and leg $l+1$ will have its fare reduced on leg l by the bidprice of leg $l+1$ for the purposes of leg optimization. After all fares have been adjusted, the optimization method, LDP in this case, is used on each leg, and ODIFs are made available if and only if the value of the ODIF is greater than the sum of the time/capacity bidprices of all of the legs traversed by said ODIF. When displacement adjustment is used in conjunction with LDP and no bucketing, or grouping, of fare classes is performed, the resulting method is known as Unbucketed Dynamic Programming (UDP).

3.2.4 Class-Based ProBP and UDP for Continuous Pricing

The ProBP and UDP methods discussed thus far have both been used in a traditional, “fixed” pricing structure. In such a pricing structure, the only fares that can be offered are the fares that are filed and pre-assigned by the airline. However, both of those methods can be relatively easily adapted into a continuous pricing structure, one in which fares are quoted from a continuous range of prices rather than fixed, pre-prescribed prices, in an unrestricted fare.

Functionally, Class-Based Continuous ProBP and UDP are very similar to their traditional class-based counterparts described in the previous section, until their fare quotation step. Both can use Q-Forecasting to scale historical booking to an equivalent number of lowest fare class bookings. Even though the historical bookings will have fare values from a continuous range of prices instead of fixed price points, the WTP formula for Q-Forecasting described previously can still be used for scaling. Forecasting, detruncating, and repartitioning of the Q-Forecasts can then occur as in the traditional class-based cases. Likewise, the fare adjustment and seat allocation and bidprice optimization steps for the class-based continuous methods are identical to those for the traditional class-based methods. It is at this point, the fare quotation step, where the class-based continuous and traditional class-based methods diverge.

One effect of removing restrictions is that there is essentially only one fare offered at a time, even in a traditional fixed pricing scheme. While any fare from the lowest available to the highest fare filed are available at any given time, a lack of restrictions means that any rational passenger will never buy any fare but the lowest offered. However, given that both ProBP and UDP are bidprice control methods, in a traditional fare structure, this one fare being offered may not be the fare that, according to the optimizer, will maximize revenue. Hypothetically, offering fares at the bidprice should be a better strategy than traditional class-based RM in a given time frame (given the fare structure), since the bidprice would be raised if it were below the optimal fare or lowered if it was higher. This offering of fares at the bidprice is what both Class-Based ProBP and UDP for Continuous Pricing do, although, if marginal revenue fare adjustment was used, the offered fare is instead the sum of the bidprice and marginal revenue (in order to prevent extreme results, both class-based continuous methods are subject to a bound which never allows a quoted fare lower than the lowest fare filed in a market).

Consider the previous example of traditional Class-Based ProBP. If the input parameters were instead to be used for Class-Based Continuous ProBP, the process would proceed exactly as it previously did until arriving at the final bidprices of \$73.13 and \$113.65, with a sum of \$186.78 for AAA–CCC. These bidprices are marginal revenue adjusted bidprices, and so they must be untransformed by adding back in the marginal revenue for each path in order to generate a fare to quote. Since the marginal revenue adjustments were \$69, \$115, and \$144 for AAA–BBB, BBB–CCC, and AAA–CCC respectively, a seat on AAA–BBB is offered for \$142.13 (\$73.13 + \$69), a seat on BBB–CCC is offered for $228.65 ($113.65 + $115), and a seat on AAA–CCC is offered for $330.78 ($186.78 + $144).

3.2.5 Classless Revenue Management

While Class-Based RM for Continuous Pricing does provide the airline with more flexibility in terms of price offerings in an unrestricted fare structure, there is an issue that the "classes" used in forecasting and optimization are arbitrary. A truly classless method for continuous pricing, where forecasting and optimization are also independent of any filed fare structure, would theoretically be better than the above mentioned class-based methods. One issue that must be addressed by any classless method is how to handle the fact that forecasted demand is no longer separable by fare class value. In terms of forecasting, this problem is addressed as it was for the class-based continuous methods: by using the Q-forecasting framework to convert any booking from where it occurred on the WTP curve into a Q-Booking. From there, two classless optimization methods have been developed in recent MIT PODS Consortium research: one based on ProBP and one based on UDP.

3.2.5.1 Classless ProBP

On an individual itinerary level, Classless ProBP seeks to optimize the fares in each time frame by determining what offered fare maximizes expected revenue based on the sell-up curve used in the forecaster at each time frame. At the same time, it also seeks to modify this optimization by accounting for the displacement that a passenger booking in one time frame will have on passengers booking in others. It does this using some of the concepts from ProBP, namely using an EMSRc bidprice and iterating said bidprice until convergence, but it uses these methods over time frames rather than over fare classes. The exact algorithm, Classless ProBP, developed by Hopperstad, used to do this is as follows.

In each time frame t, the following equation is used is used to determine the fare that maximizes expected revenue in said time frame, f_t^*, based on the probability of a passenger who is willing to pay the base fare, QFare, selling up to the quoted fare, while also accounting for the displacement of bookings at other time frames, represented as the EMSRc bidprice:

$$f_t^* = \underset{f}{\operatorname{argmax}}\big(P(\mathrm{WTP}_t \geq f) * (f - \mathrm{EMSRc})\big)$$

Where the probability of sell-up from the base fare to the quoted fare in time frame t is:

$$P(\mathrm{WTP}_t \geq f) = \min\left[1, e^{\frac{\ln(0.5)(f-\mathrm{QFare})}{(\mathrm{FRAT5}_t-1)(\mathrm{QFare})}}\right]$$

In this case, the FRAT5 value is whatever was used for the given time frame in the forecaster. Like in ProBP, EMSRc initializes as zero in the first iteration, and the value from the previous iteration is used in future iterations.

After the optimal fare given the current EMSRc value is found for each time frame, the demand at the "optimal" fares in each time frame must be determined. From the forecasting step, a Q-Demand mean and standard deviation (μ_{t_Q} and σ_{t_Q} respectively) should be known for each time frame. The demand mean and standard deviation at the quoted fare for each time frame (μ_t^* and σ_t^* respectively) can be determined by scaling from these Q-Demands using the WTP equation.

$$\mu_t^* = \mu_{t_Q} P(\mathrm{WTP}_t \geq f^*)$$

$$\sigma_t^* = \sqrt{\sigma_{t_Q}^2 \, P(\mathrm{WTP}_t \geq f^*)}$$

Once the optimal fares and their respective demand means and standard deviations have been determined for each time frame, EMSR-based logic can be applied to the time frame fares to determine the EMSRc value for each leg, given the time frame fares (this is similar to calculating EMSRc values for the purposes of proration in the Class-Based ProBP framework, except that for Classless ProBP the fares used are based on time frame rather than on path and class). After each iteration, the EMSRc value for each leg is checked against the past iteration's EMSRc value. If their difference falls within a convergence criterion for all legs in the airline's network, then the fares are considered optimal, and the fare from the current time frame is quoted to passenger (Classless ProBP, like the class-based continuous RM methods, is also, in the experiments documented in this thesis, bounded between a highest and lowest quotable fare). If the difference does not fall within said criterion, the process repeats for another iteration.

As an example of the Classless ProBP algorithm, consider a flight with three time frames and five capacity remaining. At a base fare of \$100, the forecaster for this flight has forecast an average ten demand with a standard deviation of five in each of the three remaining time frames. Additionally, the forecaster has been set to assume FRAT5s of 1.5, 2.5, and 3.0 for the remaining time frames. For the sake of this example, a convergence criterion of \$10 will be assumed (that is, the EMSRc value change between iterations must be less than $10 for the algorithm to quit). The EMSRc bidprice initializes as zero, and after the first iteration, the fares and demands are shown in Table 3-11.

TF	Optimal Fare	Demand Mean	Demand Standard Deviation
1	$100.00	10.00	5.00
2	$216.40	5.84	3.82
3	$288.54	5.20	3.61

Table 3-11: Iteration 1 of Classless ProBP Example

The EMSRc bidprice from the first iteration is found to be $219.07. Obviously, the difference between $219.07 and $0.00 is greater than $10.00, so Classless ProBP iterates with an EMSRc bidprice of $219.07 in the optimization step. The result of iteration 2 can be seen in Table 3-12.

TF	Optimal Fare	Demand Mean	Demand Standard Deviation
1	$291.21	0.71	1.33
2	$435.48	2.12	2.30
3	$507.61	2.43	2.47

Table 3-12: Iteration 2 of Classless ProBP Example

The EMSRc bidprice as found after iteration 2 is $237.73. Once again, as $|\$237.73 - \$219.07| = \$18.66 > \$10.00$, the convergence criterion is not met, and the algorithm advances to iteration 3, which uses an EMSRc bidprice of $273.73.

After four iterations, the bidprice is found to be $234.39 having not yet converged. Iteration 5, with optimal fares found using said $234.39 bidprice, can be seen in Table 3-13.

TF	Optimal Fare	Demand Mean	Demand Standard Deviation
1	$306.52	0.57	1.19
2	$450.79	1.98	2.22
3	$522.93	2.31	2.40

Table 3-13: Iteration 5 of Classless ProBP Example

In this case, the bidprice is found to be $226.40. As $|\$226.40 - \$234.39| = \$7.99 \leq \$10.00$, Classless ProBP now terminates and quotes a fare of $306.52.

Classless ProBP's algorithm can be scaled to a network algorithm by modifying the EMSRc term in the fare optimization equation so that it includes all legs m traversed by origin-destination itinerary j in time frame t.

$$f_{j_t}^{*} = \underset{f}{\operatorname{argmax}}\left(P\left(\mathrm{WTP}_{j_t} \geq f\right) * \left(f - \sum_{m \in L_j} \mathrm{EMSRc}_m\right)\right)$$

Each time frame ODIF is then prorated across the legs it traverses using the ProBP formula in order to carry out the calculation of new EMSRc values. A visualization of the network Classless ProBP process is shown in Figure 3-1.

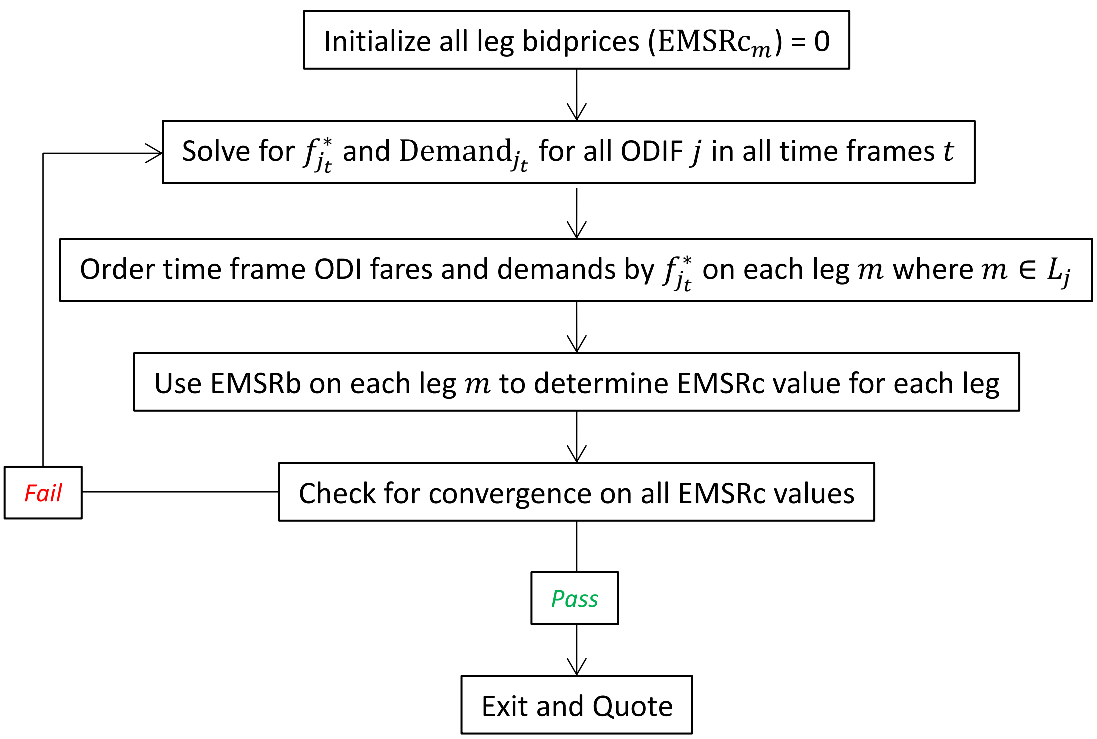

Figure 3-1: Classless ProBP Flow

3.2.5.2 Classless UDP

The other method that has been developed within the MIT PODS Consortium for Classless RM is based on UDP. As with Class-Based UDP, Classless UDP uses a network linear program and then a dynamic program on the individual legs. The use of a dynamic program, in theory, allows Classless UDP to determine the optimal fare at each time frame accounting for expected bookings and revenue later in the booking process, which is similar in purpose to the iterative EMSRc calculation for Classless ProBP. Unlike Classless ProBP, Classless UDP does not require an iterative process, and, since it is a dynamic program, should require less frequent reoptimization than Classless ProBP.

The dynamic program used by Classless UDP is similar to the Lautenbacher dynamic program used in Class-Based UDP. Instead of attempting to determine the probability a passenger will buy into a fare class, however, the dynamic program for Classless UDP seeks to pick a fare in each time slice that maximizes revenue over the entire booking window. The DP, for each time slice n, is constructed as follows:

$$U_n(x) = \max_f \left(p_{Q_n} \left(P\left(\mathrm{WTP}_j \geq f\right)\left(f + U_{n-1}(x+1)\right) + \left(1 - P\left(\mathrm{WTP}_j \geq f\right)\right) U_{n-1}(x) \right) + \left(1 - p_{Q_n}\right) U_{n-1}(x) \right), \qquad n \in j$$

$$P\left(\mathrm{WTP}_j \geq f\right) = \min\left[1, e^{\frac{\ln(0.5)(f - \mathrm{QFare})}{(\mathrm{FRAT5}_j - 1)(\mathrm{QFare})}}\right]$$

Boundary conditions (same as Lautenbacher DP):

$$U_0(x) = \begin{cases} 0, & x \leq C \\ -r_1(x - C), & x > C \end{cases}$$

In the above DP, p_{Q_n} is the probability of the arrival of a passenger willing to pay at least the base fare, QFare. The fare quoted at each time slice for Classless UDP is equal to whatever f maximizes the expected revenue to come at each time slice and number of bookings, $U_n(x)$, with this once again being based on a prospective arrival being willing to sell-up from the lowest fare to the quoted fare. The boundary conditions, as with Class-Based UDP, are designed ensure that no more revenue is expected at the final time slice before departure and that bookings, x, are never allowed to exceed capacity, C. As with the other methods for continuous pricing, classless UDP is also constrained to a maximum fare (r_1) and a minimum fare (the base fare) in the experiments documented in this thesis.

As with Class-Based UDP, the dynamic program for Classless UDP does not scale well to application across an entire network. A similar solution is employed for Classless UDP as was employed for Class-Based UDP, where a network-wide linear program (LP) is performed to determine the displacement costs of each leg. However, since Classless UDP does not have filed fares to feed into the LP, it first solves the following equation for every OD path i to determine a fare.

$$f_i^* = \text{argmax}(\text{WTP} \geq f_i) * f_i$$

Similar to Classless ProBP, a mean demand can be found for this fare by multiplying a demand at some base fare by the sell-up rate from the base fare to f_i^*. This demand and fare can then be used to solve the LP from Class-Based UDP and determine displacement costs for the leg-level DP above.

3.3 Conclusion

In this chapter, forecasting and seat allocation methods pertinent to the development of continuous pricing algorithms were described. First, forecasting for unrestricted fare structures, including Q-Forecasting and marginal revenue fare adjustment were discussed. Then, the Class-Based and Classless seat allocation optimization methods that will be further investigated in Chapter 5 were described. The next chapter will describe the simulation software and parameters used to test these seat allocation optimization methods.

Chapter 4: Simulation Methodology

Having established forecasting and optimization models for continuous pricing, the method used to test them will be discussed in this chapter. Specifically, there are two main components to this methodology. Firstly, there is the simulation software, known as the Passenger Origin Destination Simulator (PODS), which incorporates several features into its simulation methods not typically used in revenue management (RM) research. Secondly, there are the network and the input parameters that were used to test each RM method.

4.1 PODS Software Overview

PODS is an application that simulates the passenger booking process up to flight departure date for a hypothetical network with multiple airlines, each using their own fare rules and revenue management systems. Included within in it are two main components: the airlines' RM systems and the passenger choice model. The airline RM systems can implement what has been discussed in the previous chapters: forecasting, optimization, and fare quotation. While the optimization portion of PODS's RM system is fairly standard in terms of RM research simulators (although it does incorporate many different revenue management methods and parameters into the same program), the forecasting component is different. Most simulation software used in the literature allows the optimization method to know with complete certainty the distribution of demand instead of implementing a method that attempts to determine from previous data what the demand will be. For example, Gallego & van Ryzin (1997) and Zhang & Lu (2013) both test their continuous pricing models in simulators, but there is no forecasting component to either of their tests. Most of the other major component of PODS, the passenger choice model, is unique as far as RM research is concerned.

4.1.1 Passenger Choice Model

Most RM research involving seat allocation optimizers assume the optimizers are given a known demand or, in the case of dynamic programming, might give basic considerations to the probability of a passenger appearing with a willingness-to-pay (WTP). The passenger choice model in PODS, however, generates individual passengers that all behave rationally and with their own WTPs, preferences for which paths they prefer to use to get from their origin to their destination, and disutilities when restrictions are a factor. These preferences are generated from pre-assigned distributions, which are typically described by different parameters for leisure passengers, who generally arrive earlier, are willing to pay less, and have a low sensitivity to restrictions, and for business passengers, who typically arrive closer to departure, have a higher WTP, and have a higher sensitivity to restrictions. The arrival rates assumed for these passengers over the booking process are shown in Figure 4-1.

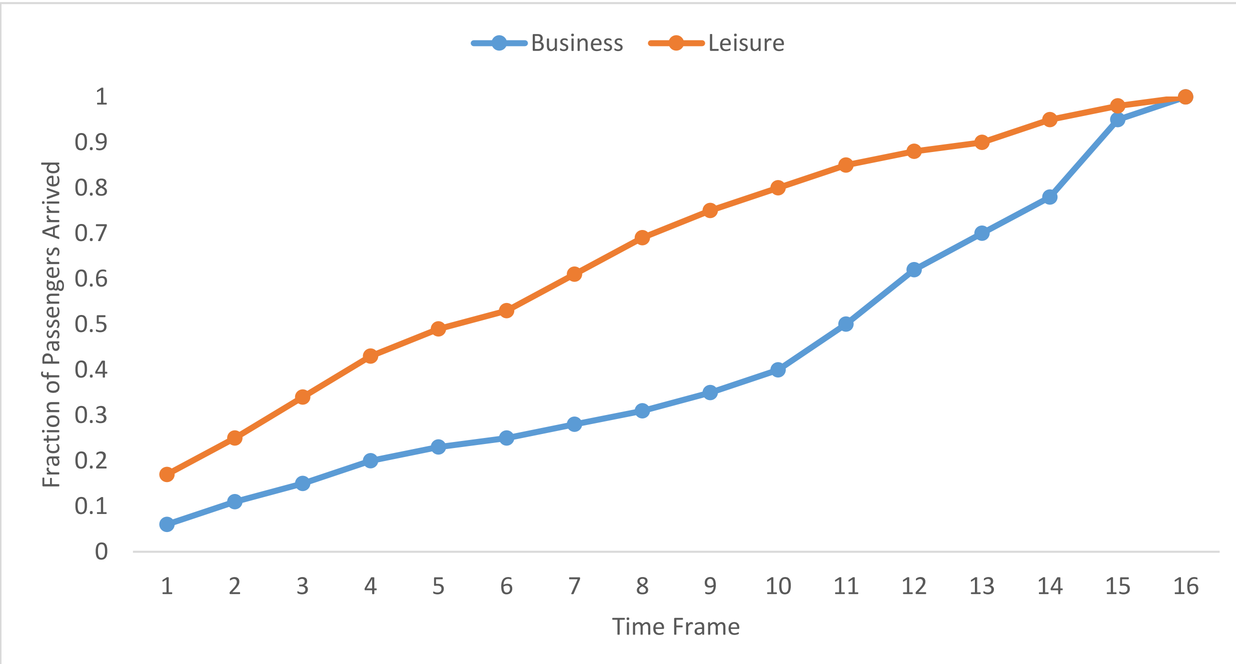

Figure 4-1: Arrival Rate of Business and Leisure Passengers over the Booking Process

The distribution of behavior of each passenger type is decided by several parameters. For every market, each passenger type has a base fare which it is assumed that all passengers of said type will be willing to pay. At this base fare, each passenger type is also given a base demand, which is the average expected demand for that passenger type at the base fare. Typically, the base fare for leisure passenger is set to be identical to the lowest fare offered for the O-D pair while the fare for the business passenger is set 2.5 times higher. The base demands are typically set to favor leisure passengers by roughly a two-to-one ratio.

Every passenger generated of type i (business or leisure) has also has a maximum WTP. The distribution of this WTP is defined by the following equation (defined by the probability a passenger is at least willing to pay f):

$$P(\text{WTP} \geq f)_i = \min\left[1, e^{\frac{\ln(0.5)(f - \text{base fare}_i)}{(\text{emult}_i - 1)(\text{base fare}_i)}}\right]$$

It may be noted that this equation takes on the same form as the sell-up estimation calculation from the forecasting section. It may be said that the above equation defines the sell-up rate from the base fare of passengers generated by the simulation. The "emult" value is a FRAT5 value, but, instead of estimating the sell-up rate for the forecaster, it defines the sell-up rate for the passenger generator. In order to reflect the higher WTP and lower sensitivity to price increases of business passengers, the emult is typically set to be higher for business than for leisure passengers. An example of business and leisure WTP curves are shown in Figure 4-2.

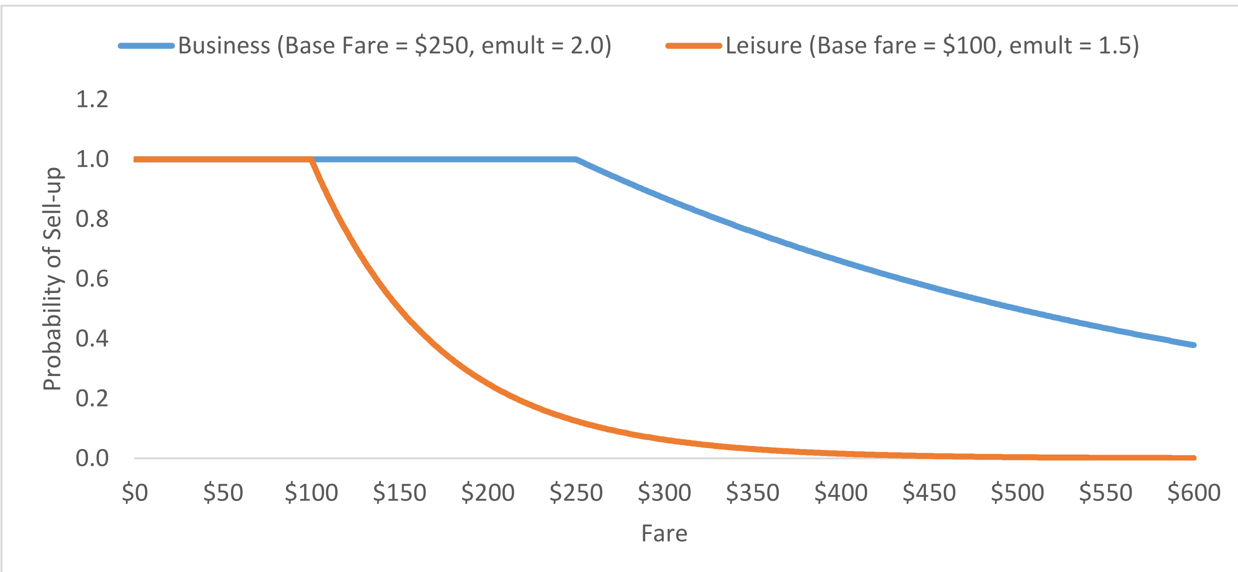

Figure 4-2: Demonstration of Sample Business and Leisure WTP Curves

During a single run of the simulation process, the base demand for each passenger type in each is modified by three numbers randomly generated from a normal distribution, one for each system, market, and passenger type, to determine for a departure day whether demand for these three categories is high or low. The modified value then represents the mean expected demand for that departure day for that passenger type in that market. Each passenger is then assigned to arrive in a given time frame is then generated with WTP properties defined by a stochastic selection from their willingness-to-pay curve. In unrestricted fare structures, like the ones used to test the various methods for continuous pricing, a passenger, given two equivalent paths, will always buy the lowest fare available, as long as the price is below their maximum WTP. If no fare lower than their maximum WTP is available, the passenger will not buy anything.

Perhaps the most important result of this passenger choice is that it allows for the effects of passenger choice and competition to be observed. While there are provably optimal solutions for RM in monopolistic settings (with known stochastic demand), once competition is introduced, those optimal strategies may generate less revenue than strategies that had been theoretically weaker in a single carrier framework. For example, in a two-airline competitive case, it would be possible that one airline could slightly undercut an airline using the "optimal" monopolistic strategy, gain a substantial slice of bookings, and end up generating far more revenue than its competitor.

4.1.2 Simulation Process

In terms of how the PODS software actually works, each "trial" typically consists of 600 "samples", with each sample containing the entire booking process of passenger choice and RM for a single departure day. Each departure day takes place within a network, which may have several airlines operating in many markets. The presence of multiple airlines results in competitive effects driven by the passenger choice model. Of the 600 samples, the first 200 only exist to generate a historical bookings database and are then discarded. In order to replicate RM systems, each sample's booking process is divided into 63 days, with the first day being when the booking process begins, and those days are lumped into 16 time frames of unequal length, which decrease in length as the day of departure approaches. The time frames used in the PODS software are designed to replicate the data collection points used by real world airlines in their revenue management systems.

Over the course of the simulation process, passengers are generated and make decisions about whether or not they can afford the available fares currently offered by for their desired origin and destination as well as which ODIF is preferable, which are then recorded by the RM system in the form of bookings traditionally by fare class or, in the case of continuous pricing models, by fare amount. It is important to note that what type of passenger, leisure or business, made each booking is not reported to the RM system, although such data is reviewable by the user after the simulation is complete. The RM system then uses its forecasting, optimization, and fare quotation components to send fare availability back to the passenger choice model, and this cycle repeats until departure day (daily for the Probabilistic Bidprice methods and at each time frame for the dynamic programming methods). This process is shown in Figure 4-3.

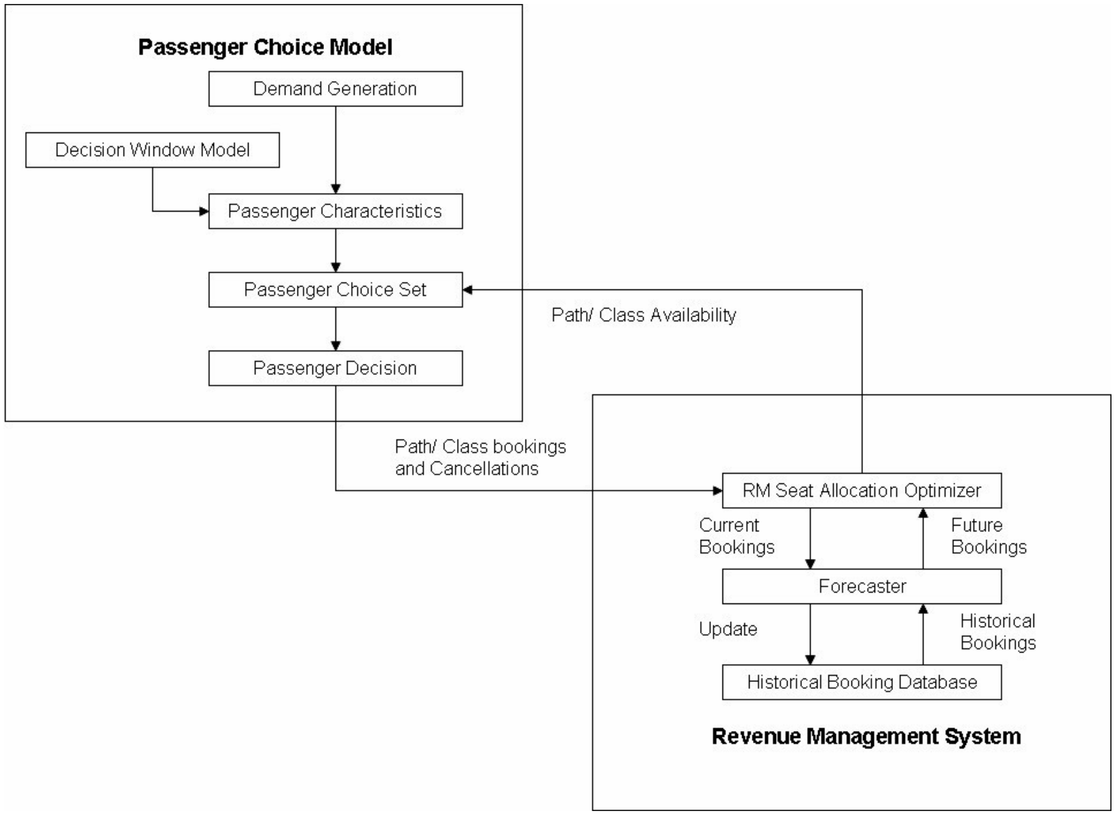

Figure 4-3: PODS Structure (PODS Primer)

The data from each of the 400 samples not discarded is averaged together to complete one trial. For smaller networks, there are typically five trials used in a test with PODS, while for larger networks, such as the one used in this thesis, there are typically two trials used as their larger size reduces expected variation in results.

4.2 Network D6 with an Unrestricted Fare Structure

The network used in all of the tests in this thesis is Network D6. It was designed specifically for use with PODS and is a generic network, as it is not designed to resemble any real world airline network, although the airports and distances between them are based on actual United States cities. D6 is a two-airline network, with each airline connecting a hub with 40 spoke cities. A map of Network D6 is shown in Figure 4-4.

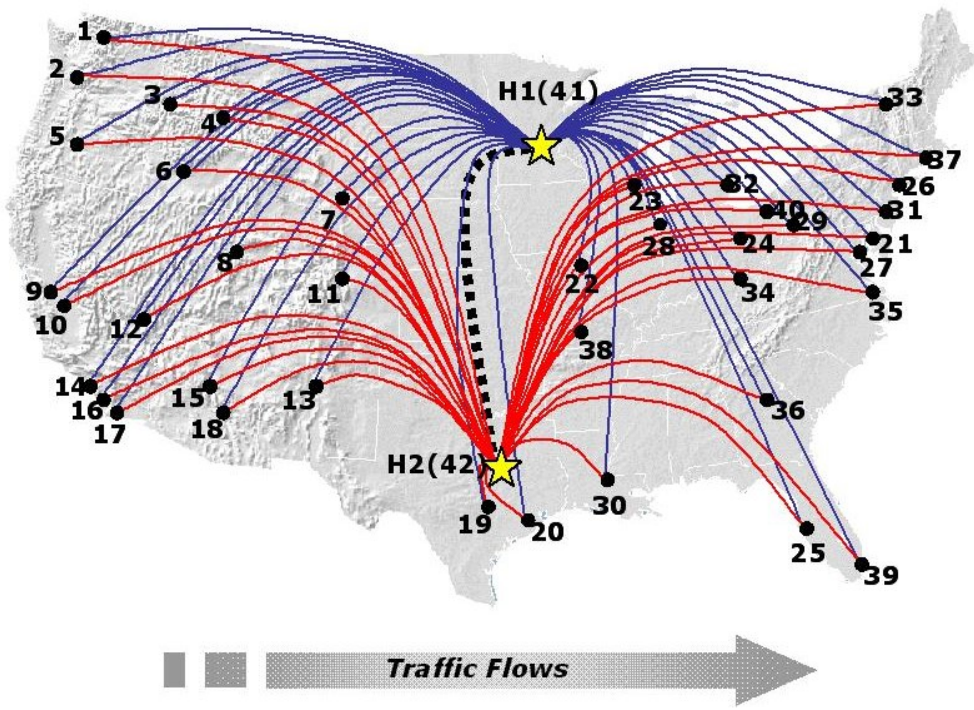

Figure 4-4: Map of Network D6

Note that, in the figure above, traffic can originate at either of the hubs or any of the spoke cities from 1 through 20. Likewise the destination of the traffic can be to either of the hubs or to the spoke cities 21 through 40. Each airline therefore offers legs on 42 unique routes (20 routes from the western spokes to its hub, 20 routes from its hub to the eastern spokes, and then to and from its hub to the other airline's hub), and each route has three flights per departure day. This yields a total of 141 individual legs operated by each airline per departure day. Additionally, there are a total of 482 possible ODs over the entire network, with many of the ODs having multiple possible paths owing to there being multiple routes and three flights a day on each route. A summary of the Network D6 fare levels can be seen below in Table 4-1.

Class	1	2	3	4	5	6
Average Fare	$412.85	$293.34	$179.01	$153.03	$127.05	$101.06
Minimum Fare	$188.33	$136.83	$87.58	$76.39	$65.19	$54.00
Maximum Fare	$742.52	$514.82	$297.02	$247.52	$198.02	$153.00

Table 4-1: Network D6 Fare Structure Summary

There are two major advantages to using Network D6: its size and its relative symmetry. In terms of its size, Network D6 is one of the largest generic networks used in PODS testing. While the network is still much smaller than most real-world networks, it does allow for the inducement of network effects. At the same time, these networks effects are kept also kept from being too complicated, as both airlines in D6 use a pure hub-and-spoke model, and each only have one hub. Perhaps more important than its size, however, is D6's relative symmetry. Both airlines have the same number of flights to the same spoke cities. The only difference between the two airlines is which hub they operate. This reduces the risk of one airline having an inherent network advantage over the other, thereby eliminating such an advantage from distorting experiments strictly focusing on the advantages granted by using particular revenue management models.

It is also worth noting that the fare structures used by both airlines in the experiments for this thesis are completely unrestricted and offer the same fares for the same markets. As previously stated, the continuous pricing tested in this thesis only generates a single fare, and, as a result, cannot work within a restricted fare structure. In order to allow for a fair comparison, even the traditional class-based RM was therefore tested in an unrestricted fare structure. The use of unrestricted fare structures requires the use of Q-Forecasting and fare adjustment by the airlines in the simulation (for reasons discussed in Chapter 3). Past work with Network D6 has indicated that the following three FRAT5 curves (A, C, and E) used by the simulated airlines for Q-Forecasting and fare adjustment typically yield the best revenue results, with FRAT5 C, the moderate of the three, typically generating the best revenue results of the three.

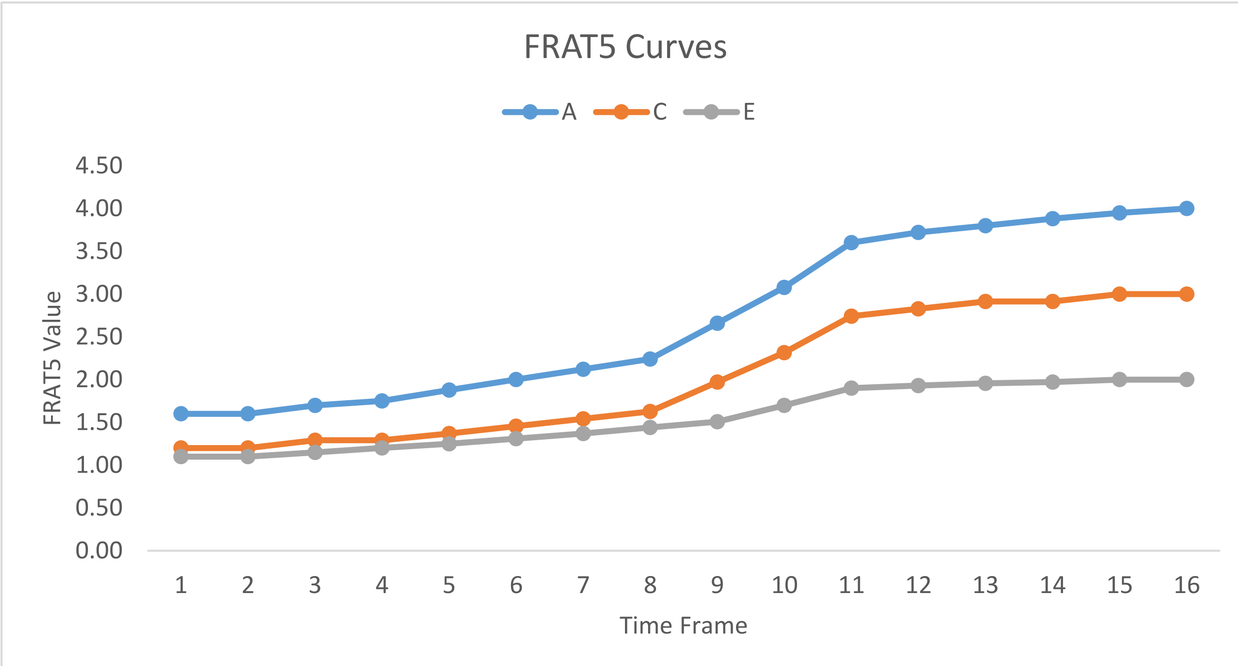

Figure 4-5: FRAT5 Curves Commonly Used for Network D6

4.3 Summary

The PODS software gives several advantages over most previous RM simulation methods, notably its use of a forecasting frame work and a passenger choice model. As a large network with multiple-leg paths and competing airlines, Network D6 takes full advantage of both the forecasting and passenger choice model used in PODS, and could be used to get detailed results about the continuous pricing methods detailed in Chapter 3. In the next chapter, the results of testing the described methods of continuous pricing will be discussed.

Chapter 5: Results of Continuous Pricing Simulations

Having established in the previous chapter the methods of continuous pricing tested and the simulation environment used for testing, this chapter discusses how each method performs. Initially, ProBP-based and UDP-based methods will be considered separately, before the ProBP and UDP methods will then be compared to each other.

The process for testing each group of methods was the same. First, a baseline was established for each method type with a pre-existing traditional class-based revenue management (RM) method. This was done by first determining which of the three previously discussed FRAT5 Curves (A, C, and E) generates the greatest revenue when used with the traditional class-based method being considered. The selected FRAT5 was then used with the traditional class-based method with an increasing number of fare classes. Fare classes were added to Network D6 by adding one, two, or three fares in between the pre-filed fares.

Fares were inserted at prices equidistant from the pre-filed fares and other added points. Thus, when one fare class was added between each set of pre-filed fares, then the value of the fare added between fare class n and fare class $n+1$ (where fare classes are indexed in descending order) was equal to $\frac{f_n+f_{n+1}}{2}$, when two fare classes were added the new fares were of values $\frac{2f_n+f_{n+1}}{3}$ and $\frac{f_n+2f_{n+1}}{3}$, and when three classes were added, $\frac{3f_n+f_{n+1}}{4}$, $\frac{f_n+f_{n+1}}{2}$, and $\frac{f_n+3f_{n+1}}{4}$. The reason for testing the traditional class-based methods in this way is that, while not making them continuously priced, this does give the class-based methods far more price granularity. As increased price granularity is one of the major theoretical advantages for continuous pricing, this would show if increasing fare classes is a good approximation for continuous pricing that still allows the use of a traditional class-based RM system. In addition to establishing how revenue generation changes for traditional class-based methods with additional fare classes, a discussion of the underlying mechanics causing these changes, such as how changes in average paid fares, bookings, bidprices, and forecasts effect final revenue, will also be presented. All of the baseline class-based tests were performed in symmetric scenarios, where both airlines in Network D6 use the same RM system parameters.

After examining the trends of a traditional class-based RM method, the corresponding class-based continuous method was tested in much the same manner. The first set of tests for class-based continuous methods was in symmetric scenarios. Once again, a FRAT5 curve was selected based on which curve maximized revenue in the baseline Network D6. Whichever curve was selected was used for experiments which again increased the number of fare classes used for optimization, and the fare class count was increased in the same way as before. As with traditional class-based RM, changes in underlying mechanics as the number of fare classes increases, as well as the effects of these changes, will be considered. Additionally, the class-based continuous method will be compared to its traditional class-based counterpart. In addition to symmetric tests, the class-based continuous methods were also tested in asymmetric scenarios. In these asymmetric tests, one airline in Network D6 used class-based continuous RM while the other used traditional class-based RM, with this being done in order to simulate the effect of one airline changing its RM method while the other one remains using its traditional class-based method.

After discussing the class-based methods of either ProBP or UDP, experiments with the corresponding classless methods will be discussed. Once again, a FRAT5 was first selected. Since there is no effect of adding fare classes to consider for classless RM, the performance of classless methods of RM will be discussed in the context of their corresponding class-based methods, and, as with the class-based continuous cases, asymmetric tests will be considered along with symmetric tests.

Once the ProBP and UDP methods have been discussed and analyzed independently, the performance of ProBP and UDP methods will be analyzed with respect to each other.

5.1 ProBP Results

This section will discuss the results of experimentation on ProBP. As previously stated, this section will start with traditional Class-Based ProBP, then Class-Based Continuous ProBP, and finally Classless ProBP. All of the experiments were simulated in Network D6, and all tests with ProBP were re-optimized daily throughout the booking process (as is typical with non-dynamic program bidprice control methods).

5.1.1 Traditional Class-Based ProBP in Symmetric Competition

To determine which FRAT5 curve generated the highest revenues with traditional Class-Based ProBP, curves A, C, and E (Figure 4-5) were tested in the standard six fare class unrestricted Network D6 (described in Section 4.2). The FRAT5 curve selection experiments, like all experiments undertaken with Class-Based ProBP, utilize Q-forecasting (described in Section 3.1.2) and marginal revenue fare adjustment (described in Section 3.2.1). The results of testing traditional Class-Based ProBP with inputs of FRAT5 curves A, C, and E are shown in Figure 5-1.

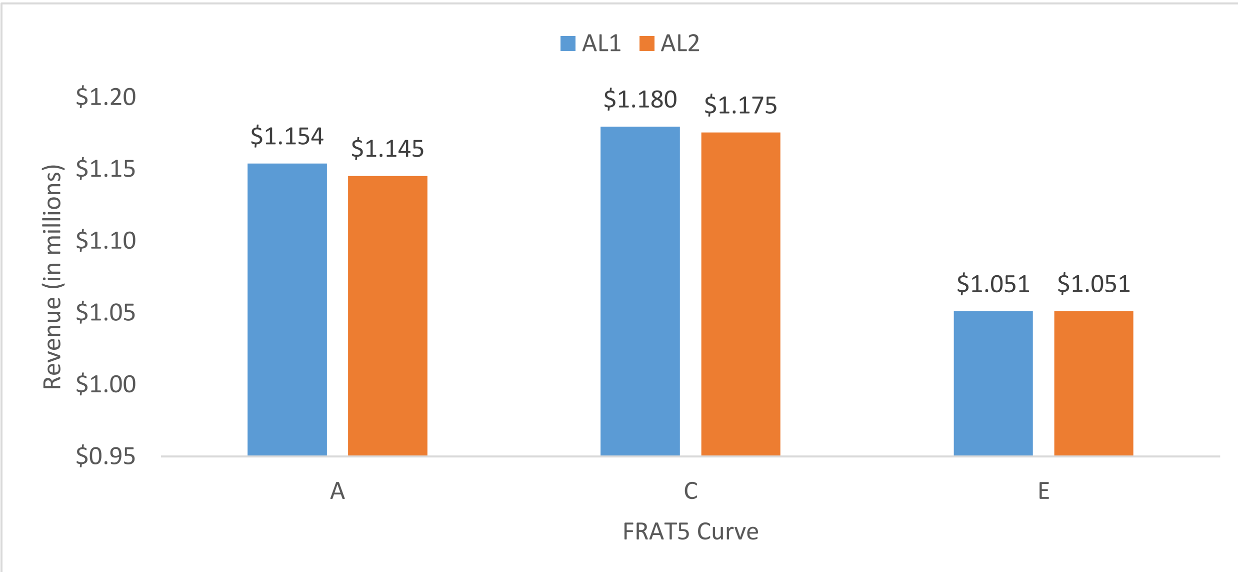

Figure 5-1: 6-Class Traditional Class-Based ProBP Revenue

FRAT5 curve C, the moderate estimate of passenger willingness-to-pay (WTP) of the three, generates the most revenue. Additionally, both Airline 1 (AL1) and Airline 2 (AL2) produce similar amounts of revenue, which is to be expected in symmetrical RM system tests in a near-symmetrical network like D6. As a result, the baseline FRAT5 curve for experiments concerning traditional Class-Based ProBP was curve C.

As previously stated, adding more fare classes should give traditional Class-Based ProBP more granularity in both its optimization and its fare quotation. In turn, this should, combined with the forecasting resulting from using FRAT5 C, increase revenue for traditional Class-Based ProBP. As shown in Figure 5-2, this is, in fact, what happens.

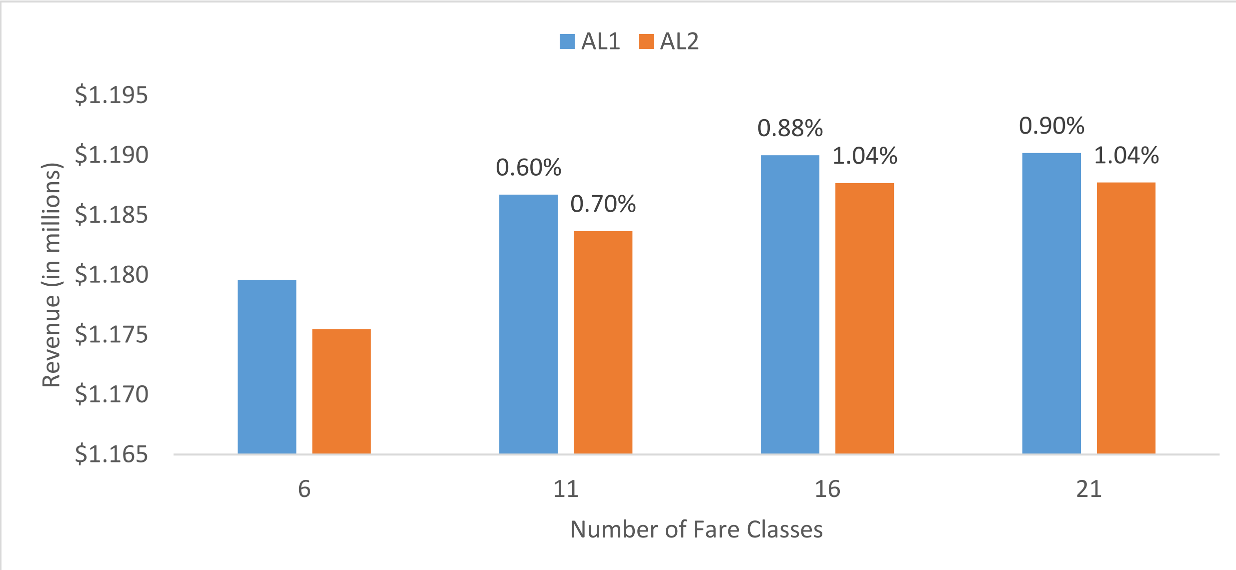

Figure 5-2: Traditional Class-Based ProBP Revenue (percent change in revenue from 6-class experiment)

The revenue increases for traditional Class-Based ProBP with additional fare classes are substantial. Adding five more fare classes increases revenue by over half of a percent. While adding more fare classes does continue to increase the total revenue generated, the amount of revenue generated by each additional fare class diminishes with each extra class. This would seem to indicate that additional fare classes do improve traditional Class-Based ProBP revenue by granting it greater granularity in fare generation. This may be confirmed by observing the average paid fare by time frame (TF) for each number of fare classes.

Figure 5-3 shows how the adding of price points helps improve Class-Based ProBP's revenue. With six fares, the average paid fares are similar for all fare classes until TF 9. Between TF 9 and 10, however, the 6-class case average paid fare jumps up far more rapidly than the cases with more classes. Referring to the Network D6 average fare for each fare class (Table 4-1), it is shown that, in between TFs 9 and 10, the average paid fare crosses between the average class 3 and class 2 prices. Because traditional class-based RM methods can only offer fares at pre-defined price points, traditional Class-Based ProBP can only offer fare class 2 or 3 and not any price in between. Adding fare classes alleviates this problem, and a much more steady increase in average paid fares is shown in Figure 5-3 for the 11-, 16-, and 21-class cases.

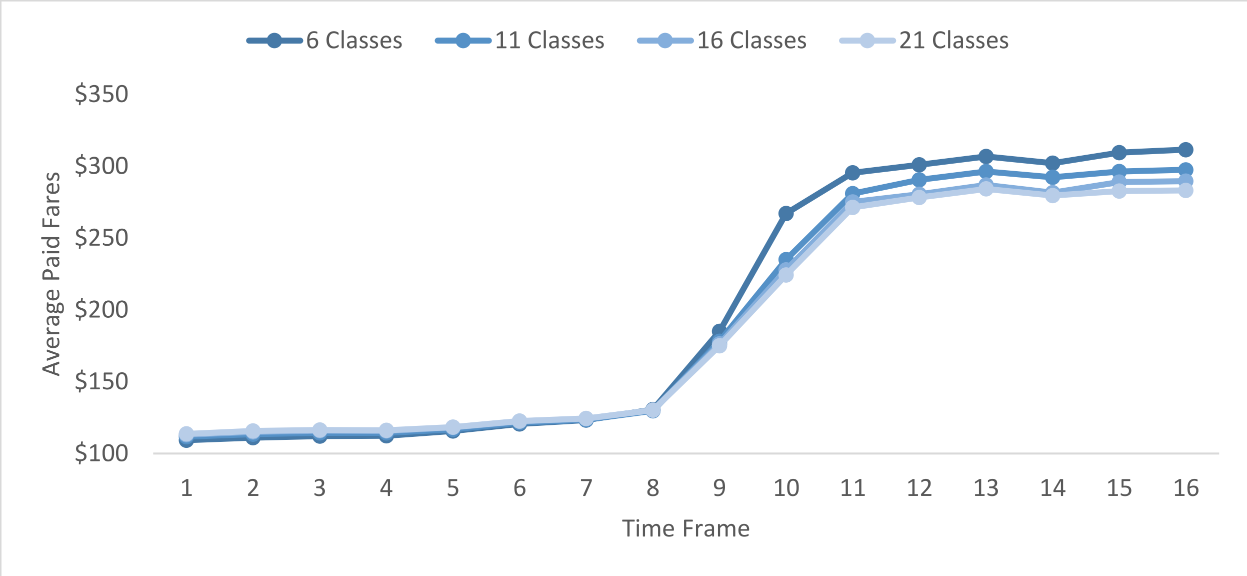

Figure 5-3: Airline 1 Traditional Class-Based ProBP Average Fares

The effect of more fare classes allowing more granularity in fare quotation is illustrated by average bookings by TF (Figure 5-4). As a result of not having too high of fares in the later TFs, the 11-, 16-, and 21-class cases end up having much higher bookings in later TFs. Having more bookings in later, more-expensive TFs increases the forecasts and bidprices in early TFs (Figure 5-5 and Figure 5-6) in an RM feedback effect, captured by the PODS simulator.

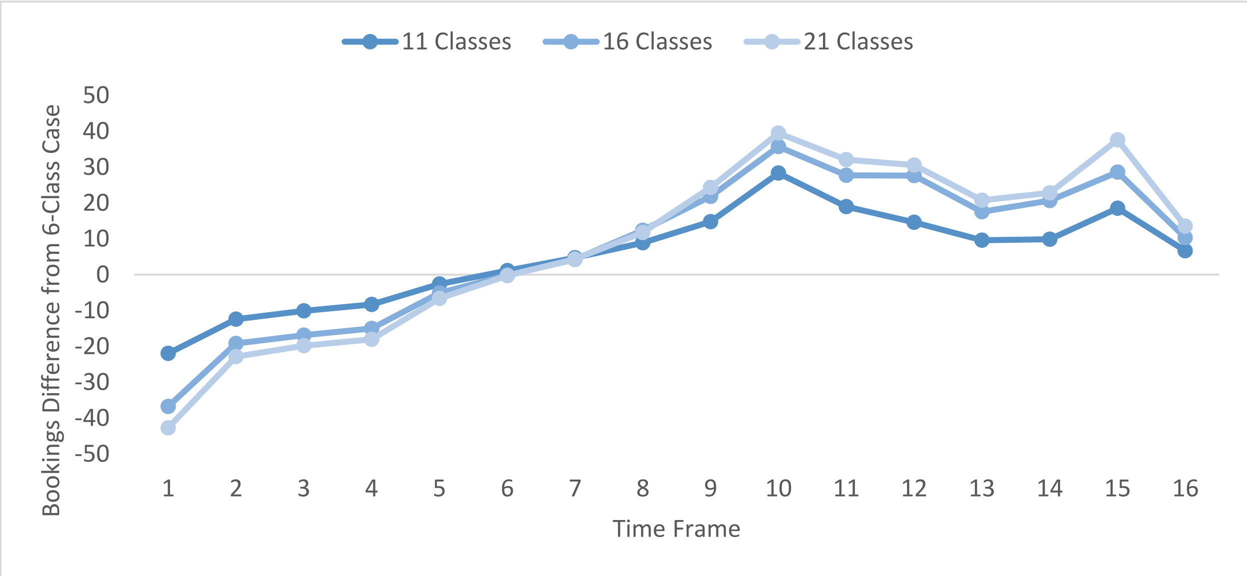

Figure 5-4: Airline 1 Traditional Class-Based ProBP Bookings Difference from 6-Class Case

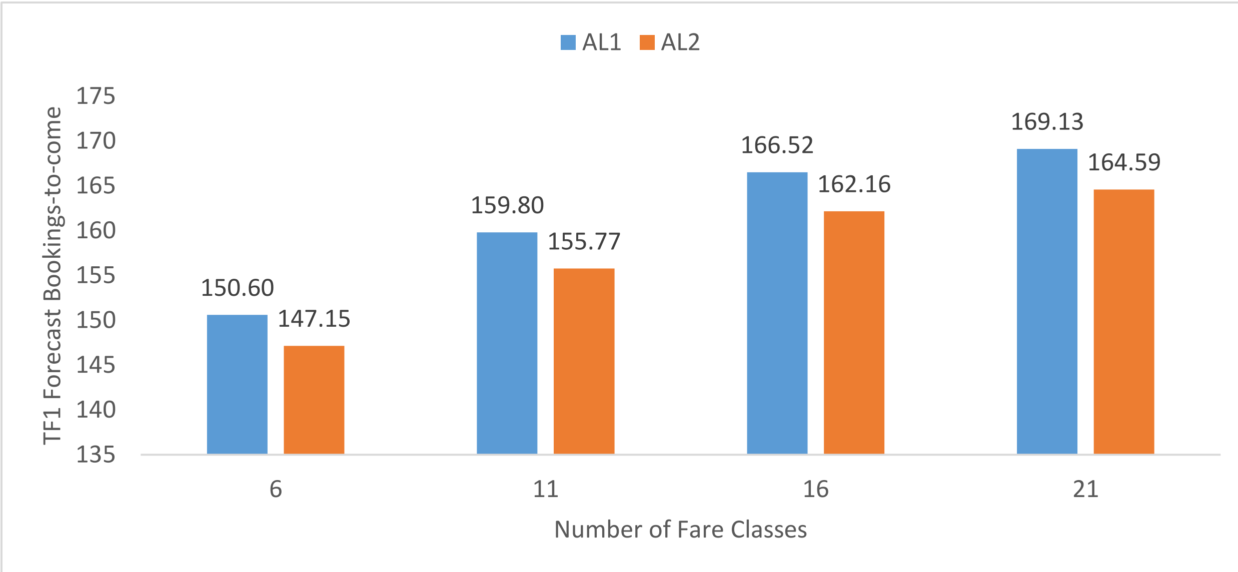

Figure 5-5: Traditional Class-Based ProBP Time Frame 1 Forecast Bookings-to-come

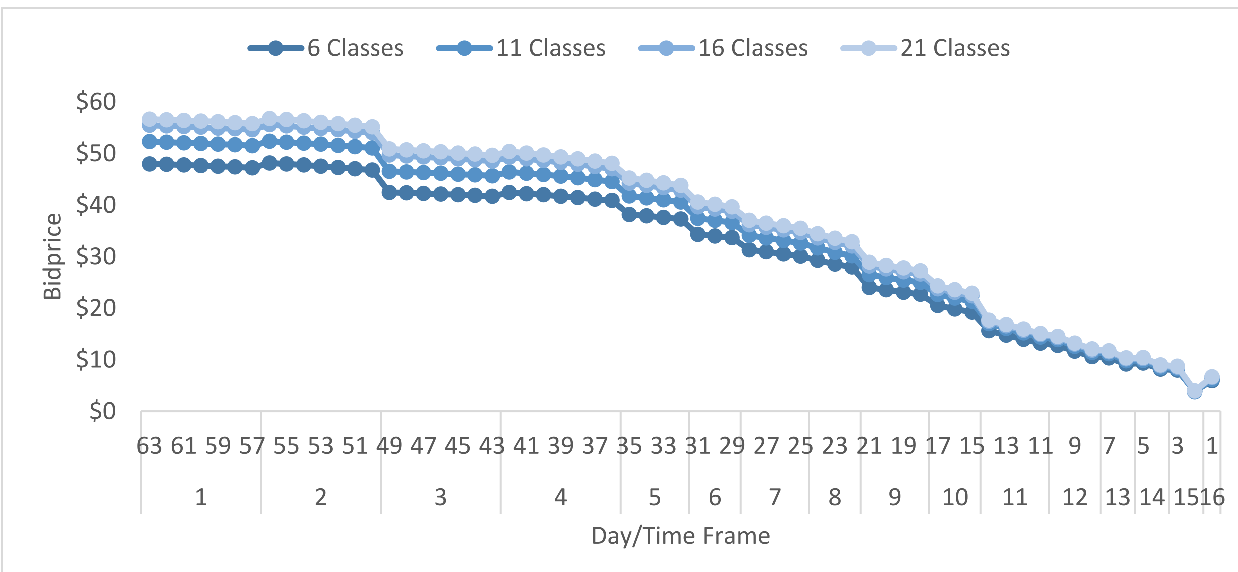

Figure 5-6: Airline 1 Traditional Class-Based ProBP Bidprices

These increased forecasts and bidprices then result in the higher fares in earlier TFs shown in Figure 5-3, which results in fewer, less valuable early TF bookings (Figure 5-4) which further encourages bookings to be shifted later to when they are more valuable.

5.1.2 Class-Based ProBP for Continuous Pricing in Symmetric Competition

Having established the performance of traditional Class-Based ProBP, Class-Based Continuous ProBP in symmetric competition will be examined in this subsection. The tests on the

Class-Based Continuous ProBP method were identical to those on traditional Class-Based ProBP, starting with a "best" FRAT5 determination followed by experiments with different numbers of fare classes. Although Class-Based Continuous ProBP already has unlimited price granularity, the optimization step may benefit from the additional fare classes, as increasing the number of fare classes more closely aligns the seat protection optimization process with unlimited price granularity provided by continuous pricing.

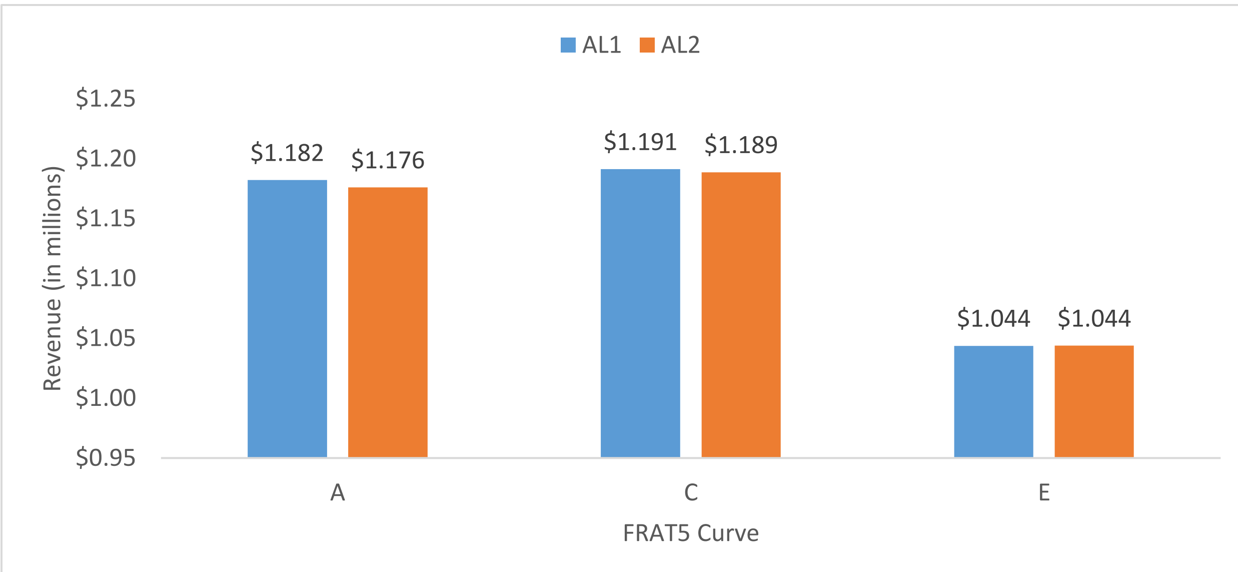

Figure 5-7: 6-Class Class-Based Continuous ProBP Revenue

Figure 5-7 shows that, as for traditional Class-Based ProBP, Class-Based Continuous ProBP revenue is maximized in the 6-fare class case with airline estimates of WTP equal to FRAT5 C. As such, FRAT5 C was used in all Class-Based Continuous ProBP tests.

Figure 5-8: Class-Based Continuous ProBP Revenue

Hypothetically, increasing the number of fare classes could increase Class-Based Continuous ProBP's revenue in much the same way that it increased traditional Class-Based ProBP's. A finer partition of the forecast for the optimization step could help the effectiveness of the optimization step be less dependent on the fare structure being used, although using too fine of a partition could render forecast partitions so small as to affect their forecast accuracy. Figure 5-8 shows that revenue actually decreases slightly for Class-Based Continuous ProBP as fare classes are added. The likely reason for this is shown in the forecasts (Figure 5-9).

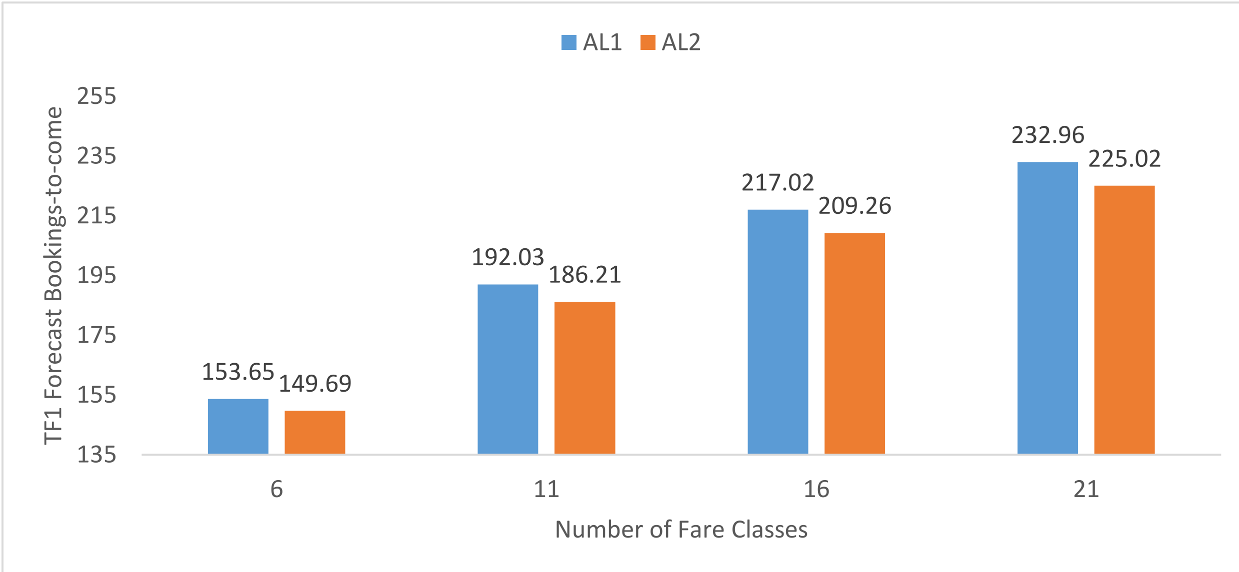

Figure 5-9: Class-Based Continuous ProBP Time Frame 1 Forecast Bookings-to-come

Class-Based Continuous ProBP forecasts increase very rapidly between the 6-class and 21-class case. This, in turn, causes bidprices to rise and bookings to drop in early TFs (Figure 5-10 and Figure 5-11).

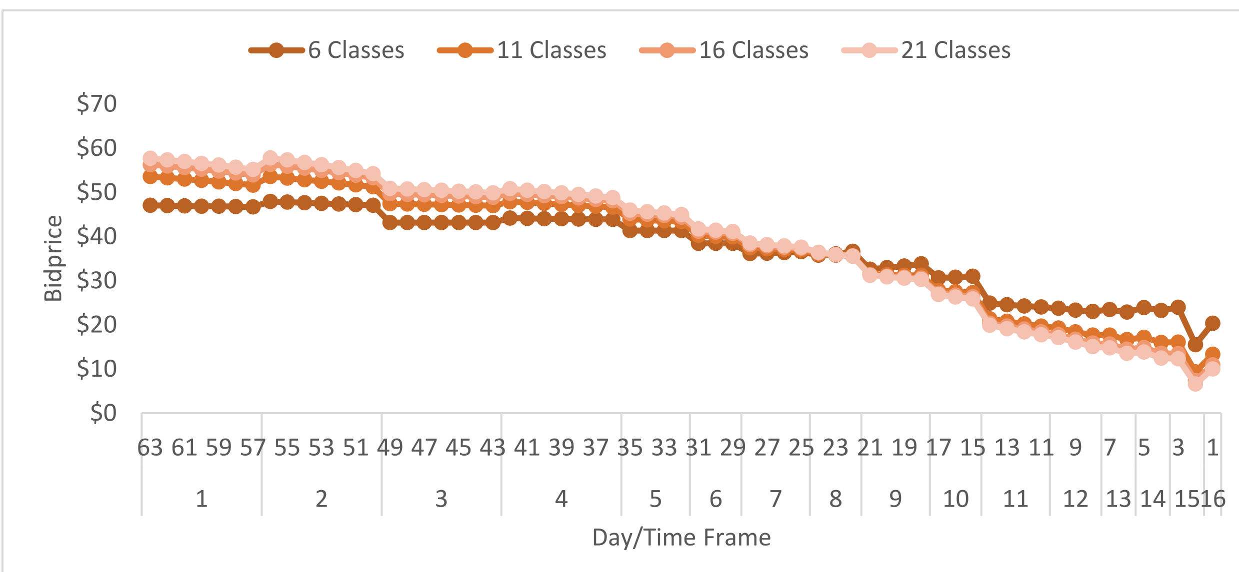

Figure 5-10: Airline 1 Class-Based Continuous ProBP Bidprices

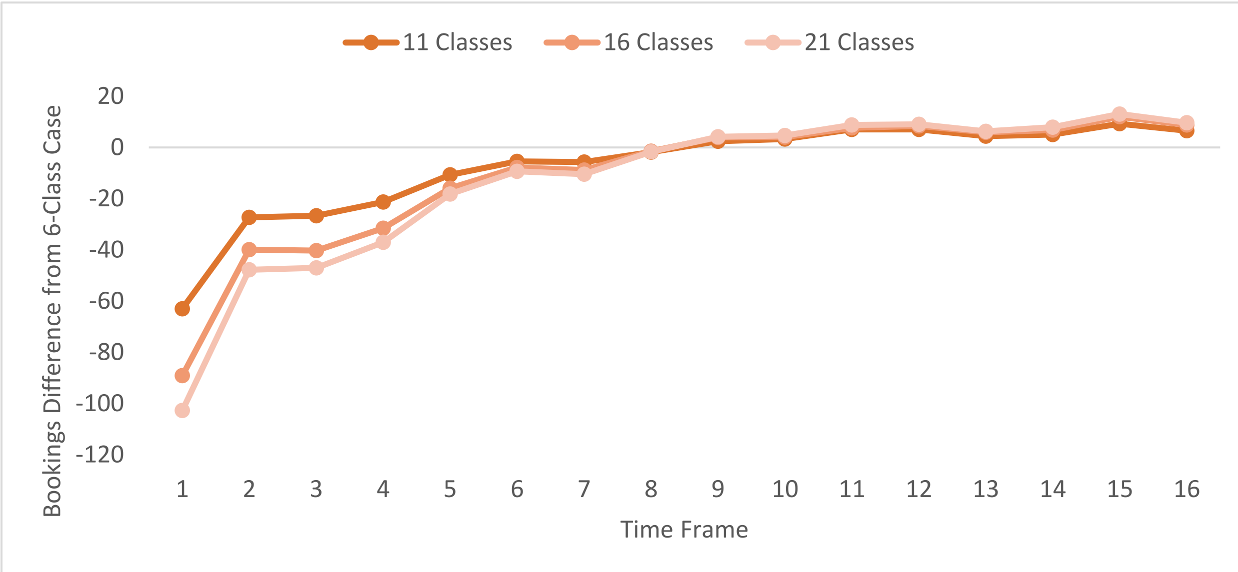

Figure 5-11: Airline 1 Class-Based Continuous ProBP Difference in Bookings from 6-Class Case

The reason for the forecast explosion is not in the Class-Based Continuous ProBP seat protection optimization algorithm, but instead is a result of a difference in the traditional and continuous pricing methods' forecasters. As stated in Section 3.1.2, a limit, known as the "XSCL" limit, is typically employed on how many "Q-Bookings" each actual booking can be scaled to. However, this limit was not employed on the bookings scaling for the continuous pricing methods. The more precise optimization resulting from the extra fare classes only increases the number of bookings in later TFs and exacerbates this issue.

5.1.2.1 Comparison: Traditional and Continuous Class-Based ProBP in Symmetric Competition

The purpose of examining continuous pricing methods is to see how they compare to existing, traditional "fixed" methods. In this subsection, traditional Class-Based ProBP and Class-Based Continuous ProBP's performances will be compared. Figure 5-12 shows a comparison of revenue performance for traditional and continuous Class-Based ProBP using 6, 11, 16, and 21 fare classes.

Figure 5-12: Airline 1 Traditional or Continuous Class-Based ProBP Revenue (percent change in revenue from switching from traditional to continuous class-based)

Observing the revenue results, a few things become apparent. Firstly, the contrasting nature of how adding fare classes affects traditional and continuous Class-Based ProBP is shown. With six fare classes, Class-Based Continuous ProBP generates nearly 1% more revenue than traditional Class-Based ProBP. However, the combination of more fare classes increasing traditional ProBP revenue and decreasing Class-Based Continuous ProBP revenue results in Class-Based Continuous ProBP generating slightly less revenue than traditional Class-Based ProBP with 21 fare classes. It is also worth noting, however, that the continuous method with six fare classes generates more revenue than the traditional method with 21.

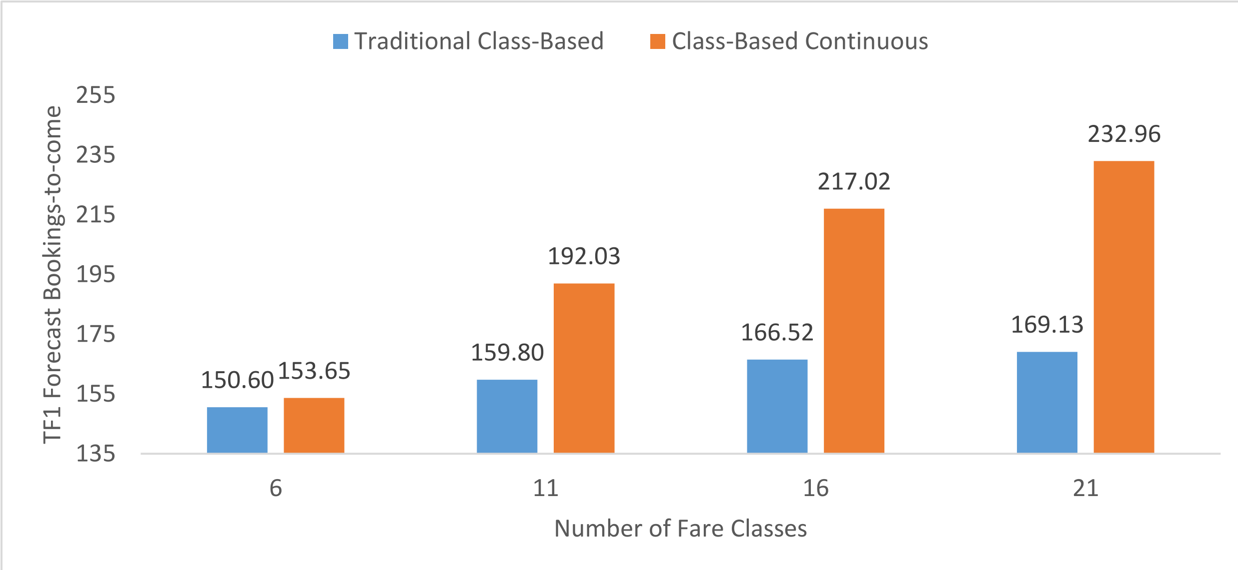

Figure 5-13: Airline 1 Traditional or Continuous Class-Based ProBP Time Frame 1 Forecast Bookings-to-come

As previously stated, the reason for Class-Based Continuous ProBP's revenue decrease is likely the scope to which its forecasts increase as fare classes are added. The effect of the lack of a scaling limitation on Class-Based Continuous ProBP forecasts is clearly shown in Figure 5-13. It may be that instituting the XSCL limit on Class-Based Continuous ProBP would keep its forecasts similar to traditional Class-Based ProBP and cause the two methods to converge to even amounts of revenue.

Theoretically, if they were using identical forecast scaling limits, the traditional and continuous Class-Based ProBP's algorithms should converge as the number of fare classes used approaches infinity, as having an infinite number of fares to quote is the same as continuously pricing (it is worth noting, as previously stated, that forecast partitions would become meaningless to the optimization step with too many fare classes). If the prior speculation is correct, however, convergence may occur within a relatively small number of fare classes. If that were, in fact, the case, there would be few advantages granted by using Class-Based Continuous ProBP in cases with symmetric competition. While the lack of a scaling limit on Class-Based Continuous ProBP may disguise quick convergence, it may be possible to detect such a convergence by observing the underlying mechanics of the two class-based methods as they add fare classes.

One underlying metric that may show evidence of rapid convergence is the average paid fares by TF. Figure 5-14 shows that, for six fare classes, traditional class-based and Class-Based

Continuous ProBP sell similar fares in TFs 1 to 9 before drastically separating at TF 10 (this is likely as a result of the aforementioned jump in price between fare class 2 and 3). When observing the 21 fare class experiments (Figure 5-15), however, the difference in average paid fares for the two methods is substantially reduced, with only a small difference in average paid fares in the later TFs, which may be explainable owing to differences in early TF forecasts owing to Class-Based Continuous ProBP's lack of a scaling limit.

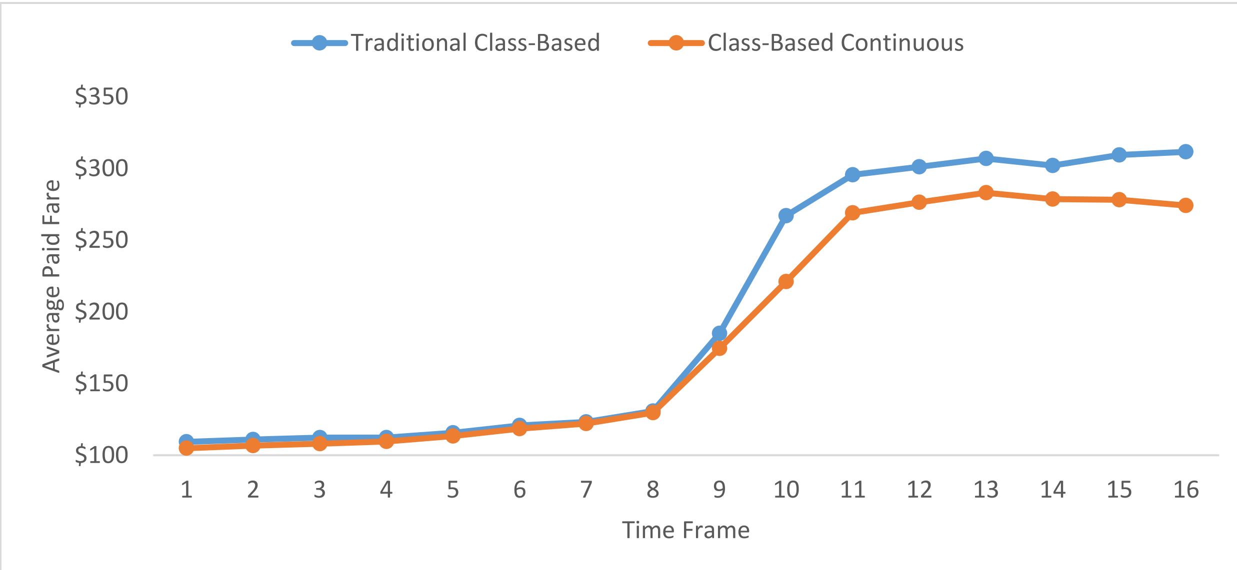

Figure 5-14: 6-Class Airline 1 Traditional or Continuous Class-Based ProBP Average Fare

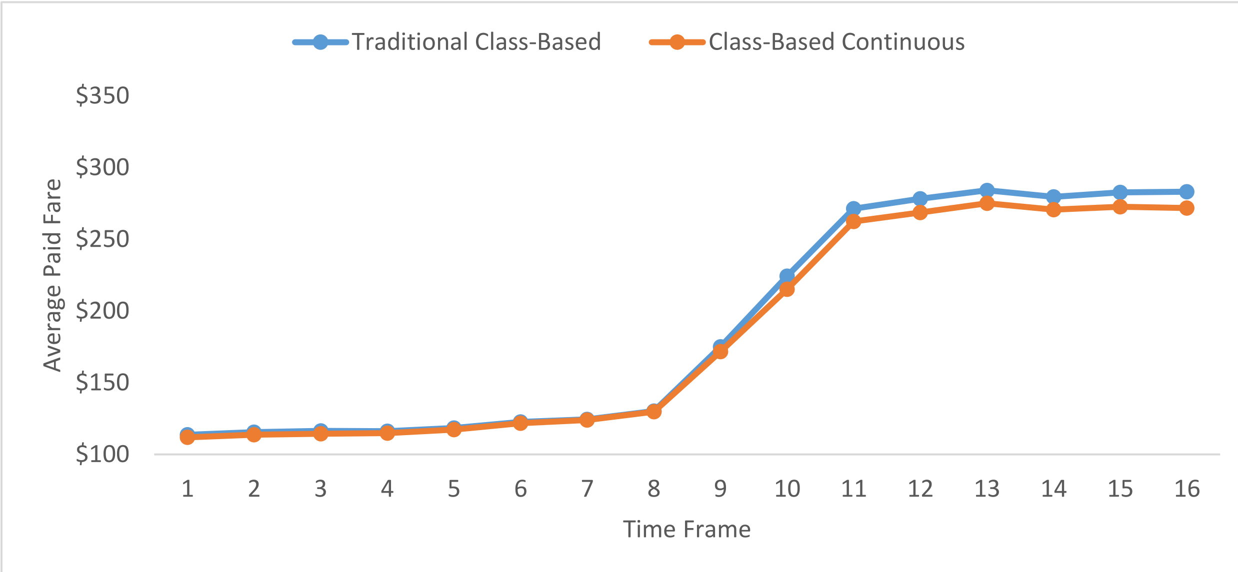

Figure 5-15: 21-Class Airline 1 Traditional or Continuous Class-Based ProBP Average Fare

As average paid fares converge, bookings by TF do as well. With six fare classes, the lower Class-Based Continuous ProBP fares result in slightly higher bookings for Class-Based Continuous ProBP in all TFs (Figure 5-16). With 21 fare classes, the similarities in paid fares also cause bookings to be very similar in all TFs (Figure 5-17).

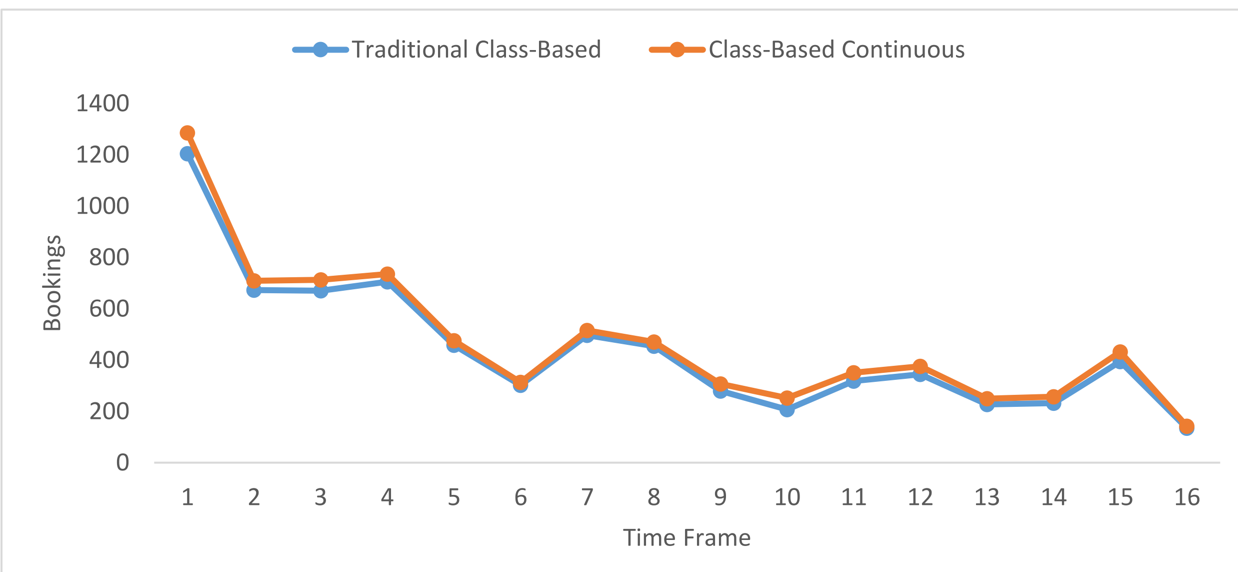

Figure 5-16: 6-Class Airline 1 Traditional or Continuous Class-Based ProBP Bookings

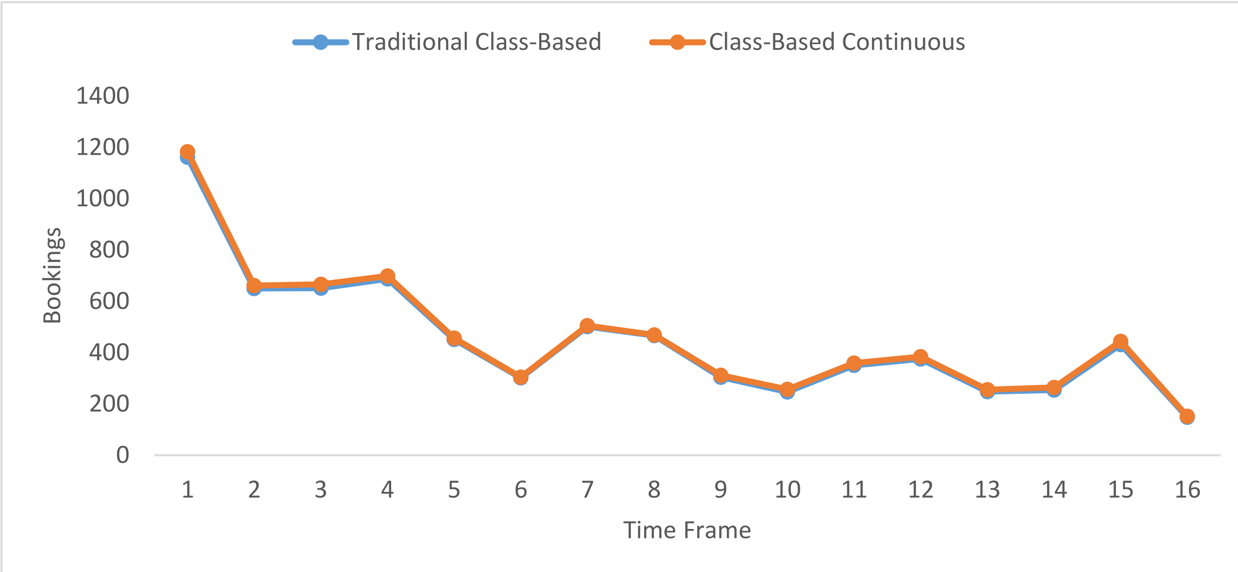

Figure 5-17: 21-Class Airline 1 Traditional or Continuous Class-Based ProBP Bookings

One final parameter that may be able to indicate convergence is comparing the bidprices for traditional and Class-Based Continuous ProBP. For the 6-fare class case, bidprices are similar for both ProBP methods before diverging in later TFs (Figure 5-18), while, as with average paid fares and bookings, bidprices are very similar in the final TFs for the 21-fare class case (Figure 5-19).

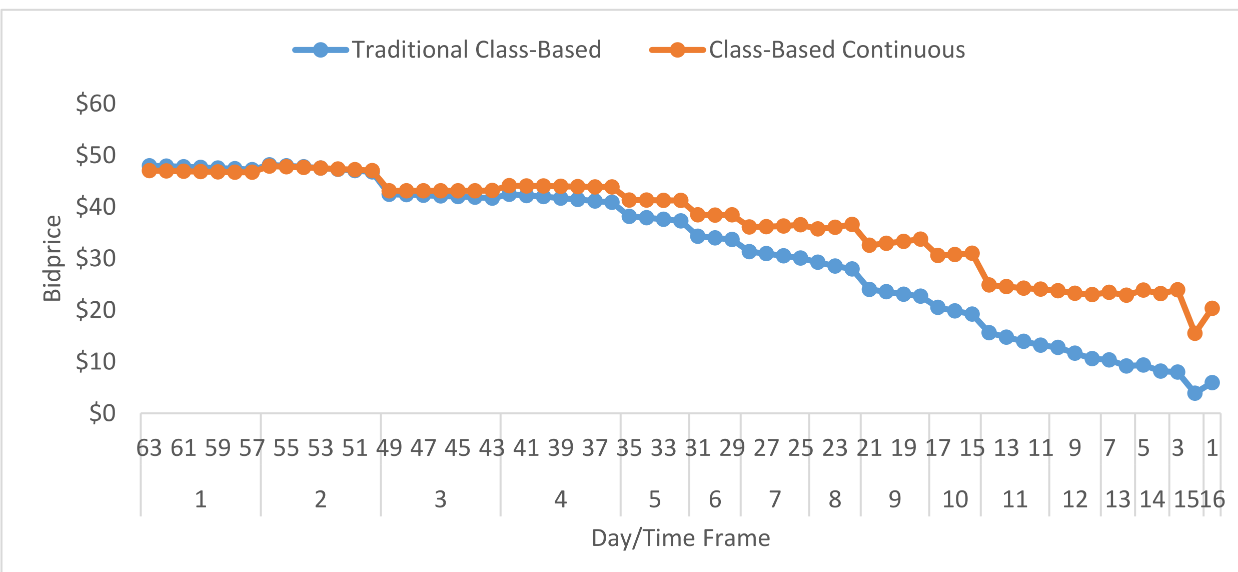

Figure 5-18: 6-Class Airline 1 Traditional or Continuous Class-Based ProBP Bidprices

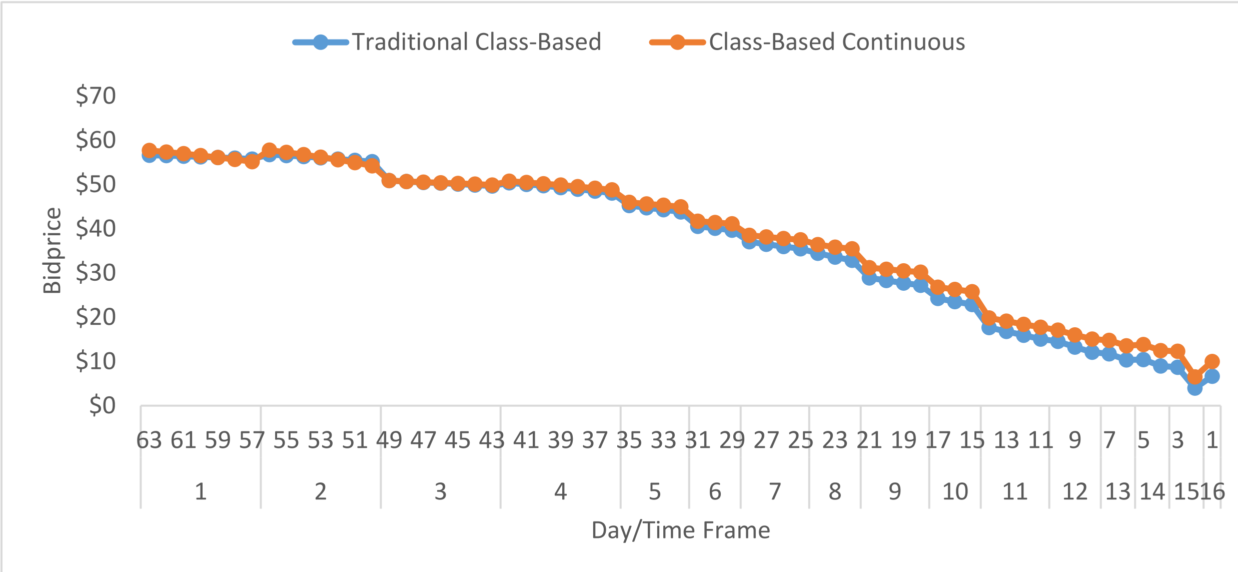

Figure 5-19: 21-Class Airline 1 Traditional or Continuous Class-Based ProBP Bidprices

The average paid fares, bookings, and bidprices for traditional class-based and Class-Based Continuous ProBP give strong indications of a rapid revenue convergence between the two methods. As a result, particularly within the context of New Distribution Capability (Westermann, 2013) where airlines could hypothetically have any number of fare classes, the gains from Class-Based ProBP for continuous pricing as indicated by symmetric competition tests are likely not large enough to warrant the design of a fare quotation system capable of using continuous pricing. However, Class-Based Continuous ProBP may still have a competitive advantage over traditional Class-Based ProBP that may be detectable in asymmetric tests.

5.1.3 Class-Based Continuous vs. Traditional Class-Based ProBP in Asymmetric Competition

While performance of an RM method in a symmetric network is a good metric for measuring the theoretical performance of said method, symmetries such as the ones previously tested are very rare in real-world airline markets. At the simplest level, it is very unlikely that two airlines would switch their methods to Class-Based Continuous ProBP at the same time. To reflect these facts, experiments with asymmetric competition were performed, and their results described in this subsection.

In these tests, AL1 from Network D6 switches from traditional ProBP to Class-Based Continuous ProBP while AL2 continues to use traditional ProBP. Initially, both airlines use FRAT5 C and fare adjustment. As before, they were tested with 6, 11, 16, and 21 fares.

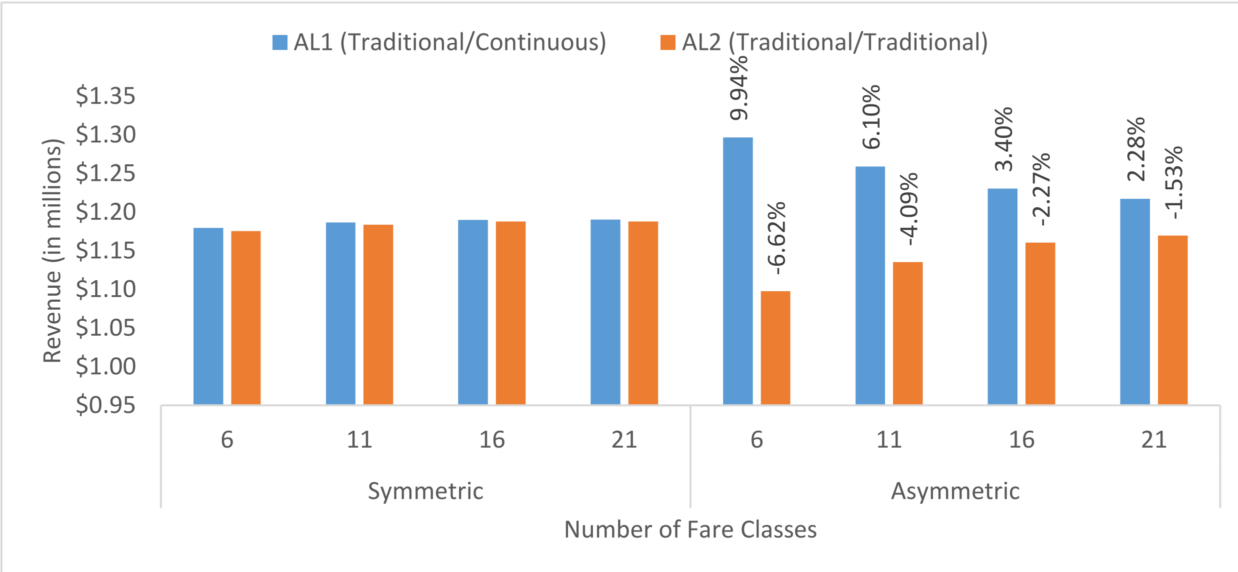

Figure 5-20: Continuous vs. Traditional Class-Based ProBP Revenue (percent change in revenue resulting from Airline 1 switching from traditional to class-based continuous)

The revenue results of the asymmetric tests (Figure 5-20) are very clear. AL1 has a competitive advantage over AL2 regardless of the number of fare classes used by the two airlines, although this advantage does decrease as fare classes are added. The fact that the advantage decreases as fare classes are added gives the first clue to the likely reason for this, which is that Class-Based Continuous ProBP has more granularity to the fares it can offer than traditional Class-Based ProBP. While this advantage diminished to the point of nonexistence in the symmetric competition experiments as the extra granularity only gave Class-Based Continuous ProBP the ability to better calculate passengers' WTP, in a scenario with asymmetric competition, this granularity gives AL1 the ability to undercut AL2, which, as shown in Figure 5-21, allows AL1 to generate more bookings in later TFs where passengers have a higher WTP (and where the gaps in traditional class-based prices become larger).

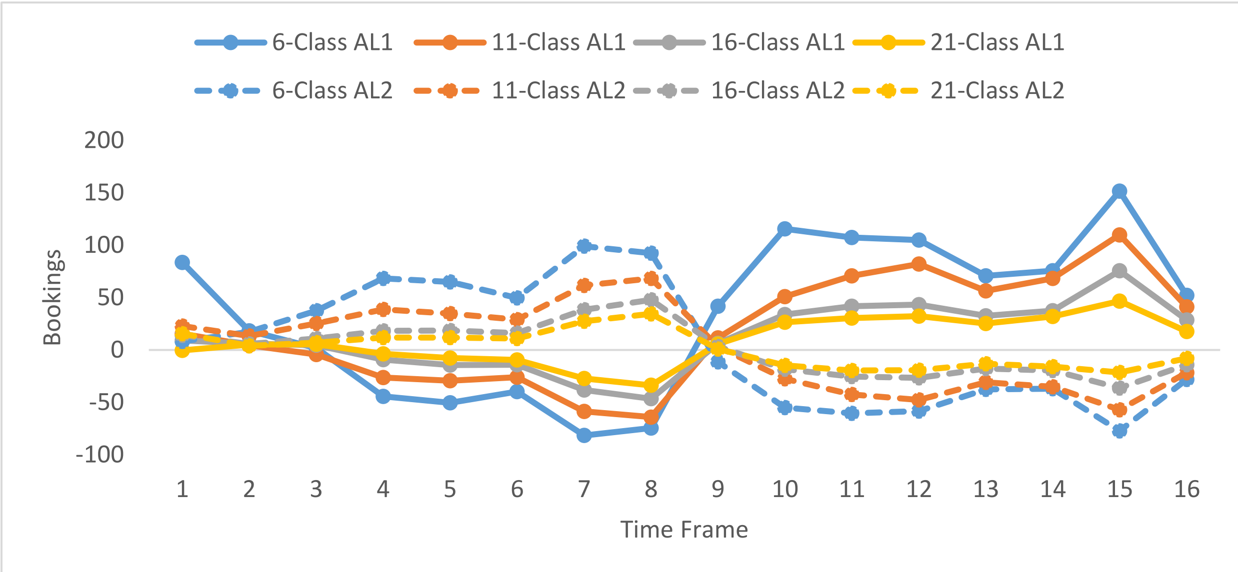

Figure 5-21: Continuous vs. Traditional Class-Based ProBP Change in Airline Bookings from Symmetric to Asymmetric Experiments

Class-Based Continuous ProBP is better able to capture higher fare bookings in later TFs (note that bookings are more useful than average paid fares when considering asymmetric tests as passengers who choose to buy will always take the lowest fare available). As an effect of Q-Forecasting, these bookings increase the overall forecast for Class-Based Continuous ProBP (Figure 5-22), which causes it to be more aggressive in early TFs, hence the fact that it has lower bookings in many of the early TFs.

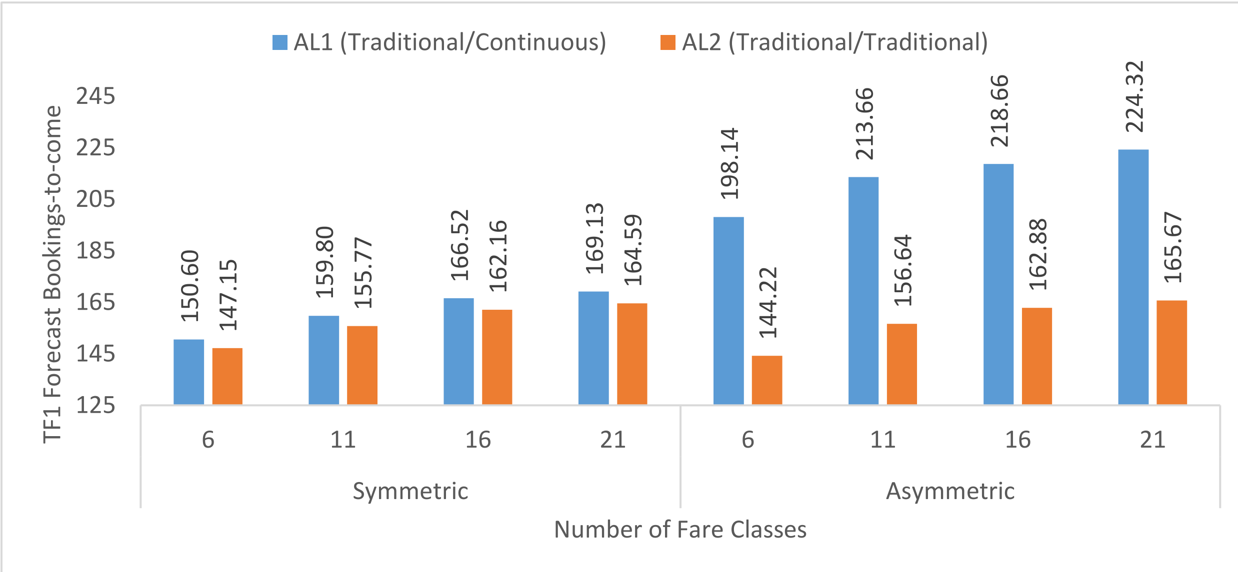

Figure 5-22: Continuous vs. Traditional Class-Based ProBP Time Frame 1 Forecast Bookings-to-come

The fact that AL1 using Class-Based Continuous ProBP generates more revenue by underpricing then leads to the question of whether the traditional Class-Based ProBP using AL2 can recapture much of that revenue by lowering its own FRAT5 and attempting to undercut Class-Based Continuous ProBP. Lowering the FRAT5s used by Q-Forecasting and fare adjustment will result in smaller forecasts and reduced fare adjustment, both of which will cause ProBP to be less aggressive about closing fare classes. This can be checked by lowering the FRAT5s used by AL2 in the previous asymmetric test to the FRAT5 E curve (Figure 4-5).

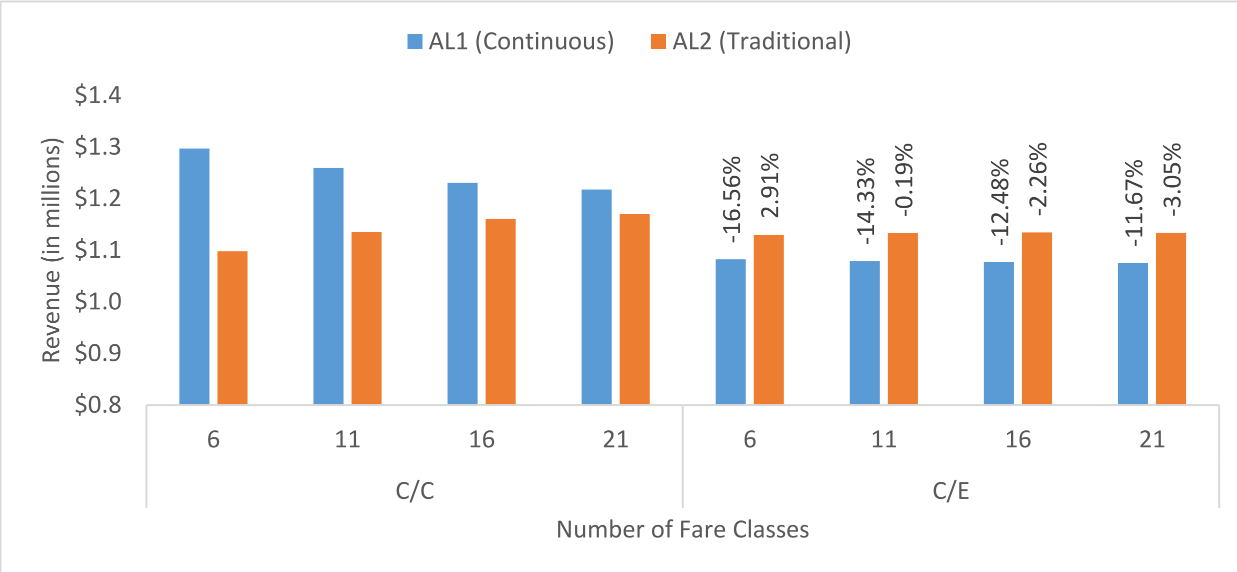

Figure 5-23: Continuous vs. Traditional Class-Based ProBP with Different FRAT5s Revenue (percent change in revenue resulting from Airline 2 switching FRAT5 curve from C to E)

While AL2, using traditional Class-Based ProBP, can drive down revenue for AL1, which uses Class-Based Continuous ProBP, this comes at the expense of its own revenue in all but the 6-fare class case (Figure 5-23). Intuitively this makes some sense, as this undercutting of AL1 can only be achieved by offering fares that the symmetric baseline tests indicate should really be too low. Once again, this can be checked by observing bookings by TF (Figure 5-24).

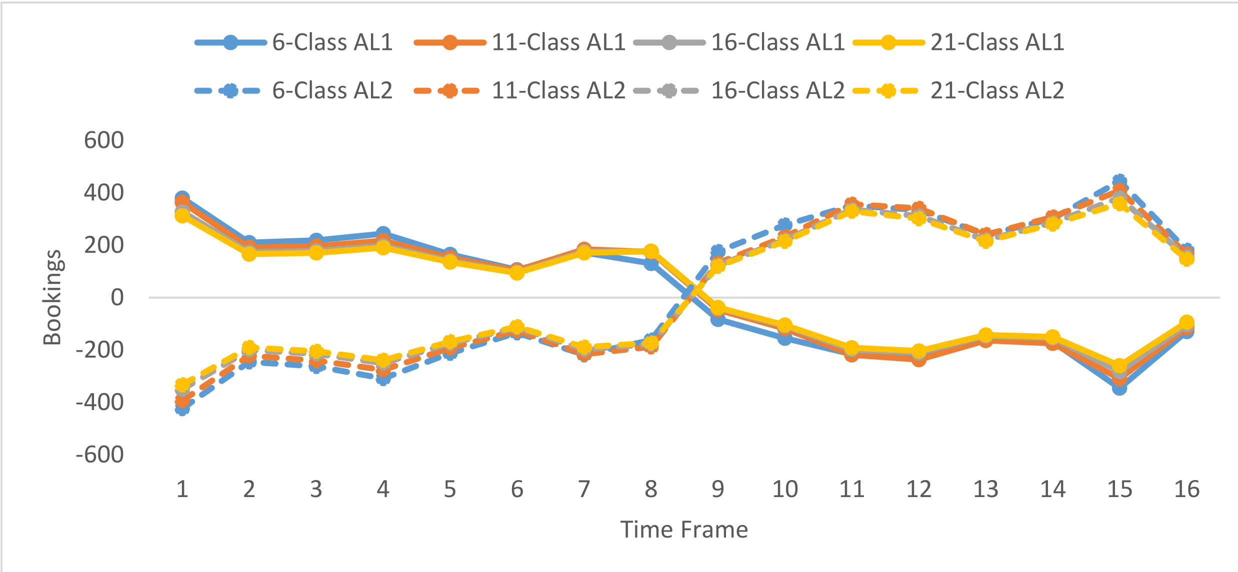

Figure 5-24: Continuous vs. Traditional Class-Based ProBP Change in Bookings from the Traditional Class-Based Airline 2 Switching to FRAT5 E

The bookings by TF may seem oddly distributed, as, despite AL2 being less aggressive, it is AL1 that ends up with far more bookings in the early TFs. The reason for this lies in the fact that lowering FRAT5 does more than just reduce how much bookings are scaled by the Q-Forecaster; it also lessens fare adjustment. With FRAT5 E causing much smaller fare adjustment in later TFs, fare classes that might have had negative adjusted fares in later TFs with FRAT5 C (and would thus automatically be closed) are opened. At the same time, Class-Based Continuous ProBP in its optimization step is still using FRAT5 C and seeing those classes closed by the fare adjuster. As a result, the traditional Class-Based ProBP using airline offers lower fares in later TFs. Since these fares are much lower than the fares Class-Based Continuous ProBP had offered when both airlines used FRAT5 C, however, AL2 ends up mostly decreasing AL1's revenue instead of increasing its own.

The asymmetric tests show that, even though they perform similarly as the number of fare classes increases in independent symmetric tests, Class-Based Continuous ProBP has a competitive advantage over traditional Class-Based ProBP by being able to sell fares at values between those associated with filed fare classes. While the airline using traditional Class-Based ProBP can counter by reducing its FRAT5 curve, this does little to help the traditional Class-Based ProBP using airline and instead mostly lowers the revenue of the Class-Based Continuous ProBP using airline. Said Class-Based Continuous ProBP using airline would also, hypothetically, be able to counter this move by reducing its own FRAT5.

5.1.4 Classless ProBP in Symmetric Competition

Having established the performance of Class-Based ProBP methods, the next item to examine is the performance of the Classless ProBP algorithm in symmetric competition scenarios. As with Class-Based ProBP, the first step was to select a FRAT5 curve to use in the symmetric experiments, and, as with both Class-Based ProBP methods, FRAT5 C was found to maximize revenue in the symmetric tests (Figure 5-25).

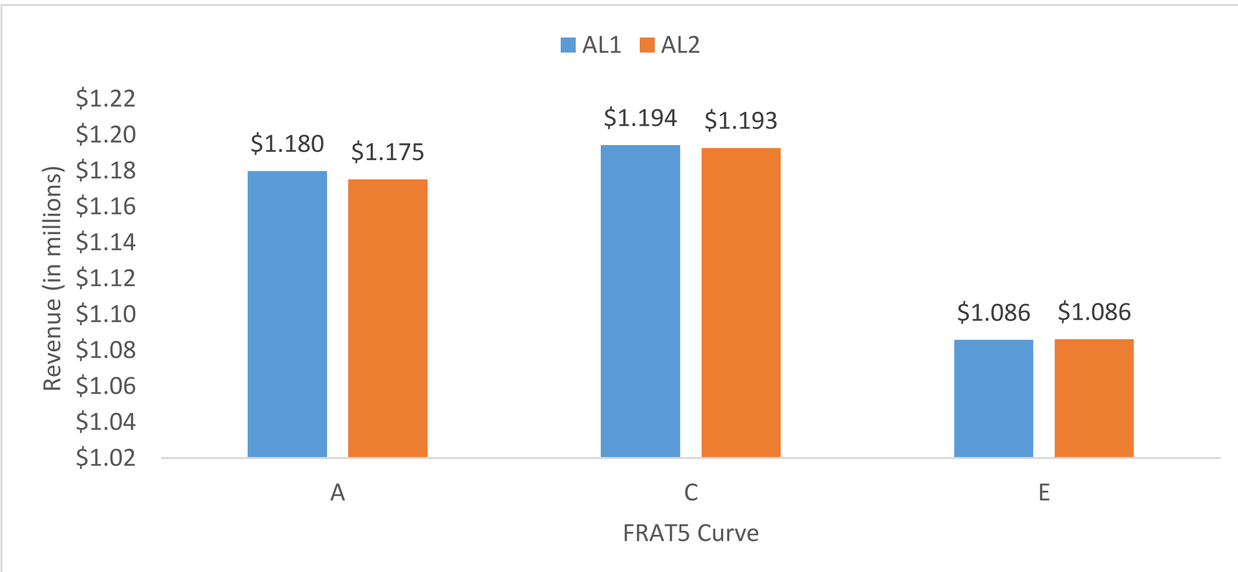

Figure 5-25: Classless ProBP Revenue

Since Classless ProBP obviously uses no fare classes in its optimization process, there is no effect of adding more fare classes to discuss. Instead, its symmetric competition performance will be compared with the class-based methods' performances, starting with traditional Class-Based ProBP.

5.1.4.1 Comparison: Traditional Class-Based and Classless ProBP in Symmetric Competition

Unlike with Class-Based Continuous ProBP, there was no theoretical expectation that the results of traditional Class-Based ProBP would converge to those of Classless ProBP as the number of fare classes used by traditional Class-Based ProBP approached infinity. In fact, the Classless ProBP algorithm addresses the forecast partitioning feasibility issues that would occur if the number of fare classes used by traditional Class-Based ProBP was very large. The results of the simulation bear this out, with Classless ProBP getting an additional revenue increase on top of the revenue gained by traditional Class-Based ProBP for adding fare classes (Figure 5-26).

Figure 5-26: Traditional Class-Based or Classless ProBP Revenue (percent change in revenue from 6-class experiment)

While the additional revenue gain for Classless ProBP is not particularly large (about another 0.34% for AL1), it is still clearly an additional gain over the traditional class-based experiments, for which revenue increases diminish as more and more fare classes are added. This revenue increase likely results from the fact that, similar to Class-Based Continuous ProBP, Classless ProBP has unlimited granularity in its fare offerings, while, at the same time, it uses an optimization algorithm specifically designed to take this unlimited granularity into account. The effect of this extra granularity is demonstrated both in how Classless ProBP's average fares (Figure 5-27) and bookings (Figure 5-28) by TF compare to traditional Class-Based ProBP's.

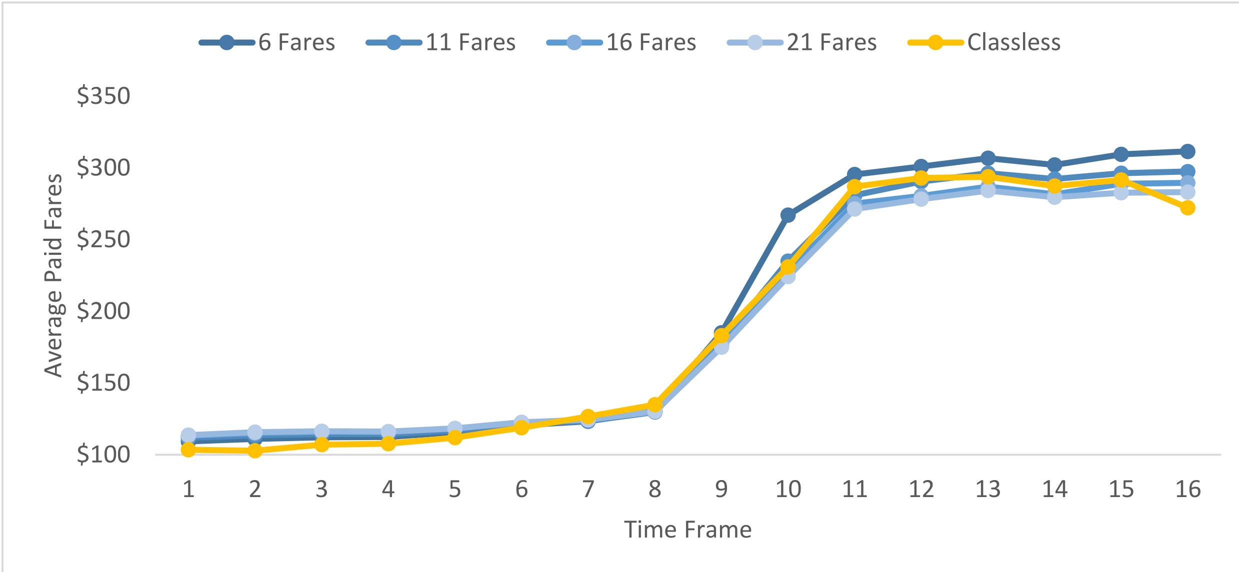

Figure 5-27: Airline 1 Traditional Class-Based or Classless ProBP Average Fares

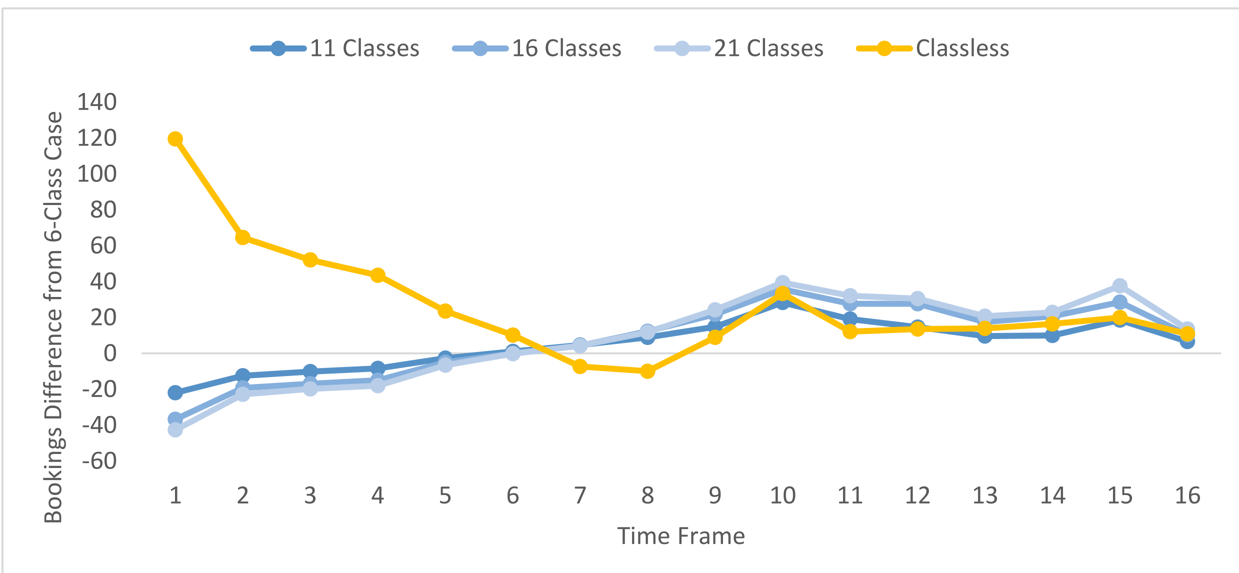

Figure 5-28: Airline 1 Traditional Class-Based or Classless ProBP Bookings Difference from 6-Class Case

Classless ProBP is able to sell somewhat lower fares than traditional Class-Based ProBP in early TFs. This causes a substantial increase in early TF bookings. At the same time, bookings and fares in later TFs are similar to those of 11-fare class traditional Class-Based ProBP.

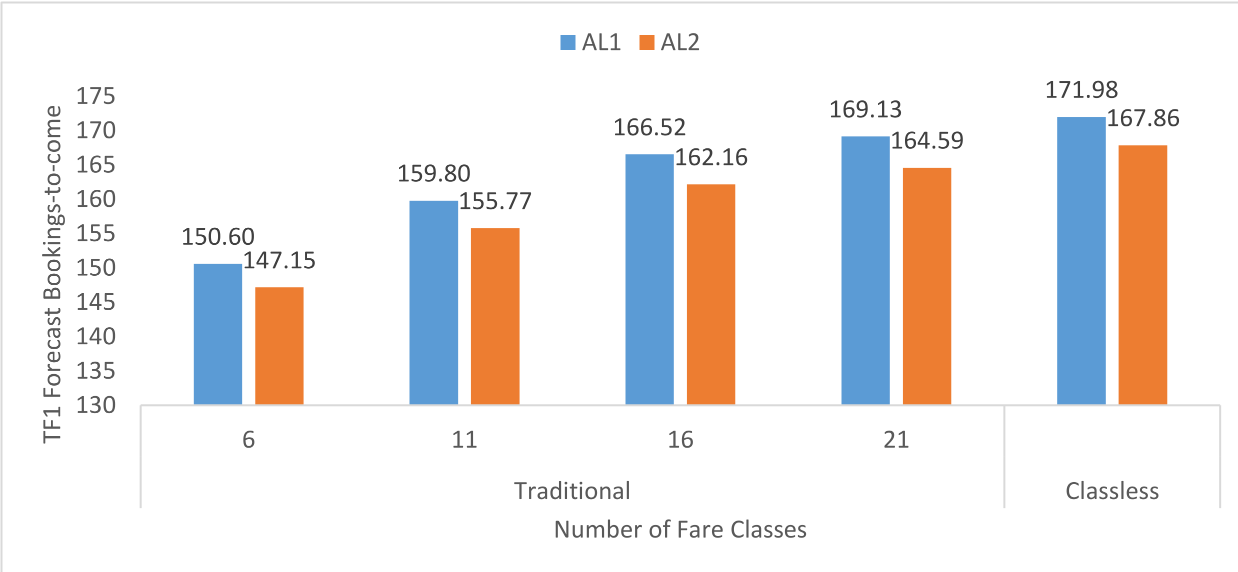

Figure 5-29: Traditional Class-Based or Classless ProBP Time Frame 1 Forecast Bookings-to-come

With respect to how it compares to traditional Class-Based ProBP, the last item worth noting for Classless ProBP is its forecasts. Similarly to Class-Based Continuous ProBP, the forecasting method for Classless ProBP is not limited as to how much bookings are scaled. Unlike Class-Based Continuous ProBP, however, Classless ProBP forecasts remain within a range close to traditional Class-Based ProBP forecasts (Figure 5-29), likely as a result of Classless ProBP gaining lower value bookings in early TFs, which are scaled less by Q-Forecasting than the higher value bookings gained by Class-Based Continuous ProBP.

5.1.4.2 Comparison: Class-Based Continuous and Classless ProBP in Symmetric Competition

In symmetric competition, Classless ProBP clearly increases revenue over traditional Class-Based ProBP. In this subsection, results of symmetric competition experiments for Classless ProBP will be compared with those for Class-Based Continuous ProBP to determine whether Classless ProBP gains a revenue advantage from using an optimizer specifically designed for continuous pricing.

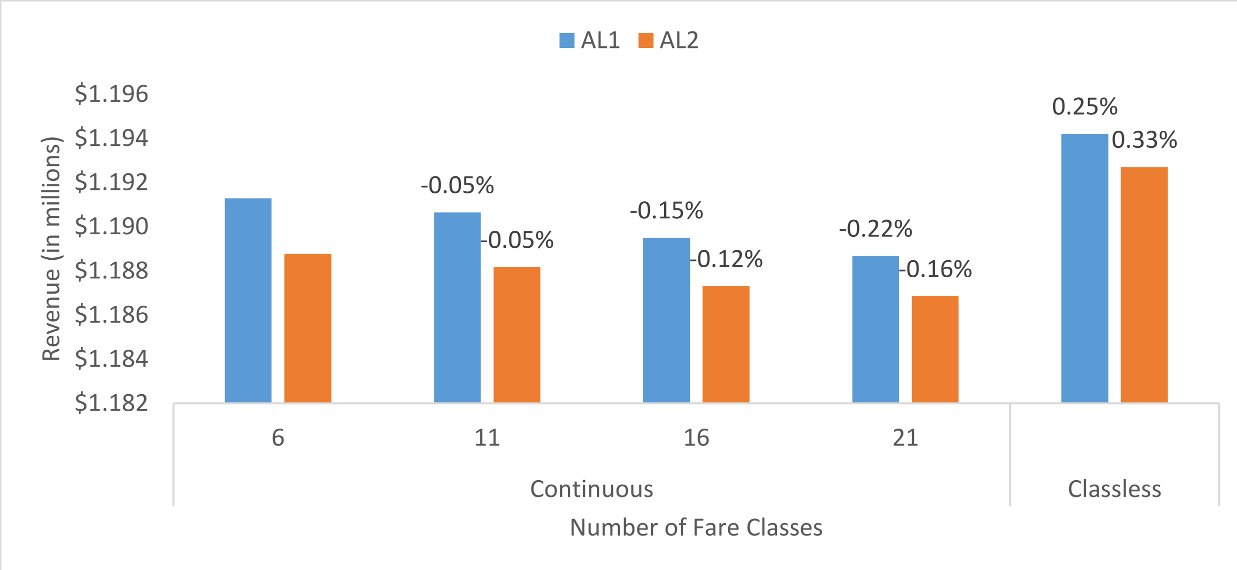

Figure 5-30: Class-Based Continuous or Classless ProBP Revenue (percent change in revenue from 6-class experiment)

Symmetric Classless ProBP generates more revenue than symmetric Class-Based Continuous ProBP, no matter the number of fare classes for Class-Based Continuous ProBP (Figure 5-30). While both Classless ProBP and Class-Based Continuous ProBP have unlimited granularity in terms of what fares they may quote, Classless ProBP does have an optimization algorithm designed to make use of this fact. The algorithm differences are shown by again observing average paid fares (Figure 5-31) and bookings (Figure 5-32) by TF.

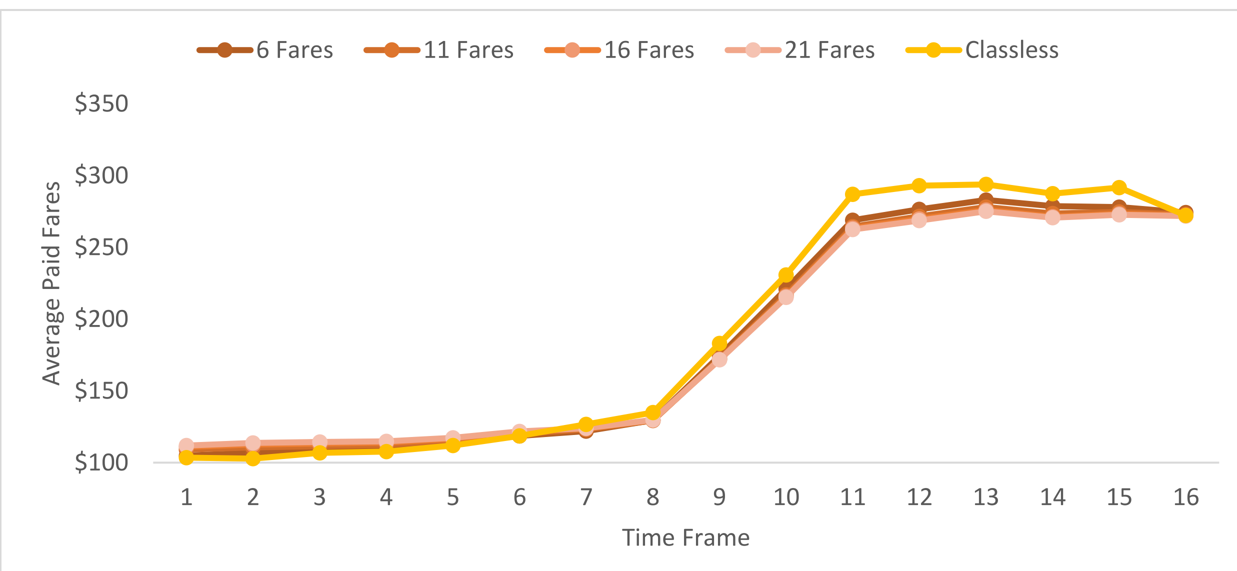

Figure 5-31: Airline 1 Class-Based Continuous or Classless ProBP Average Fares

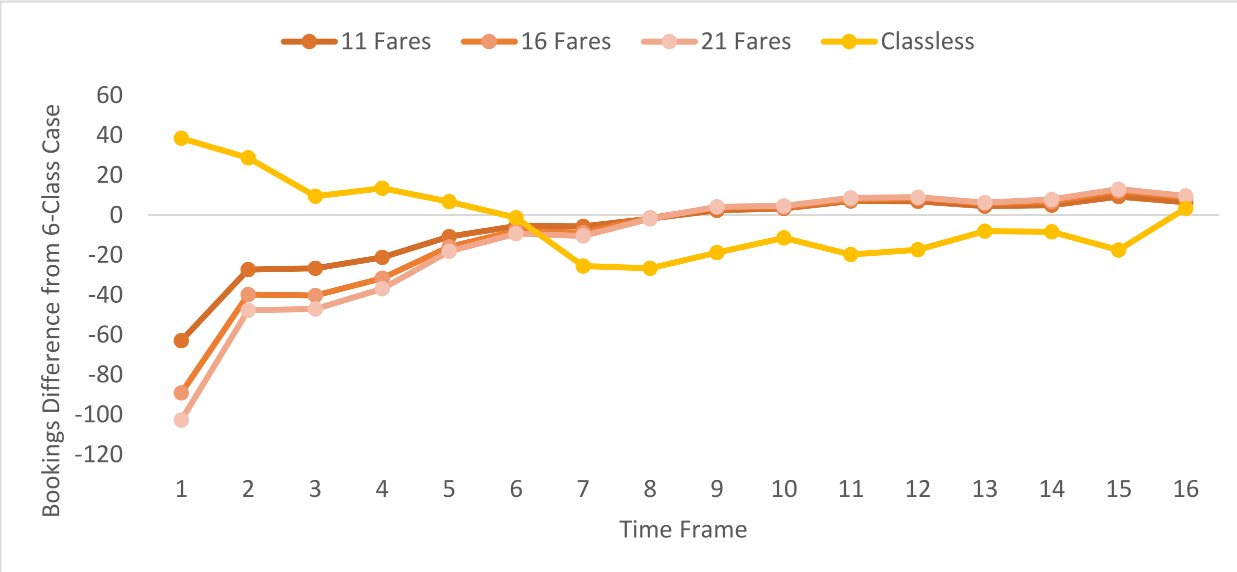

Figure 5-32: Airline 1 Class-Based Continuous or Classless ProBP Bookings Difference from 6-Class Case

While the average fare and bookings differences between Classless ProBP and Class-Based Continuous ProBP are not as stark as they were between Classless ProBP and traditional Class-Based ProBP, Classless ProBP still sells lower fares in early TFs and ends up with more bookings. While, unlike with traditional class-based, this does result in Classless ProBP having fewer bookings in later TFs, enough revenue is generated as a result of its more efficient algorithm in order outperform Class-Based Continuous ProBP.

5.1.5 Classless vs. Traditional Class-Based ProBP in Asymmetric Competition

As previously mentioned, symmetric tests do help establish behaviors for theoretical applications, but asymmetric tests are far more likely to mimic real-world conditions. Thus, as with Class-Based Continuous ProBP, the effects of Network D6 AL1 switching to Classless ProBP while AL2 remains using Traditional Class-Based ProBP will be examined. Both airlines again used FRAT5 C, while only AL2 uses fare adjustment (as fare adjustment is not applicable for classless RM).

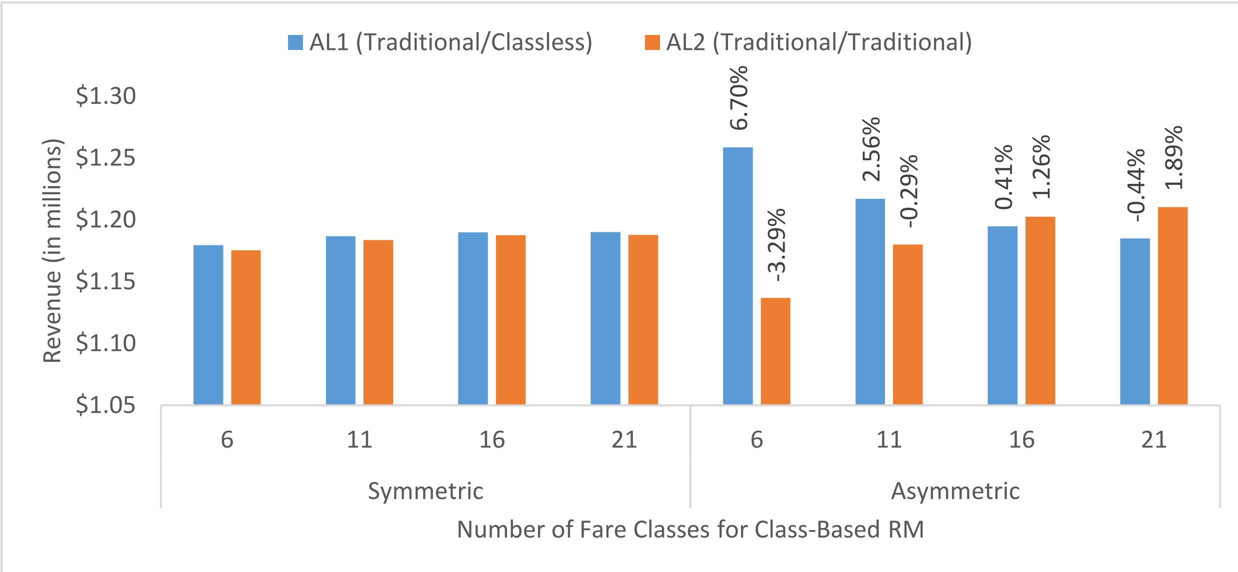

Figure 5-33: Classless vs. Traditional Class-Based ProBP Revenue (percent change in revenue resulting from Airline 1 switching from traditional to classless)

The revenue results for asymmetric application of Classless ProBP (Figure 5-33) differ somewhat from those of Class-Based Continuous ProBP. While the classless AL1 does outperform AL2 when the latter only uses 6 or 11 fare classes, AL2 generates more revenue as it adds more fare classes. This is most likely explainable by observing the bookings by TF for AL1 and AL2. It was previously established that Classless ProBP increases revenue over the six-fare traditional Class-Based ProBP baseline in symmetric competition experiments by increasing its early bookings without too negatively impacting its later bookings. Traditional Class-Based ProBP, meanwhile, added revenue by decreasing its early bookings and adding later, more valuable bookings. It is possible that these forces pulling early TF bookings in opposite directions is what causes the revenue effects shown in Figure 5-33.

Figure 5-34: Classless vs. Traditional Class-Based ProBP Change in Airline Bookings from Symmetric to Asymmetric Experiments

Figure 5-34 shows that this opposing pull is in fact what causes the revenue effects previously shown. When AL2 uses six fare classes, it tends to offer lower fares in early TFs (once again, bookings are a better metric of comparing average offered fare in asymmetric tests than average paid fare as only the lower offered fare will ever be purchased). AL1's Classless ProBP algorithm can, therefore, siphon off bookings from AL2 without risking taking too many and blocking out higher paying passengers in later TFs. As AL2 adds fare classes however, it becomes more aggressive about protecting higher value fare classes, resulting AL1 taking too many lower fare bookings in early TFs and preventing itself from having space for higher-value, later TF passengers.

5.1.6 Summary of ProBP Continuous Pricing Experiments

Both Class-Based Continuous ProBP and Classless ProBP show some potential for increasing revenue generated over traditional Class-Based ProBP. However, their potentials for increase are different and both come with weaknesses.

Class-Based Continuous ProBP does increase revenue over the traditional class-based baseline in 6-fare class symmetric tests as a result of it allowing more granularity in the fares it can offer. This is particularly helpful for Class-Based Continuous ProBP in later TFs, when the spacing between the filed fares in Network D6 becomes much larger and Class-Based Continuous ProBP is able to offer fares that fall in between the gaps. However, this advantage diminishes when traditional Class-Based ProBP is used with more fare classes added. With symmetric competition, this diminishes the advantage Class-Based Continuous ProBP has of more precise pricing, and most indicators show that traditional Class-Based ProBP and Class-Based Continuous ProBP performances converge with 16 or 21 fare classes. However, with asymmetric competition, the price granularity advantage once again becomes important for Class-Based Continuous ProBP, as it is often able to undercut traditional Class-Based ProBP in later TFs, increasing its bookings of valuable passengers and generating more revenue.

Classless ProBP is in many ways the opposite of Class-Based Continuous ProBP. In symmetric competition experiments, its specifically designed optimization algorithm helps it better select fares that are more likely to increase revenue. However, it does this by increasing how many low-fare passengers it books without substantially decreasing how many high-fare passengers book. In asymmetric competition experiments, this can cause a problem for Classless ProBP, as traditional Class-Based ProBP generates more revenue when it focuses on obtaining more later TF higher-priced bookings (which adding fare classes enables), and this fact may end up causing Classless ProBP to take low-fare bookings at the expense of the aforementioned high-fare later bookings.

5.2 UDP Results

In this section, the continuous pricing methods for UDP will be examined and discussed in the same way as those for ProBP. First, baseline metrics will be discussed for traditional Class-Based UDP, followed by discussions of symmetric and then asymmetric competitive impacts of Class-Based Continuous UDP. This will be followed by analysis of experiments on Classless UDP, which was also tested in symmetrical and asymmetrical competition. In this section, the results of UDP tests will only be compared with the results of other UDP tests. A comparison of the ProBP and UDP methods will be undertaken in the next section.

5.2.1 Traditional Class-Based UDP in Symmetric Competition

As with ProBP, it is best to compare the continuous pricing methods for UDP relative to a baseline of traditional Class-Based UDP. As before, the first step was to determine which FRAT5 curve to use for testing traditional Class-Based UDP, and Figure 5-35 shows that, once again, the best curve for use is FRAT5 C.

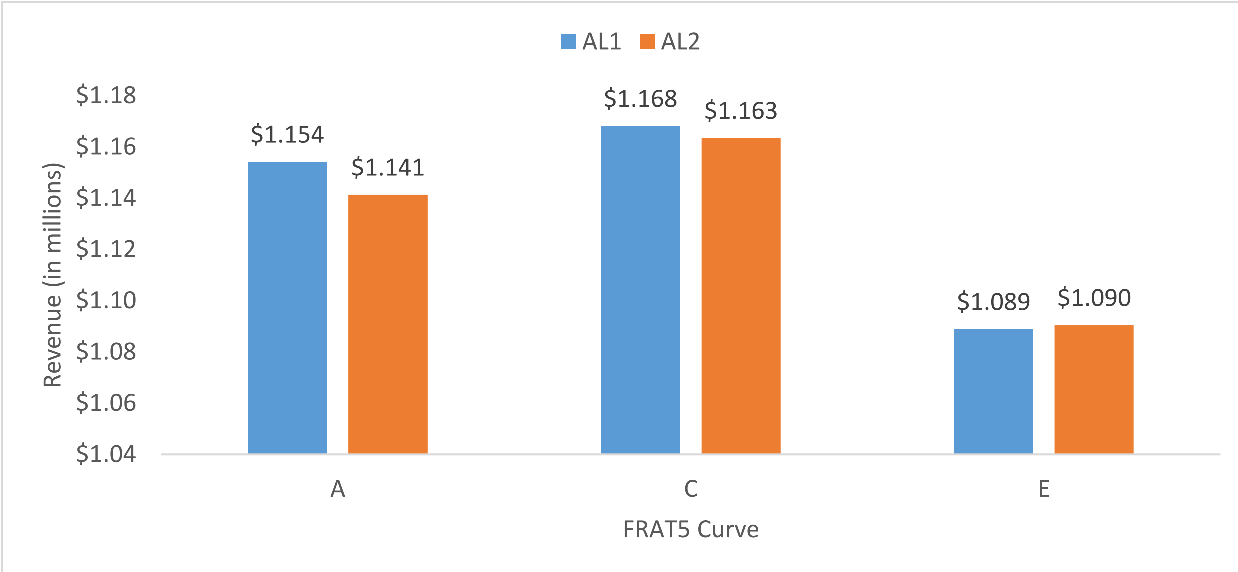

Figure 5-35: 6-Class Traditional Class-Based UDP Revenue

As with ProBP, the revenue results from traditional Class-Based UDP and Class-Based Continuous UDP hypothetically could converge in terms of how much revenue they generate as the number of fare classes used by either of them approaches infinity. As a result, like traditional Class-Based ProBP, traditional Class-Based UDP was tested with 6, 11, 16, and 21 fare classes using FRAT5 C and fare adjustment.

Figure 5-36: Traditional Class-Based UDP Revenue (percent change in revenue from 6-class experiment)

Figure 5-36 shows that traditional Class-Based UDP, like traditional Class-Based ProBP, gains revenue with additional fare classes and that these revenue gains diminish as more fare classes are added. As with ProBP, observing the average fares (Figure 5-37) and bookings (Figure 5-38) by TF for traditional Class-Based UDP shows that increasing the number of fare classes provides the method with more price granularity, particularly in later TFs, which shifts bookings from the earlier, lower-WTP TFs into the later TFs where passengers have greater WTP.

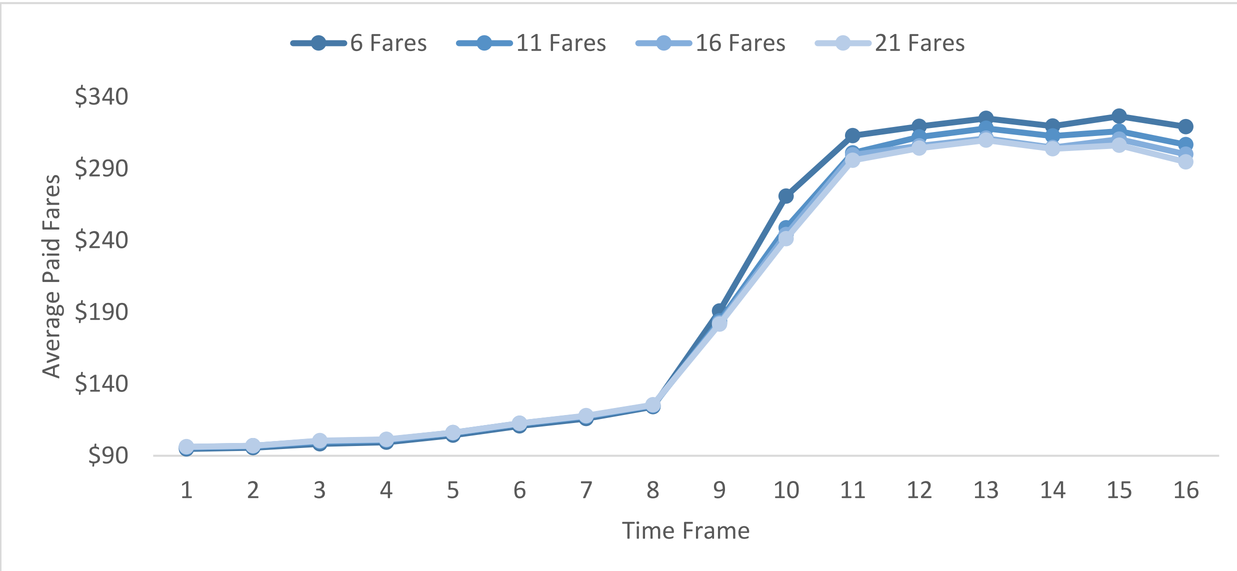

Figure 5-37: Airline 1 Traditional Class-Based UDP Average Fares

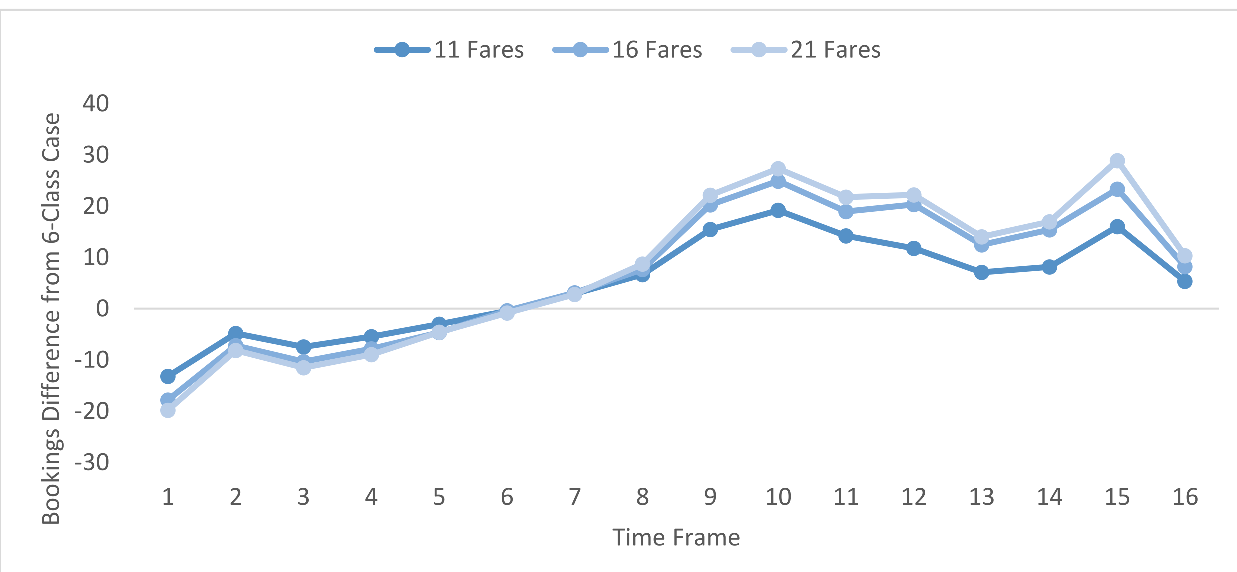

Figure 5-38: Airline 1 Traditional Class-Based UDP Bookings Difference from 6-Class Case

As was the case with traditional Class-Based ProBP, forecasts also increase with additional fare classes (Figure 5-39) as a result of the shift of bookings into higher priced fare classes in later

TFs. As with ProBP, this also increases the UDP bidprices in early TFs (Figure 5-40), which results in offering higher fares in earlier TFs (best observed in the fact that having more fare classes leads to fewer early TF bookings).

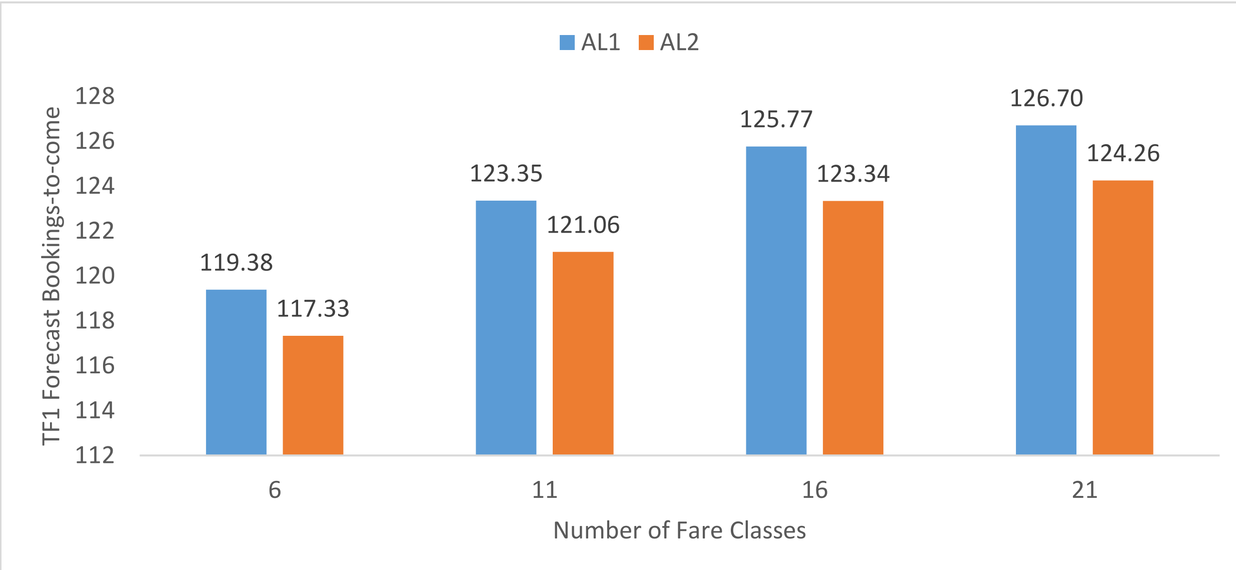

Figure 5-39: Traditional Class-Based UDP Time Frame 1 Forecast Bookings-to-come

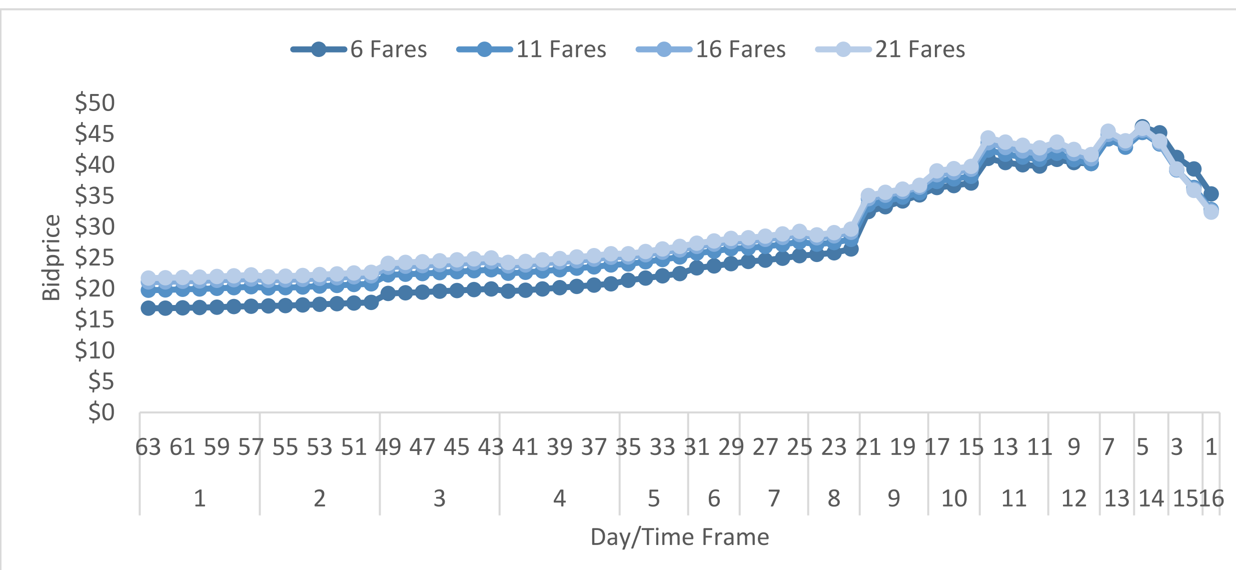

Figure 5-40: Airline 1 Traditional Class-Based UDP Bidprices

5.2.2 Class-Based UDP for Continuous Pricing in Symmetric Competition

Class-Based Continuous UDP will first be considered as fare classes are added, and then its symmetric competitive environment results will be compared with those of the symmetric competition tests of traditional Class-Based UDP. Once again, the first step is to select a FRAT5 curve that maximizes revenue for Class-Based Continuous UDP in a six-class fare structure, and that curve is once again curve C, although the difference in revenue between curves A and C is quite small (Figure 5-41).

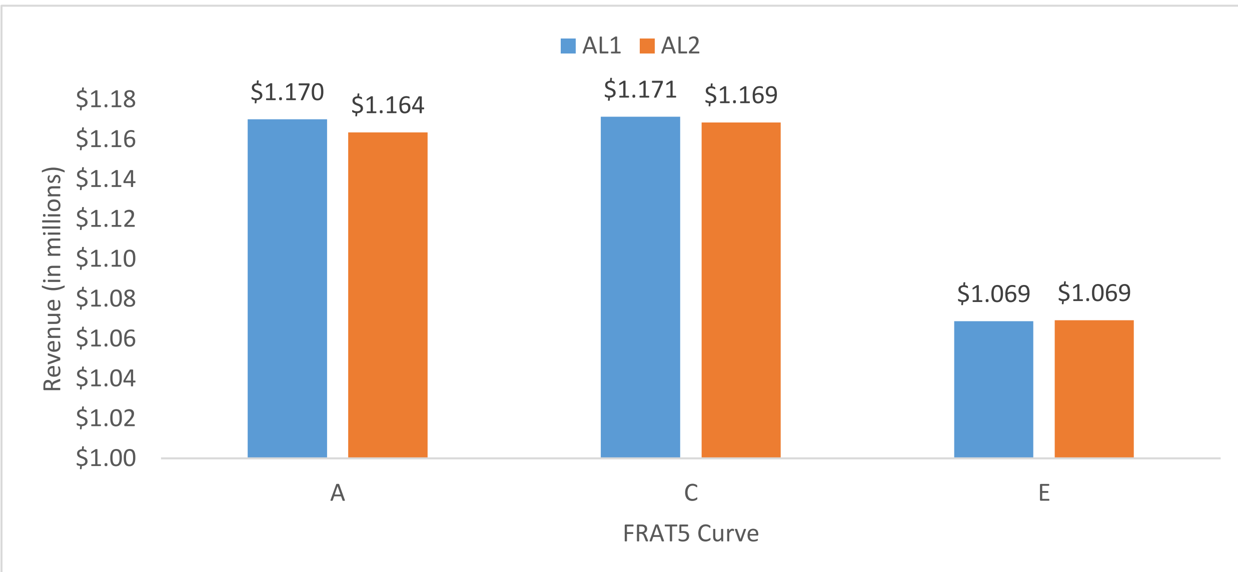

Figure 5-41: 6-Class Class-Based Continuous UDP Revenue

While Class-Based Continuous UDP does not gain the advantage of increased fare granularity from adding fare classes, doing so could serve to render Class-Based Continuous UDP optimization algorithm more precise. Experiments with 6, 11, 16, and 21 classes for Class-Based Continuous UDP show that revenue, does, in fact increase with additional fare classes (Figure 5-42). As with traditional Class-Based UDP, revenue increases from the additional fare classes diminish as the number of fare classes increases.

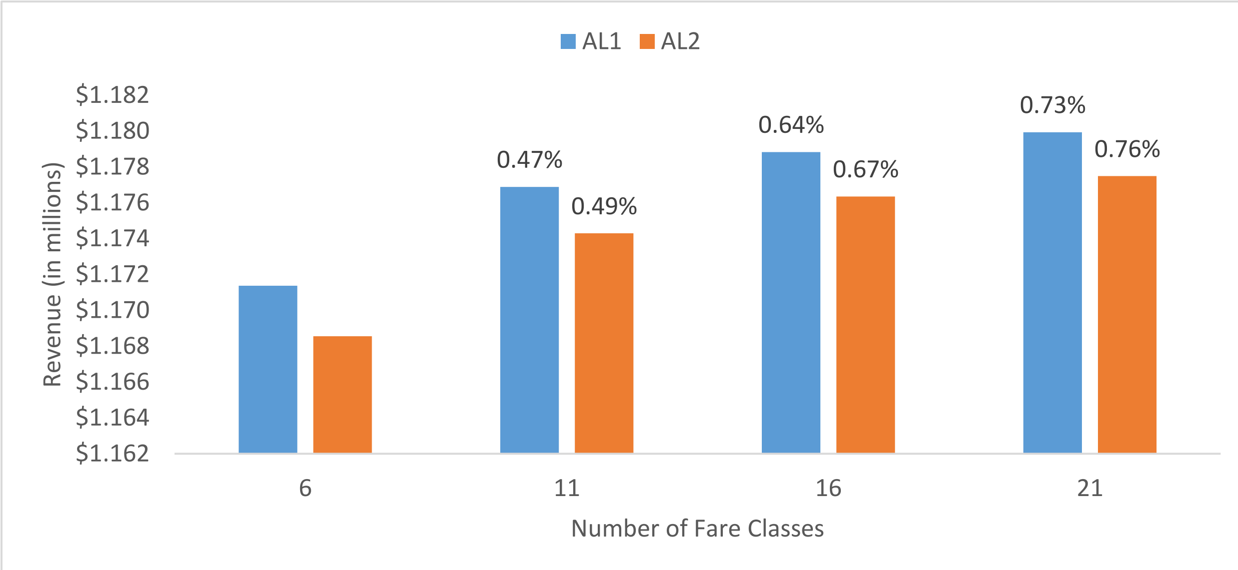

Figure 5-42: Class-Based Continuous UDP Revenue

Class-Based Continuous UDP's revenue increases are in stark contrast to the revenue decreases Class-Based Continuous ProBP caused when the number of fare classes increased (Figure 5-8). The reason for those revenue decreases was found to likely be that forecasts were allowed to and did grow unchecked as a result of a difference in forecasting methods used by traditional and continuous class-based RM methods. While Class-Based Continuous UDP uses the same forecasting method as Class-Based Continuous ProBP (and does not have any limit on how much its bookings are allowed to be scaled), Figure 5-43 shows that the forecasts only increase a modest amount for Class-Based Continuous UDP as fare classes are added.

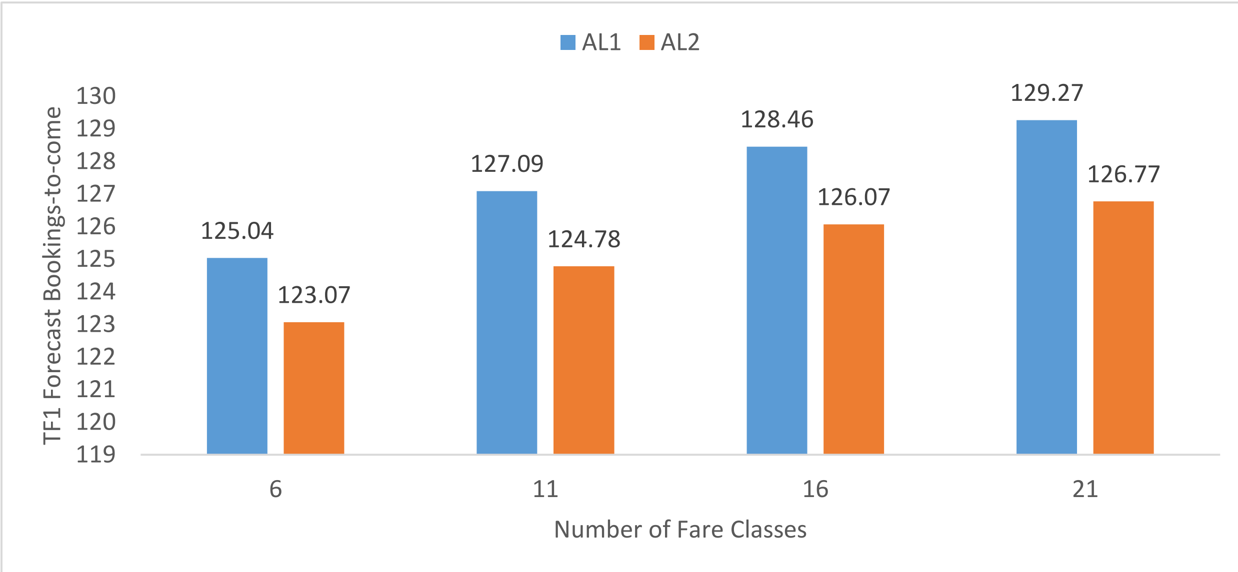

Figure 5-43: Class-Based Continuous UDP Time Frame 1 Forecast Bookings-to-come

The most likely explanation for the differences in forecast behavior for Class-Based Continuous ProBP and UDP is that, in order for the dynamic program used by UDP to work properly, the assumed forecast demand distribution for UDP must be a Poisson distribution, a limit which is not imposed upon ProBP. As the mean-to-variance ratio of demand used in the PODS simulator (and in the real world) is typically greater than 1 (the mean-to-variance ratio for a Poisson distribution), UDP will almost always assume a forecast demand variance smaller than the actual variance of demand. An RM optimizer that is more certain of late booking high fare demand will be more likely to make lower fare classes available early in the booking process. As lower fare class bookings contribute less to a Q-Forecast, the forecast demand mean (which is said Q-Forecast) will always be smaller for UDP than ProBP, particularly in unrestricted fare structures like the ones used in these experiments.

The effects of lower forecasts are readily apparent when observing the bidprices (Figure 5-44). The higher forecasts for the cases with a greater number of fare classes increase bidprices in early TFs, but this increase diminishes as fare classes are added. This causes what was observed with both traditional Class-Based ProBP and traditional Class-Based UDP, with earlier TF fares being higher causing lower bookings in early TFs and then more bookings in later TFs (Figure 5-45).

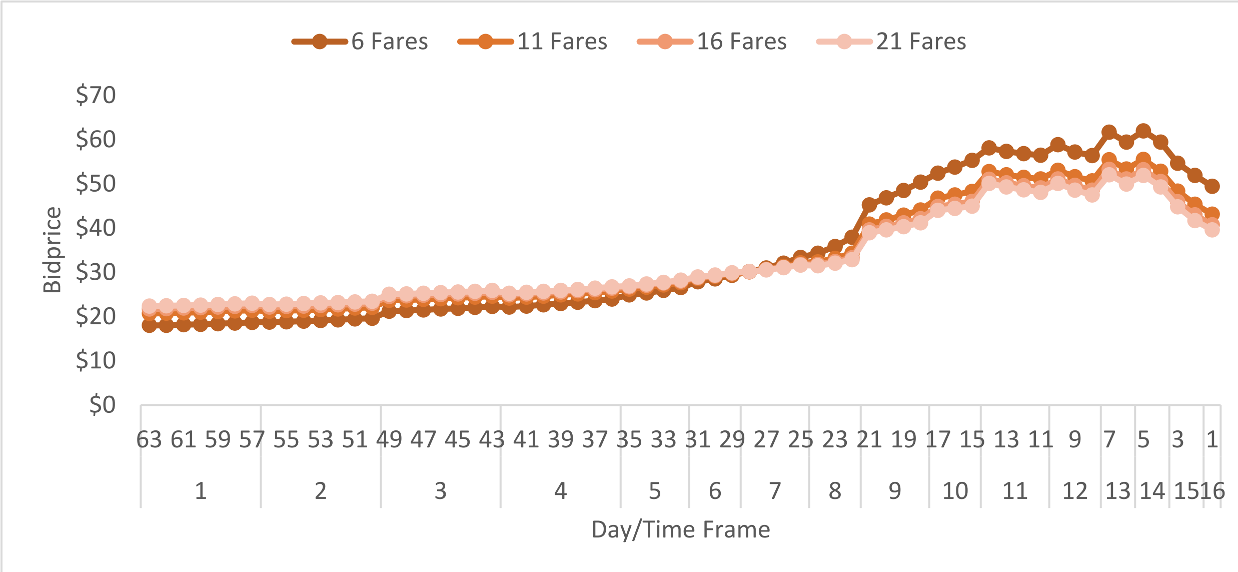

Figure 5-44: Airline 1 Class-Based Continuous UDP Bidprices

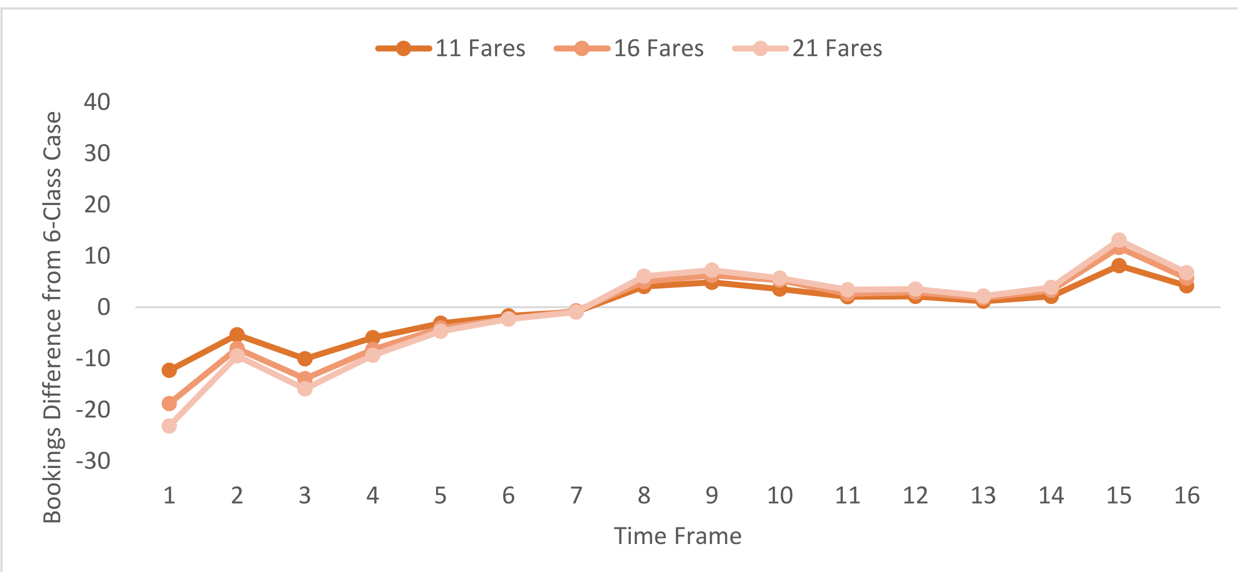

Figure 5-45: Airline 1 Traditional Class-Based UDP Bookings Difference from 6-Class Case

5.2.2.1 Comparison: Traditional and Continuous Class-Based UDP in Symmetric Competition

Having established that both traditional and Class-Based Continuous UDP saw increased revenue when fare classes were added, their revenue performances can be compared to each other. Figure 5-46 shows that, although adding fare classes increased revenue for both methods, traditional Class-Based UDP saw far greater gains and generated more revenue with 11, 16, and 21 fare classes. While traditional Class-Based UDP saw greater revenue improvements with the increase of the number of fare classes from 6 to 11 and then again from 11 to 16 fare classes, the gap actually narrows between traditional class-based and Class-Based Continuous UDP when moving from 16 to 21 fare classes. It is, therefore, possible that traditional Class-Based UDP leads to more rapid gains from additional fare classes, but that the incremental effect of these gains diminishes more quickly for traditional Class-Based UDP than for Class-Based Continuous UDP.

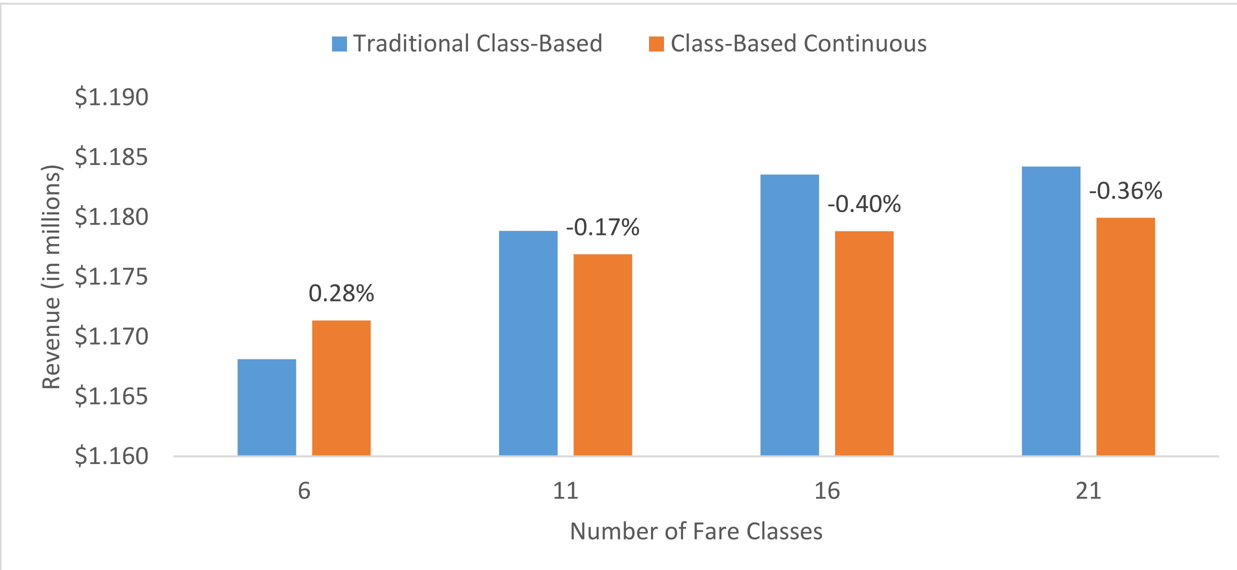

Figure 5-46: Airline 1 Traditional or Continuous Class-Based UDP Revenue: (percent change in revenue from switching from traditional to continuous class-based)

Average paid fare is one metric that shows signs of convergence for traditional and continuous Class-Based UDP. In the 6-fare class case, both the traditional class-based and the class-based continuous methods have similar average paid fares through the first eight TFs. Once the gaps in the fixed fares increase in later TFs, however, class-based continuous and traditional Class-Based UDP fares separate as Class-Based Continuous UDP's fare granularity advantage sets in (Figure 5-47). For 21 fare classes, however, the traditional class-based and Class-Based Continuous UDP average paid fares, while still slightly different in later TFs, are much closer together (Figure 5-48).

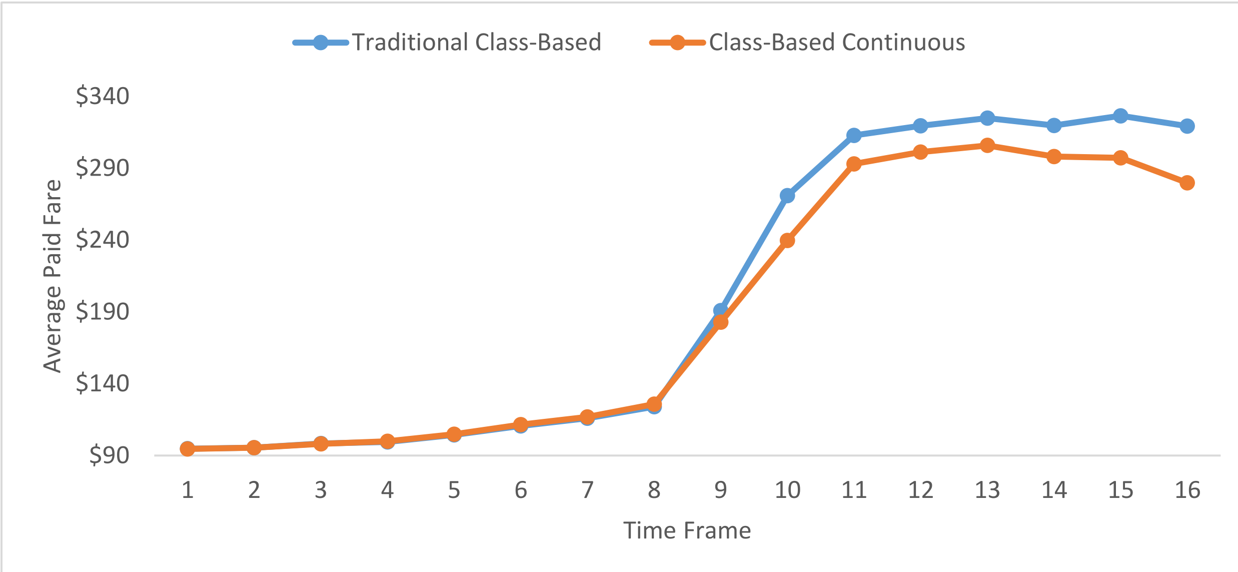

Figure 5-47: 6-Class Airline 1 Traditional or Continuous Class-Based UDP Average Fare

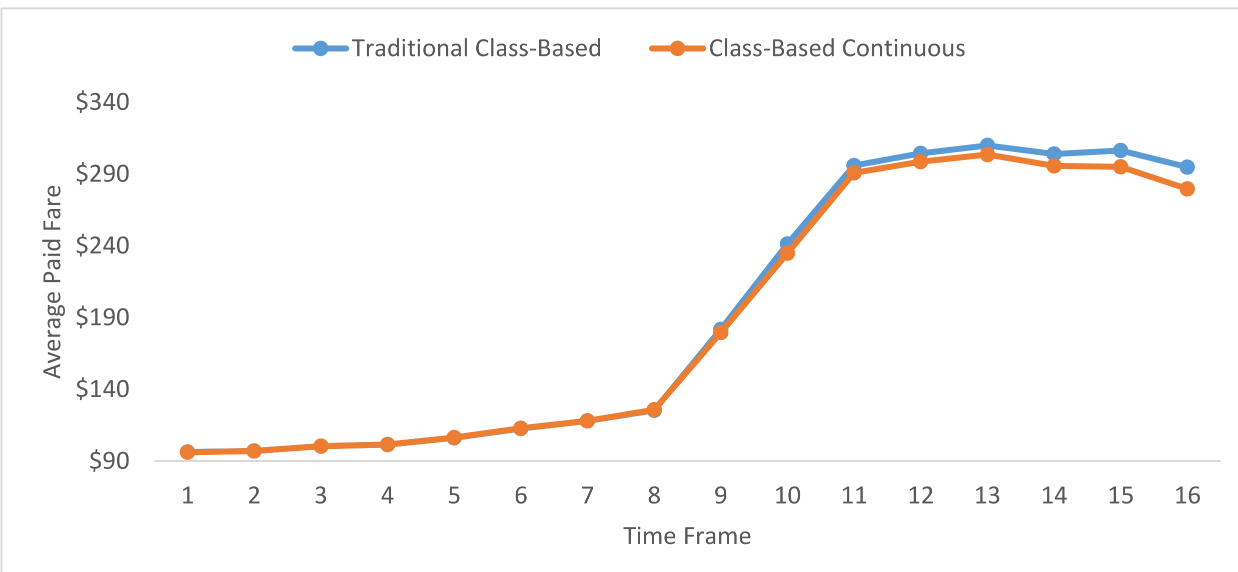

Figure 5-48: 21-Class Airline 1 Traditional or Continuous Class-Based UDP Average Fare

Bookings by TF is another metric that seems to indicate convergence for the Class-Based UDP methods. While bookings are already very similar in the six-fare class case, there is a small amount of separation in TFs 9–15, which are also the TFs when the Class-Based Continuous UDP's fare granularity becomes an advantage (Figure 5-49). In the 21-fare class case, however, there is almost no separation at all as the traditional Class-Based UDP method is provided far more fare granularity by the extra classes and is able to sufficiently cover passengers WTPs in later TFs (Figure 5-50).

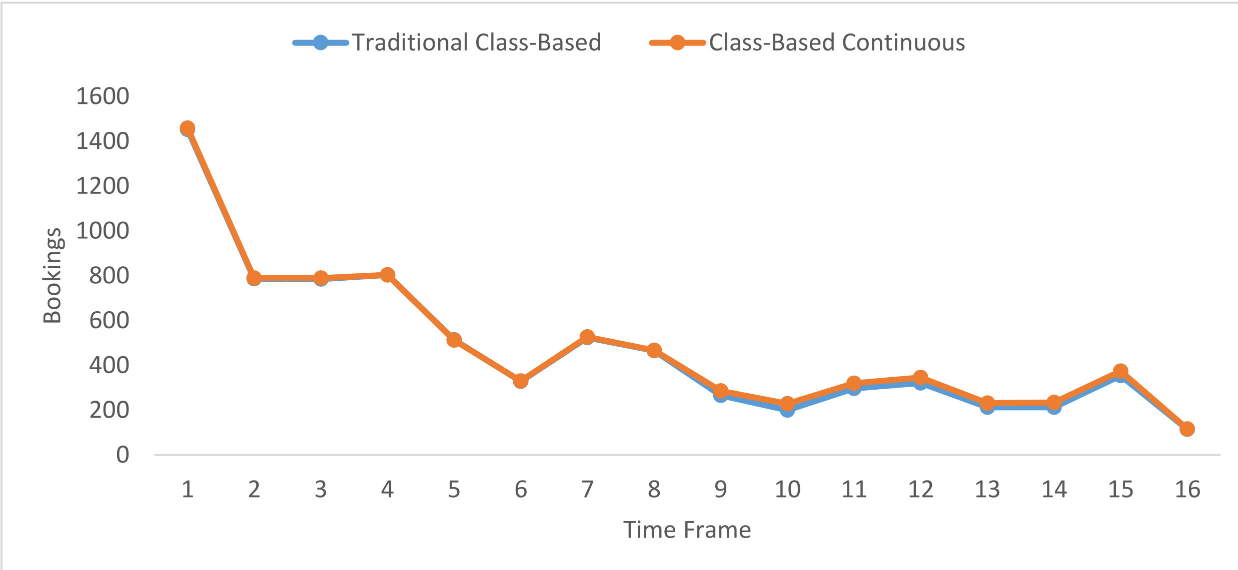

Figure 5-49: 6-Class Airline 1 Traditional or Continuous Class-Based UDP Bookings

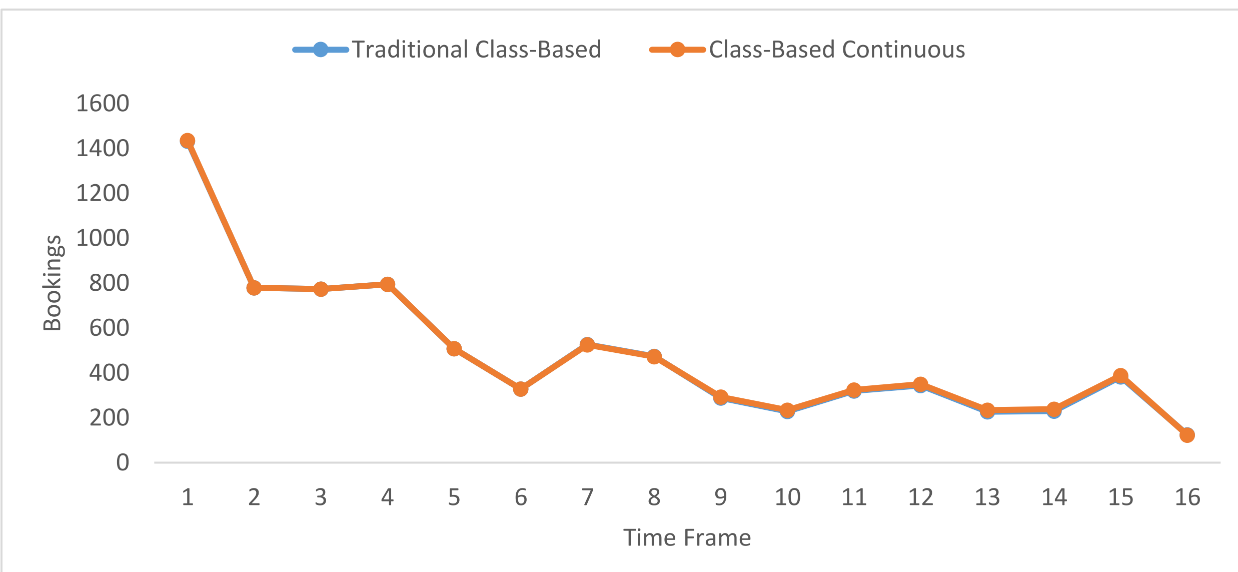

Figure 5-50: 21-Class Airline 1 Traditional or Continuous Class-Based UDP Bookings

5.2.3 Class-Based Continuous vs. Traditional Class-Based UDP in Asymmetric Competition

Experiments with ProBP showed that, while the difference between a continuous class-based RM method and a traditional class-based method could be limited with enough fare classes in symmetric competition, asymmetric competition could allow the class-based continuous method to undercut the traditional method in terms of fares offered. More granularity allowed the airline using the class-based continuous method to have fares lower than the traditional class-based method-using competitor's fares. This was also tested for UDP. As before, these experiments took place in in Network D6, with AL1 switching from traditional Class-Based UDP to Class-Based Continuous UDP, while both airlines continued to use FRAT5 curve C and fare adjustment. Figure 5-51 shows that the results for UDP are much the same as they were for ProBP, with AL1 gaining a revenue increase that decreased as both airlines added fare classes.

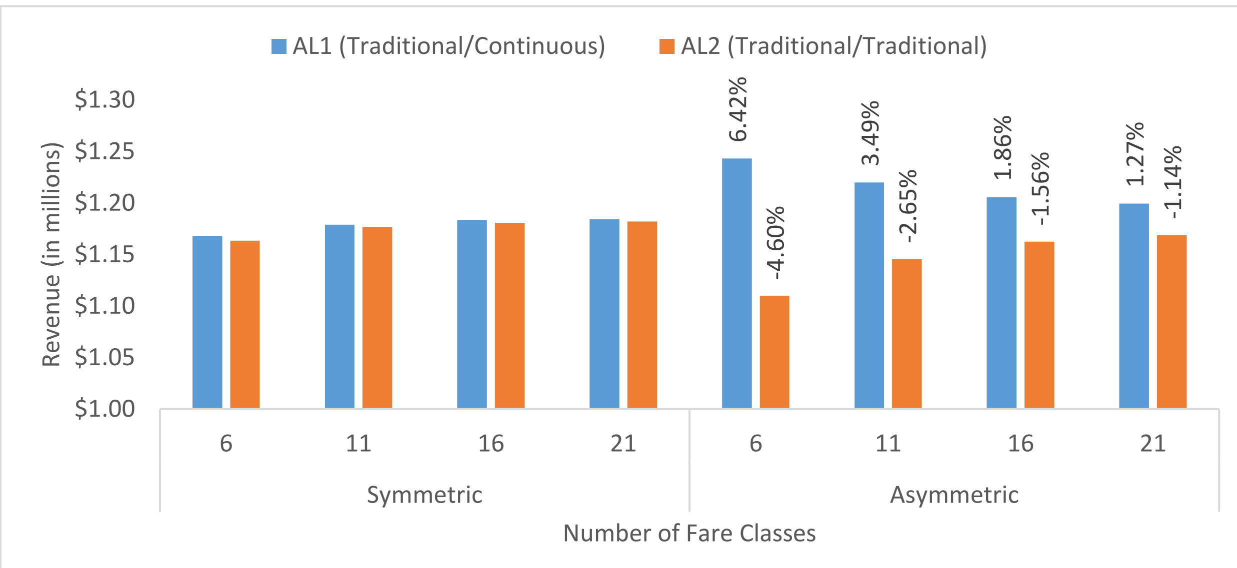

Figure 5-51: Continuous vs. Traditional Class-Based UDP Revenue (percent change in revenue resulting from Airline 1 switching from traditional to class-based continuous)

The similarities of the behavior of revenue indicate very strongly that Class-Based Continuous UDP, like Class-Based Continuous ProBP, undercuts its traditional class-based competitor in terms of price in later TFs and picks up more higher-WTP passengers. This is confirmed by observing the booking data (which is a better indicator than average paid fare of which airline offered a lower fare in an asymmetric test), which, like with ProBP, shows AL1 giving up early TF bookings and picking up more later TF bookings (Figure 5-52). Additionally, this effect also diminishes as AL2 adds fare classes and gains pricing granularity in later TFs.

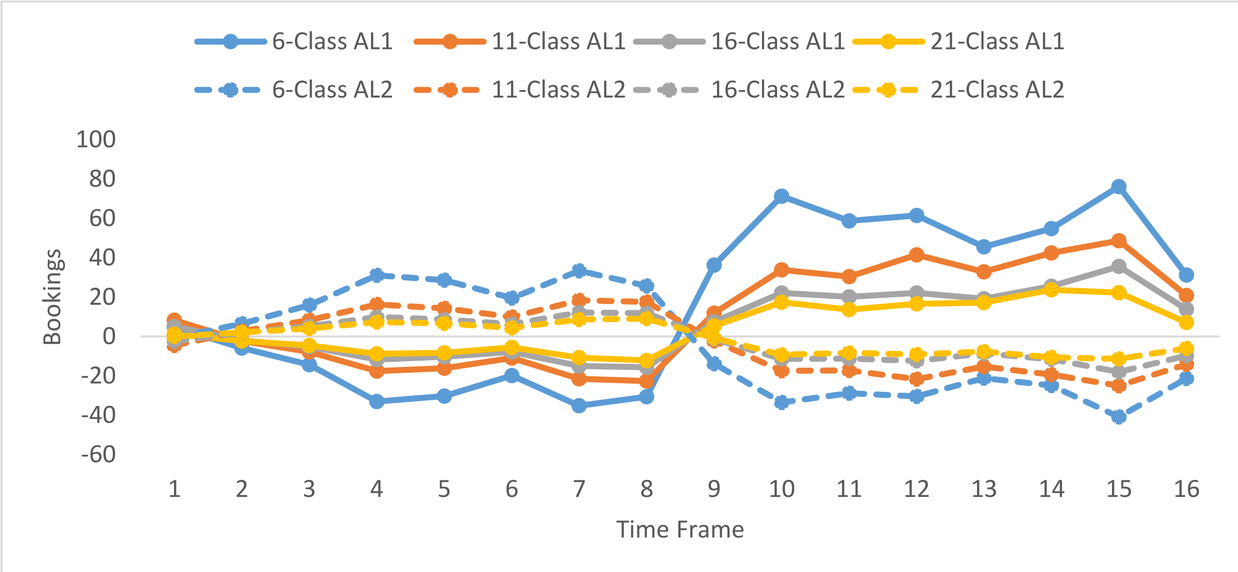

Figure 5-52: Continuous vs. Traditional Class-Based UDP Change in Airline Bookings from Symmetric to Asymmetric Experiments

As with ProBP, the question with UDP becomes for AL2 whether it can increase its revenue by simply using a less aggressive FRAT5 curve and undercutting AL1. The asymmetric Class-Based UDP tests were, therefore, rerun with AL2 instead using FRAT5 E. The results of AL2 lowering its FRAT5 curve swing revenue heavily in favor of AL2, with this advantage decreasing as fare classes are added (Figure 5-53).

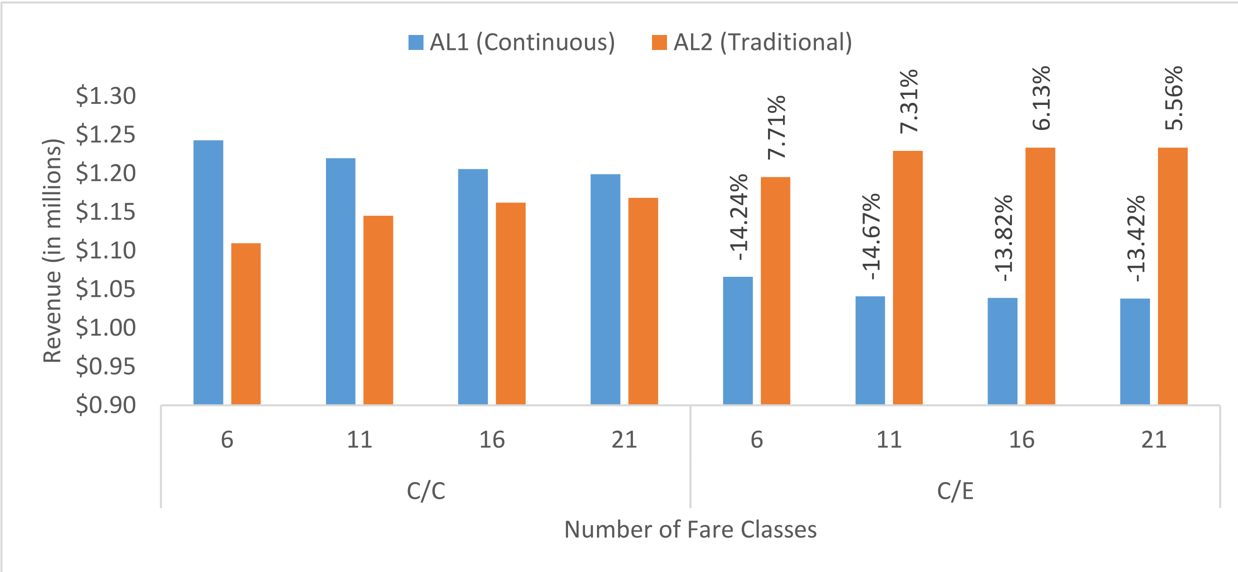

Figure 5-53: Continuous vs. Traditional Class-Based UDP with Different FRAT5s Revenue (percent change in revenue resulting from Airline 2 switching FRAT5 curve from C to E)

Observing the bookings by TF data, the reason for this revenue swing becomes very clear. AL2 not only recaptures some of the later TF bookings from AL1 by lowering its FRAT5 curve, but also captures a substantial portion of AL1's original later TF bookings (Figure 5-54). As with ProBP, this recapture likely occurs due to AL2 allowing fares to be open in later TFs. AL1, as a result of fare adjustment, never allows these fares to be open.

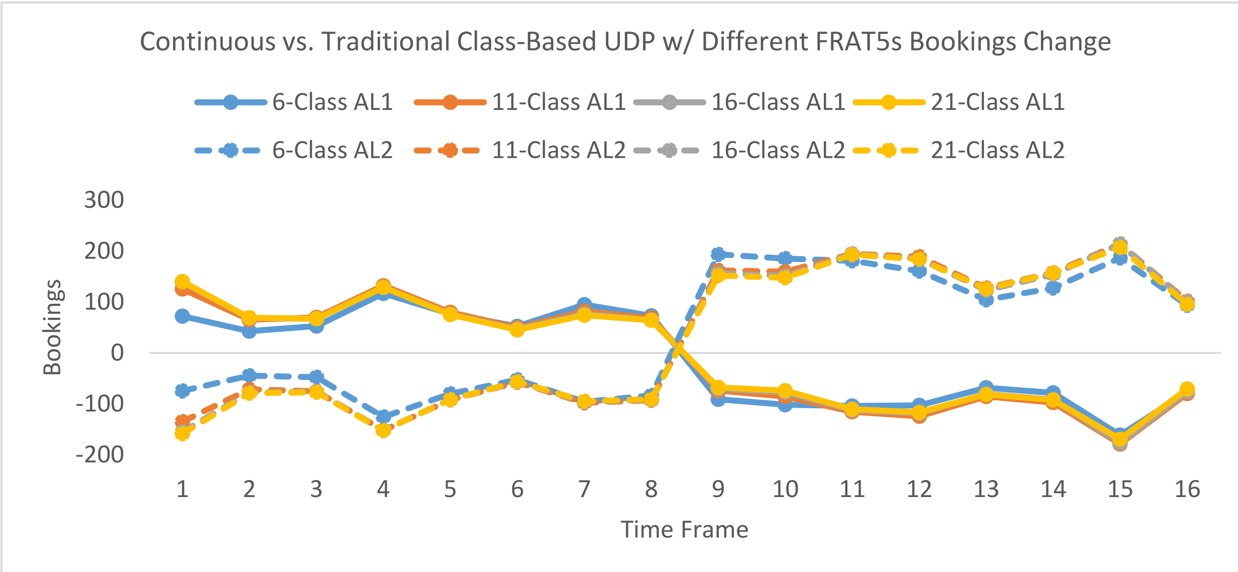

Figure 5-54: Continuous vs. Traditional Class-Based UDP Change in Bookings from the Traditional Class-Based Airline 2 Switching to FRAT5 E

It is important to note that Class-Based Continuous UDP's potential to be undercut is not a complete indictment of its ability to outperform traditional Class-Based UDP in asymmetric competition. In the previous example, AL1 could very easily lower its FRAT5 curve and recapture many of the bookings lost to AL2. What is important to note is that, in asymmetric unrestricted competition, whichever airline offers the lower fares in later TFs will generally end up generating more revenue.

5.2.4 Classless UDP in Symmetric Competition

The final method for continuous pricing to be discussed is Classless UDP. The tests for Classless UDP were the same as the tests for Classless ProBP, starting off with determining a FRAT5 curve that maximized revenue in symmetric competition. As with all of the other continuous pricing methods, FRAT5 C was found to maximize Classless UDP's revenue with symmetric competition in Network D6 (Figure 5-55). As the method is classless, there were no and could be no tests concerning the effect of adding fare classes to Classless UDP, so the first set of comparisons to be made will be of comparing traditional Class-Based UDP to Classless UDP with symmetric competition.

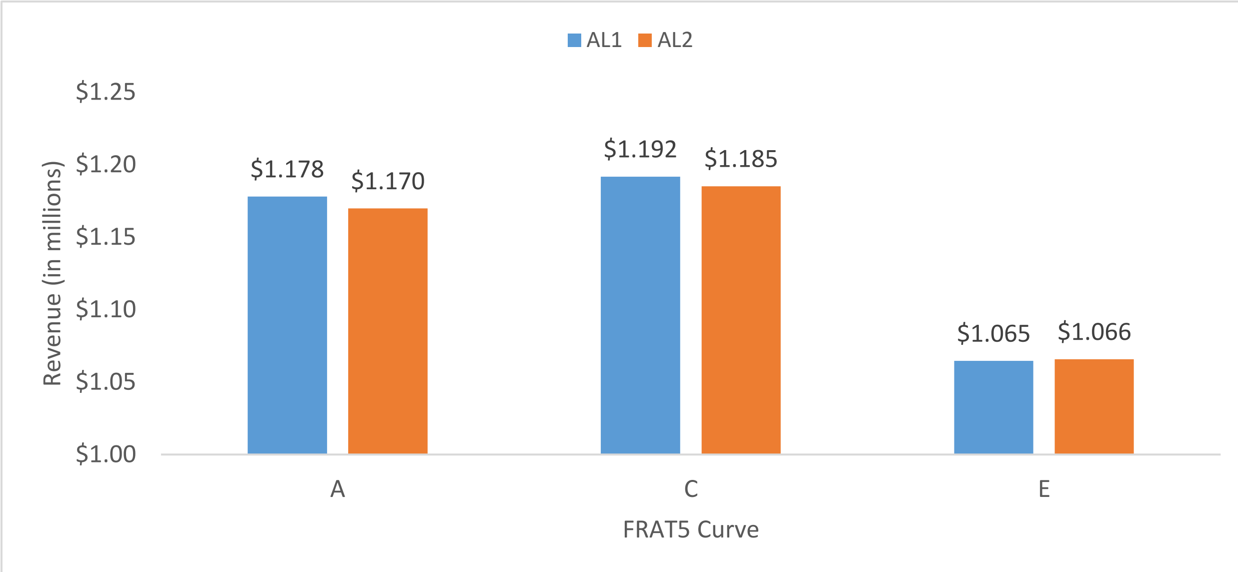

Figure 5-55: Classless UDP Revenue

5.2.4.1 Comparison: Traditional Class-Based and Classless UDP in Sym. Networks

As was the case with Classless ProBP, no convergence for traditional class-based and Classless UDP was expected as the optimization algorithms for each are different. This was confirmed by the revenue results, which show that Classless UDP gains an increase in revenue over 21-fare class traditional UDP (Figure 5-56).

Figure 5-56: Traditional Class-Based or Classless UDP Revenue (percent change in revenue from 6-class experiment)

The reason for Classless UDP's revenue increase is the same as the reason which revenue increased in most of the previous experiments: Classless UDP quotes much lower fares in later TFs and is far more aggressive about protecting seats in early TFs with higher fare quotations (Figure 5-57). The result of this is, once again, an exchange of lower-priced, early bookings for higher-priced, later bookings (Figure 5-58).

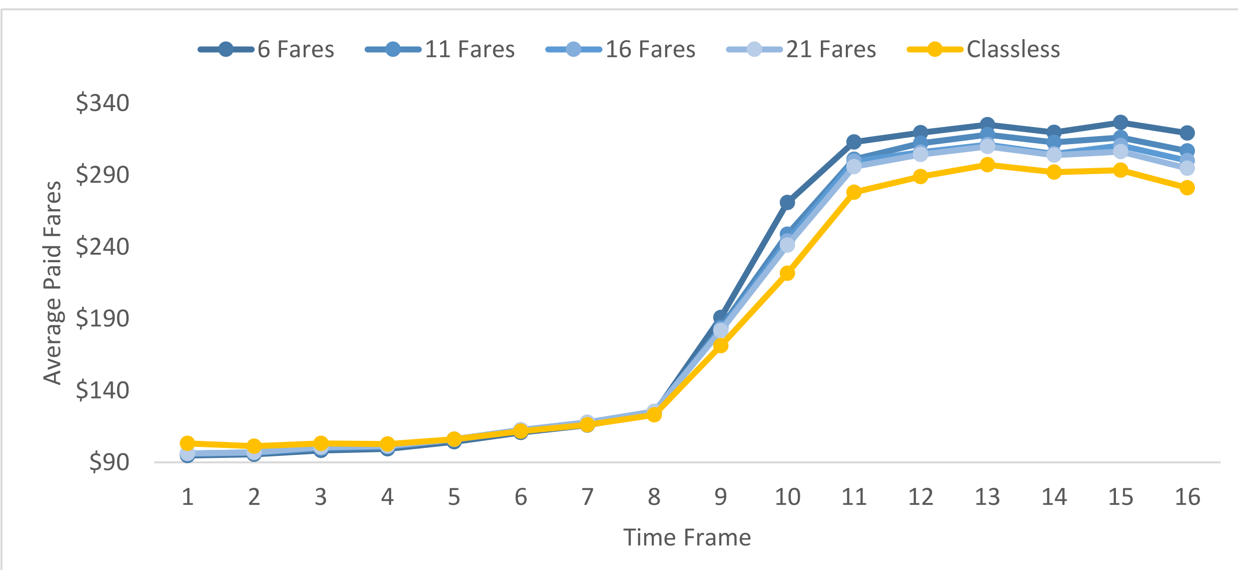

Figure 5-57: Airline 1 Traditional Class-Based or Classless UDP Average Fares

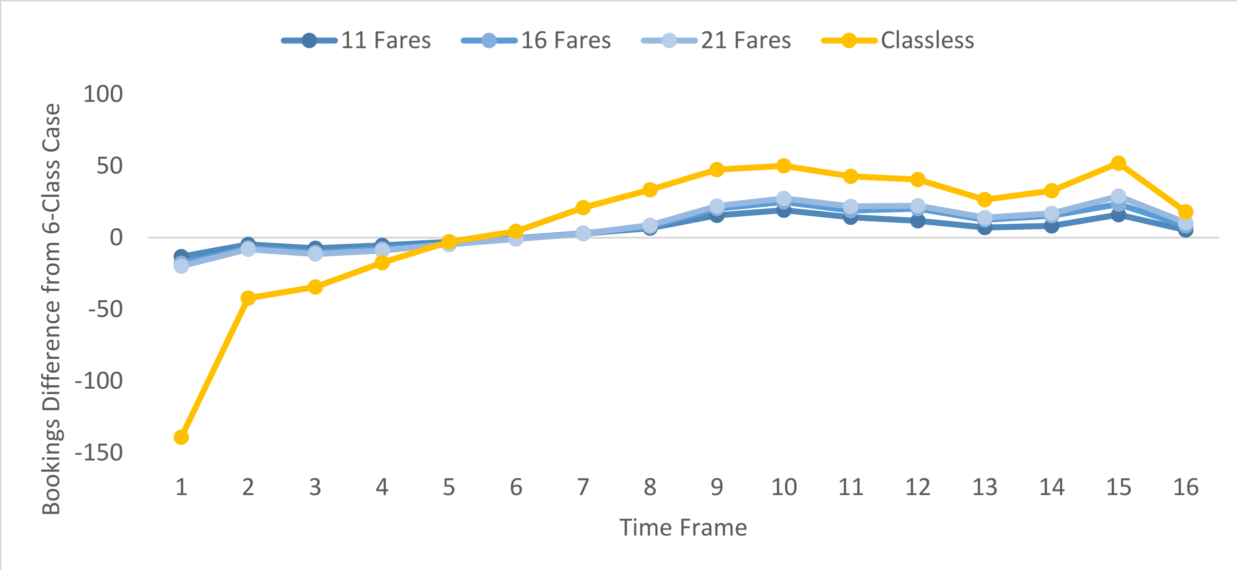

Figure 5-58: Airline 1 Traditional Class-Based or Classless UDP Bookings Difference from 6-Class Case

The reason for this change in aggressiveness is shown in the forecast of traditional Class-Based and Classless UDP, as the forecast bookings-to-come for the latter increases by about 33% over the former (Figure 5-59). While the increased forecasts are no doubt encouraged by the extra bookings in the later TFs, they are also self-supporting as they cause Classless UDP to be more aggressive in those early TFs and keep seats available in later TFs.

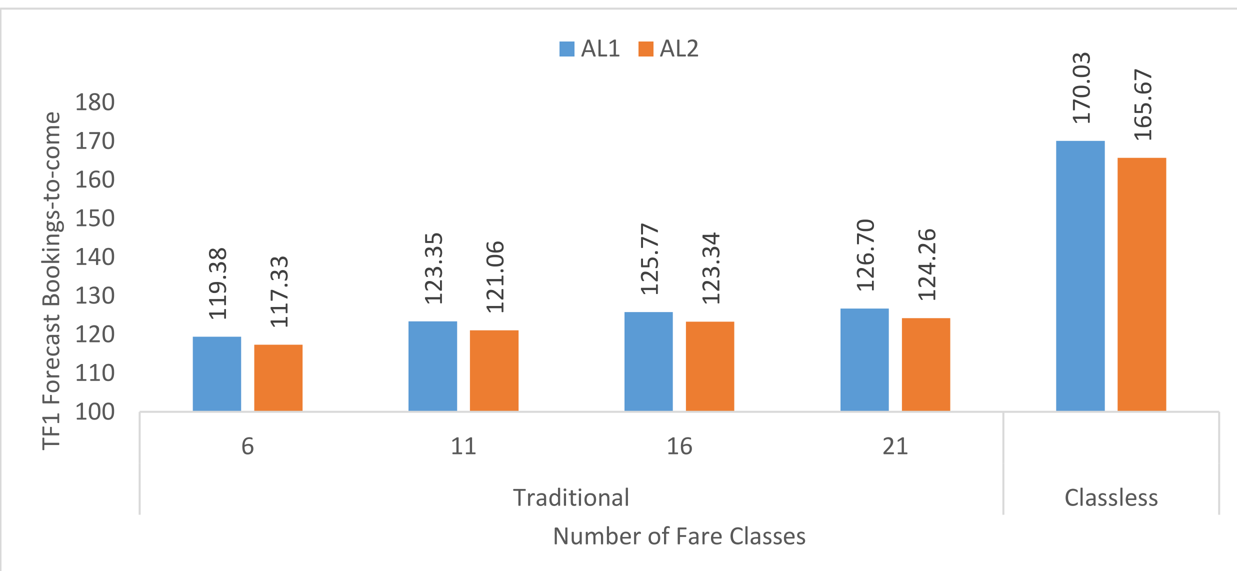

Figure 5-59: Traditional Class-Based or Classless UDP Time Frame 1 Forecast Bookings-to-come

5.2.4.2 Comparison: Class-Based Continuous and Classless UDP in Symmetric Competition

Because of the convergence of traditional Class-Based UDP and Class-Based Continuous UDP with symmetric competition, it is not surprising that Classless UDP, after outperforming traditional Class-Based UDP, substantially outperforms Class-Based Continuous UDP. For AL1, this increase constitutes an additional 1% increase over AL1 in the 21-class continuous UDP case (Figure 5-60).

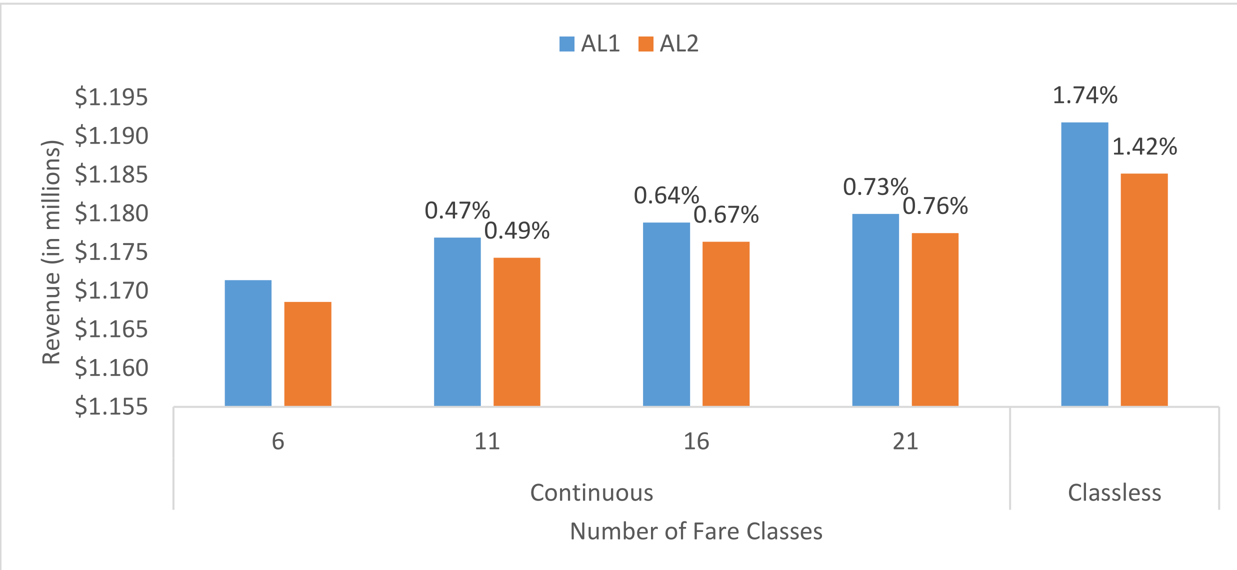

Figure 5-60: Class-Based Continuous or Classless UDP Revenue (percent change in revenue from 6-class experiment)

When considering the underlying metrics responsible for the revenue increase, the theoretical advantages of using a classless method instead of a class-based one for continuous pricing are on full display with UDP. For example, while the difference is not as big as it was between traditional class-based and Classless UDP, Classless UDP offers lower later TF fares than those of the Class-Based Continuous UDP (Figure 5-61), capturing additional later TF passengers (Figure 5-62). As a result, Classless UDP does a better job protecting against low-fare, early TF bookings, with booking counts substantially lower in those TFs than with any number of fare classes for Class-Based Continuous UDP (Figure 5-62).

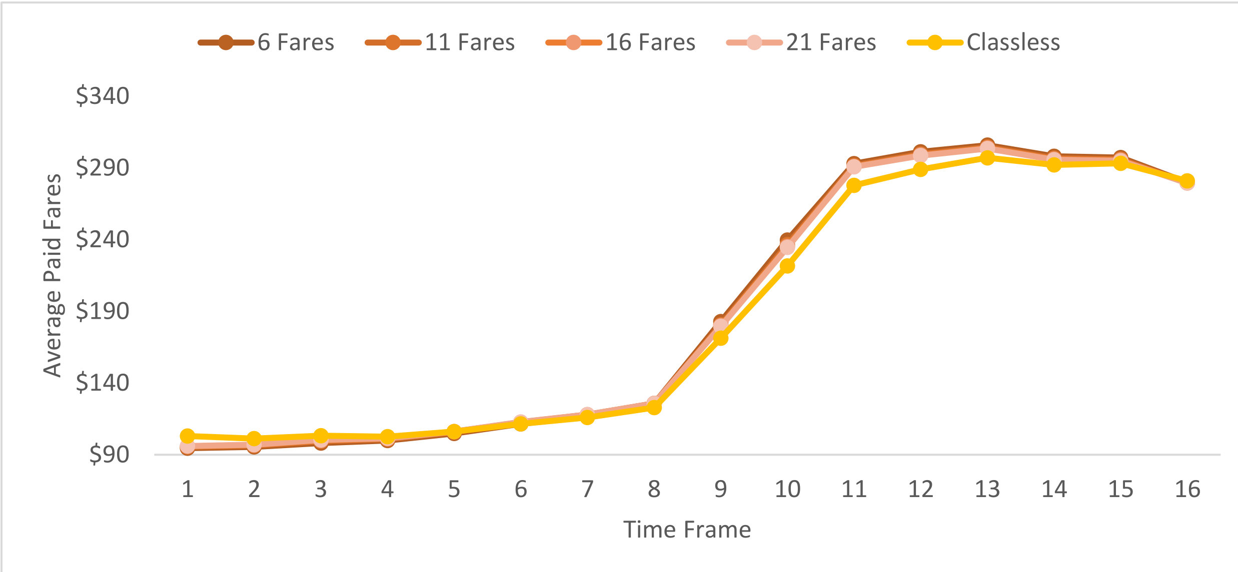

Figure 5-61: Airline 1 Class-Based Continuous or Classless UDP Average Fares

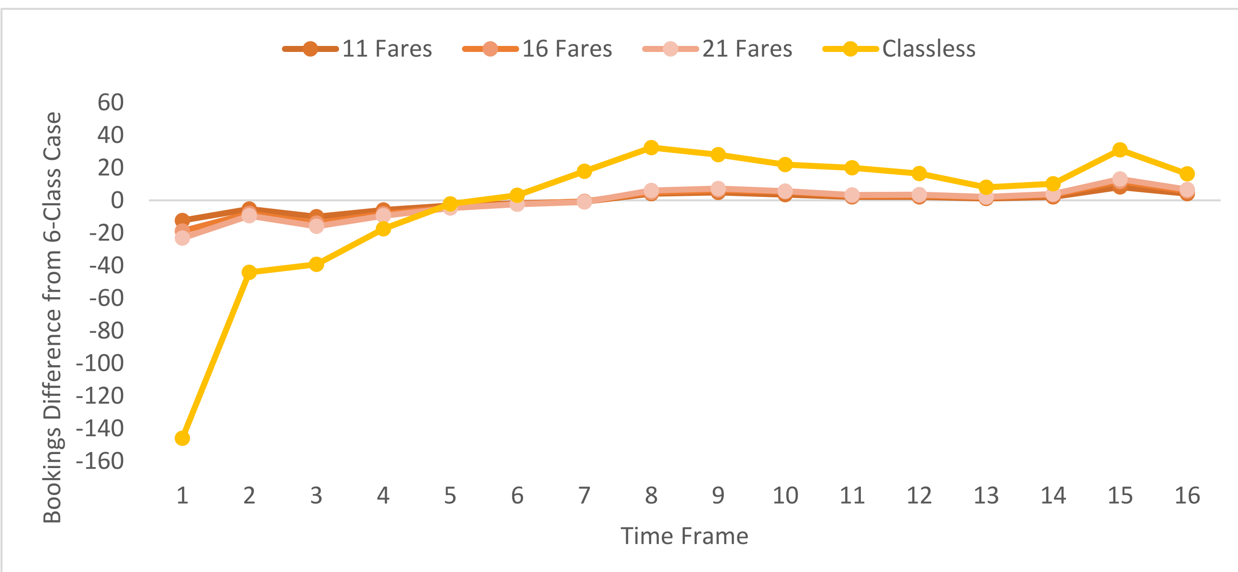

Figure 5-62: Airline 1 Class-Based Continuous or Classless ProBP Bookings Difference from 6-Class Case

As before, these better seat protection practices (for situations with symmetric competition) result from Classless UDP having substantially higher bookings than Class-Based Continuous UDP (Figure 5-63). The forecasts, along with the revenue results and the other metrics shown here, demonstrate how, with symmetric competition, Classless UDP is better for revenue than either of the methods it was derived from.

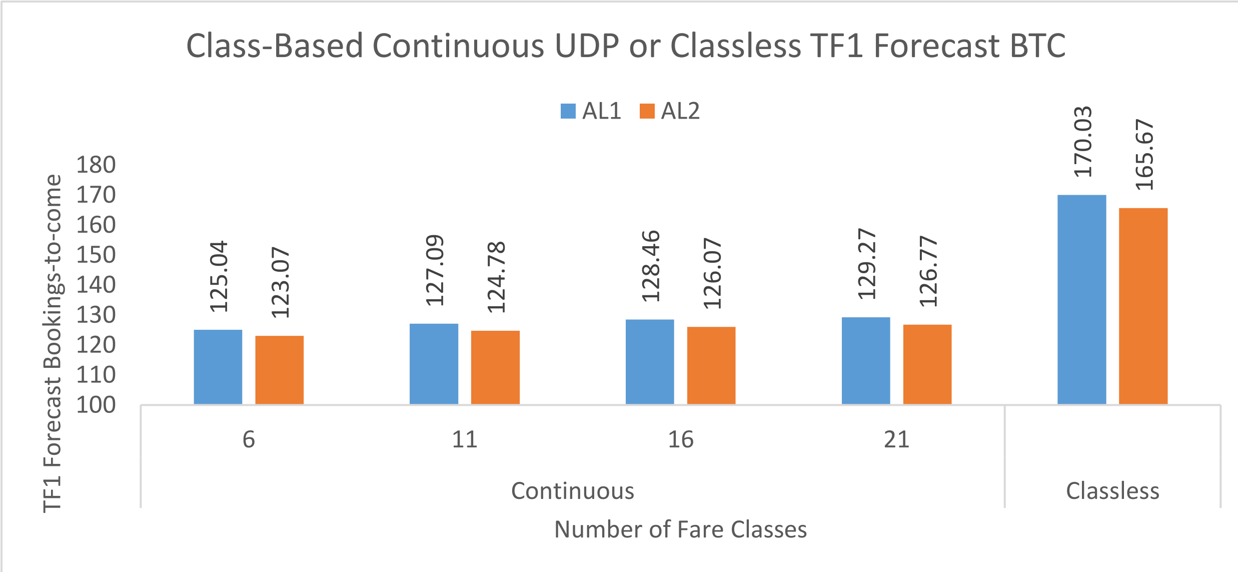

Figure 5-63: Airline 1 Class-Based Continuous or Classless UDP Bookings Difference from 6-Class Case

5.2.5 Classless vs. Traditional Class-Based UDP in Asymmetric Competition

Although Classless UDP has been shown to outperform traditional Class-Based UDP in terms of revenue when they are used in separate cases with symmetric competition, the last question about Classless UDP is how it performs against traditional Class-Based UDP in scenarios with asymmetric competition. As with previously discussed experiments on asymmetric competition, this was tested by starting from the traditional Class-Based UDP tests in Network D6 and having AL1 switch to Classless UDP. AL1, under these conditions, was found to generate more revenue than traditional Class-Based UDP with any number of fare classes, with its margin of revenue advantage decreasing as fare classes were added for AL2 (Figure 5-64).

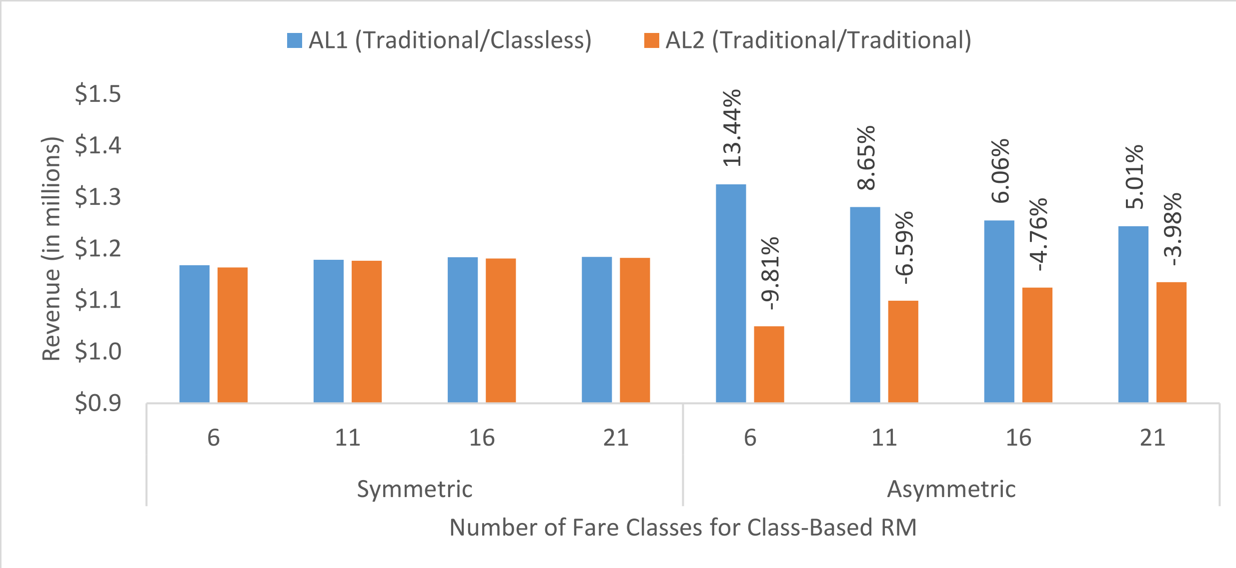

Figure 5-64: Classless vs. Traditional Class-Based UDP Revenue (percent change in revenue resulting from Airline 1 switching from traditional to classless)

As with all previously discussed cases of asymmetric competition, the reason for AL1's revenue gains is that it offers lower fares in later TFs and higher fares in earlier TFs, therefore resulting in it undercutting AL2 for the more valuable bookings which occur in the later TFs (Figure 5-65).

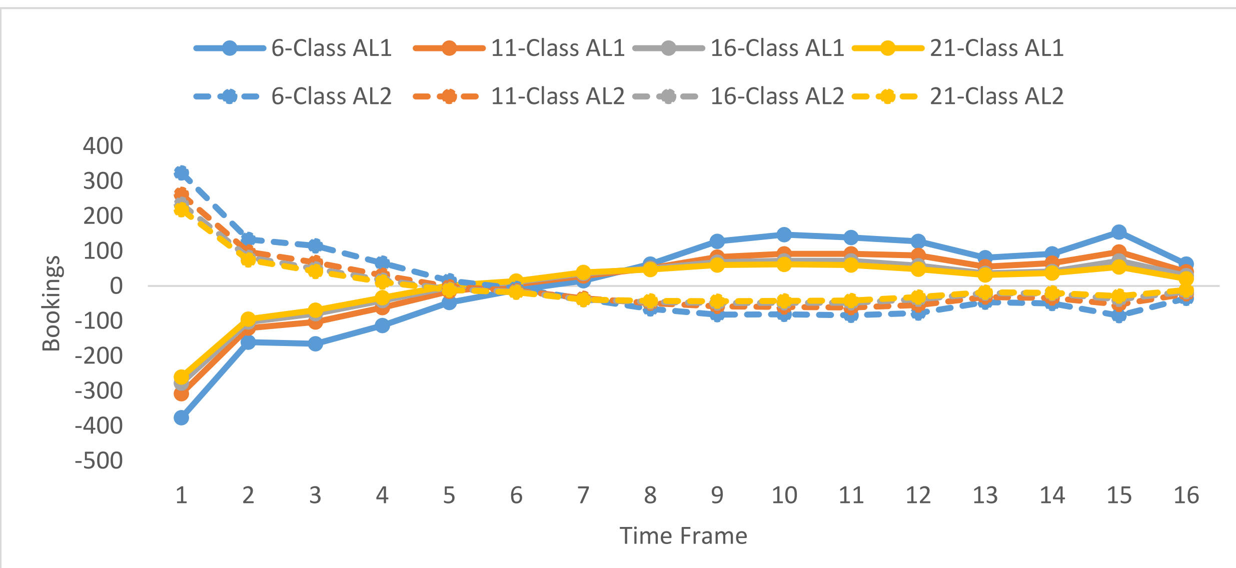

Figure 5-65: Classless vs. Traditional Class-Based UDP Change in Airline Bookings from Symmetric to Asymmetric Experiments

5.2.6 Summary of UDP Results

As with ProBP, the UDP methods for continuous pricing both showed revenue gains over traditional Class-Based UDP. In symmetric competition tests, Class-Based Continuous UDP showed signs of converging with traditional Class-Based UDP and lost any revenue advantage with just 11 fare classes. However, in asymmetric competition tests, Class-Based Continuous UDP was able undercut the prices offered by traditional Class-Based UDP in later TFs to gain additional higher-priced bookings. Classless UDP, on the other hand, showed revenue gains in both the symmetric and asymmetric experiments as a result of it being able to better protect for the aforementioned more valuable bookings in later TFs.

5.3 Comparison of ProBP and UDP in Symmetric Competition

With the advantages and disadvantages of Class-Based Continuous and Classless ProBP and UDP over their baselines established, the final question becomes whether the ProBP or the UDP methods generate more revenue and why. This will be answered by taking data from the previous ProBP and UDP symmetric competition experiments documented in this chapter and comparing the results to each other. Figure 5-66 shows that said answer is definitive: in the simulations documented in this thesis, the ProBP method outperforms the corresponding UDP method for all symmetric competition cases tested in terms of revenue generated.

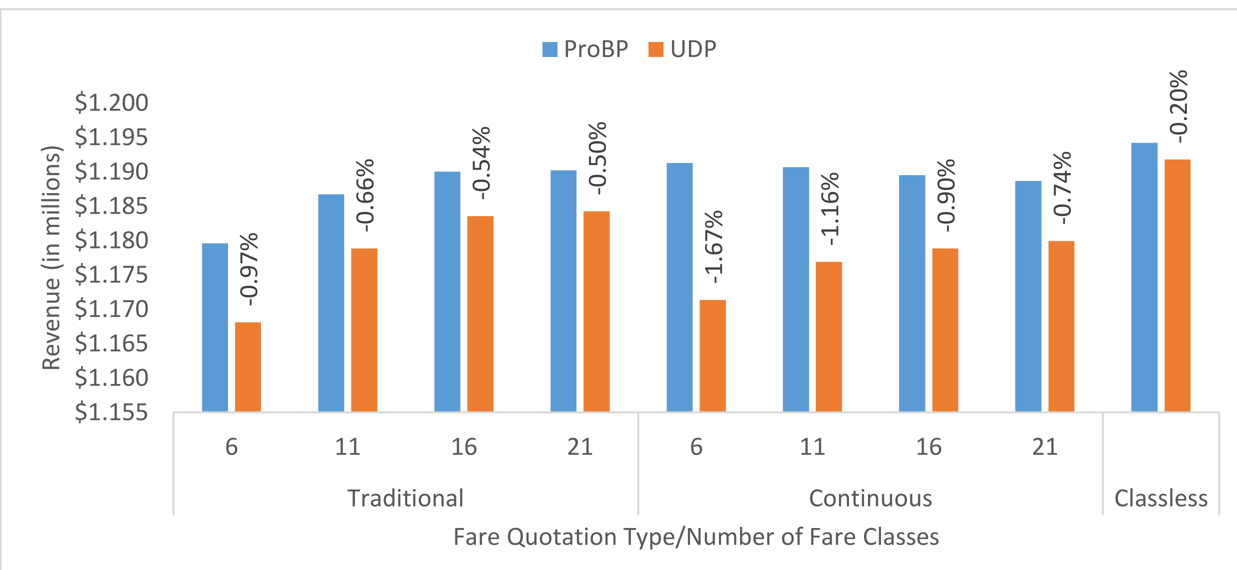

Figure 5-66: Airline 1 ProBP or UDP Revenue (percent change in revenue from between ProBP and UDP)

To explain why the ProBP methods generate more revenue than the UDP methods, the 16-fare class cases will be considered, as this is the number of fare classes where the effect of adding fare classes levels off for traditional Class-Based ProBP, traditional Class-Based UDP, and Class-Based Continuous UDP (obviously, Classless ProBP and Classless UDP do not use fare classes and may be compared directly to each other). With 16 fare classes, class-based continuous, then traditional class-based, and then finally the classless case show the biggest revenue drops for UDP compared to ProBP (Figure 5-67).

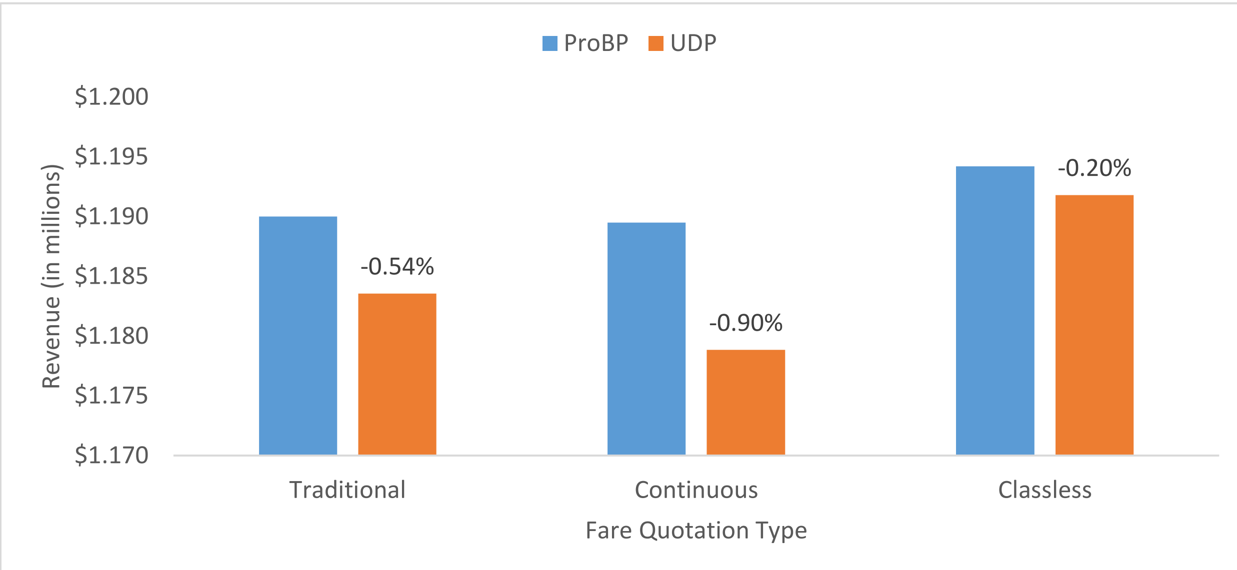

Figure 5-67: 16-Class Airline 1 ProBP or UDP Revenue (percent change in revenue from between ProBP and UDP)

The first place to look for differences between ProBP and UDP is the forecasts each method generates. As has previously been shown, higher forecasts are linked to higher bidprices and more aggressive protection in early TFs, which increases overall revenue. The forecast data for 16-class ProBP and UDP correlates with the revenue performance of each method with a larger difference in ProBP and UDP forecasts corresponding to a larger difference in revenue between ProBP and UDP methods (Figure 5-68).

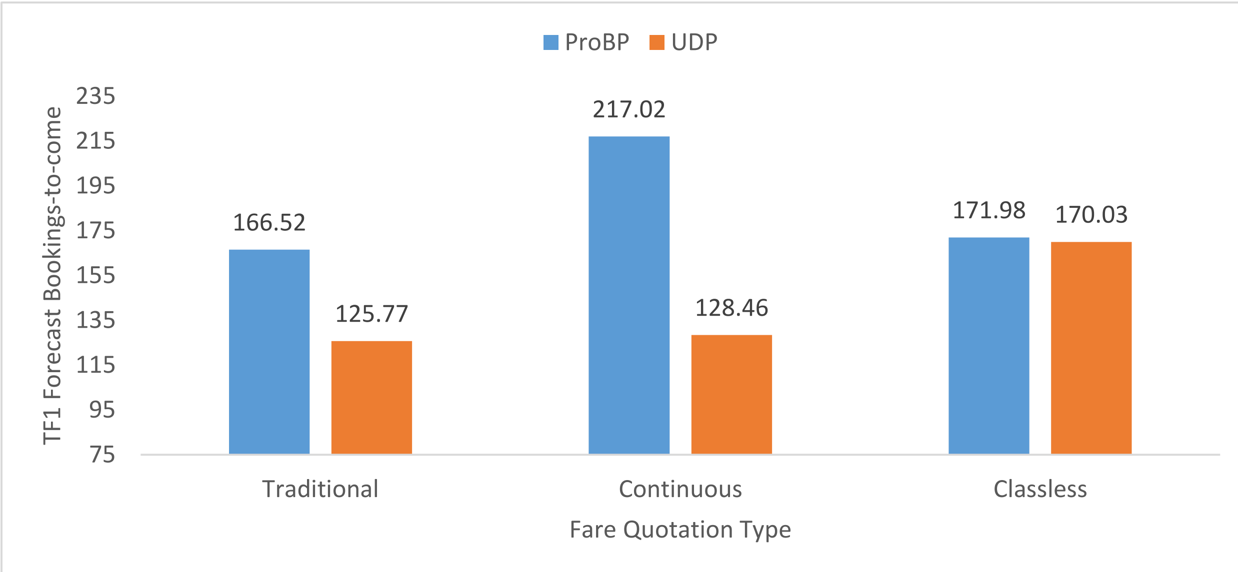

Figure 5-68: 16-Class Airline 1 ProBP or UDP Time Frame 1 Forecast Bookings-to-come

It is worth noting, however, that the difference between Classless ProBP and Classless UDP forecasts is very small, almost to the point of being negligible. While the revenue difference for Classless ProBP and UDP is also not very large, it is at least noticeable. The combination of these two facts suggests that Classless ProBP may have an algorithmic advantage over Classless UDP that goes beyond interaction between the optimization algorithm and the forecaster. The most likely reason for this goes back to UDP's previously mentioned Poisson distribution assumption for arrival of demand.

Given the same forecast, both Classless ProBP and Classless UDP determine what fare to offer in the present TF while protecting for the future. As shown in Section 3.2.5, Classless ProBP's algorithm relies on maximizing expected revenue contribution, which is calculated using forecasted demand means and variances, while Classless UDP uses a dynamic program, which does use demand means but assumes a lower Poisson variance of demand. As previously mentioned, demand in the PODS simulations (and in the real world) has a variance greater than that of the Poisson distribution, which in turn causes Classless UDP to calculate a lower optimal fare than Classless ProBP given the same demand.

Assuming the forecast is accurate (or, at the very least, unbiased), the fact that Classless ProBP accounts for variance makes the optimal fare it calculates more likely to be closer to the "true" revenue maximizing fare. The average paid fare data shows that this is the case, with Classless UDP fares being slightly lower than Classless ProBP fares in TFs 1–12 (Figure 5-69).

Traditional and continuous Class-Based ProBP and UDP, meanwhile, show the typical results of low forecasts on average paid fares of selling many lower fares in early TFs and higher fares in later TFs (Figure 5-69).

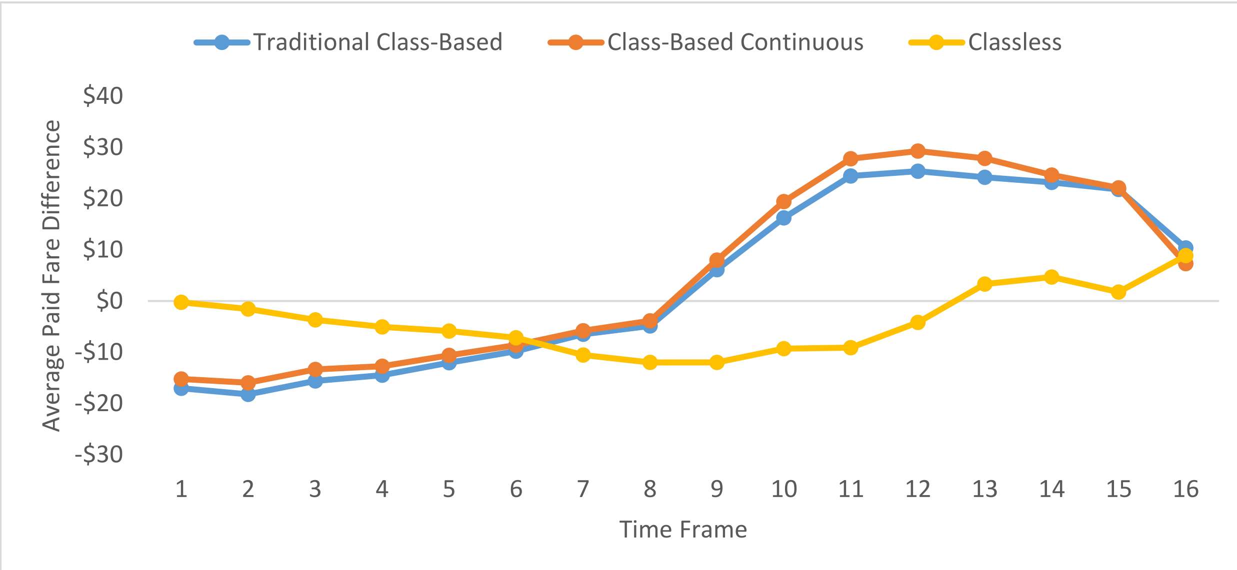

Figure 5-69: 16-Class Airline 1 Average Fare Difference Between ProBP and UDP

The data for bookings by TF shows a similar story to the one offered by the average paid fare data. Traditional and continuous Class-Based UDP's lower fares result in much higher early TF bookings at the expense of more valuable later TF bookings (Figure 5-70). Classless UDP, meanwhile, generates bookings that are generally slightly higher than Classless ProBP's, but these extra bookings are not enough to offset the extra revenue earned per booking by Classless ProBP (Figure 5-70), likely a result of Classless UDP being further off from the "true" revenue maximizing fare.

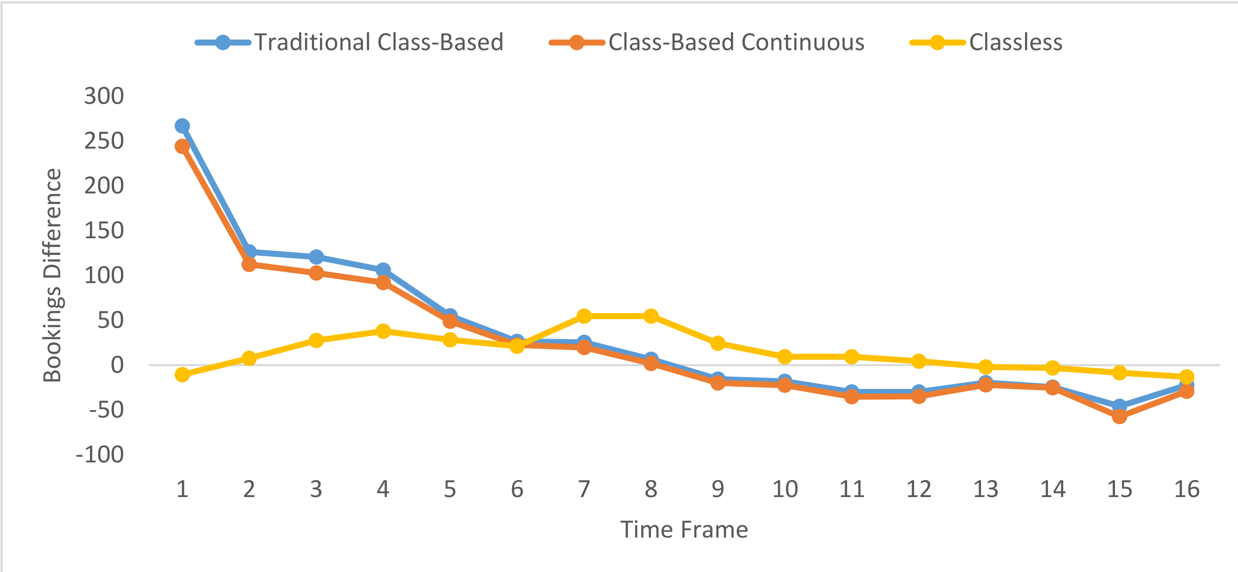

Figure 5-70: 16-Class Airline 1 Bookings Difference Between ProBP and UDP

5.4 Summary

In this chapter, the results of testing continuous pricing methods for ProBP and UDP were documented. For each ProBP and UDP, traditional class-based RM with a varying number of fare classes in a network with symmetric competition was first tested to establish a baseline. This was followed by symmetric and asymmetric competitive tests of class-based continuous RM, also with a varying number of fare classes. Finally, symmetric and asymmetric competitive tests of classless RM methods were also discussed. Across the continuous pricing experiments, one prevailing pattern emerged: the ability to protect seats for later TF bookings was the best indicator for increased revenue.

In symmetric tests, it was found that, for both ProBP and UDP, class-based continuous RM had an advantage over traditional class-based RM with six fare classes (0.99% for ProBP and 0.28% for UDP). Added granularity of fare quotation allowed the class-based continuous methods to offer lower fares in later TFs, generating more revenue. With 21 fare classes, however, this advantage diminished to the point of non-existence, as the increased number of price points for the traditional class-based method allowed the bidprices and fares of the two methods to converge. In asymmetric competitive tests, however, it was found that the fare granularity afforded to class-based continuous methods once again had an effect, as the class-based continuous method was able to undercut the traditional method in later TFs.

Classless ProBP and UDP both outperformed their class-based counterparts in symmetric competition. Classless ProBP gained 0.34% and 0.25% more revenue over the maximum earned by traditional Class-Based and Class-Based Continuous ProBP respectively, while Classless UDP gained 0.64% and 1.00% over the maximums earned by traditional Class-Based and Class-Based Continuous UDP. Both classless methods were better able to optimize fares for continuous pricing, as both of them were designed to function in such a setting. Classless ProBP was the only method that increased revenue by a means other than by protecting for bookings in later TFs; instead, Classless ProBP took more bookings in early TFs without substantially impacting its later TF bookings. Classless UDP, however, did follow the previously mentioned pattern of protecting for bookings in later TFs. Both of the aforementioned booking trends held true in asymmetric tests of the classless methods. Classless ProBP typically undercut its traditional Class-Based ProBP competitor in early TF without losing too many bookings in later TFs, while Classless UDP typically protected for later booking periods. While Classless UDP's method resulted in a revenue advantage regardless of the number of classes used by its competitor, Classless ProBP could end up in a revenue disadvantage (making 2% less than an airline using traditional Class-Based ProBP with 21 fare classes) if its competitor had enough fare classes to take higher value bookings from Classless ProBP in later TFs.

The final comparisons were between corresponding ProBP and UDP methods with symmetric competition. For all cases, it was found that ProBP generated more revenue than UDP by about 0.54% for 16-class traditional Class-Based RM, 0.90% for 16-class Class-Based Continuous RM, and 0.20% for Classless RM. For the class-based cases, this was due in large part to the Poisson distribution assumption leading to lower forecasts for UDP, which caused the UDP methods to be less aggressive in early TFs. As a result, the ProBP methods protected more seats for later bookings and generated more revenue. For Classless ProBP and UDP, the forecast volumes were almost identical, yet Classless ProBP generated slightly more revenue. This was also linked to the Poisson assumption making Classless UDP slightly less aggressive about the fares it generated in early TFs, even with identical forecasts, which was enough for Classless ProBP to generate slightly more revenue.

Chapter 6: Conclusions

The introduction of New Distribution Capability (NDC) seems likely to change the way airlines will be able to sell their tickets. Among these changes is the removal of the requirement for an airline to sell its tickets within the confines of fare classes. There is hope that this will allow the implementation of continuous pricing, where fares offered by airlines are taken from a range of prices rather than from fixed price points. Theoretically, continuous pricing should allow airlines to increase their revenues by matching the willingness-to-pay of their passengers, particularly in the unrestricted fare structures that some airlines have been forced to adopt by the introduction of low-cost carriers.

Several methods of implementing continuous pricing have been proposed, for example, Gallego & Van Ryzin (1994 & 1997) and by Zhang & Lu (2013). Most of these proposals have made critical assumptions that would make them difficult to adopt for real world purposes. Most importantly, these assumptions did not take into account competition or uncertain estimates of demand. It was, therefore, useful to develop continuous pricing methods that can account for these real-world parameters.

6.1 Thesis Objectives and Models Overview

The purpose of this thesis was to introduce revenue management (RM) methods for continuous pricing designed to account for the above mentioned uncertainties and to test those methods in a simulator that included components such as competition, passenger choice, and forecasted demand. Four methods of continuous pricing, developed by Hopperstad within the MIT passenger origin-destination simulator (PODS) consortium, were tested in the PODS software to observe how each of the methods performed under simulated real-world conditions. Two of the methods developed solved for static solutions that require seat protection reoptimization throughout the booking process and two of the methods were dynamic programming solutions (which were re-optimized nonetheless in the experiments to account for forecasting inaccuracies).

Since all continuous pricing methods developed result in a single fare quote for any origin-destination itinerary at any point in time, a requirement for all continuous pricing methods developed was that they would be employed in unrestricted fare structures As a result, Q-Forecasting was used as the forecasting method for all methods in order to prevent spiral-down from occurring. Spiral-down is a known effect of unrestricted fare structures resulting from the non-independence of demand across fare classes. When spiral-down occurs, passengers do not buy the fare class representing their maximum willingness-to-pay and instead buy the lowest available fare class, which causes forecasts based on historical bookings to skew over time towards lower-value fare classes. These lower fare class forecasts in turn cause seat protection optimization methods to protect fewer high-value seats, which only causes further buy down (Cooper et al, 2006). Q-Forecasting counters spiral-down by transforming all historical bookings into equivalent lowest fare class bookings, which takes into account the aforementioned non-independence of demand (Belobaba & Hopperstad, 2004).

Two of the continuous pricing methods tested, one with "static" seat protection optimization and one with "dynamic" seat protection optimization, were based on existing class-based seat protection optimization methods. The "static" method utilized probabilistic bidprice (ProBP) optimization. First discussed by Bratu (1998), ProBP is a method that addresses issues related to network effects of RM by iteratively prorating path fares according to probabilistic bidprices of the legs the paths traverse until the bidprices converge. In doing so, it accounts for the contribution of a path fare to the network. While, for traditional Class-Based ProBP, the bidprices from the final iteration are used to determine the minimum fare class that will be offered at a given point in time, for Class-Based Continuous ProBP, these bidprices are used to directly quote a fare. The reasoning behind this is that, in an unrestricted fare structure, the bidprice (plus fare adjustment) is the fare where expected revenue decreases if said fare is either increased or decreased.

The "dynamic" method used the already existing unbucketed dynamic programming (UDP) optimization method available in PODS. UDP is a combination of a network linear program displacement algorithm (Williamson, 1992) and a dynamic program for RM introduced by Lautenbacher (1999). Unlike the "static" methods, the dynamic program is designed to determine an optimal seat protection solution for any remaining capacity at any time during the booking process. When used with fixed price points, the dynamic program, like ProBP, determines bidprices as cutoffs for minimum fare classes that will be open at a given point in time. When used for continuous pricing, these bidprices are also used to determine quoted fares.

Both Class-Based Continuous ProBP and UDP have the drawback of using seat protection optimization methods designed to work in a fare structure with fixed price points. As a result, two new classless seat protection optimization algorithms, one derived from ProBP and one derived from UDP, were also described and tested. The ProBP-derived method was designed by extending the idea of iteratively calculating revenue to protect seats for legs based on contribution by calculating optimal fares for each individual time frame, or period before departure, on the same leg accounting for the contribution of each time frame. Classless UDP, meanwhile, was designed by substituting the fare class-based dynamic program of Class-Based UDP for an algorithm that instead calculated fares based on what fare would maximize revenue at the current point in time while protecting for the future demand in the booking window.

Experimentation on all four methods was performed using the PODS software. The PODS software is a stochastic simulator that provides two key advantages over many other simulators for examining the practicality of RM concepts. Firstly, it implements several components of an RM system beyond seat protection optimization, most importantly forecasting, which allows examination of how RM methods will function in real-world conditions where demand distributions must be estimated from previous historical booking data. Secondly, the PODS software implements a passenger choice model, which allows passengers to choose between multiple options offered by one or more airlines, incorporating the effects of competition.

To examine the continuous pricing methods, each method was tested in a relatively large two-airline network with nearly symmetric route competition. By using such a network, it was possible to observe the network effects of each method tested while minimizing any competitive advantages, outside of RM method, one airline might have over the other. For both ProBP and UDP, a baseline was first established using the class-based version of each method with traditional, fixed fare quotation. To account for the impacts of increased fare granularity that the continuous pricing methods may have had over traditional RM methods, both traditional methods were tested with a varying number of fare classes. Experimenting with adding fare classes to the traditional class-based methods also served to determine if adding fare classes was an effective heuristic for continuous pricing. Each continuous pricing method was then tested with symmetric competition, where both airlines used the same RM method, and asymmetric competition, where the competitor airline used the corresponding (ProBP or UDP) traditional class-based method. As with the traditional class-based methods, both class-based continuous methods were tested with a varying number of fare classes for their seat protection optimization step (classless RM, not using fare classes, could not have its number of fare classes increased).

6.2 Results Summary

For traditional Class-Based ProBP, adding five fare classes to the baseline of six fare classes increased total revenue generated by about 0.65% with the incremental gains diminishing with further additions; the total revenue gain with 21 fare classes leveled off at about 0.97%. This was, as expected, due to the increased pricing granularity that the additional fare classes offered. With the extra granularity, the airlines were less likely to be constrained to offering fares much higher than the fare that the calculated bidprices would indicate was revenue maximizing.

In symmetric competition experiments, Class-Based Continuous ProBP, with six fare classes, generated a revenue increase of almost 1% over traditional Class-Based ProBP also using six fare classes in the baseline symmetric experiments. Additionally, the revenue generated by Class-Based Continuous ProBP with six fare classes was always greater than the revenue generated by traditional Class-Based ProBP in all symmetric competition experiments, regardless of the number of fare classes used by the latter. This was found to be the result of increased price granularity offered by continuous pricing, which allowed Class-Based Continuous ProBP to have greater accuracy in its quoted fares in later time frames (when the gaps between filed fares in the fare structures used in the network became larger), generating more bookings. However, experiments showed that increasing the number of fare classes used by the seat protection optimization component of Class-Based Continuous ProBP resulted in a decrease in revenue. This was found to most likely be a result of a difference in how Q-Forecasting scaling rules for the fixed pricing and continuous pricing methods worked, resulting in additional fare classes causing Class-Based Continuous ProBP forecasts to grow very rapidly. More importantly, however, the traditional and continuous Class-Based ProBP methods' revenues appeared to converge as fare classes were added, indicating that, with symmetric competition, the two methods could become essentially equivalent with 21 fare classes.

Asymmetric competition tests involving the use of each Class-Based Continuous ProBP and traditional Class-Based ProBP by only one of the two airlines revealed, however, that the apparent convergence of the two methods only applied in cases of symmetric competition. With asymmetric competition, Class-Based Continuous ProBP increased revenue for the implementing airline by 9.94%, while the competitor saw its revenue decline by 6.62% in terms of generated revenue when each airline used six fare classes. The Class-Based Continuous ProBP-using airline maintained an advantage for all numbers of fare classes tested, although the advantage did diminish as fare classes were added, having a 4% advantage only receiving a 2.28% revenue increase when each airline used 21 fare classes. While, in a case of symmetric competition, the effect of the difference between the granularities of the methods became negligible, when the two methods were in competition with each other, the increased granularity for Class-Based Continuous ProBP allowed the airline using it to undercut the airline using traditional Class-Based ProBP in terms of fare quoted in the later time frames, where willingness-to-pay of passengers is higher.

While it was shown to be possible for the Traditional Class-Based ProBP airline to respond by being less aggressive with its forecasting, in turn undercutting the Class-Based Continuous ProBP airline, this often came at the expense of much revenue for both airlines. The results of the asymmetric experiments seem to indicate that, in an unrestricted fare structure, Class-Based Continuous ProBP offers revenue advantages that cannot be achieved using traditional Class-Based ProBP

Classless ProBP, in symmetric competitive experiments, generated more revenue than either Class-Based ProBP method (0.34% over the maximum revenue for traditional Class-Based ProBP and 0.25% over the maximum revenue for Class-Based Continuous ProBP). Classless ProBP has the granularity advantage exhibited by Class-Based Continuous ProBP, and it has the additional advantage over Class-Based Continuous ProBP in that its seat protection optimization method optimizes for a continuous range of fares rather than fare classes not used for fare quotation. As a result, the revenue generated by Classless ProBP was increased by offering lower fares that would generate increases in early time frame bookings without seriously impacting later time frame bookings.

Asymmetric competition experiments, however, showed that, with a sufficient number of fare classes, an airline using traditional Class-Based ProBP could generate more revenue than a Classless ProBP competitor. While using Classless ProBP resulted in revenue advantages in the symmetric competition experiments by adding early time frame bookings without displacing later time frame bookings, this was not the case with asymmetric competition. Instead, as traditional

Class-Based ProBP had fare classes added, it became more aggressive about protecting for later time frame passengers, taking the later time frame bookings that Classless ProBP had previously still received in the symmetric experiments. Thus, while Classless ProBP may offer theoretical revenue advantages, it would require caution when implementing.

Adding fare classes to traditional Class-Based UDP was found to have the same effect as with traditional Class-Based ProBP. Increased granularity of fares quoted led to better protection of seats, particularly in later time frames, increasing revenue. As with traditional Class-Based ProBP, traditional Class-Based UDP also saw diminishing revenue returns with each additional fare class, as each extra fare class added a diminishing increment of price granularity.

In experiments with symmetric competition, Class-Based Continuous UDP with six fare classes saw a 0.28% revenue increase over traditional Class-Based UDP with a low number of fare classes. As with ProBP, Class-Based Continuous UDP's increased granularity caused better protection of seats in later time frames. Use of Class-Based Continuous UDP with additional fare classes, however, had different results than those of ProBP. Whereas revenue decreased for ProBP, it increased for UDP. Although Class-Based Continuous ProBP and Class-Based Continuous UDP both used the same forecast scaling rule, adding fare classes to UDP, which generally has lower forecasts than ProBP in unrestricted fare structures, did not see uncontrolled forecast growth.

Although the revenue for Class-Based Continuous UDP did increase with additional fare classes, traditional Class-Based UDP saw greater revenue increase and, with a sufficient number of fare classes (using 11 classes total in these experiments), generated more revenue than Class-Based Continuous UDP. With a sufficient number of fare classes beyond that (using 21 classes total in these experiments), however, Class-Based Continuous UDP would gain more incremental revenue, suggesting that, with symmetric competition, the two methods would converge.

Class-Based Continuous UDP in an asymmetric competitive environment was mostly the same as Class-Based Continuous ProBP. Increased fare granularity for Class-Based Continuous UDP, which was not substantially different from traditional Class-Based UDP with more than 11 fare classes in symmetric competition, gave the Class-Based Continuous UDP-using airline the ability to undercut the traditional Class-Based UDP-using airline, particularly in later time frames where it was able to allow more passengers to book. As with ProBP, the traditional Class-Based UDP-using airline could respond by being less aggressive with its forecast scaling, causing it to offer lower fares and undercut the Class-Based Continuous UDP using airline. Unlike with ProBP, this always resulted in a revenue increase for traditional Class-Based UDP but was still contingent on the Class-Based Continuous UDP-using airline not responding in kind. As with ProBP, Class-Based Continuous UDP showed signs of having an advantage over its traditional class-based counterpart in unrestricted fare structures.

As with Classless ProBP, Classless UDP, in symmetric tests, generated more revenue than either of its class-based counterparts. As with ProBP, Classless UDP was able to take advantage of the fact that its seat protection optimization method was specifically designed to generate fares for quotation with no classes. The mechanisms of Classless UDP's revenue increases, however, were different. Unlike Classless ProBP, Classless UDP was more aggressive than its Class-Based counterparts about protecting seats for later time frames when passengers arriving have a higher willingness-to-pay, an effect similar to the one that occurred when adding fare classes for traditional class-based or using class-based continuous methods.

The difference in how Classless UDP increased its revenue in symmetric experiments caused it to still see revenue increases in asymmetric competition-focused experiments. Unlike Classless ProBP, which lost revenue in asymmetric competition if the competitor airline used a great enough number of fare classes as a result of no longer being able to protect seats for later time frames, Classless UDP was already more aggressive about protecting seats for later time frames, and thus maintained its revenue advantage even as its traditional class-based competitor got more aggressive in its protections.

In comparing ProBP and UDP with symmetric competition, it was found that for traditional class-based, class-based continuous, and classless RM, the ProBP-based method generated more revenue. With 16 fare classes, revenue generated by traditional Class-Based UDP was found to be 0.54% lower than traditional Class-Based ProBP, and Class-Based Continuous UDP revenue was found to be 0.90% lower than Class-Based Continuous ProBP. For both of the class-based methods, it was found that the fact that UDP generates lower forecasts than ProBP causes it to be less aggressive about protecting seats for later time frames, causing a decrease in revenue. For the classless methods, it was found that, although ProBP and UDP had nearly the same forecasts, Classless UDP likely generated 0.20% less revenue than Classless ProBP because of UDP's incorrect assumption of a smaller variance of demand, a result of UDP assuming that demand follows a Poisson distribution. The smaller demand variance would cause Classless UDP to quote fares slightly lower than what Classless ProBP would quote. As ProBP optimized based on the variance of demand as calculated by the forecaster, it would show greater accuracy, which resulted in its slight revenue advantage.

6.3 Suggested Future Research Directions

From what was determined in this thesis, there are four suggested potential directions for further research. Two of these recommended directions are relatively minor, with the first of these recommended directions being applying the forecasting scaling limit for the continuous pricing methods. This should prevent the rapid forecast growth seen with Class-Based Continuous ProBP, and better help prove whether or not convergence between the class-based methods occurs, and would allow slightly more even comparisons between the fixed pricing and continuous pricing methods. The other minor direction for future research would be to test Classless RM for Fixed Price Points (the top right box in Figure 1-1). While methods fitting into this category would not be theoretically optimal, they could potentially allow an implementation of Classless RM not dependent on the use of a distribution system configured for continuous pricing.

The first of the two more substantial directions for future research would be to test the methods of continuous pricing in a larger network less focused on symmetry and with more airlines. Such a network would be better able to model the real world conditions in which a continuous pricing method would be expected to operate. Nearly all the tests in this thesis could be repeated in some way in such a network, but with competitor airlines possibly having different responses to the continuous pricing-using airline's decision to switch to said continuous pricing method.

The final possible direction of further research is by far the most ambitious, and, unlike the three other three potential directions for research, would involve creation of new algorithms. Despite the potential the continuous pricing methods showed in this thesis for increasing revenue generation in an unrestricted fare structure, they are not implementable in differentiated fare structures. As restrictions remain a primary method of separating demand by willingness-to-pay within a single window of time, a next step would be the development of methods for multi-fare continuous pricing. Any such method, whether class-based or classless, would be far more complex than the methods developed here, as multi-fare continuous pricing would require not only optimizing each fare with respect to forecast willingness-to-pay but also with respect to each other. One idea for developing such a method would be to adapt concepts currently used in fare family RM, as fare families already offer multiple products with each having multiple price points.

This thesis has described and tested multiple methods of continuous pricing designed to function in real world RM environments. With some in the airline industry hoping that the upcoming New Distribution Capability will allow the complete elimination of fare classes, the findings described in this thesis provide guidance with respect to RM methods that could make such a concept a possibility.

Bibliography

Belobaba, P. P. (1987). *Air travel demand and airline seat inventory management* (Unpublished doctoral dissertation). Massachusetts Institute of Technology. Retrieved from http://hdl.handle.net/1721.1/14800

Belobaba, P. P. (1992). Optimal vs. Heuristic Methods for Nested Seat Allocation. In *AGIFORS Yield Management Study Group Meeting*. Brussels, Belgium.

Belobaba, P. P., & Diwan, S. (2010, May). Performance of Dynamic Programming in Airline RM Systems. In *AGIFORS Revenue Management Study Group Meeting*. New York.

Belobaba, P. P., & Hopperstad, C. (2004, March). Alternative RM algorithms for unrestricted fare structures. In *AGIFORS Reservation and Yield Management Meeting*. Auckland, New Zealand.

Belobaba, P. P., & Liotta, N. J. (2017, November) *A New Approach to Classless RM*. Presented at the meeting of PODS Consortium, Cambridge, Massachusetts.

Belobaba, P. P., Odoni, A. R., & Barnhart, C. (2016). *The global airline industry* (2nd ed.). Chichester: John Wiley & Sons.

Bratu, S. (1998). *Network value concept in airline revenue management* (Unpublished master's thesis). Massachusetts Institute of Technology. Retrieved from http://hdl.handle.net/1721.1/9939

Brumelle, S. L., & McGill, J. I. (1993). Airline Seat Allocation with Multiple Nested Fare Classes. *Operations Research, 41*(1), 127-137. doi:10.1287/opre.41.1.127

Cléaz-Savoyen, R. L. (2005). *Airline revenue management methods for less restricted fare structures* (Unpublished master's thesis). Massachusetts Institute of Technology. Retrieved from http://hdl.handle.net/1721.1/32439

Cooper, W. L., Homem-De-Mello, T., & Kleywegt, A. J. (2006). Models of the Spiral-Down Effect in Revenue Management. *Operations Research, 54*(5), 968-987. doi:10.1287/opre.1060.0304

Fiig, T., Isler, K., Hopperstad, C., & Belobaba, P. P. (2010). Optimization of mixed fare structures: Theory and applications. *Journal of Revenue and Pricing Management, 9*(1-2), 152-170. doi:10.1057/rpm.2009.18

Gallego, G., & Van Ryzin, G. (1994). Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons. *Management Science, 40*(8), 999-1020. doi:10.1287/mnsc.40.8.999

Gallego, G., & Van Ryzin, G. (1997). A Multiproduct Dynamic Pricing Problem and Its Applications to Network Yield Management. *Operations Research, 45*(1), 24-41. doi:10.1287/opre.45.1.24

Gorin, T. O. (2000). *Airline revenue management: Sell-up and forecasting algorithms* (Unpublished master's thesis). Massachusetts Institute of Technology. Retrieved from http://hdl.handle.net/1721.1/9245

Lautenbacher, C. J., & Stidham, S., Jr. (1999). The Underlying Markov Decision Process in the Single-Leg Airline Yield-Management Problem. *Transportation Science, 33*(2), 136-146. doi:10.1287/trsc.33.2.136

L'Heureux, E. (1986). A New Twist in Forecasting Short-term Passenger Pickup. In *26th AGIFORS Annual Symposium Proceedings*, 248-261.

Littlewood, K. (1972). Forecasting and Control of Passenger Bookings. *12th AGIFORS Symposium, 4*(2), 95-117.

Tam, C. W. (2008). *Airline revenue management based on dynamic programming incorporating passenger sell-up behavior* (Unpublished master's thesis). Massachusetts Institute of Technology. Retrieved from http://hdl.handle.net/1721.1/44680

Weatherford, L. R., & Pölt, S. (2002). Better unconstraining of airline demand data in revenue management systems for improved forecast accuracy and greater revenues. *Journal of Revenue and Pricing Management, 1*(3), 234-254. doi:10.1057/palgrave.rpm.5170027

Westermann, D. (2013). The potential impact of IATA's New Distribution Capability (NDC) on revenue management and pricing. *Journal of Revenue and Pricing Management, 12*(6), 565-568. doi:10.1057/rpm.2013.23

Williamson, E. L. (1992). *Airline network seat inventory control: Methodologies and revenue impacts* (Unpublished doctoral dissertation). Massachusetts Institute of Technology. Retrieved from http://hdl.handle.net/1721.1/13174

Wittman, M. D. (2018). *Dynamic pricing mechanisms for airline revenue management: theory, heuristics, and implications* (Unpublished doctoral dissertation). Massachusetts Institute of Technology.

Wittman, M. D., & Belobaba, P. P. (2018). Dynamic pricing mechanisms for the airline industry: A definitional framework. *Journal of Revenue and Pricing Management, 18*(2), 100-106. doi:10.1057/s41272-018-00162-6

Zhang, D., & Lu, Z. (2013). Assessing the Value of Dynamic Pricing in Network Revenue Management. *INFORMS Journal on Computing, 25*(1), 102-115. doi:10.1287/ijoc.1110.0488