

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Collaborative Resource Allocation Strategies for Air Traffic Flow Management

Permalink

<https://escholarship.org/uc/item/4f07r35r>

Author

Kim, Amy Miyoung

Publication Date

2011

Peer reviewed|Thesis/dissertation

Collaborative Resource Allocation Strategies
for Air Traffic Flow Management

By

Amy Miyoung Kim

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Civil and Environmental Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Mark Hansen, Chair

Professor Samer Madanat

Professor Rhonda Righter

Spring 2011

Abstract

Collaborative Resource Allocation Strategies for Air Traffic Flow Management

by

Amy Miyoung Kim

Doctor of Philosophy in Engineering – Civil and Environmental Engineering

University of California, Berkeley

Professor Mark Hansen, Chair

The Airspace Flow Program (AFP) is a traffic management initiative that aims to mitigate delays arising from en route capacity constraints, by holding flights on the ground and metering their flow through constrained regions. It has been successful in controlling traffic with reasonable delays, but the procedures must be improved upon to handle future projected demands. This dissertation explores a future AFP concept that incorporates structured user preference inputs in a centrally-managed assignment of ground delays as well as reroutes. However, the commercial aviation industry is very competitive, and airlines are reticent to share detailed operational information without clear benefits in return. This idea must be taken into consideration when designing flight resource assignment strategies, as we address a very fundamental question in transportation service provision: how does a central authority allocate resources efficiently when they are uncertain about what users want?

A modeling framework was developed to evaluate and compare allocation strategies, under differing assumptions regarding the information that traffic managers may (or may not) have about airline flight costs. We introduce several resource allocation strategies that feature different allocation rules and route preference inputs requested of flight operators. We assess the total user-cost impact of each allocation strategy through a simple generalized cost function that represents the cost of delay for each flight caught in the AFP. This flight cost function consists of two parts: the first is a simple representation of operator flight cost characteristics that the traffic managers have adopted, while the second consists of a flight's routing preferences that are not captured by the first, and therefore are privately known to each flight's operator but not to traffic managers unless the information is offered. In one allocation strategy, users are asked to provide their private information in exchange for the employment of a resource allocation rule that rewards flights for early submission. In another, users are not asked to provide their proprietary information but, in turn, are not offered the same "guarantee" regarding resource allocations, i.e. flights have less idea about what resources they will receive in the allocation. We identify some basic properties about the relationships of the

assignment schemes' total flight cost efficiency results, under changing assumptions about the quality of the central decision maker's knowledge about the flight operators' route preferences. Numerical examples illustrate situations where sacrificing a system-optimal allocation rule for a sequential one (First Submitted, First Assigned, or FSFA), in order to obtain and utilize flight operators' private information about route preferences, will result in more user cost efficient resource allocations. The examples also illustrate situations where the opposite is true, i.e. it is more efficient to use a system-optimal allocation without private route preference information. Also, it is found that the existing Ration-by-Schedule algorithm performs poorly in the context of our modeling framework.

The above results do not consider the effects of gaming behavior. As a result, we also study the gaming and truth-telling behaviors of flight operators in response to the competition for limited en route resources. It is demonstrated that operators are incentivized to provide untruthfully high inputs in schemes where a system-optimal allocation rule is used. Using a gaming analysis, it is also demonstrated that the operators' equilibrium submission strategies in the FSFA scheme can vary significantly depending on the conditions of the AFP and their private information. However, it is also shown that flights are highly incentivized to submit at the very beginning of the FSFA planning period, which is favorable for air traffic flow management (ATFM) planning and coordination.

Dedicated to my mom and dad.

Table of Contents

Table of Contents	ii
List of Tables	v
List of Figures.....	vi
List of Acronyms	vii
Acknowledgements	viii
1. Introduction	1
1.1 Review of Current ATFM Practices	2
1.1.1 Flight Rerouting	2
1.1.2 Ground Delay Program (GDP).....	3
1.1.3 Airspace Flow Program (AFP).....	5
1.1.4 Airline Flight Planning and Operating Costs	5
1.2 En Route Resource Allocation.....	6
1.3 Literature Review.....	10
1.3.1 Traffic Assignment.....	10
1.3.2 Air Traffic Flow Management Models	11
1.3.3 Equity and Collaboration.....	12
1.3.4 Research Contributions	13
2. Flight Cost Model	15
2.1 Evaluation Framework.....	15
2.2 User Flight Cost Model.....	16
2.3 Heterogeneous Flights	18
2.4 Unknown Cost Component.....	18
2.5 Concluding Remarks.....	19

3. Resource Allocation Schemes	20
3.1 Stated Route Preference	20
3.1.1 Full Information Optimal (OPT)	22
3.1.2 First Submitted, First Assigned (FSFA).....	23
3.2 Parametric (PO)	24
3.3 Properties	26
3.4 Numerical Examples	29
3.5 Model Approximations	35
3.5.1 First Submitted, First Assigned	35
3.5.2 Parametric.....	37
3.6 Concluding Remarks and Future Work	40
4. Airline Behavior under Competition	42
4.1 Introduction.....	42
4.2 Truth-telling Properties of System-Optimal Resource Allocation.....	43
4.3 Competition in FSFA	45
4.3.1 Utility by Rank in Submission Order	47
4.3.2 Utility by Submission Time	49
4.3.3 Competitive Model.....	52
4.3.3.1 Assumptions and Payoff Function	52
4.3.3.2 Equilibrium Submission Strategy	56
4.3.3.3 Sensitivity tests	58
4.3.4 Model Extensions	64
4.4 Other Gaming Behaviors and Future Work	64
5. Other Resource Allocation Schemes	66
5.1 Introduction.....	66
5.2 Hybrid Stated Route Preference Model	66
5.3 Ration-by-Schedule (RBS)	69
5.4 Future Investigations.....	71
6. Conclusions and Future Research	74

Bibliography	77
Appendix A: <i>Stochastic properties of the allocation schemes</i>	81
A.1) Expected value of the minimum of two iid normal random variables	81
A.2) Some relationship properties	82
Appendix B: <i>Properties of the Parametric Approximation</i>.....	83
B.1) Aircraft distribution based on cost parameters.....	83
B.2) Convexity	84
B.2.1) Homogeneous flights	84
B.2.2) Heterogeneous flights	85
Appendix C: <i>Chapter 4 calculations</i>.....	86
C.1) Untruthful α submissions, 3-route system-optimal case.....	86
C.2) Equilibrium FSFA airline preference submission strategy	87
C.3) Submission strategy approximation	89
Appendix D: <i>Ration-by-Schedule</i>	90

List of Tables

Table 3.1	Toy Model Allocation	26
Table 3.2	Example supply parameters.....	30
Table 3.3	Demand Scenario 1, Parametric vs. FSFA results.....	32
Table 3.4	Demand Scenario 2, Parametric vs. FSFA results.....	33
Table 3.5	Demand Scenario 3, Parametric vs. FSFA results.....	34

List of Figures

Figure 1.1	Causes of NAS delay (% of total operations), 2008.....	2
Figure 1.2	Use of GDPs for en route capacity constraints.....	4
Figure 1.3	Current resource allocation process	7
Figure 1.4	Illustration of possible allocation outcomes	8
Figure 1.5	Proposed modeling and analysis framework.....	9
Figure 2.1	Model airspace geometry and select parameters	15
Figure 3.1	Delay thresholds	21
Figure 3.2	Total cost efficiency relationship of allocation schemes.....	28
Figure 3.3	Demand Scenario 1 total cost results	31
Figure 3.4	Assignment of flights to routes by increasing α values.	38
Figure 4.1	Untruthful versus truthful system-optimal solutions.....	44
Figure 4.2	Sample average utilities by submission rank	47
Figure 4.3	Average utility by submission rank, range of σ (Demand Scenario 1).....	48
Figure 4.4	Utility loss through imperfect information.....	51
Figure 4.5	Submission cost function $C(q)$	55
Figure 4.6	Example equilibrium submission strategies	57
Figure 4.7	Equilibrium submission strategies, t versus p	59
Figure 4.8	Equilibrium submission strategies (t , hours), p versus h	60
Figure 4.9	Equilibrium submission strategies (t , hours), $h = 1.0$	61
Figure 4.10	Equilibrium submission strategies (t , hours), $p = 1.0$	62
Figure 4.11	h_0 for decision-making and predictability	63
Figure 5.1	Example of Hybrid allocation scheme performance.	68
Figure 5.2	Demand Scenario 1 total cost results including RBS.....	70

List of Acronyms

AFP	Airspace Flow Program
ATFM	Air Traffic Flow Management
CDM	Collaborative Decision Making
DTA	Dynamic Traffic Assignment
EDCT	Expected Departure Control Time
FAA	Federal Aviation Administration
FCA	Flow Constrained Area
FCFS	First Come, First Serve
FSFA	First Submitted, First Assigned
FSM	Flight Schedule Monitor
GDP	Ground Delay Program
NAS	National Airspace System
OAG	Official Airline Guide
RBD	Ration-by-Distance
RBS	Ration-by-Schedule

Acknowledgements

Deciding to return to Cal to do the PhD was the easy part; completing it was much more difficult. There are many people that have made this journey a deeply rewarding and humbling experience. When I think about all that has aligned to make it happen, I feel very blessed.

Thank you to all the professors that I have taken classes from and that have served on my various exam committees. Professors Madanat and Righter: you are both excellent teachers and I only wish I could have done better in your classes, because the material was great. I would also like to thank the students of NEXTOR, and several others in the TE program. In addition to other things, for hours they sat through my various presentation prep sessions, offering healthy skepticism as well as incredibly thoughtful suggestions, which not only reflect intellectual and stylistic aptitude but a sincere desire to help a fellow student. I can only wish this kind of peer environment for others. Thank you to my dear friend Jing for her daily support. Our cultural differences are hilarious.

Although I have been in Mark's hair for almost five years now, his brilliance, insight, creativity, and thoughtfulness do not fail to impress me still. I cannot aptly express my gratitude for his support in all areas of my academic life, and for handling my scatterbrain, lack of confidence, and intermittent tears with grace. He has given me countless hours of critique and support, and allowed me to tease him with good humor. Even about his disco-era skiing style, which admittedly has been updated with recent equipment purchases. The truth is, I can barely keep up with him. My mother once mentioned that one of the best things that happened to my father as a grad student was to have Dr. Gourishankar as his advisor. I think I can say the same about Prof. Hansen. Thank you for being an exceptional mentor, boss, ski buddy, friend. Thank you for being on my side.

To my patient and understanding husband Stephen: there is no way to adequately express my gratitude for your daily sacrifices and acts of love and kindness. Thank you MSN instant messenger for keeping us in touch throughout our twenties. If your name could be on the cover of this document I would be honored to share it with you.

Thank you to my brother for his honesty, to my grandparents for their tough Korean love. To my mom: as a single immigrant parent starting over with nothing, you somehow managed to provide us every opportunity to succeed in this world and explore the limits of our capabilities. You remind me through example that age is nothing but a number, and that one should never stop learning, growing, being independent, and doing things that are scary. I am SO proud of you. I hope this document is of a quality that reflects even just a tiny smidgen of your excellence.

Thank you to the UCTC, PATH, FAA, Eisenhower Fellowship, and of course NEXTOR for financial support through this degree.

1. Introduction

Adverse weather frequently and severely impacts flight operations in the National Airspace System (NAS). In addition, with the growth in demand projected for the NAS over the next 20 years, weather and traffic-induced delays are anticipated to increase under the current system (Bureau of Transportation Statistics, 2007). Air traffic flow management (ATFM) programs are used to reduce the scale and cost of disruptions to flight operators. One such initiative is the Airspace Flow Program (AFP), which facilitates resource allocation decisions when en route capacity/demand imbalances exist. In the AFP, flights are held on the ground at departure airports in order to control their flow through capacity constrained airspace regions. The AFP was first implemented in 2006 in the northeastern airspace of the United States, and has proven to be successful in increasing efficiency and reducing flight delays. However, projected increases in traffic demand will eventually limit the benefits derived from the currently employed AFP. A procedure to more effectively utilize available airspace capacity must be incorporated into the AFP to handle these future demand levels.

This research proposes a more comprehensive, centralized, and collaborative resource allocation concept within a future AFP. In particular, it proposes a highly structured user preference input component. It builds and compares alternative resource allocation strategies that employ rerouting combined with ground delay to minimize the impacts of AFP initiatives on users of the NAS. These strategies differ with respect to the mechanisms and philosophies of resource allocation, and the inputs requested of users. We develop a modeling framework through which these strategies are represented, evaluated and compared. The objective of this work is to determine how the strategies perform in comparison to one another under different assumptions about airline utility as well as potential gaming behaviors.

This chapter contains a review of current en route ATFM practices, a brief description of flight planning, an introduction to the en route resource allocation problem studied in this research, and a literature review. Chapter 2 introduces the modeling framework and flight cost model. Chapter 3 describes the three basic resource allocation schemes and some properties of these schemes, and illustrates their performance through several numerical examples. Chapter 4 explores the gaming and truth-telling behaviors of the airlines when subject to each allocation scheme. Chapter 5 discusses several other resource allocation schemes. Chapter 6 contains conclusions and a summary of on-going and future work.

Throughout this dissertation, “operator” or “user” will be used to refer to NAS users such as commercial airlines and general aviation aircraft. “Traffic manager” will refer to the agent responsible for allocating resources. In the U.S. context these would be

traffic management specialists at the Federal Aviation Administration's (FAA's) Air Traffic Control System Command Center (ATCSCC).

1.1 Review of Current ATFM Practices

En route air traffic congestion occurs daily in the National Airspace System (NAS) due to capacity/demand imbalances caused by various phenomena. Of these phenomena, severe weather and excessive demand are the dominant contributors to traffic congestion, and combined were the source of over 85 percent of all NAS delays in 2008 (see Figure 1.1). Air traffic flow management (ATFM) programs are employed to reduce the scale and cost of disruptions to flight operators. One such initiative is the Airspace Flow Program (AFP), which facilitates resource allocation decisions when en route capacity/demand imbalances exist. In the AFP, flights are held on the ground at departure airports in order to limit traffic flow through capacity constrained airspace regions. The AFP was first implemented in 2006 in the northeastern region of the United States, and has proven to be successful in reducing en route flight delay (Federal Aviation Administration, 2007).

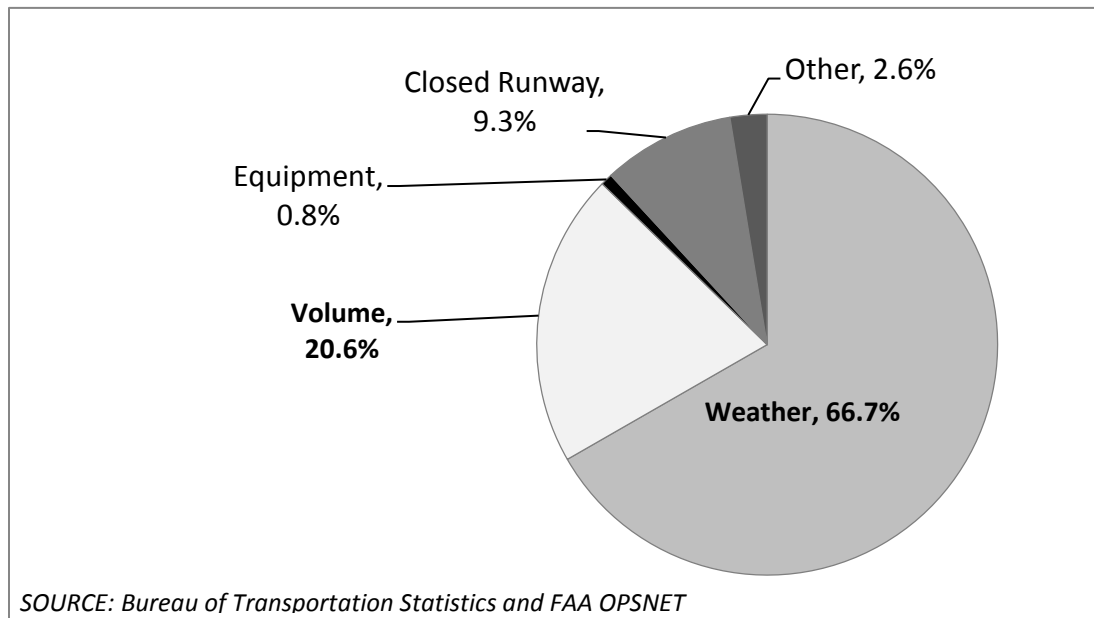


Figure 1.1 Causes of NAS delay (% of total operations), 2008

This section will provide an overview of past and current strategies for en route air traffic flow management and the AFP.

1.1.1 Flight Rerouting

Flight rerouting due to severe en route weather and traffic congestion is performed in both strategic and tactical air traffic flow management today. It is manually intensive as it

requires close coordination between several traffic management units. Consequently, FAA traffic managers typically select reroutes from a standard set compiled in the National Playbook¹, basically employing a “one size fits all” approach (Wilmouth & Taber, 2005) without input from the operators. Airlines also have the option of rerouting their own flights before and after departure, subject to traffic managers' approvals. They will often exercise this option to avoid being assigned routes they consider undesirable or long ground delays due to postponed departure times. However, if rerouting decisions were completely in the hands of the traffic managers, where airlines could not propose their own reroutes or voice any type of preference, these ATFM initiatives would likely be more inefficient and highly suboptimal. Concepts that propose more collaboration in rerouting have existed since the early 2000s (Ball, Futer, Hoffman, & Sherry, 2002); these concepts describe a more structured approach to coordination between traffic managers and operators.

1.1.2 Ground Delay Program (GDP)

A significant improvement to NAS air traffic management began in the mid-1990s with the Collaborative Decision Making (CDM) program. CDM is a program that aims to improve the technological and procedural aspects of air traffic management, by improving information exchange between government and industry. The first major application of CDM was to Ground Delay Programs (GDPs) (Chang, Howard, Oiesen, Shisler, Tanino, & Wambsganss, 2001). When an airport has reduced arrival capacity due to severe weather at or near an airport, a GDP holds flights destined for that airport on the ground at their origin airports to meter demand. CDM information exchange between operators and traffic managers drastically enhanced the effectiveness of GDPs in correcting demand/capacity imbalances and reducing delays (Hoffman, Hall, Ball, Odoni, & Wambsganss, 1999). CDM is employed in GDPs through the Ration-By-Schedule (RBS) algorithm. RBS allocates constrained airport resources by assigning flights delayed departure times in the order by which they are scheduled to arrive at the constrained airport. RBS also stipulates that airlines have ownership of their arrival slots (or, scheduled arrival times) at the constrained airport, and relinquish ownership only by choosing to vacate them. As airlines are free to swap and cancel flights² without losing slot ownership, they are incentivized to share their up-to-date flight schedule information with FAA traffic managers. RBS is a well-accepted allocation scheme that has been in use in ground delay programs since the mid-1990s (Hoffman, Burke, Lewis, Futer, & Ball, 2005). Prior to CDM and RBS, resource allocations were made using the Grover Jack algorithm. Airlines did not “own” their slots, and would lose ownership by canceling or swapping flights. They had no incentive to report updated flight schedule information, and as a result slots would often go unused. Clearly, CDM and RBS resulted in great efficiency improvements in ground delay programs (Hoffman, Burke, Lewis, Futer, & Ball, 2005).

¹ <http://www.fly.faa.gov/PLAYBOOK/pbindex.html>

² These actions are subject to a set of rules that prohibit airlines from assigning certain flights to certain routes, even if they desire to.

GDPs are very effective in dealing with reduced arrival capacity when it is caused by inclement weather near the destination airport. However, before 2006, GDPs were also employed to address capacity issues in the en route airspace. This could be inefficient, ineffective, and inequitable. We illustrate how, using the example in Figure 1.2 below.

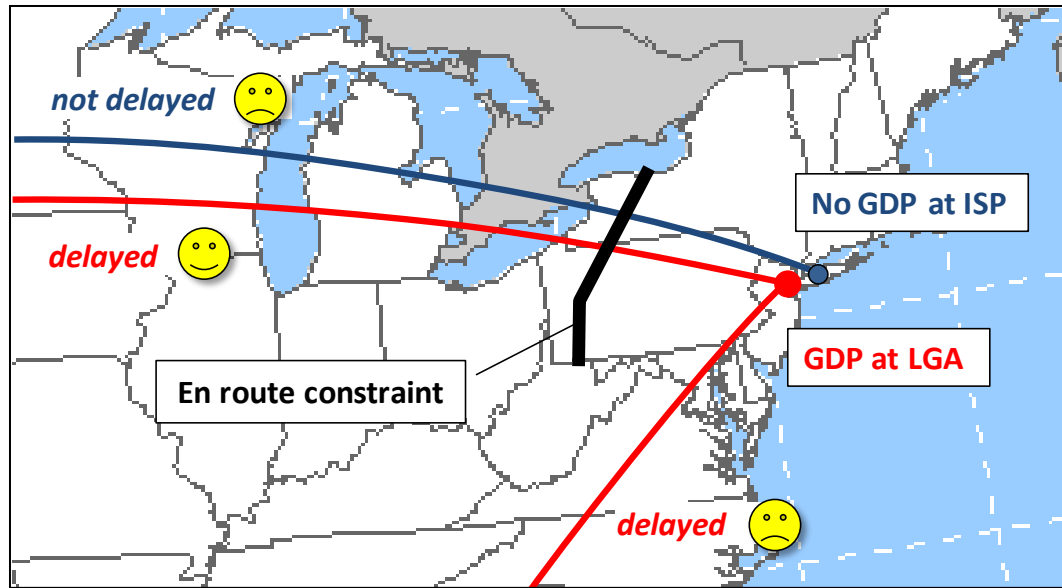


Figure 1.2 Use of GDPs for en route capacity constraints

The above figure identifies a flow constrained region in the northeastern United States en route airspace, depicted by the thick black line. It could represent a line of thunderstorms or a cordon through a corridor of heavy traffic. In fact, the constrained area shown above is a relatively common occurrence during thunderstorm season, as it is a primary corridor between the major cities on the northeastern seaboard and those further west (Chicago, Minneapolis, San Francisco, etc.). Prior to 2006, it was common to manage traffic flow through the constrained region by implementing GDPs at LaGuardia (LGA), Newark (EWR) and Kennedy (JFK) airports. However, GDPs were not commonly implemented at smaller airports such as Isip (ISP) in Long Island, New York. Under this practice, all flights destined for the New York airports with GDPs in place were assigned delayed departure times, regardless of whether or not they fly through the constrained region. However, flights destined for Isip are not delayed, although they might fly through the constrained region. It is clear that using GDPs to manage en route capacity was highly inequitable and inefficient. As a result, another CDM traffic management initiative, the Airspace Flow Program (AFP), was created and first implemented in 2006 to handle en route constraints (Federal Aviation Administration, 2008).

1.1.3 Airspace Flow Program (AFP)

The FAA will initiate an AFP to facilitate traffic assignment decisions when en route demand/capacity imbalances exist. In an AFP, the constrained airspace region and the flights scheduled to fly into this region during the period of reduced capacity are first identified. Again, constrained airspace regions include those that are experiencing undesirable weather and/or heavy demands. Most of these regions are able to accommodate a reduced amount of flight traffic; unless thunderstorms are very severe, flights are usually still permitted to fly through these regions. The reduced capacity of the constrained region is reallocated to impacted flights by assigning each flight a delayed departure time to fly their originally scheduled route. Delays are imposed only on flights that travel through the constrained airspace in question. Under the CDM rules of slot ownership, users are incentivized to provide up-to-date information about their flights to FAA traffic managers. Users will either accept these new departure times, or reject and reroute around the constrained airspace (subject to traffic managers' approval), or cancel flights altogether (Libby, Buckner, & Brennan, 2005). Slots to fly through the AFP are vacated as flights are canceled and routed out, and the schedule is compressed such that remaining flights are moved up in time. Currently, the distribution of delayed departure times combined with airline-initiated rerouting and cancellation has proven to be a sufficient improvement on GDPs in addressing en route capacity constraints (Federal Aviation Administration, 2007). However, with increasing demands and/or very severe storms, better utilization of all other available capacity options will be necessary. This is particularly true in the summer months when convective weather is most prevalent. One strategy is for FAA traffic managers to take control of allocating all aspects of en route resources, such that delayed departure times are combined with new route assignments. This would allow the FAA to manage rerouting centrally, and optimize utilization of available capacity on neighboring routes, to ultimately reduce the total delay cost of the AFP. In order to offer resource assignments that are desirable to operators, however, the FAA will require a significant level of user input. This is the main topic of this research, and is discussed in greater detail in Section 1.2.

According to 2008 data obtained from Metron Aviation, most AFPs begin after 2PM local time as airspace congestion and convective weather is more likely to occur later in the day. They typically end after 10PM when air traffic demands are low. AFPs can end earlier than their scheduled termination time, and do so about 70 percent of the time. AFPs are very rarely extended, because they are usually scheduled to end in the evening well after the time the capacity constraint is most severe.

1.1.4 Airline Flight Planning and Operating Costs

An airline's flight planning procedures typically begin about three hours before its scheduled departure (Hoffman, Lewis, & Jakobovits, 2004). Major airlines typically have sets of preferred routes that are flown under particular conditions. For a given flight, airline dispatchers (who oversee flights from planning to completion) load these routes into a program that will select the most desirable one according to their airline's

objectives for the flight. If severe weather is predicted on the selected route, the dispatcher will search for alternatives around the storm during this time. Approximately 90 to 120 minutes before departure, the airline dispatcher will file the flight plan trajectory with the FAA traffic managers. For any strategic planning activities that must occur prior to the time that flight plans are filed (for instance, an AFP) traffic managers rely on published airline schedules, historical data, or early intent flight plans from airlines when available.

An airline's operating objectives and flight cost calculations are proprietary and vary from operator to operator. As a result their preferred routes will vary as well. Traffic managers cannot assume that all flights will want to travel one route, whether it is wind-optimal or otherwise. For instance, although Northwest Airlines had a well-known preference for routes with less turbulence, timely arrival is of greater importance to Southwest Airlines. Despite these differences, however, all flight costs are clearly dependent on such things as crew time, fuel costs, maintenance costs, facility usage fees, etc. The factors contributing to flight costs differ according to whether the flight is on the ground or en route, mainly due to the fact that an aircraft's fuel costs are very high while it is in the air. En route fuel usage also varies depending on the type of aircraft flown, its speed, flight level (i.e. elevation), and other operating characteristics. Additionally, missed connections can greatly increase the cost of delay incurred either on the ground or en route. Missed connections disrupt personnel and equipment usage schedules, while missed passenger connections cause customer dissatisfaction and add rebooking workload.

1.2 En Route Resource Allocation

En route resource allocation decisions are the result of system capacity constraints, the allocation and equity principles chosen for use, and, under the CDM philosophy, up-to-date flight demand information (Figure 1.3). Flight demand information includes original flight schedules and updates about cancellations and other changes. Demand information is naturally a function of users' underlying flight cost structure. Also, the perceived quality of an allocation result will depend on the assessment metrics used to gauge its performance. Performance assessments are typically based on system efficiency measures such as total flight delay.

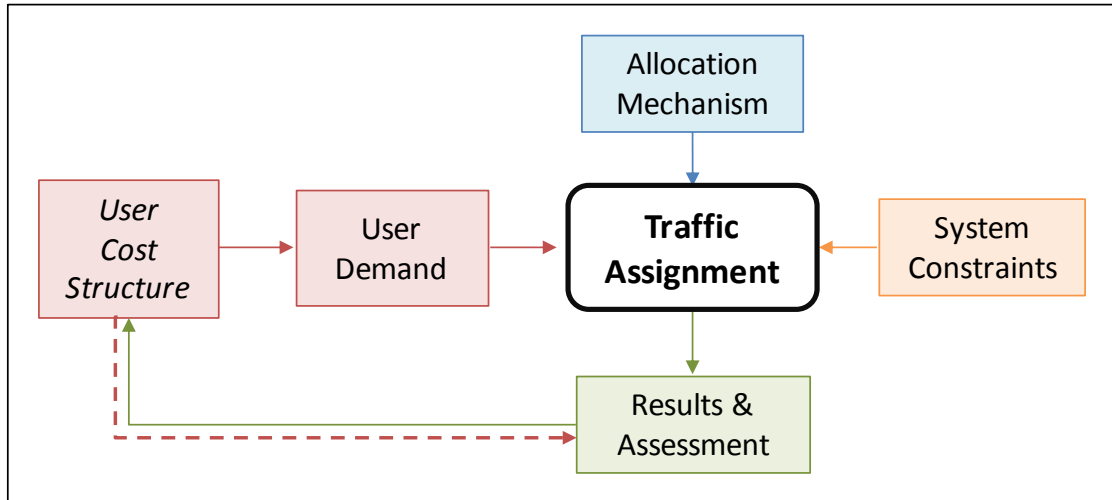


Figure 1.3 Current resource allocation process

There are many resource rationing mechanisms (uppermost box, Figure 1.3) that could be considered for use if FAA traffic managers were to take control of rerouting in an AFP. To demonstrate two possibilities, consider the example illustrated in Figure 1.4. Two flights (A and B) are planned to travel some nominal route with original departure times 0 and 5 minutes, shown in the top box. The route is completely closed due to convective weather; to accommodate these flights, departure slots on two alternative routes are offered (middle box). Suppose that flights A and B provide the traffic manager with their en route costs (in units of ground delay minutes, as discussed further below) for each route, also shown in the top box. The final cost is calculated based on the difference between the original departure time and the assigned departure time, plus the en route cost. Say traffic managers are obligated to serve Flight A first and Flight B second (Allocation 1). This kind of sequential allocation may be necessary because Flight A has some clear priority over Flight B, and expects to be served first as a result of some fairness principle. In this case, Flight A would be given Route 1 slot 1 as it is the lowest cost option available to it. Flight B would be left with Route 1 slot 2 as its best available option. The total cost of this allocation is 250. If the goal is to minimize total cost (Allocation 2), they would assign Flight A to Route 2 slot 1 and Flight B to Route 1 slot 1. The cost of this allocation is 240. Clearly the allocation results could be different if airlines submitted different cost values, or if some equity constraint, such as keeping the highest cost incurred as low as possible, were also included in an allocation.

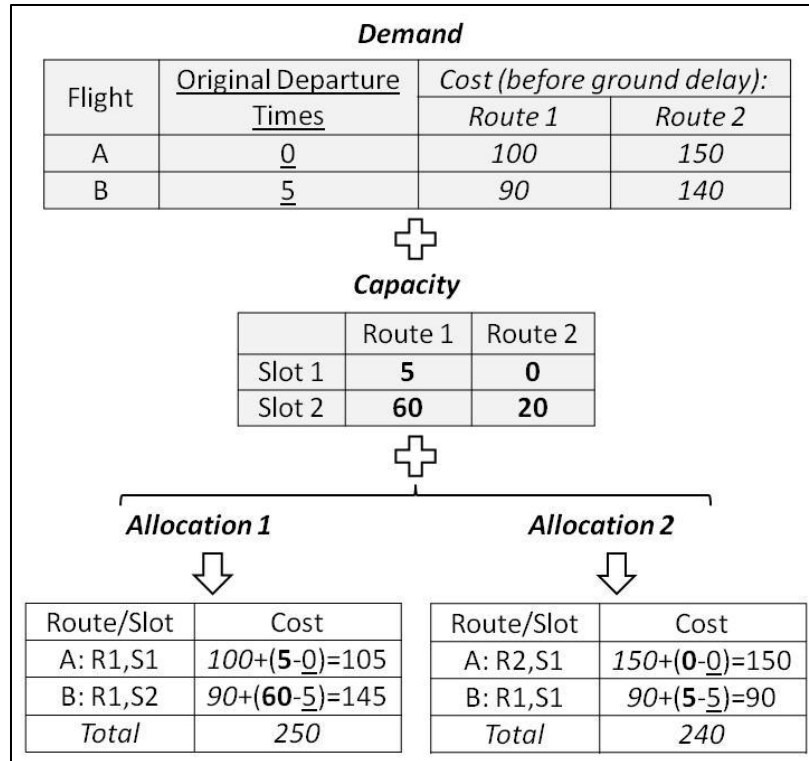


Figure 1.4 Illustration of possible allocation outcomes

We list a few of the many other allocation mechanisms that could be adopted (Ball, Futer, Hoffman, & Sherry, 2002). Traffic managers may be instructed to minimize a system cost metric such as the weighted sum, with or without consideration of flight/operator equity. Allocation could follow a “first-come, first-served” process where the ordering is based on the time of resource requests, the original schedule, or some other criterion. Real-time auctions of available resources could be employed. These are just a few of the many potential allocation mechanisms.

As stated previously, user inputs into the resource allocation process currently consist of providing up-to-date flight schedules, or demand information (Figure 1.3). We are interested in a system where airlines, in addition to providing updated demand information, provide other richer and more detailed information into the allocation process. This would ideally consist of users’ flight cost and resource preference information, and would be represented by changing the “user demand” box label in Figure 1.3 to “user demand and preference”. However, as the commercial airline industry is highly competitive, airlines are typically reticent to reveal any information about themselves that could potentially be used by competitors to gain a market advantage. Also, it is likely that additional effort will be needed to develop these user inputs. As a result, airlines must be offered incentives and rewards commensurate to the amount of information requested of them. In other words, airlines must feel that the rewards offered by traffic managers must exceed the disutility of developing and revealing proprietary information. Additionally, the competitive nature of the industry also leads naturally to gaming behavior, which is often a critical issue in the design of ATFM programs. It is not

likely that airlines will provide accurate flight cost information if they perceive that by reporting higher cost differentials they could receive more desirable resources. When designing resource allocation schemes, it is necessary to consider and address the gaming behavior that could potentially negate the benefits of these schemes.

In this research, we propose a functional form to represent user flight cost in the context of an AFP (leftmost box, Figure 1.3). We also propose three basic resource allocation schemes for AFP flights, where each consists of an allocation mechanism and a user input type. The three schemes are based on two user input types – Stated Route Preference input and Parametric input. Stated Route Preference input (Section 3.1) requires operators to supply route-level flight cost information about each alternative route. It is based on the delay thresholds concept developed as part of the Flow Constrained Area Rerouting (FCAR) Decision Support Tool by Metron Aviation (Hoffman, Lewis, & Jakobovits, 2004), which is also presented in more detail in Section 3.1. Parametric input (Section 3.2) requires flight operators to supply parameters of the flight cost function which are not route-specific, but which traffic managers can use to calculate the user costs of reroute and ground delay options available in an AFP. The first two allocation schemes employ Stated Route Preference user input. The first scheme allocates resources with the goal of minimizing total system-wide user cost (without considering equity) while the second allocates resources through a First Submitted, First Assigned (FSFA) process. In FSFA, flights are assigned the best available resources in the order they submit their input data. FSFA is just one of the many allocation mechanisms described above. It is one example of resource assignment according to a flight prioritization scheme, which could also be based on other criteria, such as flight schedule. The third assignment model uses the Parametric user input scheme and again allocates resources to minimize total system cost. Several other allocation schemes are discussed in Chapter 5, including one that employs the RBS prioritization scheme. Furthermore, the efficiency of the allocation results are assessed based on the sum of individual user costs, which depend on the user flight cost structure.

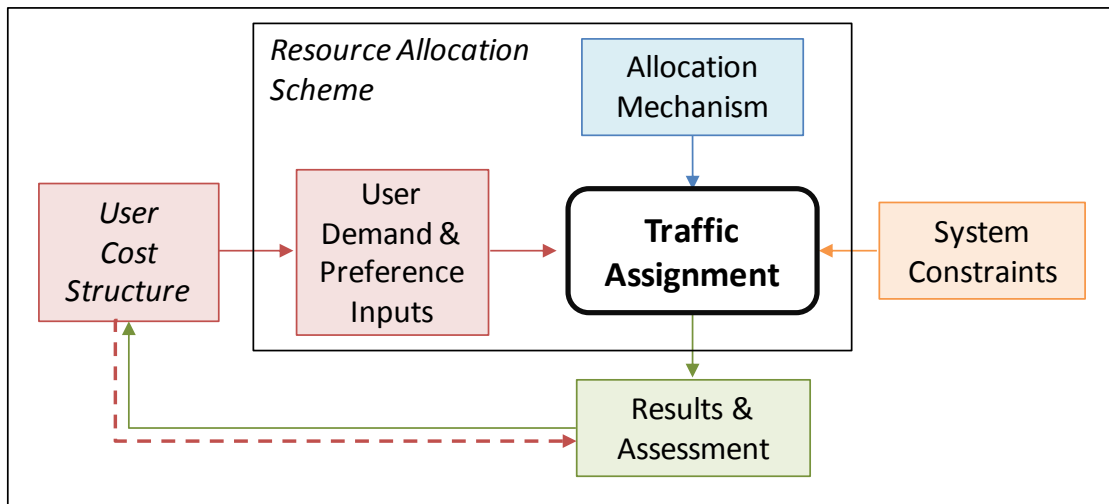


Figure 1.5 Proposed modeling and analysis framework

In this user cost-based performance assessment, allocation performance will improve when inputs that are a good representation of users' underlying cost structures are incorporated. However, the performance of a resource allocation scheme also depends on the allocation mechanism employed. We assess how the different schemes perform against one another under changing assumptions about the flights' route preferences. One of the main results of this research is a framework through which feasible user input and allocation mechanism combinations can be represented, evaluated, and compared.

1.3 Literature Review

Our review of the literature begins with a discussion about the classic traffic assignment problem, both under static and dynamic treatments, and will move to their applications to air traffic flow management (ATFM) problems. We then discuss the work that has been done to address equity and fairness issues in resource assignment for ATFM as well as a wider class of network flow problems. The final section discusses the limitations of the existing literature in addressing how a central authority like the FAA can best serve their customers using uncertain and/or incomplete information about these customers.

1.3.1 Traffic Assignment

The literature on traffic assignment is extensive and well-established. Of Wardrop's two equilibrium principles, the User Equilibrium (UE) is of greater interest in surface traffic planning, where drivers can make their own personal travel choices to minimize their travel costs. Under congested conditions, unique user equilibrium (UE) conditions exist when drivers are homogeneous. It has also been shown that for certain cases, the UE solution exists and is unique for heterogeneous driver classes (Daganzo, 1983) (Sheffi, 1985). Konishi (2004) extends Daganzo's work to heterogeneous drivers with different utility functions. Konishi shows that an equilibrium solution is unique for a general class of utility functions on a simple network. Also, Leurent (1993) accounts for driver heterogeneity through the use of continuously distributed (over the driver population) values of time, and demonstrates that there exists a unique cost-versus-time equilibrium. His work can be differentiated from a multiple user classes model (such as that analyzed by Daganzo) in that heterogeneity in his model is demand-related, not supply-related. There are many difficulties that arise in surface traffic assignment problems due to the nature of driver behavior, traffic controls and physical infrastructure characteristics on road networks. Many of these do not apply to air traffic flow management (ATFM) due to fundamental differences in "driver" behavior, the physical structure of airspace, aircraft flight patterns, and traffic management activities. However, from a traffic assignment perspective, the most significant difference is that aircraft cannot be in the airspace without permission from air traffic managers, and are always under their control. As a result, traditional ATFM models (including those discussed in the next section) have focused on system-optimal traffic assignment solutions. Dynamic traffic assignment (DTA) has also been studied extensively in the context of road and air transport networks

since the late 1970s. The problem has been addressed using optimization models, variational equality theory, and simulation, in addition to other tools (Peeta & Ziliaskopoulos, 2004). DTA has been applied to the ATFM problem mainly in reaction to changing adverse weather (Bertsimas & Stock Patterson, 2000) (Mukherjee & Hansen, 2007).

1.3.2 Air Traffic Flow Management Models

There has been much work in developing optimization models to support air traffic flow management (ATFM) decisions under disruptions. The objective of many such models is to minimize the expected system-wide cost of delay to flights. Uncertainties in capacity and weather have been extensively addressed in the single airport ground holding problem (SAGHP). This is appropriate as it is still true today that the main traffic bottlenecks in the NAS are at airports and their terminal airspace. Some of the earliest work on the SAGHP includes that by Andreatta et al. (1987) and Richetta and Odoni (1993). Ball, Hoffman, Odoni, & Rifkin (2003) show that the integer program of a stochastic version of the model can be solved using linear programming. They give consideration to its integration with the CDM paradigm. Mukherjee and Hansen (2007) extend the model further by considering a dynamic stochastic treatment of the problem. More recently, they extended their work to consider dynamic rerouting decisions in airport terminal areas (Mukherjee & Hansen, 2009).

When there is convective weather in the en route airspace, which occurs frequently in the summer months, heavily trafficked corridors can be greatly affected and cause flight delays that cause even more delays throughout the NAS. As traffic volumes increase, these en route delays become more significant. As a result, some attention has shifted to address the effects of en route congestion on ATFM. Many of these models are formulated as integer programs using discretized time intervals, as in most SAGHP models. One of the most well-known models was proposed by Bertsimas and Stock-Patterson (1998), and includes both airport and en route airspace capacity constraints. They considered ground holding, air holding and rerouting decisions in a static deterministic setting, (Bertsimas and Stock-Patterson later proposed a two-stage dynamic model (2000)). It should be noted that their cost function accounts for many aspects of operational costs including fixed costs, aircraft availability, ground delay costs, and en route flight costs. Although flight-specific air and ground hold cost ratios are included in the Bertsimas and Stock-Patterson model, it does not provide any information or assumptions about these values or how they would be obtained. Goodhart (2000) addresses the limitations of previous ATFM models, in that they assume the FAA possess all information about the flights and their costs, and have complete authority to make all decisions about these flights' departure times and routes. She proposes an ATFM model framework where decisions are made cooperatively, relying on information exchange between the FAA and flight operators. Operators are expected to provide their resource preferences to the FAA, and she approaches the problem in two ways. In the first model, she assumes that airlines offer their preferred route and departure time information to the FAA such that they can make decisions that better reflect airline preferences. In the

second model her work takes a more “bottom up” approach. She assumes that airlines have optimized their own costs and chosen their resources with respect to the NAS constraints they face, and the FAA attempts to accommodate them with respect to the capacity and equity constraints.

Jakobovits et al. (2005) formulated a variant of the GDP algorithm to schedule, reroute and airhold flights flying into and around constrained airspace. They employ a dynamic capacity estimation algorithm to forecast time-varying entry and exit points and maximum flow rates for constrained airspace. The algorithm imposes ordering schemes that align with Collaborative Decision Making.

1.3.3 Equity and Collaboration

As mentioned previously, ATFM models have traditionally aimed to optimize system efficiency in assigning resources to flights. Much of this work does not take equity into consideration, which is problematic in practice for two major reasons: a central authority like the FAA has an obligation to serve its customers (the airlines) fairly, and customers are less likely to share up-to-date information with the FAA (and potentially compromise some of their competitive advantage) without the promise of good initiatives and fair treatment. Over the last decade more attention has been given to equity in ATFM, mainly through GDPs that address airport and terminal area capacity shortfalls. One of the most significant developments involved the application of the Ration-By-Schedule (RBS) algorithm, which has been studied extensively. Vossen et al. (2003) describe a framework for equitable allocation in GDPs through RBS, illustrating its operational impacts in reducing systematic biases due to flight exemptions. Vossen and Ball (2006) also demonstrate that RBS is both efficient and equitable under certain conditions, in that the equitable solution is contained within the solution space of maximally efficient solutions. Hoffman et al. (2007) and Ball et al. (2010) introduce the Ration-by-Distance (RBD) allocation method, which is shown to be more efficient than RBS under early GDP cancellation, but less equitable. As a result of this finding, they introduce an equity-based RBD (or E-RBD) algorithm, which is a constrained version of RBD that imposes an upper bound value on a pre-defined equity metric. Glover and Ball (2010) use multi-objective optimization and investigate the use of several different objective functions to more precisely balance efficiency and equity. They compare their results to existing algorithms RBS, RBD (Ration-by-Distance) and E-RBD (Equity-based RBD).

Equity applied specifically to the problem of en route resource allocation has been studied more recently as well. Hoffman et al. (2005) discuss and evaluate the efficiency of allocation strategies tailored to the en route resource allocation problem, as an alternative to the use of GDPs in assigning en route resources. One of the ways they address the airspace resource problem is by rationing multiple resources simultaneously. Their strategy was also designed to incorporate changing user preferences and real-time operational decisions by the FAA. Pourtaklo and Ball (2009) propose an algorithm to equitably allocate airspace slots specifically within the AFP context using flight operator preference information and randomization.

Resource allocation in ATFM has also been studied to some extent from a game-theoretic and/or market-based perspective. The one-player case corresponds to a classic traffic assignment system-optimal solution, while a many-player case yields the user equilibrium solution. Haurie & Marcott (1985) formulate a non-cooperative game where players, defined by their origin-destination pairs, must send flows along a congested network to serve demand at their destination node. The cost of sending flow along a given link is a function of the flow on that link (congestion effect). They show that the Nash-Cournot equilibrium corresponds to the user equilibrium. Wie (1993) studies a dynamic extension of (Haurie & Marcotte, 1985), where each player must make decisions (to minimize their cost) about sending a fixed volume of traffic from a single origin to a single destination over a network of routes. Players make simultaneous decisions over time, which is modeled using differential game theory, and Wie establishes a dynamic game theoretic interpretation of the user equilibrium condition. He considers delay costs in addition to travel (congestion) costs, which Haurie & Marcott do not consider. The author extends his work (Wie, 1995) to account for two types of players – a user-equilibrium player and a Cournot-Nash player. The latter behaves to establish a system-optimal cost outcome.

Waslander, Raffard, & Tomlin (2008) propose a market mechanism-based approach in allocating en route resources to competing airlines, and incorporate this rationing mechanism into an ATFM model. They show that it is in the airlines' best interest to participate in the airspace resource allocation market; they can do no worse by participating than if they do not participate and allow the central decision maker to assign them resources without taking their preferences into account. The resources in this study consisted of access to airspace sectors during particular time intervals. In Waslander et al. (2008), the authors consider that airlines submit maximum lump-sum bids for resources in a market where they influence resource prices through their bids. They show that a Nash equilibrium and a bound on the worst efficiency loss exist for utility maximizing players that anticipate how their bids will affect resource prices.

1.3.4 Research Contributions

This literature review has briefly touched on the extensive body of research in transportation network analysis and most precisely, air traffic flow management models applied in the context of en route resource allocation. Over the last decade, collaborative decision making concepts have been incorporated into ATFM models, some of which go a step further to consider systems where airlines provide resource preference information to the FAA's resource allocation processes. However, these models do not consider that airlines may not be inclined to provide their preference information into the allocation process without a guarantee of fairness or efficiency in return. They also do not consider that airlines will attempt to gain advantages by gaming the system through their input information. And although many have formulated complex stochastic and dynamic models in response to uncertainty regarding how weather impacts NAS operations, they do not consider the uncertainty and incompleteness inherent in the FAA's knowledge about the airlines they serve. We address these gaps in the literature. We propose

allocation schemes that focus on how combinations of structured user preference inputs and resource rationing mechanisms can be combined to allocate constrained en route resources fairly and efficiently, particularly within the context of the AFP. Our work pairs user preference input schemes with resource rationing mechanisms, taking into account the fact that airlines will not be inclined to reveal preference inputs without being offered some type of resource guarantee from traffic managers in return for this proprietary information. One of the most significant contributions of this work is in the use of random utility theory to represent airline flight costs from the FAA's perspective. We propose a methodological framework that utilizes a random flight cost model to assess and compare these resource allocation schemes. We also look at some of the gaming issues behind our allocation schemes by considering some simple aspects of competitive behavior.

2. Flight Cost Model

2.1 Evaluation Framework

Figure 2.1 depicts the geometry of our modeling framework. Two points, or fixes, in en route airspace are connected by a nominal route, designated as such because it is the lowest cost path between the two points. Flights enter the nominal route at entry fix “A” and leave at exit fix “B”. The nominal route has sufficient capacity to serve the pre-AFP scheduled demand $D_0(t)$ between the fixes, until a constraint develops at some point along its path and lasts for some duration. The capacity of the nominal route is now reduced, and all N flights originally scheduled to use this route must be reassigned to observe the reduced capacity. All N flights are either given delayed departure times on the nominal route, or rerouted to one of $R - 1$ alternate routes and assigned a delayed departure time on their alternate route. Each alternate route r is characterized by its capacity and travel time. We characterize routes by travel time rather than distance, as favorable winds on a longer route can result in better travel times and fuel savings than a shorter route with less favorable winds. The nominal route is assumed to have the lowest travel time during the period of interest, and therefore the lowest cost of travel.

We further assume that fixes A and B are not bottlenecks, and for the purpose of this problem they can be thought of as the flights’ origin and destination. Flight trajectories upstream of fix A and downstream of fix B are not considered in this analysis.

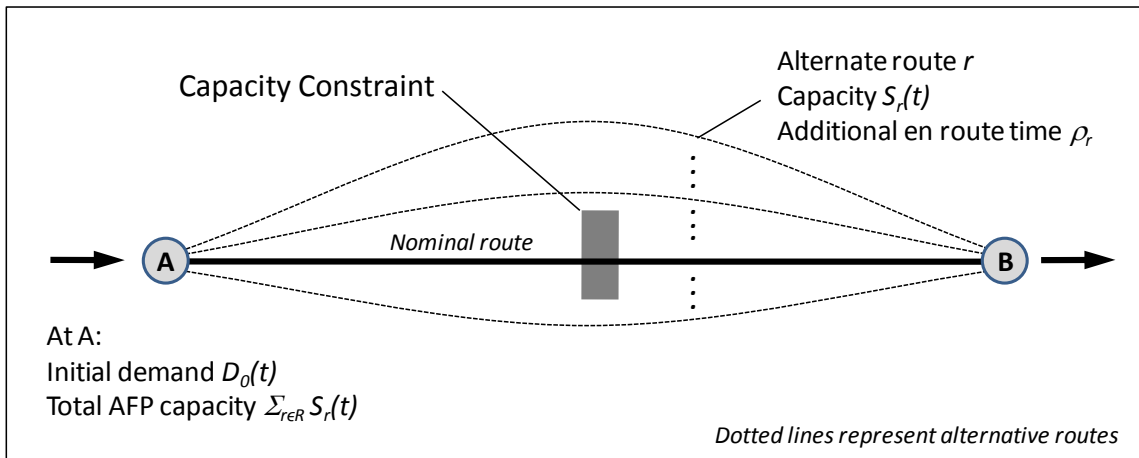


Figure 2.1 Model airspace geometry and select parameters

As mentioned in Section 1.2, FAA traffic managers typically have limited access to airline flight cost details and subsequent routing preferences. This analysis is not concerned with the costs of the airlines’ original scheduled flight plans, because we

assume that these flight plans were those most preferred under ideal conditions. Instead, this analysis focuses on evaluating the added costs associated with greater en route time and ground delay due to the AFP.

2.2 User Flight Cost Model

The flight cost function, $c_{n,r}$, represents the additional cost of flight n taking route r in an AFP. $c_{n,r}$ is a function of the additional travel time of route r compared to the nominal route (assuming that aircraft fly at desirable, fuel-efficient speeds), time spent waiting for their assigned slot on their AFP route (ground delay), and other factors. It is a generalized cost function as it is quantified in units of ground delay minutes. We assume it is the sum of the above components:

$$c_{n,r} = c_{n,r}^{en\ route} + c_{n,r}^{grd\ delay} + other_{n,r} \quad (2.1)$$

$c_{n,r}$ accounts for direct costs including additional fuel, crew time, and equipment maintenance, and indirect costs such as passenger dissatisfaction, gate time, flight coordination, the airline's satisfaction with their own specific objectives (such as minimization of en route turbulence), and others. Air holding is not included because we assume that traffic managers have perfect information about the capacity constraint location and duration, scheduled demand $D_0(t)$, and all AFP route capacities $S_1(t), S_2(t), \dots, S_R(t)$. As such, all anticipated delay can be incurred on the ground.

We assume that the term $other_{n,r}$ represents flight n 's private preferences about taking route r . These preferences are private in that they are known by the airline that operates flight n but not to the FAA traffic managers. We replace it with $\varepsilon_{n,r}$ and further identify the other components of the function as follows:

$$c_{n,r} = \alpha_n \cdot \rho_r + d_{n,r} - g_{0,n} + \varepsilon_{n,r}, \quad \varepsilon_{n,r} \sim P \quad (2.2)$$

where

- α_n is a ratio that identifies flight n 's en route time cost in units of ground delay minutes, and $\alpha_n \geq 1$;
- ρ_r is the additional flight time (in en route minutes) on route r compared to the flight time of the nominal route, $\rho_r \geq 0$;
- $d_{n,r}$ is the new departure time for flight n on route r at fix A;
- $g_{0,n}$ is the original scheduled departure time for flight n at fix A;
- $\varepsilon_{n,r}$ is a random error term representing flight n 's "other" preferences for flying route r . It represents AFP- and route-specific preferences factors that are, by definition, known to a flight's operator but not known to traffic managers or other airlines. $\varepsilon_{n,r}$ follows some distribution P .

The unit cost of airborne delay exceeds that of ground delay, so $\alpha_n \geq 1$. ρ_r is non-negative, assuming the nominal route had the shortest flying time under good weather

conditions; hence its status as the nominal route. Ground delay, or $d_{n,r} - g_{0,n}$, is non-negative as well because aircraft cannot be rescheduled to depart before their original scheduled departure time. Recall that $d_{n,r}$ and s_n are in fact departure times *from Fix A*, since we are only concerned with capacity restrictions that occur between Fix A and B.

Slots correspond to the sequence ordering as well as departure times of flights on a route, such that a flight assigned to slot i on route r is the i th flight to depart on that route. Let us index slots by increasing departure time on a given route r by $i_r = 1, \dots, X_r$. The variable X_r represents the total number of flights assigned to route r , and it follows that $\sum_{r=1}^R X_r = N$. We can identify the route to which flight n is assigned by $r(n)$, and its slot on that route by $i_r(n)$.

The AFP capacity of each alternative route r is $S_r(t)$, and it follows that the instantaneous minimum headway at time t is $S_r^{-1}(t)$. Now assume that $S_r(t)$ is constant over the duration of the AFP such that $S_r(t) = S_r$, and aircraft on route r are scheduled at constant headways g_r such that $g_r = S_r^{-1}$. By this assumption of constant headways and the mapping of flights to their assigned route and slot, the departure time of flight n on route r is $g_{r(n)} \cdot i_r(n)$. It then follows that the total cost of an AFP can be expressed as:

$$C = \sum_{n=1}^N \alpha_n \rho_{r(n)} + g_{r(n)} i_{r(n)}(n) - g_{0,n} + \varepsilon_{n,r(n)} \quad (2.3)$$

where

- C is the total cost of an AFP, measured in ground delay minutes;
- $\rho_{r(n)}$ is flight n 's additional en route time when it has been assigned to route r from the nominal route;
- $g_{r(n)}$ is the new AFP departure headway on the route r to which n has been assigned;
- $i_{r(n)}(n)$ is the slot on r to which n has been assigned;
- $g_{0,n}$ is flight n 's original pre-AFP scheduled departure time, and
- $\varepsilon_{n,r(n)}$ is a stochastic term representing n 's "other" private preferences for flying on the route r to which it was assigned.

Ground delay $(g_{r(n)} i_{r(n)}(n) - g_{0,n})$, is non-negative $\forall i, r, n$. The assumption of constant route headways g_r facilitates the formulation of analytic approximations to the resource allocation schemes introduced in Chapter 3 (Section 3.5). However, the assumption is retained for the numerical results, which are built from the exact models, because it does not detract from any insights we obtain from the analysis.

2.3 Heterogeneous Flights

We now give attention to the idea that N flights operators submit different en route cost parameters to traffic managers such that $\alpha_1 \neq \alpha_2 \neq \dots \neq \alpha_N$. We assume that cost parameters are distributed over the flight population according to some distribution, and that the N AFP flights are a representative population sample. Furthermore, if flights are ordered by increasing α , we define $\alpha(n) = \alpha_n$ to be the en route cost parameter for the n th flight, and also the n th smallest value for this parameter.

In this research, we assume that the en route cost parameter has a uniform distribution with a minimum value α_{min} and a maximum value α_{max} . This implies that α_n can be approximated as:

$$\alpha_n = \alpha_{min} + \frac{(\alpha_{max} - \alpha_{min})}{N} \cdot n \quad (2.4)$$

Equation (2.4) also facilitates the analytic approximations to the resource allocation schemes of Chapter 3. It is also used for the numerical results.

2.4 Unknown Cost Component

Recall that the stochastic term in Equations (2.1) through (2.3) represents the part of airlines' route-specific flight preferences that traffic managers have little or no information about. The stochastic term enters the objective functions at different places in the various resource allocation schemes, depending on whether or not the traffic manager has received this preference information from the airlines. The specification of the stochastic term affects the comparative performance of each allocation scheme. The strategies are evaluated over different variances of the stochastic term, where the variance is used to indicate the degree to which we have captured operator utility in the deterministic part of the cost equation.

In the following chapter, we specify that the random term is independent and identically distributed (iid) normal, with mean 0 and variance σ^2 . The independence assumption is reasonable if we trust that our cost function captures the major elements of airlines' AFP reassignment costs, both across flights and across routes. Although the (deterministic and stochastic) costs of flights operated by the same airline may be correlated, we assume here that intra-airline flight differences are so pronounced that this correlation can be ignored. In the approximations to the resource allocation schemes, we assume the error term takes an extreme value (Gumbel) form. The Gumbel distribution is characterized by its location (a) and scale (b) parameters, where the variance is determined by the scale parameter alone but the mean is determined by both a and b . In the standard Gumbel distribution, $a = 0$ and $b = 1$. The Gumbel distribution has several important properties that make it analytically convenient to use in the specification of choice probabilities and expected cost, when faced with a given set of choices (Train, 2003). In addition, the Gumbel distribution is fairly similar to the normal distribution.

2.5 Concluding Remarks

It has been demonstrated in the literature that flight costs are not a linear function of flight delay, and in fact the marginal cost of delay is increasing. This means that a small increase in delay will cost much less when a flight is experiencing low delays than if it experiences a comparable increase in delay when already experiencing higher delay levels. An increase in delay is not likely to disrupt airline operations or passengers to a great extent when delays are low. However, when delay levels are already high, that same increase is likelier to cause a service disruption that affects more of the airline's downstream operations. The delay is much more likely to cause missed connections by passengers, crew time-out problems, and delays to downstream flights due to late equipment turnaround, all of which are very costly. In the future we propose to modify the flight cost model introduced in this chapter to account for non-linearity in delay cost, and we discuss the specifics of this further at the end of Chapter 3.

3. Resource Allocation Schemes

This chapter introduces three schemes by which flights caught in an AFP might be assigned en route resources. Two resource allocation schemes are based on the Stated Route Preference input concept, while the third allocation scheme is based on the Parametric input concept. Section 3.1 introduces the concepts and formulations of the two schemes based on Stated Route Preference input, and Section 3.2 presents the Parametric scheme. In Section 3.3 we introduce a simple example to illustrate how the stochastic term of the flight cost function affects the relative performance of the models. From this example we can understand some basic properties of these models. Section 3.4 contains numerical examples, which serve to explore differences in the allocation efficiencies of each model, under a variety of demand and supply scenarios. Section 3.5 presents analytic and quasi-analytic approximations of the Parametric and the First Submitted, First Assigned (FSFA) Stated Route Preference input schemes. These approximations are formulated by relaxing the departure constraint, which stipulates that a flight's ground delay must be non-negative. We demonstrate that the approximations are quite good under certain conditions. The final section (3.6) discusses directions for future work.

3.1 Stated Route Preference

Stated Route Preference user input utilizes the delay thresholds concept of the Flow Constrained Area Rerouting (FCAR) Decision Support Tool (Hoffman, Lewis, & Jakobovits, 2004). Delay thresholds are a method by which airlines can express their privately known route preferences to traffic managers. Developed by Metron Aviation, it was designed to incorporate user preferences in an AFP resource allocation process that combines rerouting decisions with delayed departure times. The FCAR also outlines a method to facilitate information flow between traffic managers and operators.

The delay threshold concept gives operators flexibility in identifying the best resource options for their AFP-impacted flights, in a way that does not reveal explicit cost information about them. In an AFP, if a flight wants to fly the original route (whose capacity was reduced in the AFP) it must wait on the ground until its assigned AFP departure slot. There may, however, be an alternate route that is less desirable (for instance, it has a greater distance or more unfavorable winds) but has less associated ground delay. Despite the increased en route time, which is generally more expensive than ground delay, it may be less costly for the flight to take that alternate route. This would be the case if the flight avoids a long ground delay by flying the alternate route,

and/or an on-time arrival were particularly critical because of a downline connection or other factors. In the FCAR process, operators of impacted flights are asked to submit inputs that communicate these types of preferences to the FAA traffic managers.

When an AFP is announced, the flight operators are asked by traffic managers to submit route preference information to them. For each route r available in the AFP, the operator of flight n submits a delay threshold $\Delta_{n,r}$. The quantity $\Delta_{n,r}$ is assumed to reflect the airline's complete generalized cost of its flight n traveling on route r , relative to the original flight plan, without any additional ground delay costs and before ground delay is assigned. By submitting delay thresholds, airlines are providing the FAA with complete information about their route preference structure. Delay threshold values are expressed in units of ground delay minutes such that true airline costs are not explicitly revealed, and can be used directly with ground delay times to choose a route and slot.

Once traffic managers receive the delay threshold values, they will apply their adopted allocation mechanism to assign resources to flights. We illustrate a flight's costs using the example shown in Figure 3.1. Suppose a flight n had three route options in the AFP, and the flight operator submitted a delay threshold value for each route. Traffic managers will calculate the cost of n taking some slot on some route r as the sum of the delay threshold for that route and the ground delay for n to take the said slot. In the First Submitted, First Assigned scheme (introduced in Chapter 1 and to be discussed in detail shortly), once it is n 's turn for allocation, traffic managers check the slot availability on each route and determine the ground delay that flight n must take on each route: $GD_{n,1}$, $GD_{n,2}$, or $GD_{n,3}$. If the minimum cost route were assigned to this flight, it would be $\min(\Delta_{n,1} + GD_{n,1}, \Delta_{n,2} + GD_{n,2}, \Delta_{n,3} + GD_{n,3})$, or route 3 according to Figure 3.1.

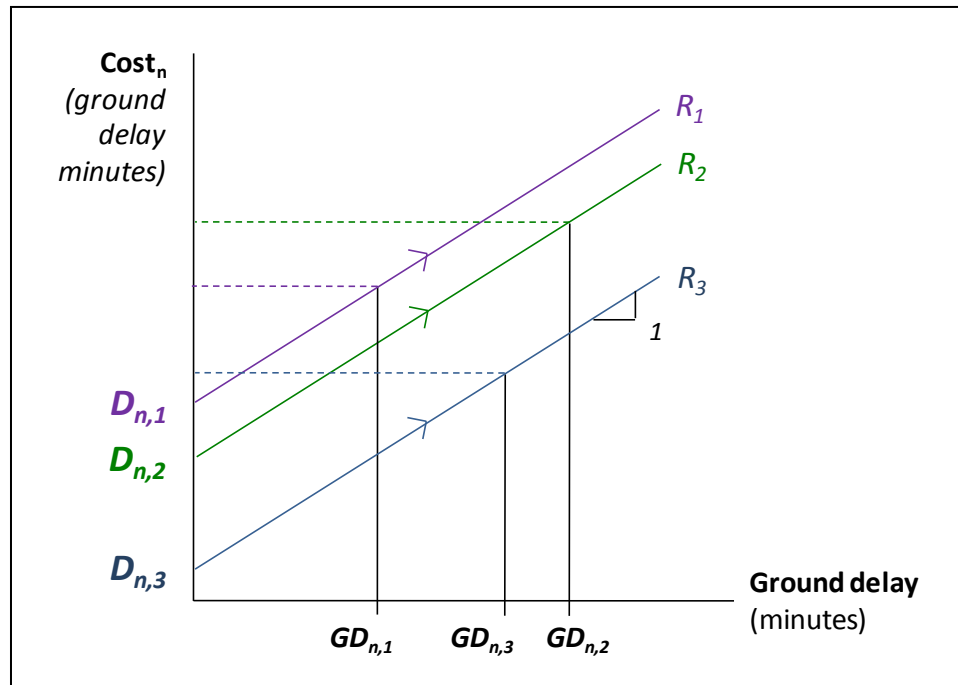


Figure 3.1 Delay thresholds

In the Stated Route Preference concept, we assume that each airline would calculate the added cost of a reassignment option using Equation (2.2). However, as indicated above, airlines do not know what slots the traffic managers have available for their flight(s) on each route, and therefore have no information about the amount of ground delay that will be assigned to their flights. As a result, airlines submit delay thresholds by which traffic managers can compare the costs of allocation – routing options and ground delay slots – for each flight. In this research, we use the flight cost model introduced in Section 2.2 to define the functional form of the delay thresholds submitted by flight operators, such that they are calculated as follows:

$$\Delta_{n,r} = \alpha_n \rho_r + \varepsilon_{n,r}, \quad \varepsilon_{n,r} \sim P \quad (3.1)$$

As discussed above, traffic managers will assign resources to each flight through a chosen allocation scheme using these delay thresholds which, again, represent complete route preference information. These inputs ensure that under any combination of ground delay slots that could be assigned to their flight, the airlines have informed the FAA about their route preferences.

We first consider a situation where an AFP has been created, and FAA traffic managers request airlines to submit their inputs before some cutoff time. Only after this cutoff time, when traffic managers have presumably received information for most of the participating flights, do they allocate resources. To represent this scheme we employ a model where the entire batch of “requests” is considered simultaneously in making optimal resource rationing decisions. We then consider a system where airlines are allocated their preferred resources on a First Submitted, First Assigned (FSFA) basis, and are thereby incentivized to submit their required inputs in a timely manner.

3.1.1 Full Information Optimal (OPT)

In the full information optimal (OPT) model, we assume that when AFP details are known and released to flight operators, traffic managers request complete route preference information from all operators with flights in the AFP, and receive them immediately upon request. Traffic managers then make system-optimal resource allocations to the AFP flights using the received route preferences. The objective is to minimize total operator cost, measured in units of ground delay minutes, without explicit considerations for flight equity. This model is highly idealized in that commercial aviation is an extremely competitive industry, and airlines are unlikely to offer such detailed information about their proprietary flight costs if they are not offered any incentives (resource or equity guarantees) in exchange for the information. They may submit untruthful and/or inaccurate information, or none at all. However, the OPT scheme yields the best system performance that can be achieved from any AFP allocation scheme, and its results can be used as a benchmark by which all other schemes are evaluated and compared.

Due to the fact that the flight operators’ complete route preference information is available for decision making through the information they offer to traffic managers, the

stochastic term of the flight cost model (representing private airline route preferences) is available to the resource allocation process, and therefore is included within the objective function. We must randomly draw these values for our numerical examples, and as a result Equation (2.3) cannot be solved analytically for OPT. Instead we formulate this model as an assignment problem, where the decision variable is a binary indicator of whether a flight takes a given slot j .

Decision variables: $x_{n,j} \in \{0,1\}$; $x_{n,j} = 1$ if flight n is assigned to slot j and $x_{n,j} = 0$ otherwise.

Objective function:

$$\min_{x_{n,j}, \forall n,j} C = \sum_{n \in F, j \in J: d_j \geq g_{0,n}} c_{n,j} \cdot x_{n,j} \quad (3.2)$$

Where

- F is the set of flights in the AFP (set F contains N total flights);
- J is the set of slots available over all routes, and each slot j is associated with some route r ;
- $c_{n,j} = \Delta_{n,r(j)} + d_j - g_{0,n}$, and is the cost of assigning flight n to slot j ;
- $\Delta_{n,r(j)} = \alpha_n \rho_{r(j)} + \varepsilon_{n,r(j)}$, and
- d_j is the departure time associated with slot j .

Constraints: $d_j \geq g_{0,n}$, $\sum_{j \in J} x_{n,j} = 1, \forall n$; $\sum_{n \in F} x_{n,j} \leq 1, \forall j$; $x_{n,j} \in \{0,1\}, \forall n, j$.

We assume that the stochastic term is distributed iid normal with mean zero and standard deviation σ , $\varepsilon_{n,r} \sim N(0, \sigma)$. Random sampling is used to generate $\varepsilon_{n,r}$ values; as a result we solve Equation (3.2) 5,000 times to generate an average result for C , for a given σ value, in the numerical examples of Section 3.4. We retain the assumption from Section 2.2 that the AFP departure headways on each route r are constant at g_r , such that if slot j is the x th slot on route r , then $d_j = x \cdot g_{r(j)}$. We also assume that pre-AFP scheduled departure headways are constant for all flights such that original scheduled departure times $g_{0,n}$ are constantly spaced. These assumptions are retained because they do not detract from what we are interested in learning about these resource allocation schemes. The constraints ensure that each flight n 's ground delay is non-negative, each flight is assigned to one slot, each slot is assigned at most one flight, and $x_{n,j}$ can only take values of zero or one.

We find C for increasing values of σ , starting at $\sigma = 0$, in Section 3.4.

3.1.2 First Submitted, First Assigned (FSFA)

In the First Submitted, First Assigned (FSFA) allocation scheme, we imagine that the details of an impending AFP have been released to the airlines. They are instructed to

submit their flight cost information sometime within the AFP planning period, which begins a few hours prior to the start of the AFP. As a result, traffic managers receive complete route preference information from flight operators in a sequence unknown beforehand. Each time an operator submits their route preference values for a flight, the FSFA algorithm identifies the best available and feasible departure time slot on each route. The flight in question is then assigned the minimum cost resource (route/slot combination) available at the time, and future requests are not considered at the time of allocation. As a result, the FSFA is a greedy allocation algorithm in that it makes a locally optimal choice when each flight submits. The FSFA allocation scheme offers flight operators a relatively clear incentive to supply their complete route preferences in a timely manner.

One of the consequences of FSFA is that it allows for early submitters to gain major rewards while late submitters can be greatly penalized. It also introduces an additional level of competition between the flight operators, in the timing of submissions. We consider these issues in Chapters 4 and 5.

The FSFA scheme can be simulated using the following recursive algorithm.

- 1) Assign α , g_0 and randomly drawn ε_r values to flights $n = 1, 2, \dots, N$. Recall that we assume $g_{0,n}$ are constantly spaced.
- 2) Order flights by increasing route preference ($\Delta_{n,r}$) submission times. At this point we cannot say what this order looks like – it could be completely random and independent of any flight or AFP characteristics. Or, it could be correlated to α , original scheduled departure time, total distance from origin to destination airport, aircraft size, etc. Whatever the order of submissions is, flights are now ordered and indexed by that ordering $m = 1, \dots, N$.
- 3) By increasing m , or $m = 1, 2, \dots, N$:
 - a. Find feasible and available slot j that minimizes $c_{m,j}$.
 - b. If $d_j \geq g_{0,m}$, then assign m to j such that the cost of the assignment is $c_{m(j)}$. Flag j as unavailable for all subsequent flights $m + 1, m + 2, \dots, N$.
 - c. If $d_j < g_{0,m}$, mark slot j as infeasible to m . Go back to (a).
 - d. Repeat (a) through (c) until $m = N$.
- 4) Find $\sum_{m=1}^N c_{m(j)}$.

Repeat (1) through (4) Z times (iteration $z = 1, 2, \dots, Z$) for each value of σ , starting with $\sigma = 0$.

3.2 Parametric (PO)

Under the Parametric resource allocation scheme, we envision that an FAA mandate would require operators to provide cost parameters for their flights to a centrally-managed FAA database. This requested parameter is the air-to-ground cost ratio α if we

assume that traffic managers have adopted the flight cost model of Equation (2.2). Operators would be encouraged to update their parameters as necessary. At the time that resource rationing decisions must be made (at some time after the AFP is announced several hours prior to its actual start time), the set of parameter values contained in the database at that time will be used to determine route and ground delay assignments for all AFP-affected flights. We assume that airlines are implicitly incentivized to provide their most up-to-date cost parameters, in order to maximize their likelihood of obtaining desirable allocations in the AFP. In addition, although it is true that detailed route cost information (such as that represented by the Δ values) is proprietary to each airline, it is assumed that the parameters of this database are high-level flight cost information typically accepted to be general knowledge. Therefore, the FAA would not need to provide additional incentives to obtain these parameters from flight operators.

By assigning flights to resources using the Parametric allocation scheme, the FAA traffic managers aim for a system-optimal total AFP flight cost solution. However, the only flight-level information that traffic managers have obtained from the operators are α values, as traffic managers have not obtained the private route preference information that would be provided through the stochastic term (as in the OPT scheme). As a result, if the stochastic term does not exhibit a high level of variance (i.e. σ is low), it means that the Parametric resource allocation can be very efficient. It also implies that the deterministic portion of the flight cost model is a good reflection of actual cost and/or an airline's preferences for routes are not highly variable or volatile. If, however, the stochastic term has a high variance (high σ), resource allocations will be less efficient. We would like to ascertain how this approach performs in comparison to the Stated Route Preference strategies as the variance of the stochastic utility – and hence the incompleteness of the traffic managers' information about flight operators' route preferences – increases.

The mathematical formulation for the Parametric scheme is identical to that of the OPT scheme, which is an assignment problem where flights are assigned to slots on routes. However, unlike the OPT scheme, the objective function consists only of the deterministic part of the flight cost function, because traffic managers only have information about deterministic costs in the Parametric resource allocation scheme.

Decision variables: $x_{n,j} \in \{0,1\}$, $x_{n,j} = 1$ if flight n is assigned to slot j and $x_{n,j} = 0$ otherwise.

Objective function:

$$\min_{x_{n,j}, \forall n,j} \hat{C} = \sum_{n \in F, j \in J: d_j \geq g_{0,n}} \hat{c}_{n,j} \cdot x_{n,j} \quad (3.3)$$

Where

- F is the set of flights in the AFP (set F contains N total flights);
- J is the set of slots available over all routes R ;

- $\hat{c}_{n,j} = \alpha_n \rho_{r(j)} + d_{n,j} - g_{0,n}$, and is the deterministic (known to traffic managers) cost of assigning flight n to slot j , and
- d_j is the departure time associated with slot j .

Constraints: $d_j \geq g_{0,n}$, $\sum_{j \in J} x_{n,j} = 1, \forall n$; $\sum_{n \in F} x_{n,j} \leq 1, \forall j$; $x_{n,j} \in \{0,1\}, \forall n, j$.

We retain the assumptions that AFP departure headways on each route r are constant at g_r , and original scheduled departure times $g_{0,n}$ are constantly spaced. The constraints are identical to those of the OPT scheme, which ensure that each flight n 's ground delay is non-negative, each flight is assigned to one slot, each slot is assigned at most one flight, and $x_{n,j}$ can only take values of zero or one.

Remember that the Parametric resource allocation is performed using only the deterministic portion of the flight cost model, which is all the flight cost information that traffic managers have. Once allocations are made, we can calculate the expected “true” cost of the allocation $E[C]$ by adding the stochastic term representing flight operators’ private route preferences. Given that we have assumed the stochastic term is distributed iid normal with mean zero and standard deviation σ , we know that

$$E[C] = \hat{C} + E \left[\sum_{n=1}^N \varepsilon_{n,r(j(n))} \right] = \hat{C}, \quad \varepsilon_{n,r(j(n))} \sim N(0, \sigma) \quad (3.4)$$

where $\varepsilon_{n,r(j(n))}$ represents the private route preference for flight n assigned to slot j , which in turn is associated with route r . To simulate C , randomly generate N values of $\varepsilon_{n,r(j(n))}$.

3.3 Properties

We construct a simple example to investigate the stochastic properties of the three allocation schemes introduced in this chapter. In this example, two flights ($N = 2$) can be assigned to one of two routes ($R = 2$) with one slot each. The cost of flight n taking route r is $c_{nr} = w_{nr} + \varepsilon_{nr}$, where w_{nr} represents deterministic costs and ε_{nr} is the stochastic term. The deterministic costs of a flight taking any route are equal such that $w_{11} = w_{12} = w_{21} = w_{22} = w$, and the total deterministic cost of any allocation is W . It follows that the stochastic terms will dictate how resources are rationed in each allocation scheme.

Table 3.1 Toy Model Allocation

Allocation	Flight 1 takes	Flight 2 takes	Total Deterministic Allocation Cost	Total Allocation Cost
1	Route 1	Route 2	$w_1 = w_{11} + w_{22} = W$	$C_1 = W + \varepsilon_{11} + \varepsilon_{22}$
2	Route 2	Route 1	$w_2 = w_{12} + w_{21} = W$	$C_2 = W + \varepsilon_{12} + \varepsilon_{21}$

Assume that the stochastic term $\varepsilon_{n,r}$ is distributed iid normal with mean 0 and standard deviation σ . Now say that Flight 1 and Flight 2 have different optimal slots, i.e. $c_{1r}^* = c_{2(-r)}^*$. The expected costs of the OPT and FSFA allocations will then be identical, $E[C_{OPT}] = E[C_{FSFA}]$. If Flights 1 and 2 have the same optimal slot, $c_{1r}^* = c_{2(-r)}^*$, then $E[C_{OPT}] < E[C_{FSFA}]$. Unless $\varepsilon_{n,r} = 0 \forall n, r$, we cannot say how $E[C_{PO}]$ will compare to the results of the other two models. In any case, we can express the expected costs resulting from the application of each resource allocation scheme as follows:

$$\begin{aligned} E[C_{OPT}] &= E[\min(C_1, C_2)] \\ &= W + E[\min(\varepsilon_{11} + \varepsilon_{22}, \varepsilon_{12} + \varepsilon_{21})] \end{aligned} \quad (3.5a)$$

$$\begin{aligned} E[C_{FSFA}] &= \frac{1}{2} \begin{cases} W + E[\min(\varepsilon_{11}, \varepsilon_{12}) + \varepsilon_{22}] & \text{if } \min(\varepsilon_{11}, \varepsilon_{12}) = \varepsilon_{11} \\ W + E[\min(\varepsilon_{11}, \varepsilon_{12}) + \varepsilon_{21}] & \text{if } \min(\varepsilon_{11}, \varepsilon_{12}) = \varepsilon_{12} \end{cases} \\ &\quad + \frac{1}{2} \begin{cases} W + E[\min(\varepsilon_{21}, \varepsilon_{22}) + \varepsilon_{12}] & \text{if } \min(\varepsilon_{21}, \varepsilon_{22}) = \varepsilon_{21} \\ W + E[\min(\varepsilon_{21}, \varepsilon_{22}) + \varepsilon_{11}] & \text{if } \min(\varepsilon_{21}, \varepsilon_{22}) = \varepsilon_{22} \end{cases} \end{aligned} \quad (3.5b)$$

$$E[C_{PO}] = W + E[\varepsilon_{1r}] + E[\varepsilon_{1(-r)}] = W \quad (3.5c)$$

Since ε_{nr} are iid normal, the moments of the maximum of N random variables are easily calculated (Bose & Gupta, 1959). The equations are summarized in Clark (1961); if $\varepsilon_1, \varepsilon_2 \sim iid N(0, \sigma)$, the expected value of the minimum of ε_1 and ε_2 can be expressed as:

$$E[\min(\varepsilon_1, \varepsilon_2)] = -\sigma/\sqrt{\pi} \quad (3.6)$$

and we can rewrite Equations (3.5a) and (3.5b) such that

$$E[C_{OPT}] = W - \sqrt{2}\sigma/\sqrt{\pi} \quad (3.7)$$

$$E[C_{FSFA}] = W - \sigma/\sqrt{\pi} \quad (3.8)$$

See Appendix A.1 for the detailed calculations.

We are interested in understanding how well an allocation scheme performs relative to the OPT scheme, which yields the most efficient total user cost solution possible under any given situation. As a result, we express the total cost results of the FSFA and Parametric schemes as ratios of OPT:

$$E[C'_{FSFA}] = \frac{E[C_{FSFA}]}{E[C_{OPT}]} = \frac{W\sqrt{\pi} - \sigma}{W\sqrt{\pi} - \sqrt{2}\sigma} \quad (3.9)$$

$$E[C'_{PO}] = \frac{E[C_{PO}]}{E[C_{OPT}]} = \frac{W\sqrt{\pi}}{W\sqrt{\pi} - \sqrt{2}\sigma} \quad (3.10)$$

Due to the fact that the total deterministic cost of any allocation is always W in this simple example, when traffic managers have perfect information about flight operators (represented by $\sigma = 0$, and therefore, $\varepsilon_{n,r} = 0 \forall n, r$) the OPT and FSFA models yield identical resource allocations and total generalized operators costs for a given set of parameters. Therefore, $E[C'_{FSFA}] = 1$ when $\sigma = 0$. Clearly in more realistic scenarios it would be the case that the deterministic results of system-optimal and greedy (such as FSFA) assignment algorithms will differ from one another; under $\sigma = 0$, a FSFA allocation will be less cost efficient than a system-optimal allocation. Also, as σ increases, Equation (3.9) increases as well.

It can be observed from Equation (3.10) that $E[C'_{PO}] = 1$ when $\sigma = 0$; the OPT and Parametric schemes yield identical resource allocations and total operator costs not only in this simple example but under any scenario at $\sigma = 0$. Recall that Parametric resource allocations do not utilize the operators' private route preference information provided through the stochastic term. As a result, as traffic managers' uncertainty about operators' private route preferences increases (represented by increasing σ), $E[C'_{PO}]$ will also increase.

Figure 3.2 displays $E[C'_{FSFA}]$ and $E[C'_{PO}]$ with respect to σ . The x -axis represents increasing values of σ as a proportion of w , where σ is the standard deviation of the stochastic term in the flight cost model and w is the deterministic cost of any flight n taking any route r (identical for all n and r). For instance, the point "0.10" on the x -axis indicates that σ is 10% of w , or $\sigma/w = 0.10$. Increasing σ represents greater variations in the flights' routing preferences.




Figure 3.2 Total cost efficiency relationship of allocation schemes

Some important properties of $E[C'_{FSFA}]$, $E[C'_{PO}]$, and their relationship to one another are illustrated in the above figure, including those discussed in the previous paragraphs. Firstly, both $E[C'_{FSFA}]$ and $E[C'_{PO}]$ are increasing convex functions of σ . $E[C'_{PO}]$ increases at a faster rate than $E[C'_{FSFA}]$, and as both equal one at $\sigma = 0$, $E[C'_{PO}]_{\sigma}$ is greater than $E[C'_{FSFA}]_{\sigma}$ at any positive value of σ .

Secondly, the slope of $E[C'_{PO}]$ is found to be 3.4 times larger than the slope of $E[C'_{FSFA}]$ at all positive values of σ . However, it is unlikely this is true under other sample constructs, as seen in the following section.

Finally, as mentioned previously, because the total deterministic cost of any allocation is always W in this simple example, $E[C'_{FSFA}]$ equals one when $\sigma = 0$. In more realistic scenarios where deterministic allocation costs differ, one can imagine that the entire $E[C'_{FSFA}]$ curve would be shifted up because FSFA does not offer a system-optimal resource assignment. However, under increasing uncertainty about user cost, we should expect that the FSFA results would also improve in comparison to the Parametric results, just as the OPT results do. Hence, in the numerical examples of Section 3.4, the FSFA total cost results cross the y -axis at a point greater than one, and therefore, the FSFA and Parametric total cost curves cross at values of σ greater than zero.

Appendix A.2 contains the calculations pertaining to the discussion of the results above.

3.4 Numerical Examples

To obtain some insight about the performances of the three allocation schemes under the FAA traffic managers' increasing uncertainty about the flights' private route preferences, represented by increasing σ , several numerical examples are presented here. Suppose N flights must be reassigned routes and departure times as part of the AFP. The nominal route remains open, but under a reduced capacity. We assume that air-to-ground cost ratios α_n are evenly distributed in $(1.5, 2.5]$ across all flights. We make this assumption because it is commonly cited in the literature that one unit of en route time is equal in cost to about two units of delay incurred on the ground. As a result, α is often assumed to take on a value of two in existing ATFM models (Mukherjee & Hansen, 2009), but we assume a distribution to capture the idea that it most likely varies across different flights. We also retain the assumption that the FSFA $\Delta_{n,r}$ submission order is random and independent of any flight or AFP characteristics. The competitive aspect of flight submission ordering is explored in Chapter 4, specifically Section 4.3.

We focus our attention to one supply scenario combined with several different demand profiles. The supply parameters are listed in Table 3.2, and the demand profiles are described below the table. The AFP will have a total of five route options, one of which includes the nominal route under a decreased capacity.

Table 3.2 Example supply parameters

Route	Capacity (aircraft per hour)	Departure Headway g_r (min) ^a	En Route Time h_r (min)	ρ_r (min)
1	24	2.5	135	35
2	20	3	130	30
3	10	6	120	20
4	12	5	115	15
5 (nominal)	7.5	8 ^b	100	0
$C_{AFP} =$		73.5		

a This is the arrival (and departure) headway at Fix A.

b Headways after AFP capacity reduction.

With the supply parameters above, we investigate three groups of demand profiles:

Scenario 1: Increasing rate of demand, D_0 ; $N = 50, 55, \dots, 95, 100$, $T = 60$ min. α_n are evenly distributed in $(1.5, 2.5] \forall n$.

Scenario 2: Fixed $D_0 = 85$ aircraft per hour; $N = 50, 55, \dots, 95, 100$, T increases. α_n are evenly distributed in $(1.5, 2.5] \forall n$.

Scenario 3: α_n are evenly distributed in $(1.5, \alpha_{max}] \forall n$, for $\alpha_{max} = 3.5$ and 5 , for the following demand profiles:

- $N = 50$, $T = 60$ min, $D_0 = 50$ airc/hr
- $N = 75$, $T = 60$ min, $D_0 = 75$ airc/hr
- $N = 100$, $T = 60$ min, $D_0 = 100$ airc/hr

Figure 3.3 contains results for demand Scenario 1, for $N = 75$ flights. The axes are identical to those of Figure 3.2, where the x -axis represents increasing values of the standard deviation of the stochastic term in the flight cost model, which in turn represents greater variation in the AFP flights' routing preferences. Specifically, each point on the x -axis represents the value of σ as a proportion of the average flight cost using OPT under perfect information, \hat{C}_{OPT} (where $\sigma = 0$); for instance, the point "0.10" means that σ is 10% of \hat{C}_{OPT} . The y -axis again represents the total cost result of each model as a ratio of the OPT total cost, or $y = C_{model}/C_{OPT} = C'_{model}$. Each point on Figure 3.3 represents the average of 5,000 simulation runs.

We can make two important conclusions from Figure 3.3. As traffic managers know less and less about the flight operators, the Parametric (PO) scheme's allocation efficiency degrades at an increasing rate with respect to σ , in comparison to the FSFA scheme. One can also observe that the PO solution is superior to the FSFA solution only when traffic managers have relatively good quality information about the operators (represented by smaller variance values, in this case, smaller than approximately $x = 0.18$). At values of x larger than 0.18, the FSFA model yields a more user cost efficient solution. This result is intuitive: when traffic managers use a more predictive flight cost model, a system-optimal allocation with some missing information is superior to a FSFA allocation with complete information. However, when the flight cost model's predictive

capabilities are lower, it may be more advisable to collect more information about the flight operators and use a FSFA allocation scheme. Identifying trade-off points such as this one is a core component of this research.




Figure 3.3 Demand Scenario 1 total cost results

Results such as those shown in Figure 3.3 could have some important policy implications. Say that traffic managers believe the predictive capability of their flight cost model is quite good in that it captures a large amount of information about the flights’ resource preferences, and other private route preferences are fairly small. In this case, traffic managers would be advised to adopt the parametric user input and system-optimal resource assignment of the Parametric (PO) allocation scheme. If traffic managers believe that their cost model’s predictive capabilities are weaker in that it captures less information about flight costs, then they would be advised to use the FSFA allocation method. If the traffic managers have no idea about the quality of their cost model’s predictive capabilities, they would be better advised to use the FSFA model to control the amount of “damage” that could be done by using the PO model at high values of σ . However, if we assume that the occurrence of information quality level is represented by some probability distribution, and traffic managers are risk neutral, we could calculate the expected cost of each strategy over all values of σ , and then choose a model.

The following table (Table 3.3) displays the results of a sensitivity analysis on the Scenario 1 demand parameters outlined above. The shaded cells contain the results of the difference between the Parametric (PO) and FSFA total flight cost results (again, as ratios of the OPT total cost results). Each row contains results from demand Scenario 1 with

flight population N identified at the very left of the table, over increasing standard deviation values. Each column represents results at the given standard deviation value listed. The darkest (red) cells at the left of the table represent those results where the Parametric total allocation costs are less than FSFA costs, at the value contained in the cell. So, for instance, when $N = 50$, at $\sigma/\hat{c}_{OPT} = 0$ the cell contains “-0.10”. This indicates that when $N = 50$, under zero error the Parametric total cost ratio is 0.10 less than that of FSFA. The opposite is true at the right side of the table (blue cells): for example, when $N = 100$ and $\sigma/\hat{c}_{OPT} = 0.4$, the value in the cell is “0.49”. This indicates that when standard deviation is 40% of the average deterministic flight cost under a system-optimal allocation, the Parametric total cost ratio is 0.49 greater than that of FSFA. Towards the middle of this table are the lightest cells, where the Parametric and FSFA results are similar (near the “crossing point” shown in Figure 3.3).

Table 3.3 Demand Scenario 1, Parametric vs. FSFA results

N	$D_0 < C_{AFP}$	$C'_{PO} - C'_{FSFA}$								
		σ/\hat{c}_{OPT}								
		0	0.05	0.10	0.15	0.20	0.25	0.30	0.35	0.40
50	y	-0.10	-0.09	-0.08	-0.06	-0.03	0.02	0.07	0.13	0.21
55	y	-0.09	-0.09	-0.07	-0.05	-0.01	0.03	0.09	0.16	0.24
60	y	-0.09	-0.09	-0.07	-0.04	0.00	0.04	0.10	0.17	0.26
65	y	-0.09	-0.08	-0.06	-0.04	0.00	0.05	0.12	0.19	0.28
70	y	-0.09	-0.08	-0.06	-0.03	0.01	0.06	0.13	0.20	0.29
75	n	-0.08	-0.08	-0.05	-0.02	0.02	0.07	0.14	0.22	0.31
80	n	-0.07	-0.06	-0.04	0.00	0.04	0.10	0.17	0.25	0.35
85	n	-0.06	-0.05	-0.03	0.01	0.06	0.12	0.19	0.28	0.39
90	n	-0.05	-0.04	-0.02	0.03	0.08	0.14	0.22	0.31	0.43
95	n	-0.05	-0.04	-0.01	0.03	0.09	0.16	0.24	0.34	0.46
100	n	-0.04	-0.03	0.00	0.05	0.10	0.17	0.26	0.36	0.49

The above table clearly shows that as the demand rate D_0 increases, the Parametric model’s performance compared to that of FSFA deteriorates, resulting in a lower σ threshold where the Parametric model cost results become higher than those of the FSFA model. As the demand rate increases, the average cost per flight increases for all models, but most rapidly for the Parametric model, followed by the OPT, and then FSFA. The difference between the Parametric and FSFA results is most likely explained in their flight assignment mechanisms. The Parametric model does not take into account the increasing stochastic term; as the average cost per flight \hat{c}_{OPT} increases with increasing D_0 , the stochastic term σ increases as well because it is calculated from \hat{c}_{OPT} . As a result, the Parametric allocation becomes more and more inefficient compared to FSFA because while the Parametric scheme does not assign resources using this increasing σ , the FSFA scheme does.

The results for demand Scenario 2, where the demand rate D_0 is constant over all flight populations, are contained in Table 3.4.

Table 3.4 Demand Scenario 2, Parametric vs. FSFA results

		$C'_{PO} - C'_{FSFA}$								
		σ/\hat{c}_{OPT}								
N	0	0.05	0.10	0.15	0.20	0.25	0.30	0.35	0.40	
50	-0.06	-0.04	-0.01	0.03	0.08	0.13	0.20	0.29	0.40	
55	-0.06	-0.04	-0.01	0.03	0.07	0.13	0.20	0.29	0.39	
60	-0.06	-0.04	-0.02	0.02	0.07	0.13	0.20	0.28	0.39	
65	-0.06	-0.04	-0.02	0.02	0.07	0.12	0.19	0.28	0.38	
70	-0.06	-0.04	-0.02	0.02	0.07	0.12	0.19	0.27	0.39	
75	-0.06	-0.04	-0.02	0.02	0.06	0.12	0.19	0.28	0.38	
80	-0.06	-0.04	-0.02	0.02	0.06	0.12	0.19	0.27	0.38	
85	-0.06	-0.04	-0.02	0.01	0.06	0.12	0.19	0.27	0.37	
90	-0.06	-0.04	-0.02	0.01	0.06	0.12	0.19	0.27	0.37	
95	-0.06	-0.05	-0.02	0.01	0.06	0.11	0.18	0.27	0.37	
100	-0.06	-0.05	-0.02	0.01	0.06	0.11	0.18	0.27	0.37	

The results indicate that the models' performances are relatively insensitive to increasing AFP durations and flight populations, given a constant demand of $D_0 = 80$ aircraft/hour and total AFP capacity of $C_{AFP} = 73.5$ aircraft/hour. The cost results of the Parametric and FSFA allocations are similar at σ/\hat{c}_{OPT} values somewhere between 0.10 and 0.15 (increasing as N increases). These results also hold true in situations where $D_0 < C_{AFP}$.

In demand Scenario 3, the total cost efficiency results for flight populations with higher maximum air-to-ground cost ratio values, α_{max} , are investigated. They are contained in Table 3.5 below.

Table 3.5 Demand Scenario 3, Parametric vs. FSFA results

				$C'_{PO} - C'_{FSFA}$									
				σ/\hat{c}_{OPT}									
N	T	D_0	α_{max}	0	0.05	0.10	0.15	0.20	0.25	0.30	0.35	0.40	
50	60	50	2.5	-0.10	-0.09	-0.08	-0.06	-0.03	0.02	0.07	0.13	0.21	
			3.5	-0.12	-0.12	-0.10	-0.08	-0.05	-0.01	0.04	0.10	0.17	
			5	-0.13	-0.13	-0.12	-0.10	-0.06	-0.03	0.02	0.08	0.14	
75	60	75	2.5	-0.08	-0.08	-0.05	-0.02	0.02	0.07	0.14	0.22	0.31	
			3.5	-0.10	-0.10	-0.08	-0.05	0.00	0.05	0.11	0.19	0.27	
			5	-0.12	-0.11	-0.09	-0.06	-0.02	0.03	0.09	0.16	0.24	
100	60	100	2.5	-0.04	-0.03	0.00	0.05	0.10	0.17	0.26	0.36	0.49	
			3.5	-0.07	-0.06	-0.03	0.01	0.06	0.13	0.20	0.29	0.40	
			5	-0.09	-0.08	-0.06	-0.02	0.03	0.08	0.15	0.23	0.33	

The results above indicate that as α_{max} increases, the FSFA model's performance compared to that of the Parametric model deteriorates, resulting in higher σ at which the FSFA scheme begins to outperform the Parametric scheme. Additionally, as α_{max} increases, the Parametric model results improve in comparison to the OPT results. We also conjecture that as flight α values increase, the significance of the en route portion of the flight cost model increases with respect to the other cost elements. Consequently, the Parametric model allocation results do not deteriorate as quickly compared to the OPT results with increasing α . The results of demand Scenario 3 suggest that the choice of a suitable allocation scheme is highly dependent on the distribution of α values over the flight population, and consequently, this choice may vary from one situation to another depending on the population of flights involved. $C'_{PO} = C'_{FSFA}$ does not change significantly as α_{max} increases.

Recall that the FSFA solutions are generated under the assumption that flights submit their preference information as soon as they know it, and this is in some independent and random order. Submission times may in fact be correlated to flight characteristics and dependent on operator beliefs regarding the time at which other flight operators will submit their preference information. Due to this competitive aspect, flights may even be inclined to submit their preference information before they are certain what it is. As a result, the assumption that flights submit their inputs randomly and independently is not likely to be accurate, and the FSFA scheme's total cost results may be better or worse than shown in these examples. We explore these FSFA flight preference submission ordering issues in Section 4.3.

3.5 Model Approximations

It was stated in Section 2.2 that a flight’s ground delay in an AFP must be non-negative, or $g_{r(n)}i_r(n) - g_{0,n} \geq 0$. We relax this constraint in order to construct analytic and quasi-analytic approximations to the Parametric and FSFA allocation schemes. We are interested in building these analytic approximations because they can be useful in better understanding the allocation schemes’ performance. For instance, these analytic formulations are used to investigate flights’ truth-telling behavior in Chapter 4. Because the error term is contained within the objective function in the OPT scheme, we do not construct an approximation to this scheme.

By relaxing the ground delay constraint to be non-binding, a flight can be assigned to a resource with a slot departure time that occurs earlier than its originally scheduled departure time prior to the AFP. Relaxing the constraint gives the FSFA and Parametric schemes some convenient properties that make them more analytically tractable. Also, the approximations are very good when pre-AFP scheduled demand is much higher than the total AFP capacity, or $D_0 \gg \sum_{r=1}^R S_r$. This situation could very well occur if capacity on the alternate AFP routes is low – they may have to accommodate their regular scheduled traffic demands in addition to the new AFP flights, or these alternate routes are weather-impacted like the nominal route. Moreover, even when $D_0 \approx \sum_{r=1}^R S_r$, numerical tests have shown that violations of the earliness constraint are small, and the approximations provide reasonable estimates of the true solutions.

For the approximations, we assume that the stochastic term of the flight cost function – representing private flight routing preferences – is distributed Gumbel. We assume this for the quasi-analytic FSFA formulation, which is described below.

3.5.1 First Submitted, First Assigned

Each time an operator submits their flight’s complete route preferences, the FSFA algorithm assigns it the best resources available at the time. This is identical to each flight choosing the minimum cost route and slot combination available when it submits its route preferences. If we assume that the stochastic term of the flight cost function is iid Gumbel with location parameter a and scale parameter b , the FSFA process can be represented using the expected received utility concept of the logit discrete choice model (Train, 2003). Each flight’s expected minimum cost and choice probabilities associated with a given set of alternatives can be determined. According to Domencich and McFadden (1975) and Ben-Akiva and Lerman (1994), the probability of agent n choosing an alternative r is:

$$P(V_{n,r}) = \frac{\exp(V_{n,r}/b)}{\sum_{j=1}^R \exp(V_{n,j}/b)} \quad (3.11)$$

where $V_{n,r}$ is the deterministic utility of option r to agent n . Assuming the flight cost function of Equation (2.2),

$$V_{n,r} = -C_{deterministic} = -(\alpha_n \cdot \rho_r + d_{n,r} - g_{0,n}) \quad (3.12)$$

Let us say that $E[c_n]$ represents the additional expected cost for flight n due to the AFP, based on the assignment options available to this flight because of the AFP. If $E[c_n^{AFP}]$ is the expected cost of the set of AFP resources available to n , and $E[c_n^0]$ is the expected cost of n 's original resource prior to AFP inception, then the difference between the two is $E[c_n]$. Given that the stochastic term is Gumbel distributed, we can represent $E[c_n]$ using the well-known expression for the change in consumer surplus that results from a change in alternatives and/or choice set (Train, 2003):

$$\begin{aligned} E[c_n] &= E[c_n^{AFP}] - E[c_n^0] \\ &= b \cdot \ln \left[\sum_{r=1}^R \exp\left(\frac{V_{n,r}}{b} + a\right) \right] - b \cdot \ln \left[\exp\left(\frac{V_n^0}{b} + a\right) \right] \end{aligned} \quad (3.13)$$

where a and b are parameters of the Gumbel distribution, and V_n^0 is the deterministic utility of n 's original resource prior to the AFP. As stated previously, we are not concerned with a flight's cost under normal operating conditions but rather the flight's additional costs due to the AFP. As a result, we set $V_n^0 = 0, \forall n, r$, and we rewrite (3.13) as:

$$E[c_n] = b \cdot \ln \left[\sum_{r=1}^R \exp\left(\frac{V_{n,r}}{b}\right) \right] \quad (3.14)$$

The location parameter a cancels out of the expression due to its inclusion in both the AFP cost and in the original cost. We now describe a recursive, quasi-analytic procedure by which the expected minimum AFP cost is found for all flights. This procedure yields approximate solutions because we calculate the expected ground delay on a given route r (based on the probability of the previous flight $n - 1$ having chosen to take route r or not) to calculate the expected utility of flight n taking r , which in turn is used to calculate the probability of n taking r . To check the use of this assumption, we compared the results of this procedure against solutions simulated using a minimum cost slot assignment algorithm. It was found that our quasi-analytic procedure consistently overestimates total AFP cost in the order of 0.3-1%. In addition, when we compared the average flight cost by submission order from the quasi-analytic versus simulation methods, the two results were extremely similar in both magnitude and trend. It was therefore concluded that the quasi-analytic procedure leads to a satisfactory approximation of the true results. The steps of the procedure are described below.

- 1) Assign an α value to each flight. Arrange flights in their order of input submission, or $n = 1, \dots, N$. At this point, we cannot say what this order looks like – it could be random, or correlated to α, n , total distance from origin to destination airport, aircraft size, etc.

2) For flight $n = 1$, we calculate $V_{1,r}$, $P(V_{1,r})$, and $E[c_1]$ using (3.12), (3.11), and (3.14) respectively, for each route r . We calculate $V_{1,r}$ under the assumption that $n = 1$ will take the first slot on the route it chooses.

3) For $n = 2, 3, \dots, N$:

- a. Determine the expected ground delay $E[d_{n,r}]$ on each route r for flight n . $E[d_{n,r}]$ is calculated based on the conditional probability that the previous flight $(n - 1)$ took r . Event “ $(n - 1)$ took route r ” is represented by B ; event “ $(n - 1)$ did not take route r ” is represented by B^c . $E[d_{n,r}]$ then becomes:

$$\begin{aligned} E[d_{n,r}] &= E[d_{n,r}|B] \cdot P(B) + E[d_{n,r}|B^c] \cdot P(B^c) \\ &= (E[d_{n-1,r}] + g_r) \cdot P(B) + E[d_{n-1,r}] \cdot (1 - P(B)) \end{aligned} \quad (3.15)$$

$P(B)$ is the probability of agent $n - 1$ taking route r , and was calculated for flight 1 in step 2 using (3.11), and flights 2,3, ... N in step 3(d) below.

- b. Find the expected utility of each alternative route r for flight n , expressed as $E[V_{n,r}] = -(\alpha_n \rho_r + E[d_{n,r}] - g_{0,n})$.
- c. Find the expected cost $E[c_n]$ using (3.14) and $V_{n,r} = E[V_{n,r}]$ found in the previous step 3(b). This calculation step leads to an approximation of the true solution because we use expected ground delay (based on the probability of the previous flight $n - 1$ choosing to take route r or not) to find the utility of flight n taking r , which in turn is used to find the probability of n taking r .
- d. Find n 's route choice probabilities $P(V_{n,r})$ using (3.11) and $V_{n,r} = E[V_{n,r}]$. This calculation step leads to an approximation of the true solution.
- e. Repeat (a) through (d) until $n = N$.

4) Find $\sum_{n=1}^N E[c_n]$.

We can perform the above calculations for different values of the Gumbel scale parameter b , where increasing b reflects increasing variance of the Gumbel-distributed stochastic term $\varepsilon_{n,r}$. Both Gumbel distribution parameters a and b are generated as functions of σ such that $b = \sqrt{6}\sigma/\pi$ and $a = 0.5772b$.

3.5.2 Parametric

To obtain analytic solutions where the decision variable is the number of flights to assign to each route r , or X_r , we solve the deterministic part of Equation (2.3) using properties that result from the non-binding schedule constraint. With the schedule constraint relaxed, flight n 's resource allocation is not affected by its original scheduled departure time $g_{0,n}$. Therefore ground delay does not influence a flight's allocation, and only the additional en route time cost influences which resource each flight receives in the allocation. It follows that flights with the highest air-to-ground ratio (α) values

should be assigned to the routes with lowest en route times, and vice versa, if the minimum cost solution is to be obtained (see Appendix B.1). We achieve this by ordering flights by increasing α values, and routes by decreasing en route times where $\rho_1 > \rho_2 > \dots > \rho_R$. If there are two routes 1 and 2 where $\rho_1 > \rho_2$, flights with lower α are assigned to Route 1 and those with higher α are assigned to Route 2. If α_1^{max} is the largest air-to-ground ratio belonging to a flight assigned to Route 1, and α_2^{min} is the smallest belonging to a flight assigned to Route 2, then $\alpha_1^{max} < \alpha_2^{min}$. This is illustrated in Figure 3.4, under the assumption that α values are uniformly distributed over the population of flights.

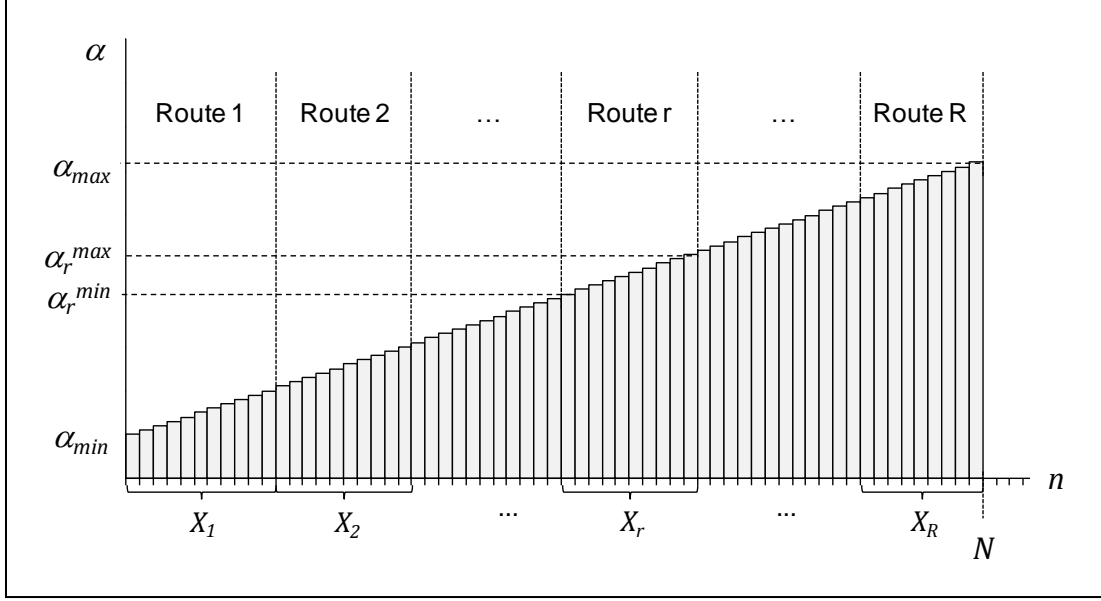


Figure 3.4 Assignment of flights to routes by increasing α values.

Out of the population of flights assigned to route r , the smallest air-to-ground cost ratio value seen within these flights is α_r^{min} , and the largest α_r^{max} . X_r is the number of flights on a route, and is also expressed as the number of flights that have air-to-ground cost ratio values between α_r^{min} and α_r^{max} inclusively. If X_r flights are ordered by increasing α values between α_r^{min} and α_r^{max} , then the total en route cost for flights on route r is expressed as:

$$C_r^{air} = \bar{\alpha}_r \cdot X_r \cdot \rho_r \quad (3.16a)$$

where

$$\bar{\alpha}_r = 0.5 \cdot (\alpha_r^{min} + \alpha_r^{max}) \quad (3.16b)$$

Based on our assumption that α has a uniform distribution with a minimum value and maximum value, using the approximation of Equation (2.4), we have

$$\alpha_r^{max} = \alpha_r^{min} + \theta \cdot (X_r - 1) \quad (3.16c)$$

where $\theta = (\alpha_{max} - \alpha_{min})/N$. Furthermore, if routes are ordered such that $\rho_1 > \rho_2 > \dots > \rho_{R-1} > \rho_R$:

$$\alpha_r^{min} = \alpha_{min} + \theta \sum_{j=1}^r X_{j-1} + \theta \quad (3.16d)$$

with $X_0 = 0$, we substitute Equations (3.16c) and (3.16d) into (3.16b) to obtain:

$$\bar{\alpha}_r = \alpha_{min} + \theta \left(\sum_{j=1}^r X_{j-1} \right) + \frac{\theta}{2} (X_r + 1) \quad (3.16e)$$

We can now rewrite Equation (2.3) for the objective function of the Parametric resource allocation scheme below, to obtain $X_r \forall r$.

Decision variables: X_r , the total number of flights assigned to route $r \in R$; $X_r \in \{0, N\}, \forall r$.

Objective function:

$$\min_{X_1, \dots, X_R} \hat{C} = \sum_{r=1}^R \left(\bar{\alpha}_r \cdot X_r \cdot \rho_r + \frac{1}{2} g_r \cdot X_r \cdot (X_r + 1) \right) - \sum_{n=1}^N g_{0,n} \quad (3.17)$$

Constraints: $\sum_{r=1}^R X_r = N$; $X_r \geq 0, \forall r$.

The first set of terms in the objective function represents the average cost of additional en route time for a flight to travel on route r , while the second and third sets of terms represents the total ground delay. The first constraint ensures that the number of flights assigned to the available routes sums to the total number of AFP flights (N), and the second constraint ensures that all route counts are non-negative. The objective function of (3.17) was checked for convexity (Appendix B.2). X_r is integer, but this is relaxed in order to obtain the analytic solution. Even if solutions are not integer, rounding (to preserve N) will still produce acceptable solutions (Richetta & Odoni, 1993). In a practical sense, the headways on each route should be designed to include some buffer space that would allow for slightly more aircraft than the capacity permits. As a result, if by rounding up X_r the route capacities were exceeded occasionally, the result would not be catastrophic.

If it is the case that $X_r = 0$ or $X_r = N$ for any r , then interior solutions to the objective function of (3.17) do not exist, and solutions lie at the boundaries. In these cases, $X_n^* = 0$ and $X_n^* = N$, respectively.

If $R = 2$, the optimal solution when $0 \leq X_1 \leq N$ is:

$$X_1^* = \frac{\alpha_{min}(\rho_2 - \rho_1) + 0.5(\theta\rho_2 + g_2 - \theta\rho_1 - g_1) + g_2N}{\theta(\rho_1 - \rho_2) + g_1 + g_2}, \quad (3.18a)$$

$$X_2^* = N - X_1^*$$

$$\hat{C}^* = \frac{\frac{1}{2}g_2N \cdot (N + 1) + \bar{\alpha}_2\rho_2N - \sum_{n=1}^N g_{0,n}}{\frac{(\bar{\alpha}_2\rho_2 - \bar{\alpha}_1\rho_1 + \frac{1}{2}(g_2 - g_1) + g_2N)^2}{2 \cdot (g_1 + g_2)}} \quad (3.18b)$$

where all terms are as described above.

Once we have the results of the Parametric resource allocation we can again find the expected “true” cost of the allocation $E[C]$ according to Equation (3.4). We now assume that the stochastic terms $\varepsilon_{n,r(n)}$ are iid Gumbel with parameters (a, b) , where the parameters are generated as functions of σ , or $b = \sqrt{6}\sigma/\pi$ and $a = 0.5772b$. When a and b are so defined, according to the central limit theorem their sum is asymptotically distributed normal with mean $\mu = 0$ and standard deviation σ , such that $E[C] = \hat{C}$. The values of the stochastic term can also be randomly drawn to obtain numerical solutions.

3.6 Concluding Remarks and Future Work

In this chapter we introduced and evaluated the performances of three resource allocation schemes that feature different preference information required from users and resource rationing mechanisms. They are designed to offer users a resource allocation method that is commensurate to the quality of information requested of them, in order to maximize the likelihood of user participation. The results confirm that when FAA traffic managers do not have a well-specified flight cost model that includes good quality information about the flights’ route preferences, using a FSFA resource allocation scheme is more cost effective. Conversely, when traffic managers do have a well-specified flight cost model, they can achieve better total cost efficiency by implementing the Parametric resource allocation scheme.

We have stated that these resource allocation schemes are designed to increase the likelihood that flight operators submit the preference inputs that are asked of them, but we have not systematically assessed whether they do. And if flight operators do indeed participate, what is the nature of their participation and the quality of their offered inputs? The results of an initial investigation into these questions are discussed in Chapter 4. Also, there are many other allocation schemes that could be considered, and we begin to explore this topic in Chapter 5.

The performances of the resource allocation schemes of this chapter may be very sensitive to the formulation of the flight cost function. At the end of the previous chapter we discussed why flight cost is not likely to be a linear function of delay as assumed by the cost function used in this research. Through empirical analyses, Xiong demonstrated in her dissertation (2010) that the cost of GDP-assigned Initial Delay (GID) can be well-represented using a piece-wise linear model. Her results show that the first 15 minutes of GID are insignificant, due to the way that airlines are required to report on their performance statistics to the FAA. GID between 45 to 90 minutes has the highest per

minute cost, and decreases after 90 minutes. She also writes that, after discussions with several airlines, the GID cost curve should actually flatten out at high GIDs, because at these levels of delay flights have already been cancelled. Hoffman et al. (2004) explicitly considered *groundcost*, *aircost*, as well as a third term – *connectcost* – in their flight cost function. *connectcost* is the cost of missing downstream connections by crew and passengers to other flights, etc. It represents both the size of a discontinuity in the flight cost function, as well as its slope (with respect to ground delay) after the discontinuity (should it be different from the slope prior to the discontinuity). The location of the discontinuity represents the ground delay after which time the flight has missed a connection, and the size represents the cost of missing that said connection.

Based on the above, we might consider two modifications to the flight cost model introduced in this chapter. We might first consider a specification such as the following:

$$c_{n,r} = [\min(m, M)]^{1+\varphi} + \varepsilon_{n,r} \quad (3.19)$$

where

- $m = c_{n,r}^{en\ route} + c_{n,r}^{grdelay}$, or the deterministic cost to flight n taking route r ;
- M is a threshold at which point the deterministic flight costs are so high that a flight will be canceled, and
- $\varphi \geq 0$.

The parameter φ is zero if total flight cost is to be a linear function of en route and ground delay costs. It is likely that $\varphi_n \neq \varphi \forall n$, and therefore φ_n may be another parameter input that could be asked of flight operators. We can consider a specification like the one above and explore its properties. An alternative approach is to approximate this missed connections concept in our flight cost function with a logistics curve. This will require a significant change to the overall structure of our cost model. Also, in order to use these formulations we must give more thought to the currency used to assess cost. For instance, what does one minute of ground delay refer to? A minute of ground delay when a flight is delayed less than 15 minutes, or when it is delayed more than an hour?

The above are two alternate flight cost function specifications we can adopt to improve its explanatory power. By using a flight cost structure that captures as much information as we know about the flight cost/delay relationship, we can in turn improve the explanatory power of our evaluation of candidate allocation schemes.

4. Airline Behavior under Competition

4.1 Introduction

The commercial airline industry is highly competitive, which leads to two main behaviors by airlines in the competition for en route resources. Firstly, airlines are reluctant to share detailed operating information about themselves with traffic managers – most aspects of operational and strategic knowledge are kept proprietary unless they receive an equal or greater (perceived) benefit in exchange. Secondly, airlines will game the system if and when they can obtain an advantage by doing so. These behaviors are a primary concern in air traffic flow management programs, where a lack of truthful information provision can undermine fairness and efficiency goals. As a result, resource allocation mechanisms should not only be designed to promote efficiency and fairness, but to incentivize airlines to provide traffic managers with truthful information.

In Chapter 3 we assumed that all flight operators provide truthful information in the resource allocation schemes introduced. In this chapter, we discuss how gaming and competition could affect flight operators' behavior, and thus, the outcomes of these schemes. We first demonstrate that in the OPT and Parametric allocation schemes, the system-optimal resource rationing mechanism encourages homogeneous airlines to provide untruthfully high preference inputs. In FSFA, airlines have an additional degree of freedom in their provision of preference information: the time at which they submit this information. Because airlines receive their highest utility option available at the time of their submission in the FSFA scheme, we conjecture that they do not have reason to submit untruthfully. The incentive to submit truthful information may be a critical advantage of the FSFA scheme.

Some numerical examples of flight utility with respect to rank in the FSFA submission order are presented. A formulation that estimates the disutility of early submission due to uncertainty regarding internal operations and NAS conditions follows. We then use the results from these utility investigations in the setup of a competition model of how airlines would behave under the FSFA allocation scheme. We illustrate the behavior of the resulting equilibrium submission time function under several numerical examples. The chapter closes with a brief discussion of other gaming behaviors and ideas for further investigation.

4.2 Truth-telling Properties of System-Optimal Resource Allocation

In both the OPT and Parametric schemes, traffic managers minimize the total user costs that are known to them, resulting in a system-optimal resource assignment. In a system-optimal assignment, the resources allocated to flights are typically of varying cost, such that some flights end up with more desirable resources than others. As a result, over time and many AFPs, flight operators will attempt to minimize the expected cost of their own allocations in our user input-based schemes by strategically altering their input behavior. We anticipate that this rational behavior will be exhibited by all flight operators, and thus lead to equilibrium conditions where no flight can expect to lower their expected assignment cost by changing their strategy.

According to the flight cost model and the two system-optimal strategies (OPT and Parametric) introduced in Chapters 2 and 3, flight operators can attempt to control their resource allocations through the information they provide to the traffic managers. We must determine whether flight operators are incentivized to submit untruthful preference information in order to reduce their expected assignment costs. We can take a first look at this issue by exploiting the properties of a basic traffic assignment analysis (Sheffi, 1985). We present a simple AFP setup with identical flights and two available routes ($R = 2$). Each route has slots spaced at identical headways (g), and route 1's en route time is greater than that of route 2 (such that $\rho_1 > \rho_2$). All N AFP flights have identical original scheduled departure times ($g_0 \approx 0 \forall n$), air-to-ground cost ratios ($\alpha^T \forall n$), and no additional unknown route preferences (such that $\varepsilon_{n,r} = 0 \forall n, r$). The results of applying the Parametric and OPT schemes are therefore identical; the system-optimal allocation using truthful air-to-ground cost ratios will result in an assignment of X_1^* flights to route 1, and $X_2^* = N - X_1^*$ flights to route 2. Equation (3.18a) gives us the values of X_1^* and X_2^* :

$$X_1^* = \frac{\alpha^T \cdot (\rho_2 - \rho_1) + gN}{2g}; X_2^* = N - X_1^* = \frac{\alpha^T \cdot (\rho_1 - \rho_2) + gN}{2g} \quad (4.1)$$

where α^T is the truthful air-to-ground ratio value for all flights, and all other parameters are as described in Chapter 2.

Figure 4.1 is a graphical representation of our flights' expected route assignment costs. The lightweight dotted lines represent the (truthful) average cost that flights can expect to incur by being assigned to a route, as a function of the total number of flights assigned to that route. Although in theory each flight is assigned to a distinct slot, after a flight is assigned to a given route, slot assignments on that route are completely arbitrary in this case. As a result, a flight can only know the expected cost of being assigned to a certain route as a function of the total flights assigned to that route. Similarly, each route's expected marginal cost curves with truthful air-to-ground cost ratios (α^T) are represented by the light solid lines. The system-optimal assignment with α^T (Equation (4.1)) is found from the point where the truthful expected marginal cost curves of routes 1 and 2 are identical, represented by point A in Figure 4.1. Since $\rho_1 > \rho_2$, it follows that $X_1^* < X_2^*$.

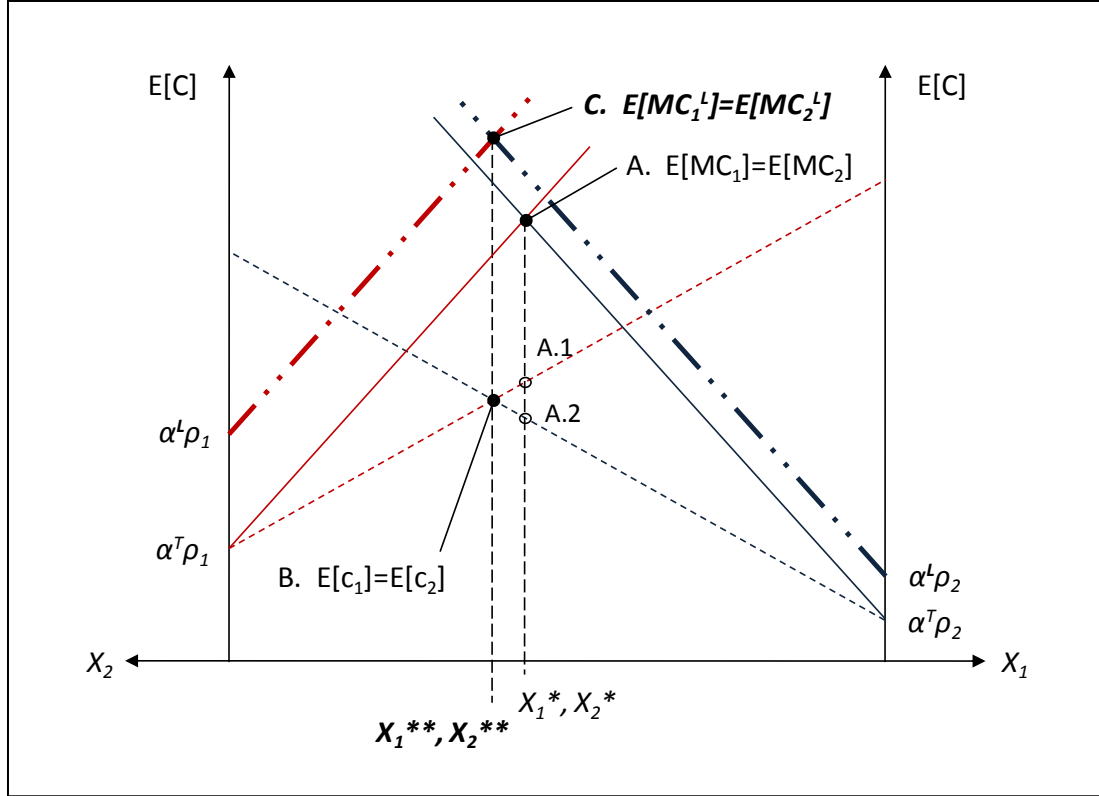


Figure 4.1 Untruthful versus truthful system-optimal solutions

Also, if one draws a vertical line at point A, one can observe that the expected cost of being on route 1 (point A.1 in the figure) is higher than the expected cost of being on route 2 at the system-optimal assignment (A.2). As a result, all flights will attempt to maximize the probability of being assigned to route 2 instead of route 1 using whatever strategic handle is available to them, which in this case are their submitted air-to-ground cost ratios. If a flight operator should submit a truthful air-to-ground cost ratio α^T , the probability of their flight being assigned to the lower cost route (route 2) is X_2^*/N . Instead, imagine they submit some untruthful air-to-ground cost ratio $\alpha^L > \alpha^T$. By submitting α^L , their flight will be assigned to the lower cost route with probability 1 if all others submit truthfully. There is clearly an incentive for flight operators to submit untruthfully high α values to traffic managers. The question is, is there a unique untruthful air-to-ground cost ratio submitted by flight operators, and what is its value at equilibrium?

Given that all flights are identical and aim to minimize their own costs, we imagine that through their submissions they will push the assignment towards a user equilibrium (UE) based on their actual cost ratios. At this UE assignment with α^T , the true expected costs of a flight being assigned to either route are equal, or $E[c_1] = E[c_2]$. At this assignment, flights have no incentive to change their inputs as they cannot lower their expected assignment costs by doing so. This user equilibrium is represented by the following expression:

$$E[c_1] = E[c_2] \Rightarrow \alpha^T \rho_1 + 0.5gX_1^{**} = \alpha^T \rho_2 + 0.5gX_2^{**} \quad (4.2)$$

X_1^{**} and X_2^{**} are the user equilibrium flight assignments to routes 1 and 2, respectively, with α^T (Point B in Figure 4.1). We can imagine that over many occurrences of this AFP, flights will submit α^L values that yield a system-optimal assignment with X_1^{**} and X_2^{**} – the assignment at the true (with α^T) user equilibrium of Equation (4.2). This system-optimal assignment is found when:

$$E[MC_1^L] = E[MC_2^L] \Rightarrow \alpha^L \rho_1 + gX_1^{**} = \alpha^L \rho_2 + gX_2^{**} \quad (4.3)$$

The value of α^L can be determined by drawing a vertical line through point B, and then finding the point along this vertical line where the marginal cost curves of Equation (4.3) intersect. It is identified as point C in Figure 4.1. We can now find an expression for α^L , the untruthful air-to-ground cost ratio submitted by flights to obtain a truthful user equilibrium solution, by solving (4.2) and (4.3):

$$\alpha^L = 2\alpha^T \quad (4.4)$$

The equilibrium (where all flights have identical expected assignment costs) is reached when all flights submit an air-to-ground cost ratio α^L that is twice that of their truthful cost ratio α^T . Given the assumption that all flight operators are rational – they make choices to maximize their own payoffs – we state the following.

Proposition 4.1. *Under the Parametric scheme and a special case of the OPT scheme (where $\varepsilon_{n,r} = 0 \forall n, r$), identical flights will submit inputs that consist of untruthful cost ratios α^L to traffic managers. These untruthful values are twice as large as the truthful values in cases with two- and three- routes.*

Proof: see above for two-route case and Appendix C.1 for three-route case.

This simple analysis has demonstrated that traffic managers will not be able to sustain a true system-optimal allocation through the Parametric and OPT schemes, as flight operators will create a user equilibrium by altering their inputs. A similar result may be true in the case of heterogeneous flights. It has been shown in past research that unique user equilibria do exist under certain conditions for heterogeneous commuters in both network and single bottleneck models (Newell, 1987) (Daganzo, 1983) (Konishi, 2004). However, there may be instabilities due to other gaming behaviors and flight characteristics due to heterogeneity, and we must perform a quantitative analysis before drawing conclusions. We leave this to future research.

4.3 Competition in FSFA

The Chapter 3 numerical examples of the FSFA allocation scheme assume that complete route preference inputs are submitted by flight operators in an arbitrary, random order. However, another level of competition is encouraged in FSFA through the additional degree of freedom airlines have in submitting their route preference information. A

flight's allocation cost may be highly dependent on its place in the preference submission order. As a result, the assumption that submission order is completely independent of flight characteristics and the AFP itself, as well as random, is likely invalid. Order could be correlated to flights' scheduled time of departure, their air-to-ground cost ratio (α_n), and/or their special route preferences ($\varepsilon_{n,r}$). We might expect earlier scheduled flights and flights with higher α_n to submit their information earlier. We might expect flights that prefer particular routes to submit earlier as well. Also, depending on the nature of competition and the conditions in an AFP, the submission process may be beneficial for ATFM. For instance, if it seems that flights are inclined to submit at the very start of the AFP planning period (which occurs several hours prior to the start of the AFP itself), traffic managers can expect to complete the majority of their AFP planning early to facilitate better coordination with other ATFM programs in the NAS. We begin investigating the effects of competition on the FSFA submission process in this part of the research.

The time that the operator of a flight n submits their complete route preference values should be the resulting balance of two opposing objectives. In order to maximize the likelihood of obtaining a desired resource before their competitors do, flight operators will want to submit their preferences early during the planning period. However, an airline's internal operations and information, as well as NAS operating conditions, can be highly volatile and change rapidly prior to a flight's scheduled and/or actual departure time. Preference information thus becomes more accurate as it gets closer to its departure time. In this regard, it is more beneficial for operators to submit as late as possible. As a result, the utility of a flight's resource assignment will depend on the resources available at the time of their preference input submission, as well as the accuracy of the preference inputs submitted to obtain a resource.

We address the following question: how does uncertainty influence an operator's decision to submit its flight's route preference information later than it would under perfect information? We first explore the characteristics of expected flight utility based on a flight's rank in the FSFA submission process. We then investigate the cost of submitting inputs under uncertainty about operating conditions during the AFP planning period. Finally, we propose a competitive model to gain some insights into airline input submission behavior during the planning period. The FSFA process differs from the Parametric and OPT schemes in that airlines can also influence their allocations through the timing of their preference submissions. However, in the FSFA scheme, airlines do not appear to have reason for submitting untruthful preferences (except to perhaps "hurt" other airlines, which we do not consider in this analysis), because the goal is to obtain their highest utility option available at the time of submission. As a result, we conjecture that gaming behavior in the FSFA scheme is manifest in the time airlines submit, rather than in the information they provide. The incentive to offer truthful information may be a critical advantage of FSFA over the system-optimal resource allocation schemes.

4.3.1 Utility by Rank in Submission Order

The cost of a flight's received resource allocation can be dependent on that flight's rank in the route preference submissions order. Here we explore flight utility characteristics as a function of submission rank. We designate x_n to be flight n 's rank in the preference submission ordering; $x_n = 10$ means flight n was tenth to submit preference inputs. The utility of flight n being x th to submit is represented by $U(x_n)$. We do not know the functional form for $U(x)$.

Figure 4.2 is a plot of $U_{avg}(x)$ for a number of numerical examples, including those presented in Chapter 3. Each curve in Figure 4.2 represents the average utility, over 5,000 simulation runs, of having been x th to submit in a given numerical example. Note that there are a total of $N!$ possible orderings of flight submissions in each example.




Figure 4.2 Sample average utilities by submission rank

For the numerical examples shown, $U_{avg}(x)$ is non-linear and decreasing with respect to x . The curves vary significantly depending on the supply and demand characteristics of an AFP. For instance, in some examples the average utility of being first is at or near zero while in one example it is approximately -20 ground delay minutes. The utility of submitting in x th place can also vary significantly from one instance (simulation run) to another: again, depending on the particular characteristics of the AFP. In these situations, $U_{avg}(x)$ is not a good representation of $U(x)$. The variance of $U(x)$ is low under situations where flights are more similar to one another in the cost model

parameters. When there are k slots that are of the highest deterministic utility to all flights in the population, and the variance of the stochastic route preference term is small, $U(1), U(2), \dots, U(k)$ have little to no variance. This is true for small k in several of the examples above.

A flight’s received resource allocation is also dependent on how that flight values the offered resources in comparison to how others value those resources. For instance, if there is a resource of high utility to two flights, the flight that submits first will “win” the high utility resource. Alternatively, two flights might also assign their highest utilities to resource sets that do not intersect, and each flight may win their desired resource regardless of their place in the submissions order.

Figure 4.3 is a plot of $U_{avg}(x)$ curves that were constructed based on a numerical example of Chapter 3 (AFP demand Scenario 1, with $N = 75$). All demand and supply characteristics of the scenario are held constant, except the variance of the stochastic error term (σ). Each curve represents the scenario constructed with a given σ value, which are indicated in the legend. The values in the legend are σ/\hat{c}_{OPT} , which is also the x -axis of Figure 3.3.




Figure 4.3 Average utility by submission rank, range of σ (Demand Scenario 1)

As σ increases, $U_{avg}(x)$ increases (the variance of $U(x)$ does as well, although this is not shown in the figure). The inflection points in both figures represent the submission rank x where slot allocations typically transition from those on one route to the next. At lower σ/\hat{c}_{OPT} values, there is less variability regarding the x values at which these transition

points occur, hence the more pronounced inflections in the curves. The utility curves that correspond to larger σ/\hat{c}_{OPT} values are smoother.

4.3.2 Utility by Submission Time

Imagine that the operator of flight n is the first to submit their preferences during the AFP planning period when the FSFA allocation scheme is used. Flight n , therefore, has all AFP slots ($1, 2, \dots, s \in S$) available to it. If n had perfect information about NAS conditions during the planning period, it would also know the true utilities (relative to flying the nominal route with no ground delay, pre-AFP) of the AFP resources to itself, or $V_{n,s} \forall s \in S$. We have specified $V_{n,s}$ using the flight cost model of Equation (2.2):

$$V_{n,s} = -(\alpha_n \cdot \rho_{r(s)} + d_s - g_{0,n} + \varepsilon_{n,r(s)}) \quad (4.5)$$

Under perfect information, it follows that n will also know which resource option is of true highest utility to itself.

$$V_n^* = \max(V_{n,1}, V_{n,2}, \dots, V_{n,s}, \dots, V_{n,S}) \quad (4.6)$$

However, n is not likely to have perfect information throughout the planning period, as NAS conditions can change rapidly in ways that cannot be predicted by airlines or traffic managers. These conditions include weather and traffic, as well as how FAA traffic managers react to these conditions. It is also true that airline internal operational situations can change rapidly, in ways that are dependent as well as independent on NAS conditions. Therefore, during the planning period, n only observes that the value of resource s to itself is $U_{n,s}$, rather than its true value, $V_{n,s}$. Flight n 's knowledge of NAS conditions over the planning period is likely to evolve, and we specify that $U_{n,s}$ is:

$$U_{n,s}(t) = V_{n,s} + \gamma_{n,r(s)}(t) \quad (4.7)$$

The stochastic term $\gamma_{n,r(s)}(t)$ represents n 's imprecise knowledge about conditions at t , and we assume it is distributed Gumbel. Remember n can only observe $U_{n,s}$, not $V_{n,s}$ nor $\gamma_{n,r(s)}$. Given that flight n only observes $U_{n,s} \forall s$ under conditions of uncertainty at time t , it can expect to obtain the following utility from the set of AFP resources available:

$$E[U_n(t)] = \sum_{s \in S} p_{n,s}(t) \cdot V_{n,s} \quad (4.8)$$

where $p_{n,s}(t)$ is the probability of choosing some resource s at time t . It can be expressed as $p_{n,s}(t) = \exp(V_{n,s}/v(t)) / \sum_{s \in S} \exp(V_{n,s}/v(t))$, where $v(t)$ is the scale parameter of $\gamma_{n,r(s)}(t)$, and indicates how much $U_{n,s}$ may differ from $V_{n,s}$. A similar probability expression was previously introduced in Equation (3.11).

Flight n 's uncertainty is likely to be greatest at the beginning of the planning period ($t = 0$) and decrease as it progresses to T , the end of the planning period and

beginning of the AFP itself. We capture this idea by assuming that the variance of $\gamma_{n,r(s)}(t)$, and therefore its scale parameter $\nu(t)$, decreases linearly with respect to t :

$$\nu(t) = p \cdot (T - t) \quad (4.9)$$

where p is a parameter that captures the overall uncertainty of conditions for the particular AFP in question. Equation (4.9) is constructed with the assumption that flight n 's information is perfect by T , such that $\nu(T) = 0$. Therefore, as the planning period approaches T , the probability of n choosing the resource of true highest utility to itself – V_n^* – approaches one. Prior to T , flight n believes that resource s is valued at $U_{n,s}(t)$, and with this information can only expect to gain $E[U_n(t)]$ with its choice. Flight n 's loss in true utility resulting from its choice under uncertain conditions at t can be expressed as:

$$L_n(t) = V_n^* - E[U_n(t)] \quad (4.10a)$$

Recall that $E[U_n(t)] \rightarrow V_n^*$ as $t \rightarrow T$, and therefore $L_n(t) \rightarrow 0$. Note that the loss function in (4.10a) assumes that all slots are available to all flights at anytime.

If $\nu(t)$ is very high, $\gamma_{n,r(s)}(t)$ is highly variable such that $U_{n,s}(t)$ is a poor reflection of its true underlying utility $V_{n,s}$, and the probabilities of choosing one resource over another become identical. As a result, at large $\nu(t)$, $E[U_n(t)]$ approaches \bar{V}_n , or the average deterministic utility to flight n of all the resources available. The maximum value that $L_n(t)$ can therefore take is $L_n^{max} = V_n^* - \bar{V}_n$. We can represent the utility loss function as a proportion of the maximum loss possible:

$$l_n(t) = \frac{L_n(t)}{L_n^{max}}, \quad l_n(t) \in (0,1] \quad (4.10b)$$

The shape of $L_n(t)$, and therefore $l_n(t)$, is highly dependent on the specification of $\nu(t)$, flight cost parameters (α_n , $g_{0,n}$, and $\varepsilon_{n,r(s)}$) and the set of available resources S . It may be convex within the planning period depending on how $\nu(t)$ is specified, or it could have an inflection point after which the function becomes concave. If only one resource is available to n , $l_n(t) = 0 \forall t$. However, if there are several resources of differing utilities, $l(t)$ is strictly decreasing and differentiable.

Figure 4.4 displays the $l_n(t)$ function for six examples.

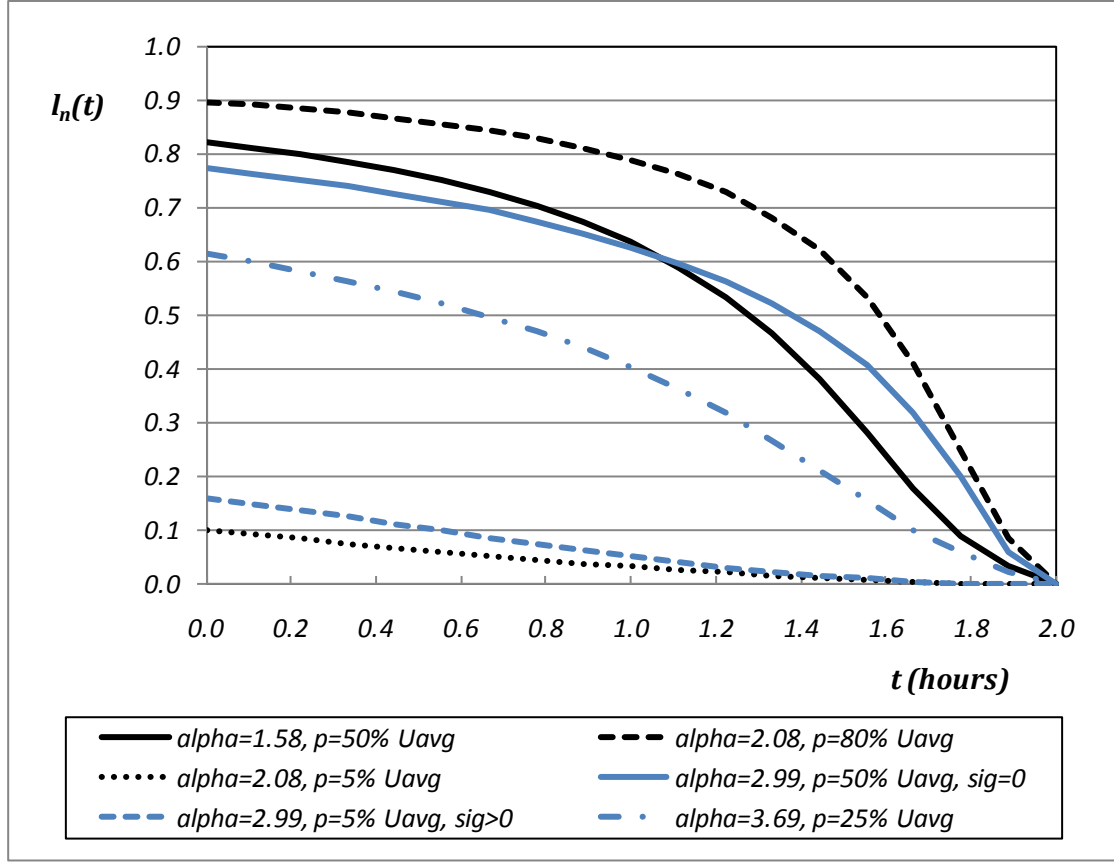


Figure 4.4 Utility loss through imperfect information

The first three examples listed in the legend above ($\alpha = 1.58, 2.08$ & 2.08) are generated for the choice set of a numerical example with $N = 40$ and $\varepsilon_{n,r} = 0 \forall n, r$. The next three examples ($\alpha = 2.99, 2.99$ & 3.69) are for the choice set of a numerical example with $N = 75$. The fourth example ($\alpha = 2.99, p = 50\% U_{avg}, \sigma = 0$) has $\varepsilon_{n,r} = 0 \forall n, r$ and the fifth and sixth examples (last two listed in the legend) have $\varepsilon_{n,r} \sim N(0, \sigma)$, where $\sigma > 0$. It can be observed that as α_n and ν (and therefore, p) increase respectively, the values and shape of the loss function within the planning period (from $t = 0 \rightarrow 2$ hours, in these examples) change significantly. The two flat (dashed) curves near the x -axis have smaller ν over the planning period. If ν were larger, or if the planning period were longer, the shape of these curves would approach that of the higher curves. In essence, these flatter loss curves correspond to the other curve values in a small time interval at high t values (near T). As the range of ν over the planning period increases, the loss curve approaches a shape where its values are near one on the y -axis for a majority of the planning period, and then decreases very quickly to zero as $t \rightarrow T$.

This section has introduced a functional form for a flight's loss in utility due to early submission in an AFP, under uncertainty regarding NAS conditions. This loss function is used in the competition model that follows.

4.3.3 Competitive Model

4.3.3.1 Assumptions and Payoff Function

The analysis of the flights' complete route preference input submission times in the FSFA competition requires some assumptions about the participating airlines and the competition process itself. Firstly, we continue to ignore the potential effects of correlation amongst flights owned by a single airline, such that each flight is considered a single non-cooperative player in the competition. This is a significant assumption in a model that is constructed for the purpose of understanding airline competition, but relaxing it will add a dimension of complexity which we would like not to consider for this first analysis. Secondly, players must submit their preference inputs sometime during the AFP planning period. Each player is informed about the resource they are allocated immediately after making their preference submission, and are not permitted to swap or modify it³.

We further assume that players are not informed as to when other players submit, what they submit, or the status of resource availability (and therefore, what allocations other players receive) at any time during the AFP planning period. The result is that players do not have any information, beyond general common knowledge, about their competitors' actions. As stated at the beginning of Section 4.3, all players in a FSFA scheme are assumed to submit truthful route preferences, because at the moment they submit, they are not competing against others for an allocation any longer. Their gaming strategies are manifest in the time at which they submit their preferences, rather than the information itself. We also assume that all players have identical flight cost functions (i.e. identical parameters α and ϵ_r) under perfect information conditions. They only differ with regards to their information uncertainty during the planning period.

There exist several candidate methods in the game theory, and in particular, the auctions literature, by which to model the FSFA competition as described above (Krishna, 2002). For this analysis we have adapted Moldovanu and Sela's (2001) analysis of player's strategies in sporting contests, where players' efforts, or "bids", are dependent on prize values, their personal ability level, and their chances of winning those prizes. We assume a three-player game setup with two prizes, one of which is worth more than the other. The third and last submitter does not win anything. This is analogous to that third player winning the resource of lowest utility, because we distinguish resources by their relative values to one another. We do not lose any information when we employ these relative utility values; the three resources constitute each player's entire choice set and we are only concerned with players' actions when faced with one choice against another in the set. Although it is certainly true that most AFPs have many more than three flights, a three-flight analysis is a starting point to gain some first insights into how players might

³ This rule is enforced to prevent players from submitting inputs at the very start of the planning period simply for the purpose of reserving a resource, any resource, with the idea that they can submit again later without cost. This behavior can lead to several major problems which we are not prepared to consider now.

behave. The analysis can be readily extended to consider larger flight populations in future work. Alternatively, this setup can be used to consider players that are airlines with groups of flights for which AFP resources must be sought.

Each player n in this competition is of a certain type that describes the degree of their informational uncertainty during the AFP planning period. Player's types are continuously distributed over the player population. Players' own types are privately known only to themselves, although they all share common information about the distribution of types within the population. Recall that players are not informed about the actions and allocations of their competitors, nor resource availability status, during the planning period. To this end, our submission competition is represented as a simultaneous incomplete information game, where each player is uncertain about their competitors' types and resulting strategies (Gibbons, 1992). We assume that all players are rational, in that they will employ strategies to maximize their own payoff with respect to what they know about themselves and their competitors. Players also believe that their competitors are rational, and believe that other players believe that they believe that they are rational, and so on. We also assume that all players' types are identically distributed, which results in a symmetric game where the expected payoff function and equilibrium submission time strategy are identical for all players. Player n 's expected payoff can be expressed as:

$$E[\pi_n] = R_1 \cdot \text{Prob}(n \text{ is 1st to submit}) + R_2 \cdot \text{Prob}(n \text{ is 2nd to submit}) - C_n \quad (4.11)$$

where

- C_n is the cost player n incurs in submitting their preference inputs at the time they choose to do so, and
- R_1 and R_2 are the utilities that players can expect to gain by being first and second in the submission order, relative to the utility of being third (and last), $R(3)$.

For example, if the expected utility of being first is $R(1)$, then $R_1 = R(1) - R(3)$. Equation (4.11) assumes that the expected utility of being in a given place in the submission order is identical for all players in a particular AFP, and is common knowledge. This follows from our assumption (stated at the beginning of this section) that all players have identical flight cost functions when they all have perfect information. Also, the average utility curves of Figure 4.2 in Section 4.3.1 demonstrate that over many AFPs, the x th submitter's expected utility will be greater than that of the $(x + 1)$ th submitter. Therefore, $R_1 \geq R_2 \geq 0$.

R_x represents the utility that a player can expect in being x th to submit, if they have good quality information. If a player n submitted very early in the planning period to be x th, n 's true expected utility from being x th will be lower than R_x , as n 's information (about NAS conditions and private airline operations) is more uncertain at the beginning of the planning period than towards the end of it. We are assuming in Equation (4.11) that the amount by which player n 's expected payoff is degraded by its uncertainty at the time of preference input submission is represented by C_n . This utility

loss due to uncertainty is assumed to be additive to the expected utility of submitting first or second. C_n is a linear function of the loss function introduced in Section 4.3.2.

$$C(q_n) = h_n \cdot L(q_n) \quad (4.12)$$

Where

- $C(q_n)$ is player n 's cost of submission, as a function of q_n ; $C_n = C(q_n)$;
- $q_n = (T - t_n)/T$, $q_n \in (0,1]$;
- h_n is a constant that distinguishes player n 's type. It determines the rate at which n 's overall informational uncertainty decreases during the planning period as $t_n \rightarrow T$, and
- $L(q_n)$ is the loss function as defined in Equation (4.10a).

A larger q corresponds to an earlier submission time, which in turn represents a costlier submission. As $L(q_n)$ is a function of q_n instead of t_n , it is increasing, $L(0) = 0$, and $L'(q_n) \geq 0$.

As mentioned above, $C(q_n)$ is the cost that player n incurs by submitting at q_n because of uncertain information. $C(q_n)$ assumes that the disutility expected from uncertainty about external NAS conditions is captured by the loss function. This loss function is identical for all players, in that they all observe the same changing information about weather, demand, and ATFM actions. However, airlines have different capabilities in regards to how they process these external uncertainties through their internal operations. Some airlines (or players, in this case) have large and/or high quality operations units that can make better decisions under uncertainty, while others do not. In fact, some airlines may not even be able to take advantage of the information about external conditions that is available. These differing (and private) intra-airline operations are captured in the term representing the players' types, or h . A player with a low h has good internal operations and information, and is able to make higher utility decisions when faced with the uncertainty represented by $L(\cdot)$. Conversely, a player with a high h has inferior internal operations and cannot make very good decisions when subject to $L(\cdot)$. We will assume all players know that h is continuously and uniformly distributed between h_{min} and h_{max} . If $h_n = h_{min}$, player n 's internal operations can handle NAS uncertainty better than other players, and their cost of submitting at some early time will be smaller than any other player. As a result, n is likely to submit well before the end of the planning period. If $h_n = h_{max}$, the opposite is true, and in fact we will assume that n would submit their inputs as close to T as possible. $C(q_n)$ assumes that n 's information is perfect by T .

Figure 4.5 displays submission cost functions $C(q)$ for two p values and three arbitrary h values. They represent a simple example where there are three resource choices in the choice set S , with utilities $V_1 = 10$, $V_2 = 4$ (and $V_3 = 0$; the relationships between these values are analogous to those between R_1 , R_2 , and R_3). The two values that p take are proportions of V_1 : $p = V_1$ or $2V_2$. Recall that p defines the value of the scale parameter ($\nu(t)$) of the stochastic term representing information uncertainty in the loss function. $\nu(t)$ was represented by Equation (4.9) and modified to be a function of q ;

$v(q) = pq$. We will use these values for some of the numerical examples illustrating the equilibrium strategy.




Figure 4.5 Submission cost function $C(q)$

$C(q)$ is a monotonic (increasing) and differentiable function. It may be strictly increasing over the strategy space of our game, $q \in (0,1]$, depending on parameter values.

Since this is a symmetric game, the functional form of the submission strategy is identical for all players and can be represented as a function of a player n 's type, or $q_n = g(h_n)$. It was also stated that players do not know their competitors' submission strategies because they do not know their types. As a result, we can assume that a player's probability of winning against one competitor is independent of the probability of winning against another. The probability that Player 1 submits before Player 2 is

$$P(t_1 < t_2) = P(q_1 > q_2) = P(g(h_1) > g(h_2)) = 1 - F(g^{-1}(q_1)) \quad (4.13)$$

For (4.13) to hold, we must assume a-priori that $g(h_n)$ is monotonic (decreasing) and differentiable, and later verify that our assumption is correct. We can also normalize all utilities (of those in the choice set S , as well as R_1 and R_2) such that the choice of highest deterministic utility is one, and then rewrite the payoff function of (4.11) as follows:

$$\begin{aligned}
E[\pi_n(q_n)] &= \frac{r_1}{L^{max}} \cdot \left(1 - F(g^{-1}(q_n))\right)^2 \\
&+ \frac{2r_2}{L^{max}} \cdot \left(1 - F(g^{-1}(q_n))\right) \cdot F(g^{-1}(q_n)) - h_n l(q_n)
\end{aligned} \tag{4.14}$$

where r_1 and r_2 are normalized utilities of having submitted first and second, respectively. Our aim is to solve (4.14) in order to obtain the submission strategy $q_n = g(h_n)$ that maximizes player n 's expected payoff, with respect to the parameters of the AFP, the information n has about its competitors, and what n knows about itself. $g(h_n)$ is the equilibrium strategy, where n cannot achieve a higher payoff by deviating from this submission time strategy.

4.3.3.2 Equilibrium Submission Strategy

It is verified that the payoff function is concave with respect to the submission strategy in $q \in (0,1]$ (Moldovanu & Sela, 2001). Depending on the parameter values, solutions (that maximize expected payoff) may lie on the boundaries of the strategy space. By rearranging terms and using the assumption that players' types are uniformly distributed according to $h_n \sim U(h_{min}, h_{max})$, Equation (4.14) is solved to find the expression:

$$\begin{aligned}
&1 - \left(\sum_{s \in S} \exp\left(\frac{v_s}{q_n p}\right) \right)^{-1} \cdot \sum_{s \in S} v_s \exp\left(\frac{v_s}{q_n p}\right) \\
&= \frac{2[(r_1 h_{max} - r_2(h_{min} + h_{max}))(\ln h_{max} - \ln h_n) + (r_1 - 2r_2)(h_n - h_{max})]}{(h_{max} - h_{min})^2} \tag{4.15}
\end{aligned}$$

where v_s are the normalized "true" slot utilities such that the most valuable slot has utility $v_1 = 1$. The above expression cannot be solved for the equilibrium submission strategy q_n in closed form, but it is possible to find solutions using Matlab or Excel. See Appendix C.2 for the derivation of (4.15). A third-order Taylor series about $q = 0.55$ was calculated, but it yielded poor approximations of the function at the boundaries of the planning period (i.e. the farthest points from $q = 0.5$), particularly with larger p values. As a result the approximation was not used.

Depending on the values of parameters $h_{min}, h_{max}, v_{s \in S}, r_1$, and r_2 , there may not exist a solution to Equation (4.15) for values of $h \in [h_{min}, h_l], h_l \leq h_{max}$. This is due to the fact that the maximum value the left side of (4.15) can take is one. If $h_n < h_l$, player n 's disutility due to uncertainty is always smaller than the utility they expect from being first or second to submit, and they will always submit as soon as they are able. In fact, all players with $h \leq h_0$ will submit as soon as the planning period begins, where $h_l \leq h_0$. All players with $h > h_0$ will want to submit at some time after $t = 0$, depending on the parameters of the submission strategy function h_{min}, h_{max}, r_1 , and r_2 . We find both h_l and h_0 numerically as they also cannot be expressed in closed form.

Once we find q_n, t_n is:

$$t_n = \max(T(1 - q_n), 0) \quad (4.16)$$

Figure 4.6 displays submission strategies for seven values of p , for a scenario where $v_s = [1, 0.5, 0]$, $r_1 = 0.8$, $r_2 = 0.8 \cdot v_2$, $h_{min} = 0.5$, $h_{max} = 1.5$, and $T = 2$. These values were chosen arbitrarily. The x -axis represents values of h from h_{min} to h_{max} , and the y -axis represents the AFP planning period, which in this case is two hours in length. It is observed that the submission time strategies expressed in t are increasing (decreasing if expressed in q) and differentiable.




Figure 4.6 Example equilibrium submission strategies

The figure shows that FSFA submission time strategies are very sensitive to p . The general uncertainty levels of AFPs can be quite different depending on storm features and behavior (if it is indeed inclement weather causing the AFP) and how traffic is managed in response to it, both of which can be represented through the parameter p . Given how much strategies can differ with respect to p , the results suggest that flight preference submission behavior may vary significantly from one AFP to the next. The figure demonstrates that the submission strategy is concave over most of the (higher) p values shown and is convex at $p = 0.25$. If it were true that the graphs represented the behavior of flight populations greater than three, when the submission strategy is more concave at higher values of p , traffic managers would observe a relatively slow arrival of

submissions after a clump of arrivals at $t = 0$. They would then observe an increasingly faster rate of arrivals until the end of the planning period.

According to the figure, when $h \approx h_{max}$, submission strategies become very sensitive in that very small increases in h result in large increases in t .

Recall that we assumed the submission strategy to be increasing and differentiable in order to determine the probabilities of being first or second in the submission order. The submission strategy is increasing within the planning period for the example shown in Figure 4.6; players that desire to submit before the planning period submit at $t = 0$. Moldovanu and Sela (2001) also prove that the bid function is strictly increasing and differentiable, and that it maximizes expected payoff.

All players with $h \leq h_0$ would expect to maximize their payoff by submitting at some time before or at the beginning of the planning period; however, it is not possible to submit inputs before the planning period begins, and thus all players with $h \leq h_0$ submit at the very beginning. As p (i.e. the general uncertainty about NAS conditions) decreases, more flights will be inclined to submit at the very beginning of the planning period (and therefore, less will submit later during the planning period); also, h_0 increases as well. This is of course assuming that h are uniformly distributed over the flight population. Again, if it were true that the graphs represented the behavior of flight populations greater than three, a smaller p would result in a large proportion of the flight population submitting as soon as the planning period begins. Traffic managers could then expect to receive more submissions at the beginning of the planning period, the majority of the AFP planning could be completed early, and the AFP could be more readily coordinated with other air traffic flow management programs in the NAS. Overall, it appears that a significant proportion of flights will submit their preference inputs at the beginning of the AFP planning period.

Another point to note is that as $h \rightarrow h_{max}$, $t \rightarrow T$ but never actually reaches it. This is because the scale parameter v is a function of t , and v enters into the denominator of the terms of the loss function. As a result, as t approaches zero, the probability of choosing the resource of highest utility V^* approaches one.

If the loss function $L(q_n)$ took on a simpler (i.e. linear or quadratic) form, we would have a closed form expression for the submission time strategy.

4.3.3.3 Sensitivity tests

The results of Figure 4.6 demonstrated that the submission strategy function is very sensitive to p values. Here we further explore the sensitivity of our strategy function to h and p , as well as v_2 values. Given that we have normalized our utility values to be between zero and one, v_1 is always one. Figure 4.7 shows the resulting submission strategies for the same example plotted in the previous figure (Figure 4.6). However, the x -axis now represents values of p from 0.25 to 3, while each curve corresponds to a value of h as labeled in the legend.




Figure 4.7 Equilibrium submission strategies, t versus p

When both p and h are high, the equilibrium submission strategy changes little with respect to changes in p and h . Conversely, at low values of both p and h , the equilibrium strategy changes at much faster rates in response to changes in h and p . This observation can be measured by the vertical distances between curves; however, it may be easier to observe in the following figure (Figure 4.8), where the contours represent differing submission strategies (or rather, the submission strategy “bins” that are demarcated by contour lines).




Figure 4.8 Equilibrium submission strategies (t , hours), p versus h

The wider contours correspond to p and h levels at which submission strategies change at a slower rate with changing p and h . As information grows more uncertain (i.e. high p and h), airlines will become increasingly motivated to submit near the end of the planning period, and strategies do not deviate greatly at this level of uncertainty. Contours become more narrow as p , h , or both, decrease. This indicates that small increases in uncertainty (at lower uncertainty levels represented by these smaller p and h values) will have a larger impact on the players' payoff functions, thereby having a greater effect on the submission strategy. The black areas represent situations where airline preference submissions should be made at the beginning of the planning period. It covers a fairly large portion of the figure, indicating that airlines should submit at $t = 0$ for many combinations and p and h . The x -axis was truncated at $h = 0.8$ (instead of at $h_{min} = 0.5$) to reduce the black space.

Figure 4.9 shows submission strategies when $h = 1$. The submission strategy was evaluated for v_2 values ranging from zero to one (which includes all the values it could possibly take), and p values ranging from 0.25 to 3. The left plot of Figure 4.9 shows submission time with respect to p , where each curve is generated using a single value of v_2 . The frontier curve is for $v_2 = 0$, and v_2 values of the curves increase in the direction of the arrow. The right plot of Figure 4.9 displays the identical example but with submission time plotted against v_2 . Each curve represents a single value of p , where p values increase in the direction of the arrow. The frontier curve represents $p = 3$.

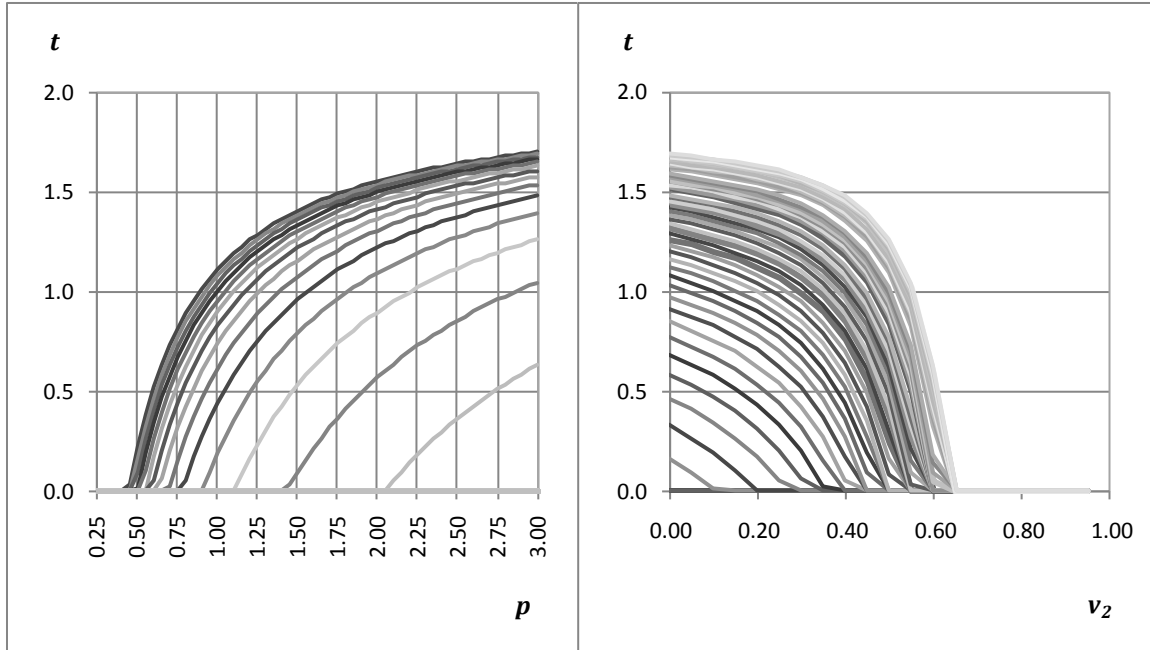


Figure 4.9 Equilibrium submission strategies (t , hours), $h = 1.0$

The figures indicate that flights will submit earlier in an FSFA allocation process as v_2 increases. This observation supports the rationale that when v_2 is larger, the total “prize” values are also larger and therefore flights have greater motivation to submit earlier. For instance, if we were to draw a vertical line at $p = 1$ in the left plot, when $v_2 = 0.5$ it is a best strategy to submit inputs at $t = 0.2$. However, if $v_2 = 0.2$ it is best to submit at $t = 1$. This is a significant difference; according to the shape of the curves these differences decrease as p increases, but they are still significant at larger values of p . More easily observed from the right graph is that for all values of p at $h = 1$, if v_2 is larger than approximately 0.65, the players’ submission strategy will be to submit at the beginning of the planning period. At smaller values of v_2 the submission strategy varies significantly with respect to the range of p values shown. However, when p is large, the submission strategy will not change significantly with a unit increase in p , at a given v_2 value. This last observation is similar to the information that Figure 4.8 imparts.

Figure 4.10 displays submission strategies when p is held constant (at $p = 1$), h ranges from $h_{min} = 0.5$ to $h_{max} = 1.5$, and again v_2 ranges from zero to one. The figure is plotted over three axes because the behavior of the submission strategy function, when plotted with respect to increasing values of h and v_2 , is more readily observed when shown in this manner.




Figure 4.10 Equilibrium submission strategies (t , hours), $p = 1.0$

It can be observed that when v_2 and h are large, the optimal strategy is to submit later than when v_2 is smaller (and h is still large). As mentioned above, more flights will be incentivized to submit earlier, and at the beginning of the planning period, when v_2 is larger. It appears that the remaining flights will submit at a more constant rate in the remaining time (again, given a uniform distribution of h) than if v_2 is smaller. One might observe the lines at the top of the figure for a given value of h (i.e. if you view v_2 on the x -axis and t on the y -axis): they are not monotonic with respect to v_2 . This behavior contradicts the idea that players are more likely to submit earlier when the prize values are higher. However, the behavior is only observed at very high h values; throughout these examples it has been shown that the submission strategy is highly sensitive at high h , and we should investigate this further to determine the cause of the non-monotonicity. The function appears to behave as expected for all other combinations of h and v_2 .

The last figure – Figure 4.11 – displays h_0 as a function of v_2 and p . Recall that if flight n has $h_n \leq h_0$ they will submit their preferences at the beginning of the planning period, and otherwise they will submit at some time later.




Figure 4.11 h_0 for decision-making and predictability

For any combination of v_2 and p , the corresponding h_0 value can be found from the figure above. For instance if $p = 1.5$ and $v_2 = 0.4$, $h_0 \approx 0.85$. Given that $h \sim U(0.5, 1.5)$, this value of h_0 corresponds to 30% of flights submitting at the beginning of the planning period (again, assuming that the results represent the behavior of >3 players). If h_0 is very high, this indicates that, for a large proportion of flight operators, the utility of being first or second is greater than the utility loss incurred by submitting early. Figure 4.11 can help airlines in making submissions decisions, and also inform traffic managers about the FSFA submissions process outcome, for all potential AFP scenarios accounted for in the figure. If airlines have access to the information contained in Figure 4.11, it can be used to quickly reduce their set of potential strategies in a particular AFP. For instance, if $h \leq h_0$, the airline knows immediately to submit at the beginning of the planning period without further analysis efforts. If $h > h_0$, they know they must submit at a later time to maximize their utility, and further analysis should be employed to determine exactly when. Traffic managers can benefit from the figure as it gives some indication of the predictability of the FSFA submissions process in a given situation. If h_0 is very high, traffic managers will expect to receive more submissions at the beginning of the planning period, such that the majority of the AFP planning will be completed early. This can help traffic managers in coordinating the AFP with other air traffic flow management programs in the NAS. It is true that early planning efforts (like early information) are more likely to become irrelevant under greater uncertainty about NAS conditions; however, we know that the submission strategy implicitly accounts for this issue as airlines are less likely to submit early when facing greater informational uncertainty.

Although not presented here, the results of our numerical investigations also indicate that a larger proportion of flights will submit at the beginning of the planning period when h_{min} is smaller.

4.3.4 Model Extensions

The results have highlighted the need for a model that represents the competitive behavior of more than three players. The model introduced in this chapter can be extended fairly easily to account for larger flight populations. Before the model is extended in this way, however, there are other issues that should be addressed with greater priority. These issues might be better addressed using an alternate auction/contest analogy. A good formulation may not require us to assume that prize values and their devaluation (due to uncertainty at the time of preference submission) are additive. Also, an alternate formulation may help us to more readily incorporate flight heterogeneity into our analysis, not only regarding intra-airline operations but also in terms of the cost parameters. Flight heterogeneity can affect the competition in several ways. Certain cost parameter values (α, ε_r) may incentivize flights to submit later or earlier in the planning period. Also, if we assume that not all flights necessarily desire the same resources because of their heterogeneity in cost parameters, and they are all aware of this fact, then they will likely submit later than the model introduced in this chapter would indicate. Valuing resources differently from other flights will dampen the competition for resources. Understanding how flight heterogeneity affects submission behavior will help us to better evaluate the efficiency of the FSFA resource rationing scheme.

It will also be helpful to revisit the assumption that player types, in regards to the quality of intra-airline information and operations, are uniformly distributed over the population, particularly if there is any empirical or anecdotal evidence available on this subject. Investigating the applicability of other distributions may give us better insight into airlines' competitive behavior in a FSFA resource allocation scheme.

4.4 Other Gaming Behaviors and Future Work

There are a number of ways by which flight operators game the current AFP and will likely game future versions. One tactic that an airline may employ is to file more flight plans through historically problematic airspace (such as the thunderstorm-prone corridors over Pennsylvania/New York and North Carolina/Virginia) than they would actually require. Under the CDM concept in which airlines "own" flight slots, by filing excess flight plans through airspace regions when resource rationing is anticipated, an airline will have reserved more slots and, therefore, more options for their flights that do require a resource. Flight plans can be modified at any time before or after a flight has departed, and as a result they have no disincentives to filing false flight plans in an AFP. In fact, one legacy carrier confirmed that they would file flights (whose optimal routes were elsewhere) into chronically constrained airspace, in order to have more slots available to them should a constraint arise. This type of behavior will impact all the allocation

schemes introduced. It might be discouraged by performing initial allocations based on information other than filed flight plans and published schedules. Potential candidates include historical flight plans, or a type of market-based system.

Flight operators may also negatively impact AFP allocation schemes not through proactive gaming behavior, but by failing to share information altogether. If a flight is scheduled towards the end of an AFP, its operator may feel that the expected utility of a wait-and-see approach is higher than the utility they could expect to receive from a premature allocation, especially when NAS conditions change rapidly and with high instability. Also, although an operator may submit information and receive an allocation for their flight, they might ultimately cancel the flight and not inform the FAA. A credit system (that extends beyond a single AFP occurrence) may discourage this behavior.

There are many ways to both improve and extend our analyses of airline truth-telling in the OPT and Parametric schemes, and preference input submission strategy in the FSFA scheme. These future research directions have already been discussed to some length in this chapter, and we continue the discussion in Chapter 6.

5. Other Resource Allocation Schemes

5.1 Introduction

In this chapter we consider several additional resource rationing schemes that employ Stated Route Preference user input. The purpose is to develop schemes that can provide greater total user flight cost efficiency and/or equity between flights. The schemes must also aim to minimize gaming behavior by flight operators, where users are encouraged to participate in the preference input process and do so with truthful information. There are many rationing philosophies that have been developed under the CDM objectives, as discussed by Vossen et al. (2003) and Ball et al. (2002), in addition to several others.

In this chapter we will discuss the Hybrid Stated Route Preference model. As the name suggests, the resource rationing mechanism consists of a combination of the OPT and FSFA resource allocation mechanisms. We will also introduce the results from an application of the Ration-By-Schedule (RBS) rationing algorithm in our model framework. RBS is the algorithm currently used to allocate airport arrival slots in Ground Delay Programs (GDPs). Investigations of both the Hybrid and RBS models are preliminary, and do not include the gaming and truth-telling aspects of flight preference input. However, potential directions for research on these aspects are discussed. The chapter ends with a brief discussion of other types of resource rationing mechanisms that can be explored as part of future research.

5.2 Hybrid Stated Route Preference Model

The two rationing schemes that employ the Stated Route Preference user input concept – OPT and FSFA – respectively represent the extremes of efficiency and reward. The OPT model minimizes total user cost, but it does not offer flight operators any clear incentives for providing expedient and truthful information. In addition, because traffic managers must collect all user preference submissions before they can allocate resources, allocation and notification may not be complete until shortly before the AFP start time. This could disadvantage flights that are scheduled at the beginning of the AFP, as they may have known their preferences much earlier on and therefore would have preferred allocation to take place earlier. Also, these flights may be required to depart only a short time after notification. Both these facts could lead to additional costs being imposed on the flight operators' operations planning. Some observations of FAA planning telcons reveal that

airlines sometimes prefer uncertain information over no information, such that they can begin planning some course of action for their flights. This implies that airlines may place a relatively high utility on early notification.

In the FSFA model, early submitters can gain major rewards while late submitters are often severely penalized. This process could be considered unfair in its allocation severity, particularly in its penalization of airlines that submit preference inputs immediately after others. Also, flight operators may have large incentives to submit very early with highly uncertain and/or inaccurate inputs as shown in the previous chapter, due to the competitive allocation process.

To address these issues, a hybrid allocation scheme that preserves the FSFA reward structure but offers greater user cost efficiency is developed. In this scheme, flight operators are instructed to submit their information whenever they choose within the AFP planning period, identical to the FSFA model. However, instead of being promised a resource assignment at the time of submission, flight operators will be informed of the pre-designated “cut-off”, or “batch allocation” times. At these times, traffic managers will perform a system-optimal allocation for all flights whose information was received after the time of the previous batch allocation. This process is repeated until all flights have been allocated resources, or the end of the planning period is reached – whichever comes first. It is clear that if a smaller number of batch allocations are used, the total user cost results of this allocation will be closer to that of the OPT model. A larger number of batches (and therefore smaller periods between allocations) will yield total user cost results that are closer to that of the FSFA model. In addition, the distribution of submission times over the AFP planning period will also have a significant effect on the resource allocation.

The Hybrid scheme requires that traffic managers decide the times at which the batch system-optimal allocations are made. Suppose the first allocation batch is at t_1 , the second is at t_2 , and so on, until the process is stopped at t_B and allocations have been made in B batches. The Hybrid scheme can be simulated as follows:

- 1) Randomly assign α values, submission times, and (randomly drawn) ε values to each AFP flight.
- 2) Order flights by these submission times, or $m = 1, \dots, N$.
- 3) For allocation batch $b = 1, 2, \dots, B$:
 - a. Define flight set F_b as consisting of all flights that made $\Delta_{n,r}$ submissions between t_{b-1} and t_b (when $b = 1, t_0 = 0$).
 - b. Run the OPT model on F_b according to Equation (5.1) below.

$$\min_{x_{n,j}, \forall n,j} C = \sum_{n \in F_b, j \in J, j \neq j_{b-1}: d_{n,j} \geq g_{0,n}} c_{n,j} \cdot x_{n,j} \quad (5.1)$$

Note that the equation specifies that flights $n \in F_b$ cannot take a slot that has been taken by a flight in the previous allocation, $n \in F_{b-1}$.

4) Repeat (1) through (4) for each value of σ and iteration $z = 1, 2, \dots, Z$.

The flights' input submission times in Step (1) could be independent of any AFP or flight characteristics, or they could be correlated to α , n , distance from origin to destination airport, aircraft size, etc. We will continue to assume that they are independent and randomly drawn from some probability distribution of our choice. We can investigate the results from several different submission time distributions, in order to gain some insight into how batch allocation times influence allocation efficiency.

The figure below shows the results of the Hybrid resource allocation scheme using different numbers of batch allocations and with two different distributions of user input submission times (uniform and normal) over the planning/submission period. The results are for an AFP scenario where there are 40 flights being assigned to slots on three routes. The batch allocations are spaced in equal time increments. In the figure, “ $b = 2$ ” represents an AFP resource allocation made using two batch allocations over an AFP planning period of one hour – one system-optimal allocation is made at 0.5 hours, and the other is made at the one hour mark. Each point on the figure represents the average of 1,500 iterations.

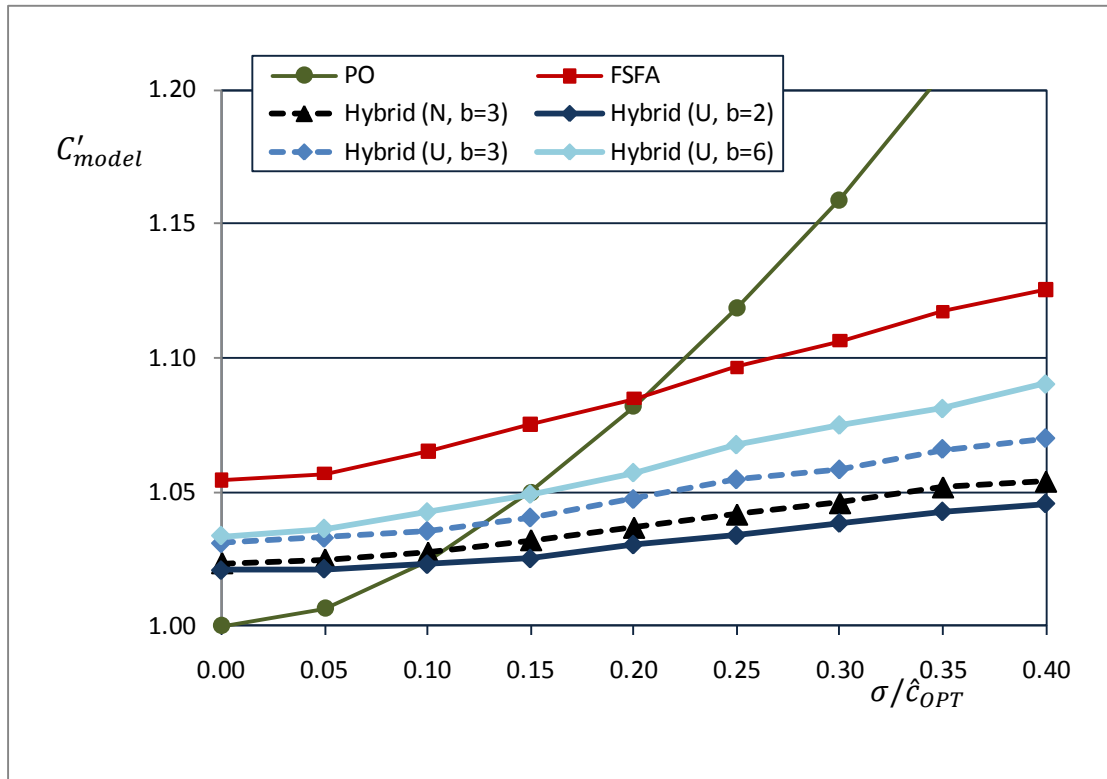


Figure 5.1 Example of Hybrid allocation scheme performance.

Clearly, if there are many allocation batches and the distribution of flight input submissions amongst the allocation batches is more or less even (i.e. uniformly distributed), the total cost outcome will approach that of the FSFA scheme. However, we can imagine that if most (or if not all) flight operators submit their inputs within a small

number of batch periods, the Hybrid allocation scheme breaks down and the process approaches that of the OPT allocation. For instance, in Figure 5.1, compare Hybrid (U, $b=3$) to Hybrid (N, $b=3$). The former represents the Hybrid model where each batch has an identical number of submissions (modeled using a uniform distribution of submissions) and there are three allocation batches. The latter represents the Hybrid model where the majority of submissions (~60-75%) are made within the second of three batch periods (modeled with a normal distribution). The latter model is more similar to OPT in that the majority of flights are assigned resources system-optimally in a single batch. Therefore, the allocation costs are closer to the x -axis (and to those of the OPT model).

The results of the FSFA gaming analysis of Chapter 4 suggest that in many cases flight operators are more inclined to submit earlier rather than later during the planning period, given the relatively high utility of submitting before their competitors relative to flight cost uncertainties. In these cases submission times might be represented using a right-skewed bell-shaped probability distribution. An analysis of flight submission times under competition, similar to that of Section 4.3, should be undertaken specifically for the Hybrid scheme.

5.3 Ration-by-Schedule (RBS)

The Ration-by-Schedule (RBS) algorithm was first introduced in Section 1.1.2 as the resource rationing mechanism currently used to allocate airport landing slots in Ground Delay Programs (GDPs). Here we consider the implications of applying RBS to determine the order by which airlines are assigned en route resources (routes and slots) in an AFP.

We imagine that when the FAA announces the start of the AFP planning period, the flights caught in the AFP would be those whose operators have filed flight plans through the affected region or, if they have not filed yet, those that have historically done so. We also assume that the operators of all AFP flights are sufficiently incentivized to provide their flights' complete and truthful Δ values, in order to participate in the allocation. The algorithm would order flights by their estimated times of arrival to Fix A (Figure 2.1), which would in turn be estimated from their pre-AFP scheduled departure times. It would then sequentially assign each flight the lowest-cost resources available to it. In the Scenario 1 example of Chapter 3 with $N = 75$ flights, it would perform as shown in Figure 5.2.

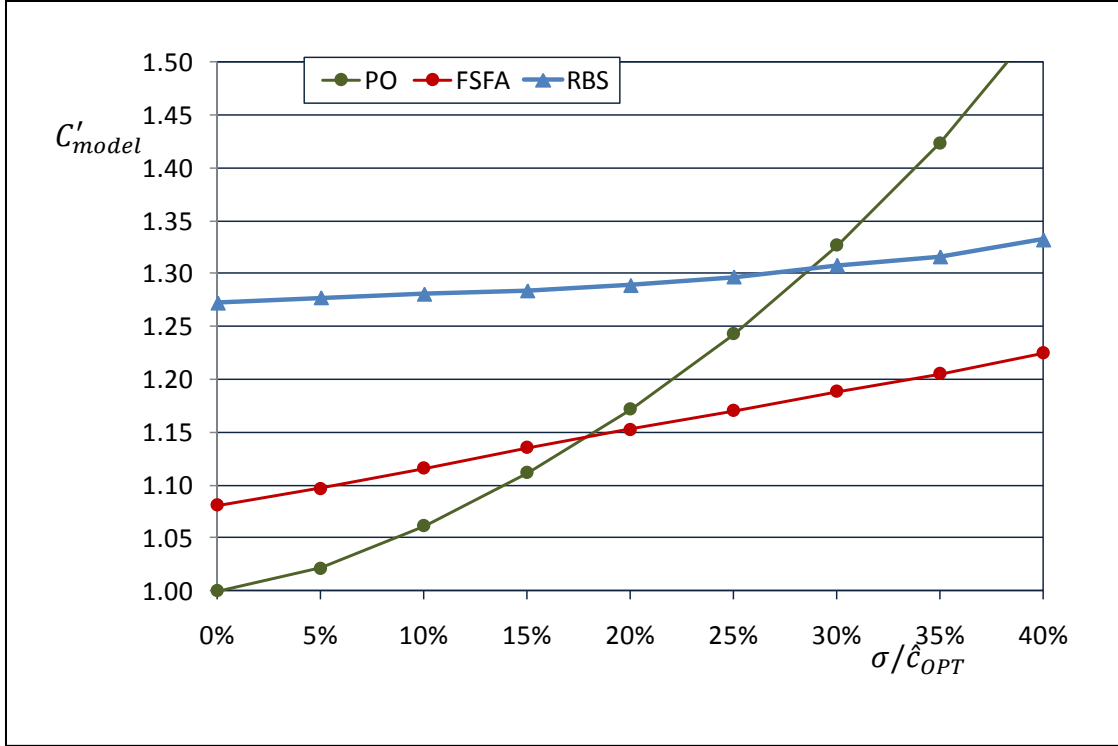


Figure 5.2 Demand Scenario 1 total cost results including RBS

The above figure demonstrates that the RBS solution, even with the assumption that all flights have provided their Δ values, is inferior to the FSFA solution generated from a random ordering of flight submissions. We can state the following:

Proposition 5.1. *Flight-slot assignments using the RBS algorithm will have equal or higher costs compared to assignments using the FSFA algorithm assuming random flight submissions.*

See Appendix D for the proof.

The RBS allocation also appears to be similar to an offline, first-fit increasing bin-packing problem. The first item (or first scheduled flight) is the “smallest” item and has much “room to fit”, in that it can take almost any slot on any route given its early pre-AFP scheduled departure time. The last item is the “largest” and has the least “room to fit”, as it cannot take a slot that is earlier than its pre-AFP scheduled departure time. RBS is, however, somewhat different from bin-packing in that the earlier scheduled flights may fit into any bin, but some fits are more costly than others. In any case, one can see that allocation in the RBS ordering will be amongst the least optimal solutions.

One of the most important points to note about Figure 5.2 above is that, no matter how suboptimal an allocation mechanism is, at some σ value it will always be more cost efficient to allocate resources with more complete information and inferior mechanisms compared to using a system-optimal assignment with incomplete information.

The RBS result shown above, in addition to its relatively poor comparative performance, is the outcome of highly optimistic assumptions about airline participation. It may again be difficult to rely on the idea that flight operators will provide complete preference information to the FAA. In some cases, airlines know that by providing accurate preference inputs they are more likely to receive a resource they desire. However, in other cases, the allocation they receive without having provided any information may be similar in cost to their allocation had they provided complete information. All operators have access to historical OAG (Official Airline Guide) schedule information, as well as all flights' airport arrival times through the Flight Schedule Monitor (FSM). As a result, operators can know their flight's "rank" or "priority" in the rationing algorithm, and can make fairly well-informed decisions about whether or not to provide their requested preference inputs. In the end, we can assume that flight operators would prefer to reveal as little operating cost information as possible, and will not offer information if they do not anticipate benefits. Clearly, we must study this behavior and understand its impacts on the RBS application to en route resource rationing.

In addition to its poor performance and optimistic assumptions about airline participation, RBS has several other practical drawbacks in its application to en route resource allocation. Firstly, en route resources are much more uncertain in nature compared to airport resources. Strategic planning for en route resources is more difficult than for airport resources, because there are typically larger sets of alternate options available due to more degrees of freedom in routing compared with landing slots. Flight plans can be changed even while a flight is en route, without discernible impact to customers. In fact, flight plans are only typically submitted to the FAA by airline dispatchers 45 minutes prior to take-off. However, airport scheduled departure and arrival times are published as part of the OAG and therefore cannot be changed so easily without disrupting customers. Secondly, RBS will certainly encourage airlines to game the allocation system by filing an excessive number of flight plans through troublesome en route regions, as described in Section 4.4. In fact, the schedule-based rationing logic may result in a stronger incentive for this behavior, particularly to those operators that have flights scheduled later in an AFP and therefore are likely to have higher cost resource allocations.

In Reverse RBS (RRBS), as suggested by the name, flights scheduled later in the AFP are assigned resources first. RRBS is demonstrated to be more efficient than RBS and FSFA (with random flight input submission times) in numerical examples. However, given that RRBS imposes smaller costs on flights scheduled later in the AFP (and therefore inherently subject to more uncertainty than earlier AFP flights), RRBS could have great inefficiencies not apparent through the setup of our numerical examples.

5.4 Future Investigations

There are many resource rationing mechanisms that could be more effective in obtaining desired levels of user cost efficiency, offer greater fairness, and discourage gaming

behavior, as well as balancing and trading off these potentially conflicting objectives. User flight cost efficiency is only one goal in the service that a governmental organization such as the FAA is obligated to provide. The FAA must also aim to serve their customers fairly, and consider other costs that are being imposed on society as a whole. To balance these different aims, some compromises must be made. Also, by providing equity, what may be compromised in explicit system cost efficiency aims can be regained through the flight operators' willingness to submit truthful route preference information. In this section we briefly introduce some potential mechanisms that employ the Stated Route Preference concept, and discuss directions for future work.

Firstly, we can investigate modifications to the greedy assignment algorithm used for the FSFA and RBS schemes. Both these schemes may benefit from the application of online algorithms and Dynamic Traffic Assignment (DTA) concepts. Both these concepts involve decision-making under uncertainty or incomplete information. Online algorithms are designed to make decisions that are as optimal as possible without knowledge of the entire input (Albers, 2003). An online algorithm consists of heuristic service rules that attempt to minimize the competitive ratio, which is the performance of the algorithm measured against that of the optimal offline algorithm (identical to C'_{model}). DTA forecasts future traffic demands and their impacts on a traffic network, which traffic managers then use to manage traffic near optimality. In the case of FSFA (or any other sequential assignment process where traffic managers must make decisions without knowing the future submissions), DTA and online algorithms may be applied to assign resources more optimally (but within some promised tolerance range of the best available resource cost, to ensure delivery of flights' complete route preferences).

If traffic managers are prepared to offer flight operators some guarantee about the maximum cost that can be imposed on any given flight(s) in the AFP, they could aim to minimize the largest cost imposed on any such flight. This could be formulated as a Linear Bottleneck Assignment Problem. However, traffic managers must be certain they are receiving truthful inputs. Also, they might want to aim explicitly for other goals in addition to user flight cost efficiency. In this research, equity has been considered implicitly in resource rationing mechanism design, one aim being to encourage user participation. We can consider incorporating explicit equity measures and constraints into the performance assessment, as well as emissions metrics and traditional performance metrics (i.e. delay). In order to choose a scheme, however, traffic managers must choose their goals a-priori. They must decide how much cost efficiency they are willing to sacrifice to gain equity benefits, equity benefits for fuel savings, etc., and therefore establish marginal rates of substitution for these measures. There are many well-established techniques in the multi-objective optimization literature to accomplish this (Miettinen, 1999). In addition, there is some existing work on the application of multi-objective optimization to balance equity and efficiency in the context of Ground Delay Programs (GDPs) (Glover & Ball, 2010). The authors investigate the use of different objective functions to more precisely balance efficiency and equity compared with RBS, RBD (Ration-by-Distance) and E-RBD (Equity-based RBD).

If the OPT scheme assignment solution is not unique, there may be a solution within the set that is more equitable by some pre-established equity measure. However, it

is likely that the OPT solution is unique given the user flight cost formulation introduced in Chapter 3. We could instead introduce an ε -constraint on the objective function (Miettinen, 1999), to specify that all other performance measures must fall below (or above) some pre-specified limit, ε . Certainly this will increase the total user cost of the AFP. However, flight operators are typically more concerned with the discrepancies between allocations rather than the total AFP cost, and solutions that provide higher but more normalized costs and more equity across operators and flights would likely be met with greater acceptance. The ε -constraint resource rationing method only works successfully when traffic managers can be certain that flight operators have submitted truthful preference inputs. Otherwise, it would encourage all flight operators to submit untruthfully high route preference inputs, thereby defeating the purpose of user participation. There are other techniques, such as goal programming, that could be explored. Gaming and truth-telling behavior must be considered for all these potential resource rationing mechanisms.

6. Conclusions and Future Research

This research has investigated how to incorporate systematic user inputs in allocating constrained en route airspace capacity, particularly within the context of the Airspace Flow Program (AFP). We have specified resource assignment schemes that feature different allocation rules and route preference inputs. In some schemes operators are offered resource allocations commensurate to the quality of information requested of them, to increase the likelihood of user participation. We have evaluated and compared the total flight cost efficiency of each assignment scheme under varying assumptions about the quality of the central decision maker's (or traffic manager's) knowledge about the flight operators' route preferences, through the use of a simple generalized flight cost function. The flight cost function consists of two parts. The first part is a simple representation of operator flight cost characteristics that the traffic managers have adopted. The second part consists of a flight's routing preferences that are not captured by the first, and therefore are privately known to each flight's operator but not to traffic managers, unless the information is offered to them. We represent the first part of the flight cost function using deterministic terms and the second part using a stochastic term. We identify some basic properties about the relative efficiencies of the three assignment models under traffic managers' increasing uncertainty about the flight operators' private route preferences. Also, numerical examples illustrated situations where sacrificing a system-optimal allocation rule to obtain the flight operators' private information about route preferences will result in more efficient resource allocations, and vice versa. If the user flight cost model is well-specified such that the stochastic route preference terms are small, traffic managers would be advised to adopt the parametric user input and system-optimal resource assignment of the Parametric scheme. If traffic managers believe their cost model does not capture sufficient information about flight route preferences, they would be better advised to use the FSFA allocation scheme. If traffic managers have little knowledge about the quality of their cost model specification, they would also be better advised to use the FSFA method, as the Parametric method becomes very inefficient when the flights' private route preferences are highly variable. We also found the performances of the assignment methods to be relatively insensitive to increasing AFP durations under a constant demand rate, but more sensitive to changes in the relative cost of airborne delay versus ground delay, and changing demand rates.

The above results were generated under the assumption that flight operators offer truthful preference inputs, and that submissions in FSFA are completely random and independent of any AFP characteristics. We know these assumptions are unlikely given the competition for resources. As a result we have also explored operator behavior in

response to the competition characteristics of the resource allocation schemes. These investigations have revealed critical issues about resource rationing schemes that rely heavily on user preference inputs and assumptions about user cost structures. We demonstrated that flight operators are incentivized to provide untruthfully high inputs in the Parametric allocation scheme. In the FSFA scheme, flight operators' submission times vary significantly depending on the conditions of the AFP and their private information, indicating that they may behave quite differently from one AFP to the next. However, the results of the gaming model also suggest that it is often optimal for flights to submit at the very beginning of the planning period in FSFA. The gaming model can be used by traffic managers to estimate when flights are likely to submit their preferences. If the number of flights that submit at the beginning of the planning period is very high, traffic managers know that the majority of the AFP planning can be completed early such that they can better coordinate the AFP with other air traffic flow management programs in the NAS. The results of Chapter 4 suggest that use of the FSFA allocation scheme may be more beneficial than the Parametric scheme. However, we must further study the relationship between the Parametric user equilibrium solution and the resulting utility of FSFA under gaming before making a definitive conclusion. We leave this to future research.

In Chapter 5, two additional allocation schemes based on Stated Route Preference input were also introduced. The Ration-by-Schedule (RBS) algorithm applied to en route resource allocation is not only difficult to implement, but it was also demonstrated to be inefficient from a total flight cost perspective. A hybrid OPT-FSFA model appeared to offer more promising results.

This research can be extended in many directions. Structuring the models and analysis on an airline rather than flight basis would better capture the behaviors of airlines that own multiple flights. We can also reassess the performances of the allocation schemes using a flight cost model specification that considers the non-linear cost of delay, as discussed at the end of Chapter 3. It would be beneficial to consider heterogeneity in the parameters of the flight cost model (α, ε_r) in the truth-telling analysis of the OPT and Parametric input schemes as well as the FSFA competition model. For the FSFA competition analysis, an alternate formulation that does not require the assumption of additive prize value devaluation might be developed. An alternate formulation may also help us to more readily incorporate flight heterogeneity into our analysis, not only regarding intra-airline operations but also in terms of the cost parameters. We also discussed ideas regarding the development of other resource allocation mechanisms. Online algorithms and dynamic traffic assignment concepts can be explored to obtain greater system efficiency from algorithms like FSFA and RBS. We can also more explicitly consider and balance equity as well as emissions and traditional performance metrics by applying multi-objective optimization techniques. These schemes should also be robust to the gaming behaviors discussed in Chapter 4.

Another aspect of this analysis will involve understanding the numerous ways that uncertainty affects allocation efficiency. In the FSFA gaming analysis of Chapter 4, we investigated how uncertainty affects airlines' input submission behavior. However, we must also understand how this same uncertainty affects the overall efficiency of the

allocation schemes. If many flight operators are incentivized to submit earlier with low quality information in order to “beat” other operators, how does this degrade the true efficiency of the FSFA scheme? Also, higher user cost efficiency may be achieved by scheduling flights to earlier AFP slots before scheduling flights to the later slots, as information changes over the course of the planning period and AFP itself, and later flights are inherently subject to more uncertainty than earlier ones. In this way, the RBS scheme may not be as inefficient as it appears. If it is true that flights with higher α values are incentivized to submit earlier, the user cost efficiency of the FSFA scheme will be higher than it has appeared in this analysis.

The concepts of this research were assessed using hypothetical scenarios. There exists a large amount of data from past and current AFPs, which not only contains information about their size and scope but also the outcomes of rerouting, cancellation, and allocation decisions. It would be helpful to use this data to not only assess the future concepts of this research on a more realistic scale, but also to compare their performance to current AFP performance, as well as better understand the airlines’ cancellation and rerouting choices in an AFP. We also continue our discussions with practitioners, in order to better understand and represent airline behavior within our modeling framework.

This research has presented future collaborative concepts for en route resource allocation, where flight operators are asked to provide very structured inputs to the resource assignment. Also, we have assessed these allocation methods from an operator cost perspective. We have employed random utility theory to account for the idea that operator preferences are not typically fully known to the traffic managers that perform resource assignments, requiring traffic managers to make assignments under uncertainty about what users truly want. This research also explored the gaming and truth-telling behavior of operators when asked to provide inputs to the resource assignment. The results from this work inform critical policy decisions regarding future FAA-airline cooperation, and the characteristics of information exchange between these parties.

Bibliography

- Albers, S. (2003). Online algorithms: a survey. *Mathematical Programming* , 97 (1-2), 3-26.
- Andreatta, G., & Romanin-Jacur, G. (1987). Aircraft Flow Management Under Congestion. *Transportation Science* , 21 (4), 249-253.
- Ball, M. O., Hoffman, R., & Mukherjee, A. (2010). Ground Delay Program Planning Under Uncertainty Based on the Ration-by-Distance Principle. *Transportation Science* , 44 (1), 14.
- Ball, M. O., Hoffman, R., Odoni, A., & Rifkin, R. (2003). A Stochastic Integer Program with Dual Network Structure and its Application to the Ground-Holding Problem. *Operations Research* , 51 (1).
- Ball, M., Futer, A., Hoffman, R., & Sherry, J. (2002). Rationing Schemes for En-route Air Traffic Management (CDM Memorandum).
- Ben-Akiva, M., & Lerman, S. R. (1994). *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press.
- Bertsimas, D., & Stock Patterson, S. (1998). The Air Traffic Flow Management Problem with Enroute Capacities. *Operations Research* , 46 (3), 406-422.
- Bertsimas, D., & Stock Patterson, S. (2000). The Traffic Flow Management Rerouting Problem in Air Traffic Control: A Dynamic Network Flow Approach. *Transportation Science* , 34 (3), 239–255.
- Bose, R., & Gupta, S. S. (1959). Moments of Order Statistics from a Normal Population. *Biometrika* , 46, 433-440.
- Bureau of Transportation Statistics. (2007). *Airline On-Time Statistics and Delay Causes*. Retrieved April 11, 2009, from Research and Innovative Technology Administration: http://www.transtats.bts.gov/OT_Delay/OT_DelayCause1.asp
- Chang, K., Howard, K., Oiesen, R., Shisler, L., Tanino, M., & Wambsganss, M. C. (2001, January-February). Enhancements to the FAA Ground-Delay Program Under Collaborative Decision Making. *Interfaces* , 31 (1), pp. 57-76.
- Clark, C. E. (1961). The Greatest of a Finite Set of Random Variables. *Operations Research* , 9 (2), 18.
- Daganzo, C. F. (1983). Stochastic Network Equilibrium with Multiple Vehicle Types and Asymmetric, Indefinite Link Cost Jacobians. *Transportation Science* , 17 (3), 282-300.

- Domencich, T. A., & McFadden, D. (1975). *Urban Travel Demand: A Behavioral Analysis*. Amsterdam: North-Holland.
- Federal Aviation Administration. (2008, May 30). *Advisory Circular 90-102A*. Retrieved February 2009, from FAA: http://www.faa.gov/documentLibrary/media/Advisory_Circular/AC%2090-102A.pdf
- Federal Aviation Administration. (2007, May 23). *Press Release – FAA Greatly Expands Air Traffic Program to Minimize Summer Delays*. Retrieved March 2009, from FAA: http://www.faa.gov/news/press_releases/news_story.cfm?newsId=8889
- Gibbons, R. (1992). *Game Theory for Applied Economists*. Princeton, New Jersey: Princeton University Press.
- Glover, C. N., & Ball, M. O. (2010). Stochastic Integer Programming Models for Ground Delay Programs with Weather Uncertainty. *4th International Conference on Research in Air Transportation*, (p. 7). Budapest, Hungary.
- Goodhart, J. (2000). *Increasing Airline Operational Control in a Constrained Air Traffic System*. University of California, Berkeley: PhD Dissertation.
- Haurie, A., & Marcotte, P. (1985). On the Relationship between Nash-Cournot and Wardrop Equilibria. *Networks*, 15 (3), 295-308.
- Hoffman, R., Ball, M. O., & Mukherjee, A. (2007). Ration-by-Distance with Equity Guarantees: A New Approach to Ground Delay Program Planning and Control. *7th Air Traffic Management (ATM) Seminar*. Barcelona.
- Hoffman, R., Burke, J., Lewis, T., Futer, A., & Ball, M. (2005). Resource Allocation Principles for Airspace Flow Control. *AIAA Guidance, Navigation, and Control Conference and Exhibit*. San Francisco.
- Hoffman, R., Hall, W., Ball, M., Odoni, A., & Wambsganss, M. (1999). *Collaborative Decision Making in Air Traffic Flow Management*. UC Berkeley. Nextor Research Report #RR-99-2.
- Hoffman, R., Lewis, T., & Jakobovits, R. (2004). *Flow Constrained Area Rerouting Decision Support Tool, Phase I SBIR: Final Report*. Metron Aviation.
- Jakobovits, R., Krozel, J., & Penny, S. (2005). Ground Delay Programs to Address Weather within En Route Flow Constrained Areas. *AIAA Guidance, Navigation and Control Conference and Exhibit*. San Francisco, CA.
- Konishi, H. (2004). Uniqueness of User Equilibrium in Transportation Networks with Heterogeneous Commuters. *Transportation Science*, 38 (3), 315–330.
- Krishna, V. (2002). *Auction Theory*. San Diego: Academic Press.
- Leurent, F. (1993). Cost versus time equilibrium over a network. *European Journal of Operational Research*, 71 (2), 205-221.
- Libby, M., Buckner, J., & Brennan, M. (2005, October). *Operational Concept for Airspace Flow Programs (AFP)*. Retrieved May 2009, from FAA Air Traffic Organization:

- http://cdm.fly.faa.gov/Workgroups/flow_eval/Operational%20Concept%20for%20Airspace%20Flow%20Programs.doc
- Miettinen, K. M. (1999). *Nonlinear Multiobjective Optimization*. Boston: Kluwer Academic Publishers.
- Moldovanu, B., & Sela, A. (2001). The Optimal Allocation of Prizes in Contests. *The American Economic Review* , 91 (3), 542-558.
- Mukherjee, A., & Hansen, M. (2009). A dynamic rerouting model for air traffic management. *Transportation Research Part B: Methodological* , 43, 159-171.
- Mukherjee, A., & Hansen, M. (2007). A Dynamic Stochastic Model for the Single Airport Ground Holding Problem. *Transportation Science* , 41 (4), 444-456.
- Newell, G. (1987). The morning commute for non-identical travelers. *Transportation Science* , 21 (2), 74-88.
- Peeta, S., & Ziliaskopoulos, A. K. (2004). Foundations of Dynamic Traffic Assignment: The Past, the Present and the Future. *Networks and Spatial Economics* , 233-265.
- Richetta, O., & Odoni, A. R. (1993). Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control. *Transportation Science* , 27 (3), 228-238.
- Sheffi, Y. (1985). *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods*. Prentice-Hall.
- Teichroew, D. (1956). Tables of Expected Values of Order Statistics and Products of Order Statistics for Samples of Size Twenty or Less from the Normal Distribution. *The Annals of Mathematical Statistics* , 27, 410-426.
- Train, K. (2003). *Discrete Choice Methods with Simulation*. Cambridge University Press.
- Vakili Pourtaklo, N., & Ball, M. (2009). Equitable Allocation of Enroute Airspace Resources. *Eighth USA/Europe Air Traffic Management Research and Development Seminar (ATM2009)*. Napa.
- Vossen, T., & Ball, M. (2006). Optimization and Mediated Bartering Models for Ground Delay Program. *Naval Research Logistics* , 53 (1), 75-90.
- Vossen, T., Ball, M., Hoffman, R., & Wambsganss, M. (2003). A General Approach to Equity in Traffic Flow Management and its Application to Mitigating Exemption Bias in Ground Delay Programs. *5th EUROCONTROL/FAA ATM R&D Seminar*. Budapest, Hungary.
- Waslander, S. L., Raffard, R. L., & Tomlin, C. J. (2008). Market-Based Air Traffic Flow Control with Competing Airlines. *Journal of Guidance, Control, and Dynamics* , 31 (1), 148-161.
- Waslander, S. L., Roy, K., Johari, R., & Tomlin, C. J. (2008). Lump-Sum Markets for Air Traffic Flow Control With Competitive Airlines. *Proceedings of the IEEE* , 96 (12), 2113-2130.
- Wie, B.-W. (1995). A differential game approach to the dynamic mixed behavior traffic network equilibrium problem. *European Journal of Operational Research* , 83, 117-136.

- Wie, B.-W. (1993). A Differential Game Model of Nash Equilibrium on a Congested Traffic Network. *Networks* , 23 (6), 557-565.
- Wilmouth, G., & Taber, N. J. (2005). *Operational Concept for Integrated Collaborative Rerouting (ICR)*. McLean, VA: The MITRE Corporation, MTR05W000053.
- Xiong, J. (2010). *Revealed Preference of Airlines' Behavior under Air Traffic Management Initiatives*. University of California, Berkeley: PhD Dissertation.

Appendix A:

Stochastic properties of the allocation schemes

A.1) Expected value of the minimum of two iid normal random variables

Since ε_{nr} are iid normal, the moments of the maximum of N random variables can be calculated (Bose & Gupta, 1959) (Teichroew, 1956). The equations are summarized in Clark (1961); the first moment of two independent normal random variables, $\varepsilon_1 \sim N(\mu_1, \sigma_1)$ and $\varepsilon_2 \sim N(\mu_2, \sigma_2)$ is as follows:

$$E[\max(\varepsilon_1, \varepsilon_2)] = \mu_1 \Phi(\alpha) + \mu_2 \Phi(-\alpha) + a\varphi(\alpha) \quad (\text{A.1})$$

where

$$a = \sigma_1^2 + \sigma_2^2 - 2\sigma_1\sigma_2\rho_{1,2}$$

$$\varphi(\alpha) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right)$$

$$x^2 = \alpha^2 = (\mu_1 - \mu_2)^2 / a^2 = 0$$

Recall that $\varepsilon_{n,r} \sim iid N(0, \sigma) \forall n, r$. It then follows that $\rho_{1,2} = 0$, $\mu_1 = \mu_2 = 0$, and $\sigma_1 = \sigma_2 = \sigma$. Equation (A.1) reduces to

$$E[\max(\varepsilon_1, \varepsilon_2)] = \sigma/\sqrt{\pi}$$

Since the normal distribution is symmetric around the mean,

$$E[\min(\varepsilon_1, \varepsilon_2)] = -E[\max(\varepsilon_1, \varepsilon_2)]$$

Recall $\varepsilon_{11}, \varepsilon_{12}, \varepsilon_{21}, \varepsilon_{22} \sim N(0, \sigma)$. Equation (3.5b) therefore reduces to

$$E[C_{FSFA}] = W + E[\min(\varepsilon_1, \varepsilon_2)] + E[\varepsilon_3] = W - \sigma/\sqrt{\pi}$$

Next, set $\varepsilon_A = \varepsilon_{11} + \varepsilon_{22}$ and $\varepsilon_B = \varepsilon_{12} + \varepsilon_{21}$. Then, we know that $var(\varepsilon_A) = var(\varepsilon_B) = var(\sum_{i=1}^2 \varepsilon_i) = \sum_{i=1}^2 var(\varepsilon_i) = 2\sigma^2$, and $\varepsilon_A, \varepsilon_B \sim N(0, \sqrt{2}\sigma)$ according to the properties of normal iid random variables. Equation (3.5a) becomes

$$E[C_{OPT}] = W + E[\min(\varepsilon_A, \varepsilon_B)] = W - \sqrt{2}\sigma/\sqrt{\pi} \quad \square$$

A.2) Some relationship properties

Define

$$E[C'_{FSFA}] = \frac{E[C_{FSFA}]}{E[C_{OPT}]} = \frac{W - \sigma/\sqrt{\pi}}{W - \sqrt{2}\sigma/\sqrt{\pi}} = \frac{W\sqrt{\pi} - \sigma}{W\sqrt{\pi} - \sqrt{2}\sigma}$$

$$E[C'_{PO}] = \frac{E[C_{PO}]}{E[C_{OPT}]} = \frac{W}{W - \sqrt{2}\sigma/\sqrt{\pi}} = \frac{W\sqrt{\pi}}{W\sqrt{\pi} - \sqrt{2}\sigma}$$

We also know that $\lim_{\sigma \rightarrow 0} E[C'_{FSFA}] = 1$, $\lim_{\sigma \rightarrow \infty} E[C'_{FSFA}] = \infty$, and $\lim_{\sigma \rightarrow 0} E[C'_{PO}] = 1$, $\lim_{\sigma \rightarrow \infty} E[C'_{PO}] = \infty$.

It is clear that as $\sigma \rightarrow \infty$, the denominators of both $E[C'_{FSFA}]$ and $E[C'_{PO}]$ decrease at faster rates than the numerators, indicating that $E[C'_{FSFA}]$ and $E[C'_{PO}]$ increase non-linearly. To verify,

$$\frac{\partial(E[C'_{FSFA}])}{\partial\sigma} = \frac{W(\sqrt{2}-1)\sqrt{\pi}}{(W\sqrt{\pi} - \sqrt{2}\sigma)^2} \in \mathbb{R}^+, \frac{\partial^2(E[C'_{FSFA}])}{\partial\sigma^2} \in \mathbb{R}^+; \sigma \geq 0$$

$$\frac{\partial(E[C'_{PO}])}{\partial\sigma} = \frac{W\sqrt{2\pi}}{(W\sqrt{\pi} - \sqrt{2}\sigma)^2} \in \mathbb{R}^+, \frac{\partial^2(E[C'_{PO}])}{\partial\sigma^2} \in \mathbb{R}^+; \sigma \geq 0$$

Both $E[C'_{FSFA}]$ and $E[C'_{PO}]$ are non-linearly increasing functions of σ . Their slopes increase with respect to σ as well. It then follows that, due to the construct of the toy model where the total deterministic cost of any allocation is always W , $E[C'_{FSFA}]_{\sigma} > E[C'_{PO}]_{\sigma}, \sigma \geq 0$. Also,

$$E[C'_{PO}]'/E[C'_{FSFA}]' = \frac{\sqrt{2}}{\sqrt{2}-1} = 3.414. \quad \square$$

Say the deterministic cost of the FSFA allocation is W_{FSFA} and that of the OPT allocation is W_{OPT} when $\sigma = 0$. Then $\lim_{\sigma \rightarrow 0} E[C'_{FSFA}] = W_{FSFA}/W_{OPT}$.

Appendix B:

Properties of the Parametric Approximation

B.1) Aircraft distribution based on cost parameters

We show that flights with higher α values should be assigned to routes with lower en route times, and vice versa, to obtain the assignment that yields the optimal total AFP flight cost solution. This is the assumption on which Equation 3.17 is formulated.

Let's assume a two-route case where

- Route 1 has a higher flying time than Route 2, such that $\rho_1 > \rho_2$;
- Slots on both routes are spaced at equal (constant) headways, or $g_1 = g_2 = g$;
- α are distributed uniformly across N flights between $(\alpha_{min}, \alpha_{max}]$, and
- $\widehat{C}_{n,1} = \alpha(n) \cdot \rho_1 + gn$, $\widehat{C}_{n,2} = \alpha(n) \cdot \rho_2 + gn$ (as per Equation (2.2))

Based on (3.17), the total cost of assignment is

$$\hat{C}_1 = X_1 \alpha_{min} \rho_1 + 0.5(\theta \rho_1 + g) X_1 (X_1 + 1) + X_2 (\alpha_{min} + \theta X_1) \rho_2 + 0.5(\theta \rho_2 + g) X_2 (X_2 + 1)$$

If we swap the flight having the highest α value assigned to Route 1 with the flight having the lowest α value assigned to Route 2, we have

$$\hat{C}_2 = X_1 \alpha_{min} \rho_1 + 0.5(\theta \rho_1 + g) X_1 (X_1 + 1) + \theta \rho_1 + X_2 (\alpha_{min} + \theta X_1) \rho_2 + 0.5(\theta \rho_2 + g) X_2 (X_2 + 1) - \theta \rho_2$$

It must be true that $\hat{C}_1 < \hat{C}_2$ for the ordering assumption of (3.17) to be true. Taking out identical terms, we have $0 < \theta(\rho_1 - \rho_2)$. Because $\rho_1 > \rho_2$, the RHS is positive, and therefore $\hat{C}_1 < \hat{C}_2$. The result clearly also holds if $R > 2$.

In the above we have assumed that the assignment and therefore X_1, X_2 do not change. If we allow that they change, we have

$$\begin{aligned} \hat{C}_1 = & (8Ng^2 + 4N^2g^2 + 4Nt^2\rho_2\rho_1 + 8\theta\alpha_{min}\rho_1\rho_2 + 8N\theta\alpha_{min}\rho_2\rho_1 - 8N\theta\alpha_{min}\rho_2^2 \\ & + 4N^2\theta g\rho_2 + 4N^2\theta g\rho_1 - 4\theta\alpha_{min}\rho_1^2 + 8Ng\alpha_{min}\rho_2 + 8\alpha_{min}^2\rho_1\rho_2 - 4N^2\theta^2\rho_2^2 \\ & - 4N\theta^2\rho_2^2 + 8Ng\alpha_{min}\rho_1 - 4\alpha_{min}^2\rho_2^2 - 4\theta\alpha_{min}\rho_2^2 - 4\alpha_{min}^2\rho_1^2 + 8N\theta g\rho_1 \\ & + 4N^2\theta^2\rho_2\rho_1 - \rho_2^2\theta^2 - \rho_1^2\theta^2 + 2\theta^2\rho_1\rho_2)/8(2g - \rho_2\theta + \rho_1\theta) \end{aligned} \quad (B.1)$$

$$\begin{aligned}
\hat{C}_2 = & (8Ng^2 + 4N^2g^2 + 4Nt^2\rho_2\rho_1 + 8\theta\alpha_{\min}\rho_1\rho_2 + 8N\theta\alpha_{\min}\rho_2\rho_1 - 8N\theta\alpha_{\min}\rho_2^2 \\
& + 4N^2\theta g\rho_2 + 4N^2\theta g\rho_1 - 4\theta\alpha_{\min}\rho_1^2 + 8Ng\alpha_{\min}\rho_2 + 8\alpha_{\min}^2\rho_1\rho_2 - 4N^2\theta^2\rho_2^2 \\
& - 4N\theta^2\rho_2^2 + 8Ng\alpha_{\min}\rho_1 - 4\alpha_{\min}^2\rho_2^2 - 4\theta\alpha_{\min}\rho_2^2 - 4\alpha_{\min}^2\rho_1^2 + 8N\theta g\rho_1 \\
& + 4N^2\theta^2\rho_2\rho_1 - 16\rho_2\theta g + 16\rho_1\theta g - 14\theta^2\rho_1\rho_2 + 7\rho_1^2\theta^2 + 7\rho_2^2\theta^2)/8(2g - \rho_2\theta \\
& + \rho_1\theta)
\end{aligned} \tag{B.2}$$

Designate the identical terms of (B.1) and (B.2) as K such that

$$\hat{C}_1 = (K - \rho_2^2\theta^2 - \rho_1^2\theta^2 + 2\theta^2\rho_1\rho_2)/8(2g - \rho_2\theta + \rho_1\theta)$$

$$\hat{C}_2 = (K - 16\rho_2\theta g + 16\rho_1\theta g - 14\theta^2\rho_1\rho_2 + 7\rho_1^2\theta^2 + 7\rho_2^2\theta^2)/8(2g - \rho_2\theta + \rho_1\theta)$$

And show that $\hat{C}_1 < \hat{C}_2$.

$$\begin{aligned}
2\theta^2\rho_1\rho_2 - \theta^2(\rho_1^2 + \rho_2^2) & < 16\theta g(\rho_1 - \rho_2) - 14\theta^2\rho_1\rho_2 + 7\theta^2(\rho_1^2 + \rho_2^2) \\
\Leftrightarrow -\theta^2(\rho_1 - \rho_2)^2 & < 16\theta g(\rho_1 - \rho_2) + 7\theta^2(\rho_1 - \rho_2)^2 \\
\Leftrightarrow 0 & \\
& < 2\theta g(\rho_1 - \rho_2) + \theta^2(\rho_1 - \rho_2)^2
\end{aligned}$$

Since $\rho_1 > \rho_2$, all terms on the RHS are positive and $\hat{C}_1 < \hat{C}_2$. \square

B.2) Convexity

The objective function of the Parametric approximation (and equivalently, OPT when $\sigma = 0$) can be represented in quadratic form:

$$\min_x \frac{1}{2}x^T \mathbf{Q}x + f^T x$$

$$\text{where } \mathbf{A} \cdot x \leq b$$

$$\mathbf{A}_{\text{eq}} \cdot x = b_{\text{eq}}$$

If \mathbf{Q} is positive semi-definite then the objective function is convex. There exists a global minimum solution x^* ($= X_1^*, X_2^*, \dots, X_R^*$) if there is at least one x that satisfies the constraints. If \mathbf{Q} is positive definite then the global minimum solution x^* is unique. A matrix \mathbf{Q} is positive definite if and only if all its eigenvalues are positive.

B.2.1) Homogeneous flights

When $\alpha_n = \alpha \forall n$, then $f = \alpha\rho_r + 0.5g_r \forall r$ and \mathbf{Q} is a diagonal matrix.

$$\mathbf{Q} = \begin{bmatrix} g_1 & 0 & 0 & 0 \\ 0 & g_2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & g_R \end{bmatrix}, \text{ where } g_1, g_2, \dots, g_n > 0.$$

Because \mathbf{Q} is diagonal, the eigenvalues of \mathbf{Q} are the diagonal values of \mathbf{Q} . The eigenvalues of \mathbf{Q} are positive, indicating that \mathbf{Q} is positive definite and the objective function is convex with respect to the number of flights assigned to each route, or $X_1^*, X_2^*, \dots, X_R^*$. \square

Flights are not uniquely matched to their original scheduled departure times g_0 for the formulation of our approximation; as a result, they are identical and therefore not uniquely defined within the objective function. This is why the solution is unique – there is only one solution x^* , or the number of flights to each route, that minimizes the objective function. Flights themselves can be interchanged between routes within this solution. As a result, if we do assume that each flight is uniquely defined by their original departure time or even some other feature, then there are many assignment solutions within the unique solution x^* .

B.2.2) Heterogeneous flights

When α_n take unique values that are uniformly distributed in $(\alpha_{min}, \alpha_{max}]$, then $f = \alpha_{min}\rho_r + 0.5\theta\rho_r + 0.5g_r \forall r$ and Q is an upper triangular matrix.

$$Q = \begin{bmatrix} \theta\rho_1 + g_1 & \theta\rho_2 & \theta\rho_3 & \theta\rho_4 & \dots & \theta\rho_R \\ 0 & \theta\rho_2 + g_2 & \theta\rho_3 & \theta\rho_4 & \dots & \theta\rho_R \\ 0 & 0 & \theta\rho_3 + g_3 & \theta\rho_4 & \dots & \theta\rho_R \\ 0 & 0 & 0 & \theta\rho_4 + g_4 & \dots & \vdots \\ \vdots & \vdots & \vdots & \dots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \dots & \theta\rho_R + g_R \end{bmatrix}$$

Where $g_1, g_2, \dots, g_n, \theta > 0$ and $\rho_1, \rho_2, \dots, \rho_R \geq 0$.

$$\text{Set } Q_2 = \begin{bmatrix} \theta\rho_1 + g_1 & \theta\rho_2 \\ 0 & \theta\rho_2 + g_2 \end{bmatrix}, Q_3 = \begin{bmatrix} \theta\rho_1 + g_1 & \theta\rho_2 & \theta\rho_3 \\ 0 & \theta\rho_2 + g_2 & \theta\rho_3 \\ 0 & 0 & \theta\rho_3 + g_3 \end{bmatrix}, \dots,$$

$$Q_{R-1} = \begin{bmatrix} \theta\rho_1 + g_1 & \theta\rho_2 & \dots & \theta\rho_{R-1} \\ 0 & \theta\rho_2 + g_2 & \dots & \theta\rho_{R-1} \\ 0 & 0 & \ddots & \vdots \\ 0 & 0 & 0 & \theta\rho_{R-1} + g_{R-1} \end{bmatrix}, Q = Q_R$$

Since $\det(Q_2), \det(Q_3), \dots, \det(Q) > 0$, Q is positive definite and the objective function is convex with a unique solution. \square

Appendix C:

Chapter 4 calculations

C.1) Untruthful α submissions, 3-route system-optimal case

According to (3.17) the total cost of a system-optimal allocation using untruthful air-to-ground cost ratio α^L is:

$$C = \alpha^L X_1 \rho_1 + 0.5gX_1^2 + \alpha^L X_2 \rho_2 + 0.5gX_2^2 + \alpha^L X_3 \rho_3 + 0.5gX_3^2$$

where $X_3 = N - X_1 - X_2$.

$$\frac{\partial C}{\partial X_1} = \alpha^L(\rho_1 - \rho_3) + 2gX_1 + gX_2 - gN = 0$$

$$\frac{\partial C}{\partial X_2} = \alpha^L(\rho_2 - \rho_3) + gX_1 + 2gX_2 - gN = 0$$

By solving for the above, we find that

$$X_1 = \frac{N}{3} + \frac{\alpha^L}{3g} \cdot (\rho_3 + \rho_2 - 2\rho_1), \quad X_2 = \frac{N}{3} + \frac{\alpha^L}{3g} \cdot (\rho_3 - 2\rho_2 + \rho_1),$$

$$X_3 = \frac{N}{3} + \frac{\alpha^L}{3g} \cdot (-2\rho_3 + \rho_2 + \rho_1)$$

The expected cost of flying any route is equal at the truthful user equilibrium with α^T . The following must be satisfied:

$$\alpha^T \rho_1 + 0.5gX_1 = \alpha^T \rho_2 + 0.5gX_2 = \alpha^T \rho_3 + 0.5gX_3$$

And we find that

$$X_2 = \frac{2\alpha^T(\rho_1 - 2\rho_2 + \rho_3) + gN}{3g} \quad \text{At the UE.}$$

We solve the two expressions for X_2 shown above and find that $\alpha^L = 2\alpha^T$, which is identical to the two-route case. \square

C.2) Equilibrium FSFA airline preference submission strategy

$P(q_1 > q_2)$ is the probability that Player 1 submits earlier than Player 2. Recall that $P(q_1 > q_2) = P(t_1 < t_2)$. If we assume a-priori that $g(h_n)$ is monotonic and differentiable, then we can say that:

$$\begin{aligned} P(q_1 > q_2) &= P(g(h_1) > g(h_2)) \\ &= P(g^{-1}(q_1) < g^{-1}(q_2)) \\ &= 1 - P(g^{-1}(q_2) < g^{-1}(q_1)) \\ &= 1 - F(g^{-1}(q_1)) \end{aligned}$$

where $q_n = g(h_n)$ is the submission time strategy for player n .

We assume that the probability of winning or losing against other players is independent. Therefore, the probability that player n submits earlier than both players is $(1 - F(g^{-1}(q_n)))^2$. The probability that player n submits earlier than one player and later than the other is $2(1 - F(g^{-1}(q_n))) \cdot F(g^{-1}(q_n))$. The payoff function for n becomes:

$$\begin{aligned} E[\pi_n(q_n)] &= \frac{r_1}{L^{max}} \cdot (1 - F(g^{-1}(q_n)))^2 \\ &\quad + \frac{2r_2}{L^{max}} \cdot (1 - F(g^{-1}(q_n))) \cdot F(g^{-1}(q_n)) - h_n l(q_n) \end{aligned}$$

Set $g^{-1}(q_n) = h_n = y_n$.

$$\begin{aligned} E[\pi_n(q_n)] &= \frac{r_1}{L^{max}} \cdot (1 - F(y_n))^2 \\ &\quad + \frac{2r_2}{L^{max}} \cdot (1 - F(y_n)) \cdot F(y_n) - h_n l(q_n) \end{aligned}$$

$$\begin{aligned} \frac{\partial E[\pi_n]}{\partial q_n} &= -\frac{2r_1}{L^{max}} \cdot (1 - F(y_n)) \cdot F'(y_n) \cdot y'_n - \frac{2r_2}{L^{max}} \cdot F'(y_n) \cdot y'_n \cdot F(y_n) \\ &\quad + \frac{2r_2}{L^{max}} \cdot (1 - F(y_n)) \cdot F'(y_n) \cdot y'_n - y_n L'(q_n) = 0 \\ \Leftrightarrow L'(q_n) &= -\frac{2r_1}{y_n L^{max}} \cdot (1 - F(y_n)) \cdot F'(y_n) \cdot dy_n \\ &\quad - \frac{2r_2}{y_n L^{max}} \cdot (2F(y) - 1) \cdot F'(y_n) \cdot y'_n \end{aligned}$$

Now we determine boundary conditions. If $h = h_{max}$ (drop the subscript), the highest uncertainty level possible, we conjecture that n will submit as late in the AFP planning period as possible, at T (or $q \approx 0$). Otherwise, when $h < h_{max}$, $q > 0$. Therefore,

$$\int_q^0 l'(x)dx = -\frac{2r_1}{L^{max}} \int_h^{h_{max}} x^{-1} \cdot (1 - F(x)) \cdot F'(x)dx$$

$$-\frac{2r_2}{L^{max}} \int_h^{h_{max}} x^{-1} \cdot (2F(x) - 1) \cdot F'(x)dx$$

We know that the operators' h take values that are uniformly distributed between h_{min} and h_{max} . If $a = h_{min}$ and $b = h_{max}$, $F(x) = \frac{x-a}{b-a}$ and $F'(y) = \frac{1}{b-a}$.

$$l(q) = \frac{2r_1}{L^{max}} \int_{h_n}^{h_{max}} \frac{1}{(b-a)x} \left(1 - \frac{x-a}{b-a}\right) dx$$

$$+ \frac{2r_2}{L^{max}} \int_{h_n}^{h_{max}} \frac{1}{(b-a)x} \left(\frac{2(x-a)}{b-a} - 1\right) dx$$

$$= \frac{2}{L^{max}(b-a)^2} \left(r_1(b \ln x - x) + r_2(2x - (a+b) \ln x) \right) \Big|_h^{h_{max}}$$

Evaluate and replace a and b to obtain

$$l(q) = \frac{2}{L^{max}(h_{max} - h_{min})^2} \left[(r_1 h_{max} - r_2 \cdot (h_{min} + h_{max})) \cdot (\ln h_{max} - \ln h) \right.$$

$$\left. + (r_1 - 2r_2) \cdot (h - h_{max}) \right]$$

Replace $l(q) = \left(1 - \left(\sum_{s \in S} \exp\left(\frac{v_s}{v(q)}\right)\right)^{-1} \sum_{s \in S} v_s \cdot \exp\left(\frac{v_{n,s}}{v(q)}\right)\right) / L^{max}$:

$$1 - \left(\sum_{s \in S} \exp\left(\frac{v_s}{qp}\right)\right)^{-1} \sum_{s \in S} v_s \cdot \exp\left(\frac{v_s}{qp}\right)$$

$$= \frac{2 \left[(r_1 h_{max} - r_2 \cdot (h_{min} + h_{max})) (\ln h_{max} - \ln h) + (r_1 - 2r_2)(h - h_{max}) \right]}{(h_{max} - h_{min})^2}$$

Recall that we assumed the submission strategy to be monotonic and differentiable a-priori. The numerical example of Figure 4.6 shows that the submission strategy is increasing through the planning period. All players that desire to submit before the planning period submit at $t = 0$. Moldovanu and Sela (2001) also prove that the bid function is strictly increasing and differentiable, and that it maximizes expected payoff.

If $r_2 = 0.5r_1$ then (4.15) becomes:


$$1 - \left(\sum_{s \in S} \exp\left(\frac{v_s}{pq}\right)\right)^{-1} \sum_{s \in S} v_s \cdot \exp\left(\frac{v_s}{pq}\right) = \frac{r_1(h_{max} - h_{min})(\ln h_{max} - \ln h)}{(h_{max} - h_{min})^2}$$

If $L(q)$ were linear with a form such as pq , we would have a closed form solution:

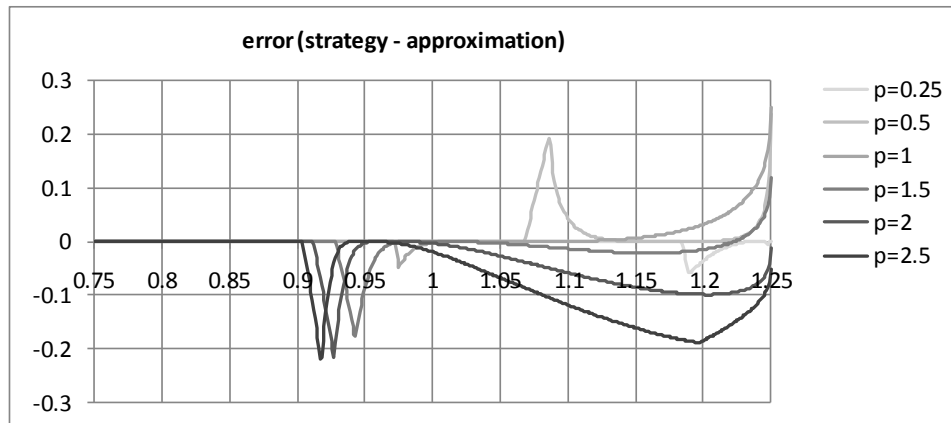
$$t = \max\left(T \left(1 - \frac{2 \left[(r_1 h_{max} - r_2 \cdot (h_{min} + h_{max})) (\ln h_{max} - \ln h) + (r_1 - 2r_2)(h - h_{max}) \right]}{(h_{max} - h_{min})^2 p}\right), 0\right)$$

C.3) Submission strategy approximation

Below is a plot of example submission strategies over six p values, with $v_s = [1,0.4,0]$, $r_1 = 0.8$, $r_2 = 0.8v_2$, $h_{min} = 0.75$, $h_{max} = 1.25$, and $T = 2$. The solid lines labeled “strategy” are exact solutions while the dashed lines labeled “T3(0.55)” are results from third-order Taylor series approximations about the point $q = 0.55$. The x -axis represents values of h from h_{min} to h_{max} , and the y -axis represents the AFP planning period.



The approximation does not provide good estimates to the expression of (4.15) near the boundaries of the planning period (i.e. the farthest points from $q = 0.55$) and at higher p values. For instance, when $p \geq 2$ (i.e. p is at least twice that of V^*) the approximation overestimates the time at which a player n would submit according to (4.15). At p values less than about one, the approximation underestimates submission times near the boundaries of the planning period. The approximations do not capture the players’ strategy to submit closer to T with a high h_n (i.e. the sharp increase in slope at high h_n).



Appendix D:

Ration-by-Schedule

Proposition 5.1. *Flight-slot assignments using the RBS algorithm will always have equal or higher costs compared to assignments using the FSFA algorithm.*

This problem may be similar/analogous to an offline, first fit increasing bin packing problem – the first item (first scheduled flight) is the “smallest” item and has a “lot of room to fit” (i.e., the first scheduled flight can “fit” into many slots). The last item is the “largest” and therefore has the least room to fit. As a result, we have some intuition that allocation in this order will be amongst the least optimal.

Proof:

Let’s assume the following are true:

- All flights are identical except for their OAG scheduled departure times, which we assume are all distinct.
- Decisions are made at some time prior to start of AFP, before the first flight’s OAG scheduled departure time (S1), and are not modified with updated airline info. In fact we will assume that all parameters (airline info, AFP characteristics) are static.

Let A1 be an assignment by RBS, and A2 an assignment by FSFA. We argue that although A1 and A2 may not be identical, they have the same total cost assignment cost. Below, we prove that the above statement is not true in all cases.

We have 2 flights, on 2 routes with 2 slots each. Flight 1’s OAG departure time is S1, and Flight 2’s is S2; $S1 < S2$ by definition. The AFP departure times for each slot on each route are listed in the left table, and the cost of taking each slot (not accounting for OAG times) are listed in the table on the right.

Departure Times			Route Costs		
Slot	Route 1	Route 2	Slot	Route 1	Route 2
1	D11	D21	1	C11	C21
2	D12	D22	2	C12	C22

Let’s assume that $D11 < D21 < D12, D22$. This means that Route 1 Slot 1’s departure time occurs prior to that of Route 2 Slot 1, and both occur prior to those of Slot 2 on either route. We now have five scenarios:

1. $S1 < S2 < D11 < D21 < D12$
2. $S1 < D11 < S2 < D21 < D12$
3. $S1 < D11 < D21 < S2 < D12$
4. $D11 < S1 < D21 < S2 < D12$
5. $D11 < D21 < S1 < S2 < D12$

In each of these scenarios we can have a) $C11 < C21 < C12, C22$ or b) $C21 < C11 < C12, C22$.

1) **$S1 < S2 < D11 < D21$**

a) **$C11 < C21 < C12, C22$**

If we allocate by RBS, Flight 1 first chooses a slot, then Flight 2. Allocation by FSFA includes either ordering.

RBS	Slot choice	Cost	FSFA	Slot choice	Cost
F1	R1 slot1	$C11 - S1$	F2	R1 slot1	$C11 - S2$
F2	R2 slot1	$C21 - S2$	F1	R2 slot1	$C21 - S1$
Total		$C11 + C21 - S1 - S2$	Total		$C11 + C21 - S2 - S1$

Also, we can say that the total cost of a system-optimal allocation is:

$$C_{SO} = \min [(C11 + C21), (C21 + C11), (C11 + C12), (C12 + C11), (C21 + C12), (C12 + C21)] - S1 - S2 = \min(A, B, C, D, E, F) - S1 - S2$$

We know that $A = B, C = D, E = F$. Because $C11 < C21, C12$, it follows that $A < C < E$. Therefore $C_{SO} = C11 + C21 - S1 - S2$.

Let's say that $A = \min(C11, C21, C12) = C11$. The cost of an RBS allocation is

$$C_{RBS} = A + (\min(C11, C21, C12) \notin A) - S1 - S2 = C11 + C21 - S1 - S2$$

And $C_{SO} = C_{RBS}$. The total costs of ordering either way are identical. \square

b) **$C21 < C11 < C12, C22$**

RBS	Slot choice	Cost	FSFA	Slot choice	Cost
F1	C21	$C21 - S1$	F2	C21	$C21 - S2$
F2	C11	$C11 - S2$	F1	C11	$C11 - S1$
Total		$C11 + C21 - S1 - S2$	Total		$C11 + C21 - S2 - S1$

Identical. \square

2) **$S1 < D11 < S2 < D21 < D12$**

a) **$C11 < C21 < C12, C22$**

RBS	Slot choice	Cost	FSFA	Slot choice	Cost
F1	C11	$C11 - S1$	F2	C21	$C21 - S2$
F2	C21	$C21 - S2$	F1	C11	$C11 - S1$
Total		$C11 + C21 - S1 - S2$	Total		$C11 + C21 - S1 - S2$

b) **$C21 < C11 < C12, C22$**

RBS	Slot choice	Cost	FSFA	Slot choice	Cost
F1	C21	C21-S1	F2	C21	C21-S2
F2	C12	C12-S2	F1	C11	C11-S1
Total		C21+C12-S1-S2	Total		C21+C11-S1-S2

Also, we can say that the total cost of a system-optimal allocation is:

$$C_{SO} = \min(C_{11} + C_{21}, C_{12} + C_{21}, C_{11} + C_{12}) - S1 - S2$$

$$= C_{11} + C_{21} - S1 - S2$$

where Flight 1 takes slot 11 and Flight 2 takes slot 21. Let's say that A is the set of slots from which Flight 1 will choose, such that $A = \min(C_{11}, C_{21}, C_{12}) = C_{21}$. The cost of an RBS allocation is

$$C_{RBS} = A + (\min(C_{21}, C_{12}) \notin A) - S1 - S2 = C_{21} + C_{12} - S1 - S2$$

Since $C_{11} < C_{12}, C_{SO} < C_{RBS}$. The RBS assignment costs more than the FSFA assignment. \square

3) **S1<D11<D21<S2<D12**

a) **C11<C21<C12,C22**

RBS	Slot choice	Cost	FSFA	Slot choice	Cost
F1	C11	C11-S1	F2	C12	C12-S2
F2	C12	C12-S2	F1	C11	C11-S1
Total		C11+C12-S1-S2	Total		C11+C12-S1-S2

b) **C21<C11<C12,C22**

RBS	Slot choice	Cost	FSFA	Slot choice	Cost
F1	C21	C21-S1	F2	C12	C12-S2
F2	C12	C12-S2	F1	C21	C21-S1
Total		C12+C21-S1-S2	Total		C12+C21-S1-S2

4) **D11<S1<D21<S2<D12**

a) **C11<C21<C12,C22**

RBS	Slot choice	Cost	FSFA	Slot choice	Cost
F1	C21	C21-S1	F2	C12	C12-S2
F2	C12	C12-S2	F1	C21	C21-S1
Total		C12+C21-S1-S2	Total		C12+C21-S1-S2

b) **C21<C11<C12,C22**

RBS	Slot choice	Cost	FSFA	Slot choice	Cost
F1	C21	C21-S1	F2	C12	C12-S2
F2	C12	C12-S2	F1	C21	C21-S1
Total		C12+C21-S1-S2	Total		C12+C21-S1-S2

Same situation for slots 2 and 3 as in scenario 1. \square

To complete the proof, do this again for flights 1 and n , 2 and n , ..., $n - 1$ and n . \square