Revealed Preference of Airlines' Behavior under Air Traffic Management Initiatives

by

Jing Xiong

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 M. Hansen, Chair
Professor Carlos F. Daganzo
Professor James L. Powell

Spring 2010

Revealed Preference of Airlines' Behavior under Air Traffic Management Initiatives

Abstract

Revealed Preference of Airlines' Behavior under Air Traffic Management Initiatives

by

Jing Xiong

Doctor of Philosophy in Engineering

University of California, Berkeley

Professor Mark Hansen, Chair

The Federal Aviation Administration uses Air Traffic Management Initiatives (TMIs) to mitigate the consequences of aviation system capacity shortfalls, for example by delaying aircraft on the ground at their origin airports. In order to make more efficient use of National Airspace (NAS) resources, reduce delay costs, and increase the flexibility of NAS users to meet their operational needs, tremendous efforts have been made to design TMIs in a manner to encourage cooperation between the FAA and airlines. Airlines are offered opportunities to make choices such as cancelling flights and increasing delays on some flights while decreasing delays on others. However, there has been little study of airlines' resulting behavior. In this dissertation, we analyze choices made by airlines in response to TMIs and attempt to infer from these key features of airlines' preference structures. Two econometric models are specified and estimated. The first model focuses on airlines' flight cancellation decisions, and the second model examines airline requests to simultaneously re-assign arrival slots and cancel flights using Slot Credit Substitution (SCS) messages.

The cancellation model captures how airlines value delays of the subject flight itself and potential delay savings of other flights in making a flight cancellation decision. Aircraft size, along with segment frequency and load factor, are all significant factors in cancellation decisions; larger, fuller, and less frequent flights are less likely to be cancelled. Somewhat surprisingly, a higher average fare is found to increase cancellation probability. Hub-bound flights are found more likely to be cancelled than spoke-bound flights. The model also confirms airlines' hedging behavior in response to TMIs by preferentially cancelling short-haul flights. In addition, a piece wise linear specification of the utility function confirms that the delay impact is non-linear. Individual airline model reveals some consistent behavior as well as some differences in how different factors enter into cancellation decisions.

The SCS model captures airlines' tradeoff behavior in dealing with flight cancellations and delays. It confirms that cancelling flights decreases airlines' utility while reducing delays increases the utility. Moreover, airlines are sensitive to the aircraft

size and average fare of flights in performing these actions. In this model, however, average fare has the expected sign. The model estimates that airlines are willing to cancel a flight if the cancellation can reduce around 100 minutes of delays on their other flights that are in the ground delay program.

To my grandparents,

Xiong Baogeng and Yao Shunxiang

Table of Contents

List of Figures iv
List of Tables v
Acknowledgements vi
Chapter 1 Introduction 1
1.1 Problem Statement 1
1.2 Dissertation Overview 2
1.3 Organization 3
Chapter 2 Background and Literature Review 4
2.1 Background 4
2.1.1 Traffic Management Initiatives (TMIs) 4
2.1.2 CDM 5
2.1.3 Airlines Practices in Schedule Disruption Management 7
2.2 Literature Review 8
2.2.1 Delays 9
2.2.2 Flight Cancellation 12
Chapter 3 Econometric Framework 14
Chapter 4 Cancellation Model 17
4.1 Introduction 17
4.2 Major Explanatory Variables Development 17
4.2.1 Delay Factors 17
4.2.2 Flight Characteristics 20
4.2.3 Segment Characteristics 21
4.2.4 Summary 22
4.3 Data Description 23

4.3.1 Data Preparation ... 23
4.3.2 Statistics of Sample Data ... 24
4.4. Model Specifications and Estimation Results ... 27
4.4.1 Baseline Model C-1 ... 27
4.4.2 Piece-wise Linear Model C-2 ... 32
4.4.3 Model C-3 with Quadratic/Interactive Terms ... 34
4.4.4 Mixed Logit Model C-4 ... 37
4.4.5 Individual Airline Model C-5 ... 39
4.5 Cancellation and Delay Tradeoff ... 43
4.6 Summary ... 45
Chapter 5 SCS Model ... 47
5.1 Introduction ... 47
5.2 SCS Mechanism ... 47
5.3 Choice Set Generation ... 48
5.4 Correlated Alternatives ... 50
5.5 Data ... 52
5.6 Explanatory Variables ... 54
5.7 Model Specification and Results ... 55
5.7.1 Standard Logit Model ... 55
5.7.2 Mixed Logit Model ... 57
5.8 Summary ... 60
Chapter 6 Conclusions ... 62
6.1 Summary of Findings ... 62
6.2 Future Work ... 63
6.2.1 Model Estimation and Validation ... 64
6.2.2 Applications ... 64
Reference ... 66
Appendix 1 Aircraft Size and Load Factor Development ... 70
Appendix 2 Average Fare Development ... 71
Appendix 3 Market Fare Development ... 73

List of Figures

Figure 2.1: Summary of Hourly Ground Delay Cost 10
Figure 2.2: Hourly Delay Cost versus Delay Duration 11
Figure 4.1: Distribution of GDP Impacted Flights among Airports, 2006 26
Figure 4.2: Distribution of GDP Impacted Flights among Airlines, 2006 26
Figure 4.3: Impact of GID on the Cancellation Decisions 30
Figure 4.4: Impact of Aircraft Size on the Cancellation Decisions 31
Figure 4.5: Impact of Average Fare on the Cancellation Decisions 32
Figure 4.6: Impact of GID on the Cancellation Decisions -- Comparison of Model C-1 and Model C-2 34
Figure 4.7: Impact of GID on the Cancellation Decisions – Comparison of Model C-1 and Model C-3 37
Figure 4.8: Impact of GID on the Cancellation Decisions --Comparison over Airlines 41
Figure 4.9: Impact of the Number of Seats on the Cancellation Decisions --Comparison over Airlines 42
Figure 4.10: Impact of Average Fare on the Cancellation Probability — Comparison over Airlines 43
Figure 4.11: GID and Delay Savings Required to Make Cancellation Probability 0.5 44
Figure 5.1: Choice Generation Algorithm 50
Figure 5.2: Cumulative Distribution of Packets over the Numbers of Alternatives 53
Figure 5.3: Willingness to Cancel 60

List of Tables

Table 2.1: Example of Schedules under CDM 6
Table 4.1: Example of How Delay Savings Metrics is Calculated 19
Table 4.2: Hypothesized Signs of Coefficients of Explanatory Variables................................ 23
Table 4.3: Number of GDPs and GDP Impacted Flights, by Airport, 2006 25
Table 4.4: Number GDP Impacted Flights, by Airline, 2006 25
Table 4.5: Summary of Definition of Variables 28
Table 4.6: Basic Statistics of All Variables in Model C-1 29
Table 4.7: Estimation Results of Model C-1 29
Table 4.8: Estimation Results of Model C-2 33
Table 4.9: Estimation Results of Model C-3 36
Table 4.10: Estimation Results of Model C-4 38
Table 4.11: Estimation Results of Model C-5 40
Table 5.1: A Hypothetical Example of SCS 48
Table 5.2: Choices Permutations 49
Table 5.3: Example of Generated Alternatives 51
Table 5.4: Distribution of Packet Sizes 52
Table 5.5: Number of Alternatives and Corresponding Packet Frequency 53
Table 5.6: Estimation Results of the Standard Logit Model 56
Table 5.7: Estimation Results of the Mixed Logit Model 58

Acknowledgements

First and foremost I would like to thank my advisor Professor Mark Hansen for his guidance in the past four years. It is him who leads me to the field of aviation research and inspires me with great ideas. His profound knowledge and insights in air transportation guide me through the challenging exploration of this research. His sense of humor makes research discussion a breeze during the intensive study. He is an advisor who not only resolves students' questions but also understand their needs. Without him, I can never finish this work. I owe so much to him.

I would like to thank my dissertation committee, Professor Carlos Daganzo and Professor James Powell for their valuable comments. I am grateful for the mentorship of Professor Adib Kanafani. I always enjoy talking to him about research, teaching, and life. I also would like to thank Professor Kenneth Train. I learned so much from his class and his book "discrete choice methods". I benefited a lot from discussing my models with Professor Joan Walker. I want to also acknowledge Dr. Robert Hoffman from Metron Aviation, who has great knowledge in air traffic management and helped me to understand the data used in this dissertation.

I am grateful to know so many smart students and make friends with them. Yu Zhang and Qin Lu took good care of me in my early PhD years. Sofia Arroyo and I spent plenty of time together taking class and chatting. I obtained tons of useful information from Edward Huang. Amy Kim gave me great comments every time she reviewed my dissertation. I was amazed by how she can understand my mind and edit my sentences to make it clearer. I would like also acknowledge my NEXTOR officemates and Chinese fellow students in this program.

Finally, I want to extend my deepest gratitude to my family in China, especially my grandparents and parents. I was not able to spend much time with them because I am studying outside of my home country. But I know no matter where I am, their love are always with me. I am very lucky to get to know Haijian Shi and marry him when I am pursuing my PhD. With his unwavering support and love, I accomplish this task.

Chapter 1 Introduction

1.1 Problem Statement

The aviation system is facing unprecedented challenges in dealing with increasing demand and unpredictable weather. The Federal Aviation Administration (FAA) forecasts that the number of air passengers will increase at a 2.7% annual rate and reach 1.1 billion by 2025, which is 1.6 times the number of passengers in 2007 [1]. Although part of the demand can be accommodated by using more large-size aircraft, high demands still impose great stress on the current infrastructure. In addition, the air transportation system is vulnerable to severe weather, which reduces both airport capacity and en-route capacity. For example, the landing capacity of San Francisco International Airport (SFO) could be 60, 45, 36 or 30 aircraft per hour depending on various weather conditions. However, at major US airports, air traffic is often scheduled close to (or sometimes even above) the maximum capacity. Therefore, a capacity drop caused by bad weather will often result in a demand-capacity imbalance, which then creates disruptions and delay. Moreover, the very competitive commercial airline market and high cost of owning commercial aircraft encourages high utilization of aircraft. Consequently, scheduled turnaround times (from the scheduled arrival time to the schedule departure time of a given aircraft) are often quite short relative to the magnitudes of the delays incurred. In these cases an arrival delay at a given airport will often result in a departure delay. Such delay propagation can make a capacity drop at a single airport affect much of the network. Congestion and delays from capacity shortfalls are costly to air passengers, airlines and the overall economy. The total cost of air traffic delays in 2007 was estimated by the Joint Economic Committee (JEC) to be as much as $41 billion to the American society [2].

Increasing capacity through physical infrastructure expansion or advanced technology implementation are two solutions to reduce delays, but come with high investment costs and environmental impacts. Mitigating the consequences of aviation system capacity shortfalls by improving the efficiency of the existing system is another alternative. Under these circumstances, Air Traffic Management Initiatives (TMIs) are developed by the FAA to manage situations where demand exceeds capacity under severe weather operations. The objective of TMIs is to shift the demand to alternative resources, such as different routes, or different times in order to minimize congestion.

The air transportation system is a complex system with many agencies involved. Each airline, as a service provider to air passengers and a direct user of the air transportation system, has their own interest: profit maximization. Furthermore, the competition among commercial airlines is high in most markets. As a result, airlines generally prefer to minimize disclosure of their flight arrangements as well as other operational and financial information. On the other hand, the FAA requires enhanced flight information and collaborative responses from airlines so as to efficiently manage the aviation system.

Tremendous efforts have been made to design TMIs in a manner that encourages cooperation and information exchange between the FAA and airlines to make more efficient use of National Airspace System (NAS) resources and to reduce operating costs. However, there is little published literature regarding attempts to understand the airlines' behavior in response to the TMIs. The airlines' underlying utility, which is believed to be essential for TMIs (and the aviation system) to be evaluated appropriately and improved accordingly, is not well understood. This dissertation proposes to infer airlines' preference structure through their collaborative responses to the FAA's TMIs.

1.2 Dissertation Overview

The Ground Delay Program (GDP), a most critical element of TMIs, provides a good environment to investigate airlines' behavior. A GDP is implemented when an airport's capacity is reduced, and consequently falls short of the demand. Flights bound for the affected airport (called the GDP airport) are delayed on the ground at their origin airport to save fuel and ensure safety. The Collaborative Decision Making (CDM) program is developed to facilitate GDP, by encouraging collaboration and information exchange between the FAA and airlines. CDM-enhanced GDP allows airlines flexibility in managing their flights based on their own operational and business objectives. In other words, airlines are given the opportunities to re-arrange and cancel flights to best utilize the arrival slots that were assigned to them at the beginning of the GDP. As a consequence, airlines are observed to reduce delays on some flights and increase delays on other flights. They also cancel flights as an alternative to incurring high delays on these flights, or to vacate slots for other flights.

Airlines' utility can be estimated through these decisions by employing discrete choice models, which are applied extensively in many other transportation domains. An example is investigating the mode choice decisions (bus, car, bike, or walk) people make when commuting to work. The commuter's choice model can help to understand how people value travel time and travel cost. In the context of this dissertation work, the airlines' reaction to the GDP will be studied to understand how airlines value different flights and the relative costs of delay and cancellation. The models can help to answer questions such as the following: Are airlines more willing to cancel a flight whose Origin-Destination (OD) pair is more frequently served? Is it true that a flight operated with a small aircraft usually experiences more delay than a flight with a large aircraft?

In this dissertation work, two sets of models are proposed and estimated to reveal airlines' preference structures. The first set of models focus on modeling individual flight cancellation decisions. The dependent variable is binary, whether a flight is cancelled or not during the GDP. Three categories of explanatory variables are developed to explain this choice. The second set of models, consider airlines' proposals to reshuffle (and cancel) a set of flights under a process known as Slot Credit Substitution (SCS) under CDM. The details of SCS will be described later. In general, the submitted proposal is considered as superior --from the airline's point of view--to other feasible arrangements for the same set of flights, allowing the process to be modeled using multinomial choice models.

Beyond the scientific interest of understanding airlines' utility, there are many other questions motivating this dissertation work. For example, what is the cost of a flight cancellation? There is very little information in the existing literature that quantifies the cost of a flight cancellation. In fact, the phenomena is ignored and understated in most aviation studies that discuss the impact of aviation congestion, by focusing solely on delays. Additionally, how can flight heterogeneity be considered in evaluating air transportation system performance? Flight delay is a major performance metric in current practice. But it is an inadequate measurement due to the exclusion of flight cancellations and little consideration of flight characteristics. The delay of a 300-seat flight is likely to have a different cost to an airline than that of a 30-seat flight. In a similar way, the impact of cancelling a short haul flight should differ from cancelling a long haul flight. Another point that raises further questions is that delay is not homogeneous. Given 15 minutes of delay on each of six flights and 90 minutes of delay on a single flight, airlines are more likely to accept the first alternative because a single high flight delay is much more disruptive to the entire network. How can the non-linear cost of delays be quantified and incorporated into the delay study? Answering these challenging questions, which are of considerable practical significance, is one of main contributions of this research.

1.3 Organization

This dissertation consists of six chapters. Chapter 2 provides a background and review of the relevant literature. Chapter 3 describes the underlying economic theory and the econometric framework for this dissertation work. Chapter 4 discusses the individual flight cancellation model. Chapter 5 describes a model that captures both flight cancellation and flight re-arrangement decisions. Chapter 6 contains a conclusion and discussions of future work.

Chapter 2 Background and Literature Review

2.1 Background

This section begins with an introduction to Air Traffic Management Initiatives (TMIs). This is followed by a detailed discussion of Collaborative Decision Making (CDM) to provide background into the development of airlines' preference structures. At the end of this section, current airlines practices in dealing with schedule disruptions are briefly discussed.

2.1.1 Traffic Management Initiatives (TMIs)

TMIs are developed by the FAA to manage flow and mitigate the cost and operational impacts of demand-capacity imbalances. TMI solutions include the Ground Stop (GS), Ground Delay Program (GDP), Airspace Flow Program (AFP), and Miles-In-Trail (MIT). A GS is implemented when an airport is experiencing extreme weather conditions (e.g. heavy snow storm) or a very high security threat. Flights that are destined to this airport are stopped on the ground at their origin airport for the duration of the GS and waited there until further notice. In this extreme case, all operations are on hold until the airport situation clears. GDP and AFP are two other major types of TMIs that are more strategic than the GS. As one of the most critical TMIs, GDPs have been in use for more than two decades and has been refined to include improvements such as CDM. It is also the most commonly used TMI in daily operations of the air transportation system. A GDP is a response to an anticipated imbalance between arrival demand and arrival capacity at an airport. Flights bound for the airport are delayed at their origin to save fuel and ensure safety by decreasing en route holding. In addition, the GDP was also used to address en-route constraints, before the AFP was initiated in 2006, by restricting flow to airports major routes to which were affected by such constraints. The AFP is a relatively new program for managing aircraft flow through capacity-reduced airspace resulting from convective weather. AFP identifies en-route constraints and develops a list of flights that are requested to fly through the constrained area. These flights are then given options to re-route out of the constrained area, or accept ground delay at the origin airport until the time of their assigned slot through the constrained region. MIT is used to assist allocation of airport arrival capacity for incoming traffic, when the capacity falls short of demand. Air traffic controllers increase the separation distance between flights coming to the airport to reduce the flow rate. The restrictions of MIT can propagate to upper streams and affect flights that are hundreds of miles away from the airport. The purpose is to avoid congestion and preserve safety in the terminal airspace.

In this dissertation, the GDP is selected to study airlines' reactions to TMIs, not only because it is a representative TMI, but also because messages exchanged between the FAA and airlines during the event of GDP are archived in a database. The detailed flight schedule updates through CDM are described next.

2.1.2 CDM

CDM was implemented in 1998 to increase the efficiency and equity of GDPs. The idea is to increase information exchange amongst all agencies involved in GDP, including the FAA, airlines, airports, and air traffic control facilities. Building a system that provides common situational awareness allows operational problems to be solved in a timely and coordinated manner [3]. CDM ensures more efficient use of reduced capacity and adds to the flexibility of NAS users in meeting their operational needs. The advantages of CDM-enhanced GDP are as follows:

1. Through CDM, airlines are able to see the "big picture" of airport operations, oversee airport total demand and capacity changes, and plan their operational schedules accordingly.
2. CDM promotes accurate flight information sharing from airlines. CDM is designed in a manner that each airline benefits from sharing schedule updates. Before CDM, airlines were reluctant to report flight cancellations and delays until the "unused" arrival slots were wasted, because they lose the slots of cancelled flights to their competitors once the cancellations were reported. CDM allows airlines to keep the slots of their cancelled flights, preventing airlines from under-reporting flight cancellations and delays.
3. CDM uses the most updated demand information, leading to a significant improvement of GDP parameters. For instance, GDP may terminate earlier than planned because many cancellations are made by airlines, dropping the demand below the capacity.
4. CDM allocates slots to flights and allows airlines to retain these slots. Airlines are free to reshuffle their flights to best utilize their allotted slots and meet their own operational needs.
5. CDM provides a platform for developing more strategies for a new process, known as Slot Credit Substitution (SCS), which will be discussed later in this section. The CDM architecture has proven to be flexible in its ability to incorporate improvements.

CDM working mechanism

The FAA Command Center monitors airport demand and capacity through the Flight Schedule Monitor (FSM) several hours in advance. When an imbalance between the number of scheduled flights and the predicted airport arrival rate is detected, a GDP is issued. FSM also provides a baseline solution to address the problem at the beginning of GDP by assigning an Original Controlled Time of Arrival (OCTA) to each flight in the order of its Official Airline Guide (OAG) published schedule arrival time. This is known as the Ration by Schedule (RBS) resource allocation process. After slots are first assigned, the airlines “own” these slots and can reassign them among their flights. Airlines’ responses and updated schedules reflecting their changes are submitted to the FAA via CDM messages. Moreover, the FAA runs inter-airline compression periodically to fill open slots that result from cancellations. An open slot is created when an airline has finished moving up all their delayed flights and reached a point where none of their flights can be fit into the open slot. Inter-airline compression ensures that open slots can be used by other airlines and thus not be wasted. The Controlled Time of Arrival (CTA) of each flight gets updated dynamically by the FAA and airlines. Every five minutes, the FAA Command Center consolidates all of these schedule updates and shares them with the airlines through the Aggregate Demand List (ADL).

The following hypothetical example (Table 2.1) demonstrates how CDM works in the event of a GDP. Suppose that, due to a morning fog problem, the FAA decides to implement a GDP from 8am to 9am at San Francisco International Airport (SFO). There are four United Airlines’ flights involved in this GDP.

Table 2.1 Example of Schedules under CDM

Flight	Origin	OAG Arrival Time	GDP Baseline Arrival Time (OCTA)	Airline Finalized Arrival Time (CTA)	Airline Delay Adjustment
UA1	SEA*	8:00	8:10	8:10	0
UA2	LAX*	8:10	8:30	9:10 (Cancel)	Cancel
UA3	ORD*	8:20	8:50	8:30	-20
UA4	DEN*	8:30	9:10	8:50	-20

*SEA is acronym name of Seattle Tacoma International Airport
*LAX is acronym name of Los Angeles International Airport
*ORD is acronym name of Chicago O'Hare International Airport
*DEN is acronym name of Denver International Airport

The FAA employs RBS to create a baseline solution to smooth the arrival traffic by allocating controlled arriving slots to each of flight; refer to the forth column of Table 2.1. United Airlines can then reassign its slots to its flights based upon its own operational needs and business objectives. The result of this reassignment is shown in column 5. Flight 1 is left at its original GDP slot. Flight 2 takes the slot originally belonging to Flight 4, and eventually gets cancelled. Flight 3 is moved to the slot that was assigned to Flight 2 and Flight 4 takes the slot vacated by Flight 3. By comparing the baseline schedule (column 4) with the final schedule created by United Airlines (column 5), the delay adjustment is shown in column 6. Flight 2 is cancelled to free up its slot for the more critical Flights 3 and 4. Given that United Airlines made this arrangement, it is assumed that it was willing to incur a cancellation to Flight 2 in order to decrease delay for flights 3 and 4.

Slot Credit Substitution (SCS)

Intra-airline substitution works well when an airline owns many flights that are destined to the GDP airport and scheduled close to each other. However, if a small airline does not offer much service, the two consecutive landing slots it possesses may be far away from each other. Thus even if the airline wanted to cancel its former flight to free the slot for its latter flight, the latter flight may not be able to fill it. Under these circumstances, inter-airline slot trading becomes necessary.

A process called Slot Credit Substitution (SCS) was incorporated into CDM in 2003 to achieve inter-airline substitution. Under a GDP, an airline can initiate communication with the FAA by sending a SCS message such as "if I (United Airlines) can get flight UA843 land at 6:30pm instead of 7:00pm, I would be willing to cancel my flight UA731." The request is then reviewed by the FAA command center. If a flight (from another airline) occupying the 6:30 slot can move to the UA731 slot, then the SCS request is accommodated by the FAA, and the schedule change is executed immediately. The straightforward trading mechanism of SCS gives airlines more flexibility to optimize their schedule and achieve their needs.

2.1.3 Airlines Practices in Schedule Disruption Management

Under GDP, the airlines' regular operations are often disrupted, which results in a state of "irregular operations" characterized by high delays, many cancellations, and a large number of disrupted passengers. Such conditions are recognized to be highly detrimental to airline profit and passenger welfare. As the objective of this dissertation is to infer airlines' valuations of cancellations and delays from the choices they make in the context of disruption management, it is necessary to first understand the setting in which these choices are made.

Airline Operation Control Center (AOCC) is in charge of the schedule disruption management under GDP. Depending on the size of GDP event, one or multiple airline dispatchers are assigned to manage the situation. Clarke [4] summarizes the state-of-practice in AOCC during the aftermath of irregular airline operations. In general, some airlines have developed procedures that are implemented manually, while other airlines use computer-aided optimization tools to facilitate the decision-making process.

However, regardless of which procedure airlines use to manage their fleet, there are some fundamental challenges they all face. Firstly, recovery from disruptions is a very complex and time-critical process. Enormous amounts of information are exchanged between the FAA and the AOCC in a very limited time. The continuous updates from changes of capacity, weather, demand and the resulting GDP revisions require dispatchers to make decisions constantly in a timely manner. In addition, AOCC must reallocate resources and execute updated schedules within an airline rapidly and efficiently. Therefore, the fast-paced nature of the work makes disruption management highly challenging.

Secondly, some important restrictions, such as crew timeouts and weather uncertainty, constrain the management procedure. In the airline industry, a very strict rule on maximum crew working hours is applied in order to ensure safety. For example, a pilot has to comply with FAA regulations regarding on-duty time, flight time and resting time. However, delaying a flight can extend the on-duty time and shorten resting time. Thus a pilot may timeout and not be available to operate the aircraft as scheduled due to the disruption. Weather also introduces more uncertainties to the problem, which makes disruption management more difficult.

Finally, although there is a trend towards the use of computer-aided optimization tools for disruption management, they are not yet fully adopted by the airline industry. The current practice of disruption management lags well behind the theoretical models. One reason is that the value of optimization tools is not fully appreciated. Unlike planning, recovery effectiveness is difficult to assess. The other reason is that practitioners often question the cost structure (e.g. delay cost and flight cancellation cost) that are used in these optimization tools, as well as the results obtained from such tools.

Regardless, optimization models attract considerable attention from the operational research community. Theoretically, to respond optimally to a disruption, an airline could minimize the total cost of flight delay and cancellation. However, the downstream effects of delayed and cancelled flights are hard to quantify and include in the objective function. In practice, the recovery problem is usually handled in four sequential steps: fleet assignment, aircraft assignment, crew assignment, and passenger assignment. Each step is an optimization process with certain objectives and constraints, and has been studied in the literature [5]. Even if the solution of each step is optimal, this does not mean that the sequential procedure will lead to globally optimum results.

A more strategic approach to handling disruptions, known as robust airline scheduling, has received increased attention in recent years. The basic idea is that when an airline plans its schedule, it should include the anticipated cost of recovery from irregular operations. For example, a high load factor and short turnaround time produces high revenue and thus is favored by the planning process. However, this also makes the system more vulnerable to disruptions. A robust airline scheduling process considers both regular and irregular operations in the objective function.

2.2 Literature Review

To the author’s best knowledge, a behavioral study about airlines’ reaction to the TMIs has not been performed. However, since flight delays and cancellations are two major occurrences resulting from the airlines’ responses, this section reviews the relevant literature about delays (in 2.2.1) and cancellations (in 2.2.2).

2.2.1 Delays

The majority of delay studies are focused on flight delays. This is because flight delay is the main performance metric in evaluating NAS operations. Also, flight delay is relative easy to define and measure. It is defined by comparing the actual arrival (or departure) time with the scheduled arrival (or departure) time. Depending on the different phase of the flight, delay can be decomposed into departure delay, airborne delay, and arrival delay; however, if delay is categorized based on where it occurs, it can be classified as ground delay or airborne delay.

The scopes of delay impact studies vary. Some studies focus on the cost to airlines only and some studies consider cost to air passengers. Other studies include airline costs, passenger costs, as well as the spill-over cost to society.

Delay cost to airlines

For the studies focusing on delay cost to airlines, using a unit cost of delay ($/hour) is a common approach. Then the total cost of delay is calculated based on the unit cost multiplied by the delay minutes. Under this approach, some studies identify different unit cost of ground delay and airborne delay (taxi delay included). Other studies treat ground delay the same as airborne delay.

The unit cost of airborne delay is usually calculated using the direct operating cost of aircraft blocking time, which includes costs of fuel, crew, maintenance, capital etc. The Air Transport Association estimates that a minute of airborne delay costs an airline $74.10 in 2008, a 23% increase from 2007 [6]. Another study done by airlines, calculates the direct operating cost to be $55 per minute in 2004 [7].

In general, the unit cost of ground delay is simply the unit cost of airborne delay minus the unit cost of fuel. With the fluctuations in fuel cost, the ratio of the unit cost of ground delay to airborne delay varies. Ratios 1:2 and 1:3 are used to simplify the problem [8] [9]. Figure 2.1 summarizes the hourly rate of ground delay cost from the literature [10]. Effects of long delays and short delays, aircraft size and airline are investigated in determining the hourly cost in these references.

Figure 2.1 Summary of Hourly Ground Delay Cost

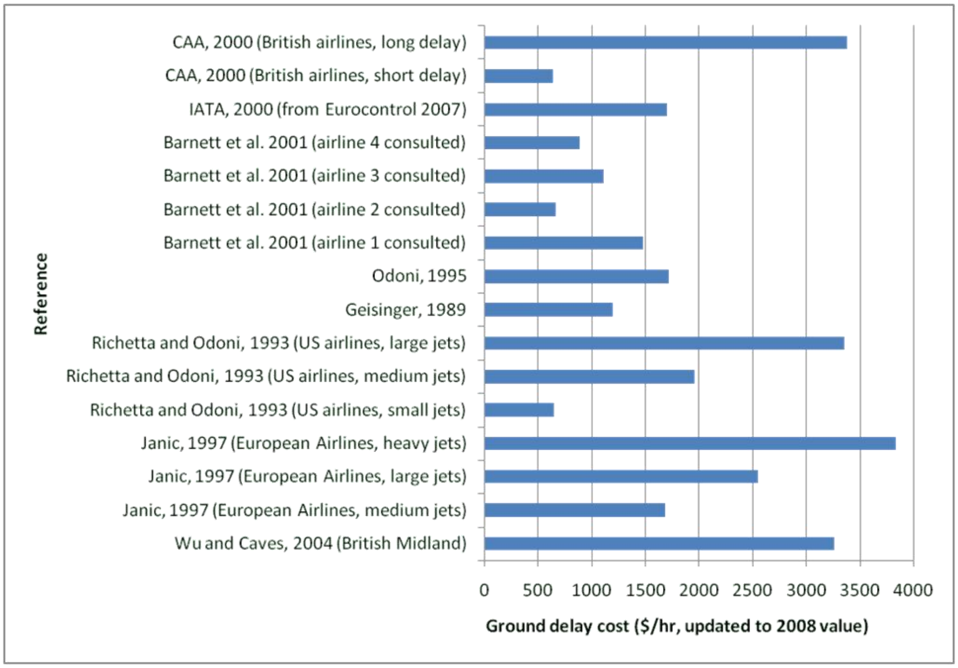

Source: *Ball et al*. "Total Delay Impact Study: A Comprehensive Assessment of the Costs and Impacts of Flight Delay in the United States."

In addition, Cook *et al.* [11] interviewed airline personnel to estimate a detailed unit cost of delay, recognizing the difference in the cost of small delay (15 minutes) versus large delay (65 minutes), aircraft type, and disaggregate locations where the delay is incurred.

The unit cost approach is intuitive and easy to use. However, m2p Consulting [12] claims that the unit cost of delay increases with delay duration, refer to Figure 2.2, because the impact of longer delays is magnified through the propagation effects. The author also states that delay costs are dependent on the airline's network, operating cost, cost of passengers' time, and the loss of a passenger's willingness to fly with the airline again. Delay propagation effects are also studied in Roger *et al*.'s paper [13]. The author proposes the concept of a Delay Multiplier (DM) and computes it by using American Airlines' actual schedule of crew and aircraft. The study concludes that the DM is the value of delay on the operating schedule as a whole. The cost of delay calculated based on flight delay only is not adequate because the downstream effect is ignored.

Figure 2.2 Hourly Delay Cost versus Delay Duration

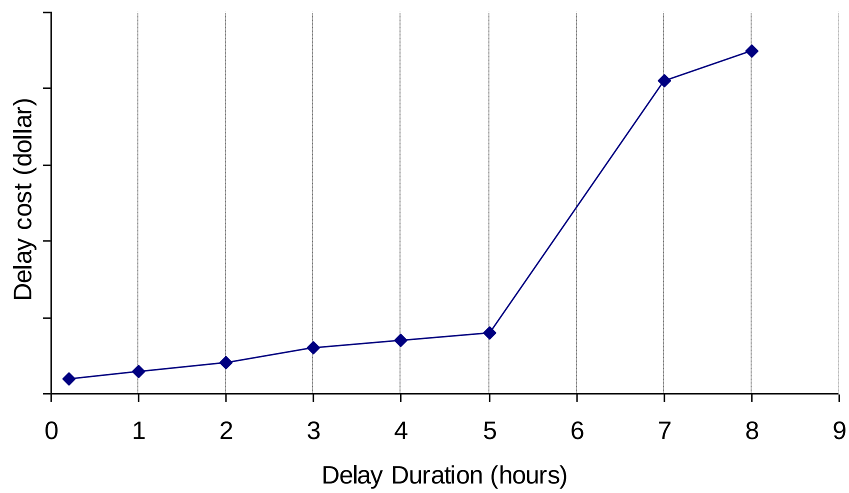

Source: m2p Consulting, "Optimized Flight Operations Delay Management"

In summary, the unit cost of delay approach simplifies the overall cost calculations and can be improved if the following issues are appropriately addressed:

1. The non-linearity of delay cost. The unit cost of delay should be a non-decreasing function of the delay, if propagation effects are to be considered.
2. Flight heterogeneity. The unit delay cost of large aircraft should be more than that of small aircraft.
3. Type of delays. The cost of ground delay and airborne delay are different.

Delay cost to air passengers

A conventional way to estimate passenger delay cost is to multiply the delay by an assumed average value of passengers' time. The Air Transport Association [14] assumes passengers' time value to be $37.18 per hour. GRA [15] investigates the value of passenger's time in air and proposes an adjustment methodology to update the value. They found that the value of time for a passenger depends on their travel purpose: $23.3 per hour for personal travelers and $40.10 per hour for business travelers.

There are also some other studies quantifying the value of passengers' time through statistical models. Morrison and Winston [16] estimates air travel demand response to delays using a logit model. He found that if the number of flights delayed more than 15 minutes is increased by 1%, the passengers' willingness to pay is reduced by $0.61 (in 1983 dollars). Forbes [17] develops a model to quantify airfare response to delays. The author found that the tradeoff between delay and fare is $1.42 per minute ($85 per hour). The author also states that a higher value resulted from the fact that there was a larger fraction of business travelers in his sample (focus on LaGuardia Airport), compared to a national sample of travelers. Alder *et al.* [18] estimates a multinomial logit model using stated preference survey data to analyze chosen and non-chosen itineraries. The tradeoff between delay and fare was found to be $22 per hour for business travelers and $6 per hour for non-business travelers.

Another challenging issue in using the value of time approach in calculating passenger delay cost is involved in quantifying passenger delay. Unlike flight delay, passenger delay is not reported anywhere in the system. For non-stop itinerary passengers, passenger delay is the same as flight delay. However for multiple-stop itinerary passengers, they are not equal. If passengers miss connections due to flight delays on previous segments, the passengers experience much longer delays than the flight delays alone, as they spend additional time waiting in the terminal for the next available flights. Disrupted passengers, defined by Cynthia and Bratu [19], are those who ultimately fly a flight itinerary other than the one originally booked. The authors found that disrupted passengers are only 3% of the total passengers, but suffered 39% of the total passenger delay. Bratu and Cynthia. [20], Wang *et al.* [21], Sherry and Donohue [22], and Zhu [23] investigate methods for computing passengers' delay.

Delay cost to overall economy

The recent work by the Joint Economic Committee estimates the cost of delays considering airline costs, passenger costs, as well as spill-over cost to the whole economy. The report concludes that the total cost of air traffic delays in 2007 is almost $41 billion to American society [2]. Of the $41 billion, $19 billion is caused by additional operating cost to airlines, and $12 billion by air passengers' loss of time and productivity. The last $10 billion results from indirect costs to hospitality industries related to air transportation.

2.2.2 Flight Cancellation

Flight delay alone cannot represent the impact of irregular operations on airlines or the degraded level of service experienced by passengers. Hansen *et al.* [24] found that schedule disruption and variability are more significant cost drivers for airlines than delay alone. Irrgang *et al*. [25] stated that irregularity or off-schedule operations cause the largest operational losses to the airlines. Diversions, cancellations, misconnections and unrealized demand are surprisingly costly. Cynthia and Bratu [19] suggest more relevant metrics from a passenger's point of view, including flight cancellation rates and the percentage of flights delayed by more than 45 minutes. Flight cancellation perturbs regular operations and carries a significant cost to air passengers and airlines. But it is often overlooked due to very limited knowledge of it.

Although airlines may know the value of a flight cancellation, this value is not revealed to the public. Some researchers, in both academia and industry, assign a rough estimate to cancellation cost based on their knowledge and experience. Sridar [26], senior scientist for Air Transportation Systems at National Aeronautics and Space Administration (NASA) Ames Research Center, states that one cancellation is equivalent to 200 to 300 minutes of delay. Metron Aviation [27] arbitrarily assigns $6000 as the cost of a flight cancellation. Later, feedback from the airline industry indicated that they considered this figure to be low. American Airlines [28] reveals that short-haul flights delayed more than 2 hours are likely to be cancelled and that long-haul flights are delayed longer before being cancelled outright. However, no quantitative studies have attempted to confirm these values.

Two studies have investigated the cause of flight cancellation in the literature. Rupp [29] studies the determinants of flight cancellations together with flight delays using a nested logit model. The author obtains individual flight delay and cancellation information between 2001 and 2003 through the U.S. Bureau of Transportation Statistics (BTS) On-Time Performance Database, and develops several categories of explanatory variables including economic variables (average revenue, load factors, etc), route competition variables (monopoly market, largely duopoly market, etc.), airport competition variables (airlines' hub, etc.), logistical variables (frequency, distance, etc.) and weather variables (rain, temperature, precipitation, etc.). The results show that route competition has little effect on cancellations or delays and airports with dominant carriers experience higher delays and cancellation rates. The author also discusses in his paper that "delays and cancellations move in opposite directions suggesting that the carrier is trading-off fewer (more) flight cancellations for more (fewer) flight delays". But his model does not capture the tradeoff between flight cancellation and flight delay.

Rupp and Holmes [30] extended Rupp's previous study by focusing solely on flight cancellations. The authors found that cancellations rates are lower on more competitive routes and higher revenue routes. They also found that airlines are less likely to cancel flights to and from their hub airport, that have high load factors and whose ODs are less frequently served. However, the study fails to consider delay as a factor in the model, which is one of the most important reasons for cancellation. Furthermore, cancelling a flight creates an opportunity to reduce delays on other flights, which will further encourage cancellation decisions. This point is also not considered in the Rupp and Holmes' model.

Chapter 3 Econometric Framework

Airlines' preference structure is investigated through the estimation of discrete choice models based on CDM archival data. By observing the actual choices made by airline dispatchers when presented with alternatives, the airlines' utility functions can be inferred. This is called the revealed preference approach to model estimation. It should be noted that during the real-time GDP, a countless number of choice situations are faced by the airlines. The first task of this research is to decompose to several situations that a reasonable number of choices can be created and used to estimate the utility. Under the circumstances, two opportunities are identified. One is to estimate airlines' cancellation utility based on their flight cancellation decisions in response to GDP. The choice is binary—a flight is cancelled, or not. The other situation is to estimate airlines' utility through SCS messages, which includes a small portion of the flights and makes the number of choices manageable.

Random Utility Theory (RUT) is employed. The utility that a decision maker i obtains from alternative a among A total alternatives, U_{ia}, can be decomposed into to two parts. V_{ia} is assumed to be known by the researcher, and is a function of variables that are observable. ε_{ia} is the unknown part, which captures the factors that affect utility but are not observable to the researcher. For example, aircraft assignment and crew assignment, which often affect airlines decision but are not observable to the researcher, are included in the unknown part. The ε_{ia} is also called error component because it is the difference between U_{ia} (real utility) and V_{ia} (the utility as estimated by researcher).

$$U_{ia} = V_{ia} + \varepsilon_{ia} \tag{3.1}$$

U_{ia}: Utility individual *i* obtains from alternative *a*

V_{ia} : Known part

ε_{ia}: Unknown part

Depending on the assumptions regarding the distribution of this unknown part, a probit model, a logit model, or other mixed models can be formulated. In this dissertation work, the logit model and mixed logit model are mainly used. The theory behind these two models can be found in Train [31].

The logit model is the most simple and widely used form of discrete choice models. By assuming that each ε_{ia} is independently and identically distributed with a Gumbel distribution, the probability of individual i choosing alternative a, P_{ia}, has a closed form:

$$P_{ia} = \frac{e^{V_{ia}}}{\sum_{j=1}^{A} e^{V_{ij}}} \tag{3.2}$$

Given Equation (3.2), the probability ratio of individual *i* choosing alternative *a* to choosing alternative *b*, is

$$\frac{P_{ia}}{P_{ib}} = \frac{e^{V_{ia}}}{e^{V_{ib}}} \tag{3.3}$$

This implies that if there are more than two alternatives, the ratio of choice probabilities for any pair of alternatives is not affected by the utility of any other alternative. This is known as the independence from irrelevant alternatives (IIA) property, and is inappropriate under some circumstances.

The limitations of the standard logit model are that it does not consider correlations among unobserved factors, and that it assumes that a single deterministic utility function applies to all choice makers—no taste variation through the population is allowed. Nevertheless, the standard logit model remains a popular tool in estimating choice preferences because it is readily interpretable and much less computationally intensive. For estimation, from a practical standpoint, the standard logit specification is a good starting point to understand utility before advancing to other more complex models.

The mixed logit model has attracted much attention in the last few decades, especially with the development of fast computers. It requires a very high level of computing resources because the choice probability is no longer closed form and is solved by simulation. By introducing a density function $f(\beta)$, the probability of individual *i* choosing alternative *a*, P_{ia}, becomes

$$P_{ia} = \int \frac{e^{V_{ia}(\beta)}}{\sum_{j=1}^{A} e^{V_{ij}(\beta)}} f(\beta) d\beta \tag{3.4}$$

The mixed logit probability is a weighted average of the logit formula evaluated at different values of β, with the weights given by function $f(\beta)$ [31]. The density function $f(\beta)$ can be specified as normal, lognormal, or some other form. The mixed logit model overcomes all the limitations of the standard logit model: it is able to consider a flexible substitution pattern among alternatives, taste variation among population, and correlations among unobservable factors.

The most straightforward interpretation of the mixed logit model is random coefficients, which introduce taste variation in the utility among populations; the variation can be captured by this density function. Revelt and Train [32] present an example where he uses the mixed logit model to explore taste variation among households in selecting efficient appliances.

An alternative interpretation of the mixed logit model is that it represents part of the error components which creates correlations among alternatives. The correlations are induced by specifying some variables that enter the error components. Utility is specified as

$$U_{ia} = \beta_{ia}{}'x_{ia} + \mu_{ia}^{'} z_{ia} + \varepsilon_{ia} \tag{3.5}$$

Where x_{ia} and z_{ia} are observable variables that related to individual i and alternative a. β is a vector of fixed coefficients and μ is a vector of random terms with zero mean. ε_{ia} is iid error component. The new error component can be seen as $\mu_{ia}^{'} z_{ia} + \varepsilon_{ia}$. By selecting appropriate variable z_{ia}, the new error components can represent the correlation patterns among alternatives. Brownstone and Train [33] demonstrate how the mixed logit model can be used to capture correlations among alternatives and account for flexible substitution patterns in forecasting vehicle choice.

Both interpretations of the mixed logit model are used in this dissertation. In Chapter 4, a mixed logit specification introduces taste variation among flights' cancellation utility. In Chapter 5, a mixed logit model specification captures the correlations among alternatives.

Chapter 4 Cancellation Model

4.1 Introduction

As discussed in Chapter 2, Under a GDP, airlines are observed to cancel flights in order to decrease the total demand, as well as free slots for other flights. In this chapter, the flight's cancellation decision is modeled using CDM archival data.

Section 4.2 describes the explanatory variables that developed for flight cancellation decision modeling. Section 4.3 discusses the data preparation and statistics. Section 4.4 presents five models, including model specifications, estimation results and discussions. Section 4.5 analyzes the tradeoff between flight cancellations and delays. The chapter is concluded in Section 4.6.

4.2 Major Explanatory Variables Development

In order to model airlines' cancellation decisions, three categories of explanatory variables are developed. The categories together with individual variables in them, and methods of calculation are described in this section.

4.2.1 Delay Factors

These variables pertain to the delays of cancelled flights and non-cancelled flights. Variables in this category include:

<u>*GDP-assigned initial delay (GID)*</u>

It is hypothesized that airlines are more likely to cancel flights that are assigned high delays in the GDP. This is because the cost of operating the flight increases with delay. In addition to the direct costs in terms of passenger, crew and aircraft idling time, long delays are more likely to affect the subsequent flights. If the delay is long enough, the airline is often forced to act to mitigate these consequences by reassigning crews and aircraft to downstream segments or rebooking connecting passengers. These actions, on the other hand, reduce the impact of cancelling the flight as compared to operating it and increase the likelihood of cancellation. Airlines are more likely to cancel a flight when the delay cost exceeds the cancellation cost. Therefore, the larger the GDP-assigned Initial Delay (*GID)* is, the more likely the corresponding flight gets cancelled instead of delayed.

GID is quantified as the difference between the baseline arrival time assigned by the FAA when GDP is implemented (known as Original Control Time of Arrival--*OCTA*) and the scheduled flight arrival time without GDP. The latter is also known as the Base Estimated Time of Arrival (*BETA*). Flight *BETA* values, effective when the GDP was issued, are saved in the database to capture the amount of arrival delay that can be attributed to a GDP.

$$GID = \min(0, OCTA - BETA) \tag{4.1}$$

Delay can be also defined as the difference between *OCTA* and Initial Gate Time of Arrival (*IGTA*) in the *OAG* schedule, published 3 months before the operation. In this dissertation, the definition based on *BETA* instead of *IGTA* is adopted because *BETA* is more up to date.

Among 167,584 sample flights in the data set, 4,130 flights (2.46%) have an *OCTA* earlier than its *BETA*. This happens sometime because *OCTA*s are assigned based on the *IGTAs*, not the *BETAs*. The *GIDs* in these cases are truncated to be zero.

Delay savings from flight cancellation

As discussed before, CDM enables individual airlines to re-use slots of cancelled flights by moving up other flights. Thus, cancelling a flight not only eliminates any delay that the cancelled flight would otherwise have incurred, but also decreases delays of other flights that can be moved up as a result. It is hypothesized that airlines consider the potential delay savings to other flights in the flight cancellation decision. Therefore metrics for such savings are developed from the archival data.

In developing these metrics, a complication, which is that the potential delay savings from a given cancellation depend on what other flights are also cancelled, should be considered. Obviously, if a given flight is cancelled, it cannot be moved up into a slot of another cancelled flight. For purposes of computing these metrics for any particular flight, it is assumed that no other flights controlled by the airline are cancelled.

An important constraint must be recognized when computing the delay savings from a cancellation. Only certain flights can be moved up into a vacant slot. The "qualified" flights are those whose Earliest Runway Time of Arrival (*ERTA*) is not too much later than the vacant slot. In the parlance of CDM, an "*x*-minute window" means that a flight must have a *ERTA* no more than *x* minutes later than a slot in order to be moved up into it. Based on discussions with CDM participants, a 20-minute window is used for this analysis. Thus, for example, a flight with an *ERTA* of 10:15 can be moved into a 10:00 slot, but one with an *ERTA* of 10:30 cannot be. The *ERTA* for each flight is reported by the airlines.

Conversely, when there is more than one "qualified" flight for a vacant slot, the possible movement is not unique any more. Therefore, the total delay savings could vary, depending on how flights are reshuffled to take advantage of the vacated slot. For purposes of this model, a vacated slot is assigned to the qualified one with the earliest *OCTA*. This assumption assures that maximum total delay savings from vacating slot through a flight cancellation [34]. On the other hand, this reassignment algorithm produces many movements, so that the saved minutes per movement could be small. Alternative reassignment strategies include only moving flights whose delay can be reduced by some minimum amount--say 15 minutes-- or even employing an optimization tool to select which flights should be moved into which slots. The imposition of the minimum could in some cases greatly reduce total delay savings, however, while the optimization approach would require situation-specific information about airline objective functions that is unavailable.

A simple example is generated to illustrate how to calculate the delay savings from a hypothetical cancellation. In Table 4.1, there are four flights from the United Airlines involved in a GDP and each flight is assigned *OCTA* by the FAA. Airline reports *ERTA* of each flight as well. Under this setting, if Flight 1 is cancelled hypothetically, what are the delay savings? The "New *CTA*" is the Controlled Time of Arrival (*CTA*) that results from reassigning flights to slots when Flight 1 is cancelled. When Flight 1 is cancelled, the arrival slot at 8:00 becomes vacant. Then Flight 2 is checked if it can take the vacant slot. Flight 2 can take the slot because its *ERTA* (7:40) is earlier than 8:00. Thus the New *CTA* for Flight 2 is 8:00 and if compare it with its *OCTA* (8:20), Flight 2 realizes a 20-minutes delay saving. Now, the slot originally to Flight 2 (8:20) is vacant. Unfortunately, *ERTA* of Flight 3 (8:50) is outside the 20- minute time window of this 8:20 slot, Flight 3 can't take the vacant slot; therefore the New *CTA* of Flight 3 is still 8:40. Next, the Flight 4 is checked if it can take the vacant slot at 8:20. It can since its *ERTA* is not later than 8:20. Therefore the New *CTA* of Flight 4 is 8:20, reducing its delay by 40 minutes. Thus, the total delay saving to Flights 2, 3 and 4 from cancelling flight 1 is 60 minutes.

Table 4.1 Example of How Delay Saving Metrics is Calculated

Flight	OCTA	ERTA	New CTA	Delay Savings (min)
1	8:00	7:00	Cancelled hypothetically	
2	8:20	7:40	8:00 (take slot pre-owned by flight 1)	20
3	8:40	8:50	8:40 (can't take the slot pre-owned by flight 2 due to 20-minute time window constraint)	0
4	9:00	8:20	8:20 (take the slot pre-owned by flight 2)	40

It is hypothesized that the greater *delay savings* that can be realized from cancelling a flight; the more likely it is to be cancelled, all else equal.

4.2.2 Flight Characteristics

Beyond the delay factors related to delay and delay savings, certain inherent characteristics of the flight itself may also influence the cancellation decision.

Flight distance

Flight distance is expected to be one such factor. It is hypothesized that, all else equal, airlines are less likely to cancel longer flights. Imagine two Flights *A* and *B* would like to land at the same time (say, 8pm) at a GDP airport. Flight *A* is a 4-hour flight (long haul) and Flight *B* is a 1-hour flight (short haul). The airline needs to make decision to cancel Flight *A* at 4pm at the latest. However, the airline can make a decision to cancel Flight *B* at as late as 7pm. Between 4pm and 7pm, situation may get better and GDP may be terminated earlier than planned. So it is reasonable for the airline to hedge by not cancelling the long haul flight, and hope that conditions at the destination airport improve so that it will not need to cancel the short-haul flight either. If there is no improvement, the short haul flight can be cancelled later. Generalizing from this example, airlines can minimize the number of unwarranted cancellations by preferentially cancelling short-haul flights.

Hub destination dummy

Most airlines' networks are hub-and-spoke, and the cancellation decision is hypothesized to be different for hub-bound and spoke-bound flights. Airlines are expected to be more willing to cancel hub-bound flights for several reasons. Airlines have more flights into their hubs, so they can generally make more use of a vacated slot of a hub-bound flight. The cost of delay into a hub may be greater because gates and other resources there are utilized more intensively. There may also be less hardship to connecting passengers if they are "stranded" at their origin because of a cancellation rather than at a connecting hub. A final reason why airlines may be more willing to cancel flights into hubs is the greater ease of obtaining replacement crews and aircraft for subsequent legs of the cancelled flight. The *hub destination dummy* variable is set to 1 whenever the flight is bound for a hub of the airline operating the flight, and 0 in all other cases.

Major airline dummy

Another factor that may affect the cancellation decision is whether the flight belongs to a major airline or one of its regional affiliates, on behalf of whom the major carriers often handle CDM. This could result in the well-known principle-agent problem in political science and economics [35], with the major favoring its own flights over those of the commuter when the two interests conflict. Thus, a *major airline dummy* variable is developed (if a flight belongs to major airlines, this variable is set to 1; otherwise, it is set to 0) to test if the principle agent phenomenon exists. Note that, while affiliate flights may also be more subject to cancellation because they employ smaller aircraft and fly shorter segments, these effects are controlled for separately.

Internal delay

The *internal delay* is defined by *BETA* minus *IGTA* and truncated to be zero when it is negative.

$$Internal\ delay = \min(\,0, BETA - IGTA) \tag{4.2}$$

BETA records the most up-to-date estimate of arrival time when GDP is implemented, while *IGTA* is the arrival time in the original *OAG* schedule, which is usually published three months in advance. Therefore the *internal delay* quantifies how much delay a flight may experience prior to imposition of the GDP. A high *internal delay* may result from some airline internal factor, such as an aircraft mechanical problem, absentee crews, and so on. Such occurrences will make a flight more likely to be cancelled even without a GDP; therefore the coefficient of *internal delay* is expected to be positive.

4.2.3 Segment Characteristics

This section describes the variables that are developed from segment information. Such characteristics pertain to the entire set of flights flying between a particular airport pair by a particular airline over some period of time, rather than the specific flight for which the cancellation is being made. In some instances segment characteristics rather than flight characteristics are included because they are inherently more relevant; in other instances segment variables are used because the corresponding flight level variables are not available.

Frequency

The airlines are expected to be more likely to cancel flights on segments with more frequent service. If there are more flights operating on a segment, it is easier to rebook passengers from a cancelled flight operating on that segment. The frequency associated with a specific flight is determined by the number of flights on the same segment, operated by the same airline, and within the same GDP.

Aircraft size

It is hypothesized that airlines preferentially cancel flights with smaller aircraft because cancelling a large flight makes rebooking passengers much harder. The aircraft type of each GDP flight is reported in the CDM archival database, but not the exact number of available seats. The number of available seats for an aircraft type varies according to its specific configuration. While this information is not available for individual flights, it is approximated by the number of seats per flight by airline, aircraft type, and flight segment. These data are available in the T100 database. Details are described in Appendix 1.

Average fare

It is hypothesized that flights with lower *average fare* are more likely to be cancelled. High fare (business) passengers' loyalty is very important to airlines. Airlines may want to avoid inconveniencing business passengers by cancelling their flights.

Ideally, the exact amount of *average fare* for each GDP flight should be used in the cancellation model. Unfortunately, airlines do not make such data publicly available. Therefore, the best a researcher can get is a rough estimate of average fare according to the Airline Origin and Destination Survey (DB1B) database. Quarterly average fare is developed by airline and flight segment. Appendix 2 documents the details. The *average fare* is a weighted average using the fare paid by non-stop passengers by the number of non-stop passengers, and the fare paid by connecting passengers by the number of connecting passengers. For the connecting passengers, the fare is allocated by the distance of the segment.

Load factor

The airlines are expected to be less likely to cancel flights on segments with high load factor. The higher the *load factor* is, the harder to rebook the passengers if a flight in the segment needs to be cancelled. The load factor of individual flight is not reported in the CDM archival database. Monthly average of *load factor* for a particular airline, aircraft type and segment is developed from T100 database. Details are also described in Appendix 1.

Market fare

The *market fare* is developed to represent the origin-destination market type of the flight involved in the GDP. The hypothesis is that airlines value premium-fare-market flights more than flights in other markets. Passengers in premium fare markets are expected to have a higher value of time, and therefore suffer greater inconvenience from a cancellation or lengthy delay. Moreover, airlines are likely to place a higher value on retaining the good will of such passengers. Thus, for instance, if the GDP airport was Chicago, a flight coming from Las Vegas would be less important compared to a flight coming from New York. This increases the probability that the Las Vegas flight would be cancelled and its slot used to reduce the delay to the New York flight. While in this example the market type is inferred from the flight origin, in most cases the categorization is not so clear-cut or obvious. Therefore the *market fare* is employed as a proxy. The *market fare* is also developed from DB1B database. It is a quarterly average of fare for all reported flights from all airlines and routes serving the market. Appendix 3 covers the development of *market fare*.

4.2.4 Summary

The hypotheses of all variables are summarized in Table 4.2.

Table 4.2 Hypothesized Sign of Coefficients of Explanatory Variables

Categories	Variable Name	Hypothesized Effect
Delay factors	GDP-assigned initial delay (GID)	Positive
	Delay savings from a hypothetical cancellation	Positive
Flight characteristics	Distance	Negative
	Hub destination dummy	Positive
	Major airlines dummy	Negative
	Internal delay	Positive
	Frequency	Positive
Segment characteristics	Aircraft size	Negative
	Average fare	Negative
	Load factor	Negative
	Market fare	Negative

4.3 Data Description

4.3.1 Data Preparation

In the CDM archival database, flights that are under control of a given GDP are assigned *OCTAs.* For each GDP impacted flight, schedules are published every five minutes or whenever there is any change--both schedule updates and the time of updates are recorded in the database. Regarding flight cancellation decisions, airlines' dispatchers are expected to continuously make choices with the most updated flight information over the entire course of the GDP. This makes decision modeling difficult since relevant factors such as delay or available alternative flights for re-booking passengers may also change over time. In addition, airlines sometimes start making cancellation decisions in anticipation of a GDP even before it is officially announced, because they can often predict that a GDP will be implemented based on weather forecast information. It is very challenging to take into account the dynamic and anticipatory nature of dispatchers' decision making.

For purposes of this analysis, it is assumed that airlines make cancellation decisions immediately after the initiation of a GDP. There are two main consequences of such an assumption.

One is that flight cancellations made before the initiation of a GDP need to be excluded because it is not possible to determine from the database whether these cancellations are in response to the GDP or some other pre-existing conditions. Also, some information regarding such pre-GDP cancellations is not available. For example, delays are not yet officially assigned to these flights through the GDP. Thus there is no way that pre-GDP cancellations can be modeled under the same utility function as cancellations made after the GDP initiation. Similarly, flights that have taken off before the initiation of a GDP should be eliminated because theoretically once flights are airborne, they can no longer be cancelled, except by means of a so-called "Diversion Cancellation". Such cancellations are rare, and it is expected that decisions regarding them are qualitatively different from those for flights that have yet to depart.

The other consequence of the assumption is that cancellation decisions are made based on the flight schedules published at the time GDP is implemented. This simplifies the dynamic decision making process and makes modeling it tractable.

4.3.2 Statistics of Sample Data

The data from all GDPs that took place in 2006 are obtained for this study, excluding the GDPs implemented in combination with other TMIs, such as ground stops. Only domestic flights are included in the model for several reasons: One is that international flights are more difficult to cancel because of the high level of coordination required at international origins and with transoceanic air traffic control. Second, the stage lengths of international flights are usually long and cancellations need to be planned well in advance. For these reasons GDP-induced cancellations of international flights are extremely rare.

This study focuses on eight airports with the largest number of GDPs in 2006, which are Hartsfield-Jackson Atlanta International Airport, Boston/General Edward Lawrence Logan International Airport, Newark Liberty International Airport, John F. Kennedy International Airport, LaGuardia Airport, Chicago O'Hare International Airport, Philadelphia International Airport, and San Francisco International Airport. Among the GPDs that took place at these eight airports, only flights from the 11 major airlines and their subcarriers are considered in the model. The 11 airlines are American Airlines, United Airlines, Airtran Airways, Alaska Airlines, Continental Airlines, Delta Airlines, Frontier Airlines, Jetblue Airways, Northwest Airlines, Southwest Airlines and US Airways.

As a result, there are 624 GDPs in our sample data, and a total of 167,584 flights subject to cancellation. The actual number of cancelled flights is 5,584, giving a cancellation rate of 3.33%. Table 4.3 lists the number of GDPs and GDP impacted flights by airports. Table 4.4 lists the number of GDP impacted flights by airlines. Similarly, Figure 4.1 shows the distribution of GDP impacted flights among airports. Figure 4.2 shows the distribution of GDP impacted flights among airlines.

Table 4.3 Number of GDPs and GDP-impacted Flights, by Airport, 2006

Airport Name	Airport Code	Number of GDPs	Number of GDP-impacted Flights
Hartsfield-Jackson Atlanta International Airport	ATL	43	20,781
Boston/General Edward Lawrence Logan International Airport	BOS	62	11,129
Newark Liberty International Airport	EWR	115	28424
John F. Kennedy International Airport	JFK	36	4,657
LaGuardia Airport	LGA	87	26,605
Chicago O'Hare International Airport	ORD	90	49,162
Philadelphia International Airport	PHL	67	14,275
San Francisco International Airport	SFO	124	12,551

Table 4.4 Number GDP Impacted Flights, by Airline, 2006

Airline	Acronym	Hubs	Number of GDP Impacted Flights
American Airlines	AAL	ORD	38,533
United Airlines	UAL	ORD, SFO	32,255
Airtran Airways	TRS	ATL	8,920
Alaska Airlines	ASA		1,114
Continental Airlines	COA	EWR	25,353
Delta Airlines	DAL	ATL, JFK	26,325
Frontier Airlines	FFT		593
Jetblue Airways	JBU	JFK, BOS	3,650
Northwest Airlines	NWA		2,793
Southwest Airlines	SWA		2,162
US Airways	USA	PHL	25,886

Figure 4.1 Distribution of GDP Impacted Flights among Airports, 2006

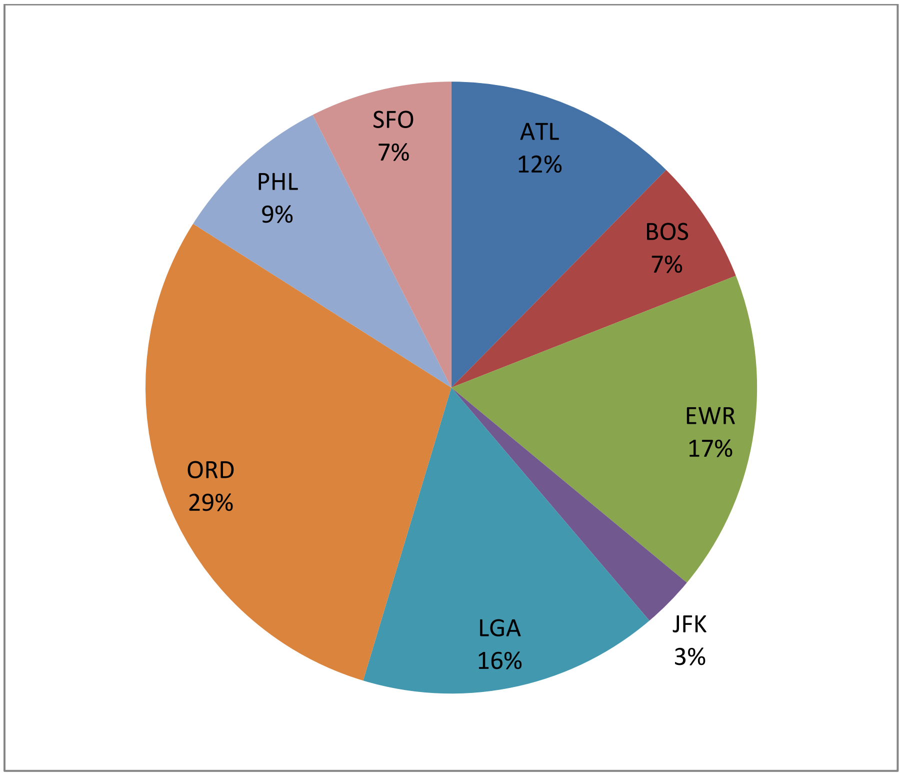

Figure 4.2 Distribution of GDP Impacted Flights among Airlines, 2006

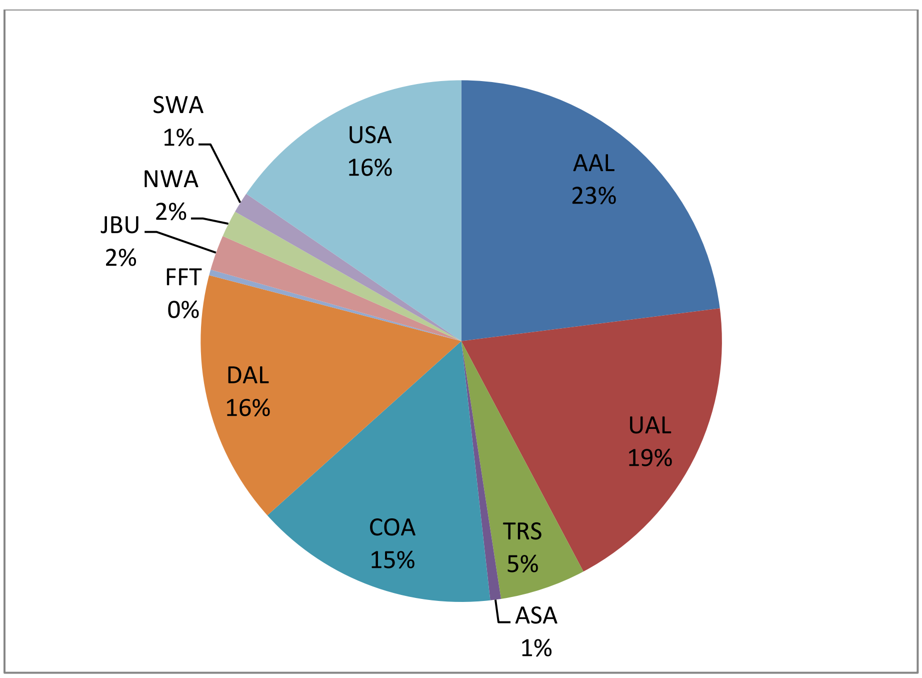

4.4. Model Specifications and Estimation Results

A choice made by the airlines' dispatchers to cancel or not cancel a flight is modeled. Many models with a variety of specifications are estimated. Among these, five are presented in this section. The first three models assume that cancellation utility is the same for all flights and all airlines. A baseline model, referred as Model C-1, is estimated using the standard logit model specification with all variables listed in Table 4.2. The second model (Model C-2) is based on Model C-1, and focuses on exploring the effect of *GID*, by employing a piece-wise linear specification. The third model (Model C-3) is also based on Model C-1 and adds quadratic and interactive forms of select variables to capture the non-linearity of these variables' impacts on the cancellation utility.

The fourth and fifth models (Model C-4 and Model C-5) relax the assumption that the same utility is shared by all flights and airlines. Model C-4 estimates a mixed logit model, allowing taste variation among flights. Model C-5 investigates the behavior of individual airlines. The specification of the Model C-1 is applied to four legacy airlines: A, B, C, and D. It is also applied to a low cost airline, L. To respect the privacy of airlines, the real names of airlines are not revealed. The details of each model are presented in the subsequent subsections.

4.4.1 Baseline Model C-1

The deterministic part of the baseline cancellation utility is formulated as a linear function in parameters as follows:

$$\beta_0 + \beta_1 * GID + \beta_2 * Delay\ saving + \beta_3 * Distance + \beta_4 * Hub\ destination\ dummy + \beta_5 * Major\ airlines\ dummy + \beta_6 * Internal\ delay + \beta_7 * Frequency + \beta_8 * Seats + \beta_9 * Average\ fare + \beta_{10} * Load\ factor + \beta_{11} * Market\ fare \quad (4.3)$$

The definition of variables are summarized in Table 4.5

Table 4.5 Summary of Definition of Variables

Variables	Definitions
GID	is GDP assigned initial delay (in minutes)
Delay savings	is the total delay savings (in minutes) to all other flights, if this flight is hypothetically cancelled
Distance	is the great circle distance (in miles) between origin and destination
Hub destination dummy	is a dummy variable to represent whether the airline operating the flight has a hub at the GDP airport; 1 for hub, 0 otherwise
Major airlines dummy	is a dummy variable that represents whether the flight belongs to a major airline or its subcarriers; 1 for major airlines, 0 otherwise
Internal delay	is the delay (in minutes) resulting from an internal airline problem prior to the GDP
Frequency	is the number of available flights on the segment by the airline in the GDP
Seats	is an estimate of the number of available seats on the flight
Average fare	is an estimate of the average fare paid (in dollars) by passengers onboard the flight
Load factor	is an estimate of the load factor on the segment for the aircraft and airline
Mark fare	is an estimate of the average fare (in dollars) to travel between two markets
$\beta_0, \beta_1 \dots \beta_{11}$	are coefficients that are estimated in the model

Statistics on all variables used in Model C-1 are summarized in Table 4.6. The standard logit model is estimated using SAS software and the estimation results are summarized in Table 4.7.

Table 4.6 Basic Statistics of All Variables in Model C-1

	MIN	MAX	MEAN	STD
GID (in minutes)	0.00	2682.00	51.25	57.83
Delay savings (in minutes)	0.00	1693.00	289.16	231.05
Distance (in miles)	54.95	4955.98	746.29	568.43
Hub destination dummy	0.00	1.00	0.63	0.48
Major airlines dummy	0.00	1.00	0.63	0.48
Internal delay (in minutes)	0.00	1434.00	2.56	17.74
Frequency	1.00	20.00	4.76	3.25
Number of seats (in seats)	19.00	348.00	110.94	52.41
Average fare (in $)	9.94	573.68	126.60	56.27
Load factor	0.00	100.00	72.28	14.13
Market fare (in $)	16.43	348.21	150.14	47.47

Table 4.7 Estimation Results of Model C-1

	Estimate	Std Error
Intercept	**-4.3372**	0.1012
GID (in minutes)	**0.0100**	0.0002
Delay savings (in minutes)	**0.0017**	0.0001
Distance (in miles)	**-0.0005**	0.0001
Hub destination dummy	**0.2058**	0.0335
Major airlines dummy	**0.1491**	0.0678
Internal delay (in minutes)	**0.0073**	0.0004
Frequency	**0.0272**	0.0048
Number of seats (in seats)	**-0.0052**	0.0007
Average fare (in $)	**0.0019**	0.0005
Load factor	**-0.0042**	0.0011
Market fare (in $)	**0.0018**	0.0004

Note: Bolded estimates are statistically significant at the 0.05 level

The results show that all the variables of Model C-1 are statistically significant. The intercept is negative, which means that airlines are generally not willing to cancel flights. The coefficients of *GID* and *delay savings* are both positive. The positive *GID* coefficient confirms that if greater GDP delay is imposed on the flight, the more likely it is that the flight will be cancelled. The positive *delay savings* coefficient confirms that a flight is more likely to be cancelled if greater delay savings can be achieved from its cancellation. *Distance* enters the model with a negative impact, as expected, which confirms that airlines hedge by cancelling shorter flights. Flights that are bound for their hub airport are more likely to be cancelled, which is also a confirmed hypothesis.

However, the hypothesis that airlines prefer to cancel flights operated by a commuter affiliate to benefit their own fleet is not supported by the estimation results. In fact, all else being equal, the flight that belongs to a major airline is more likely to be canceled than one of a regional affiliate. The estimated coefficient of *internal delay* has the expected sign. This suggests that flights disrupted as a result of an airline's internal problems are more likely to be cancelled, a phenomenon expected even in the absence of a GDP. The coefficient of *frequency* is of the expected sign, implying that airlines are more likely to cancel when more alternative flights are available. The coefficient of *aircraft size* is negative, as expected. Airlines are more likely to cancel a flight operated with a smaller aircraft. However the coefficient of *average fare* is positve. The result is unexpected and will be investigated more in this section. The *load factor* has a negative effect on cancellation utility, as expected. The higher a load factor on a segment, the more difficult it is to rebook passengers from cancelled flights. Finally, the coefficient on *market fare* is positive, which indicates that airlines are more likely to cancel flights that serve a predominantly business travel market. This may be a result of the idea that a business travel market often has more frequent service.

In addition to the above qualitative discussion of model results, three variables of particular interests-- *GID*, *aircraft size*, and *average fare*-- are explored further to see how exactly they influence cancellation decisions.

Figure 4.3 Impact of GID on the Cancellation Decisions

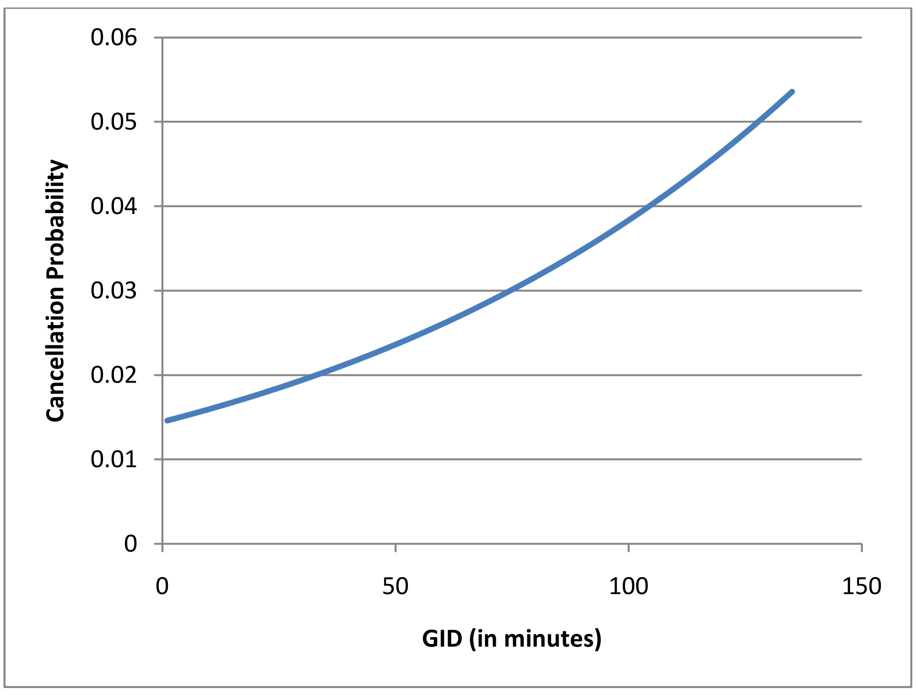

The impact of *GID* is studied by varying the *GID* only and keeping the other variables at their sample means. This is done in order to observe changes in cancellation probability due to *GID*. Figure 4.3 shows the impact of *GID* on the cancellation decision.

The range of *GID* is set from its 0.05 sample percentile of 1 minute to its 0.95 sample percentile value of 141 minutes.

Figure 4.3 shows that cancellation probability increases with the *GID*, again confirming the hypothesis. Moreover, a two-hour *GID* ,whose cancellation probability is 0.0464 as shown in Figure 4.3, makes a flight 2.4 times more likely to be cancelled than a half-hour *GID* ,whose cancellation probability is 0.0194 shown in the same figure.

In a similar way, the effect of *aircraft size* on cancellation decision is considered. Following the approach for *GID*, the cancellation probability is calculated as a function of seats only, keeping the other variables at their sample means. The range of seats is again set from its 0.05 percentile to 0.95 percentile, or 37 seats to 188 seats according to the sample. Figure 4.4 shows that cancellation probability decreases with *aircraft size*, as expected. Compared to an average sized (110-seat) flight (with cancellation probability of 0.0240 as shown in Figure 4.4), a 40-seat flight (with cancellation probability of 0.0341) is 1.4 times more likely to be cancelled. A 180-seat flight (with cancellation probability of 0.0168) is 30% less likely to be cancelled.

Figure 4.4 Impact of Aircraft Size on the Cancellation Decisions

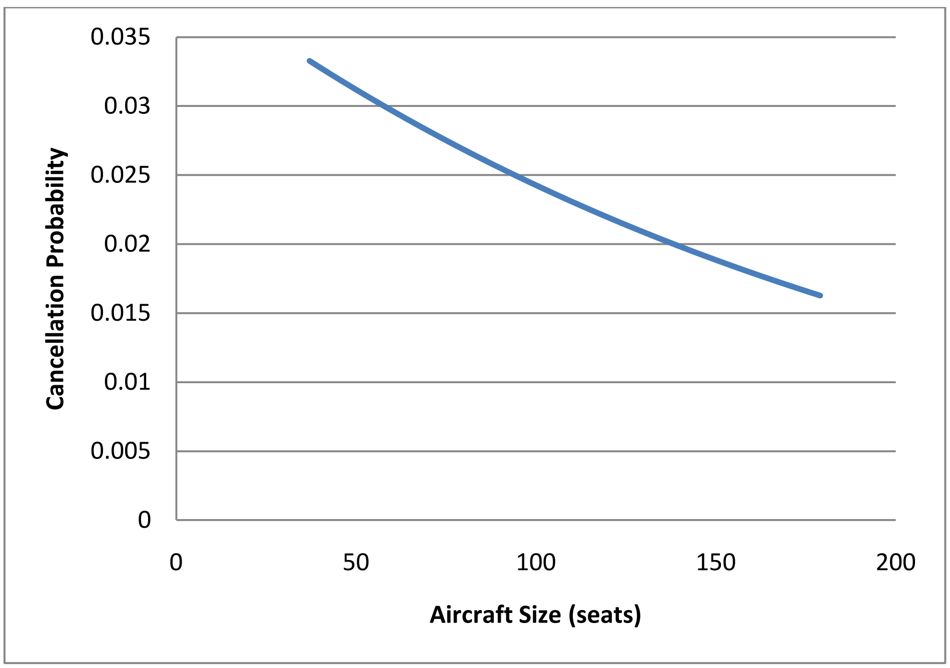

The impact of *average fare* on cancellation decisions is also investigated. Following the same procedure as above, the impact of *average fare* on cancellation probability is plotted in Figure 4.5, assuming other variables are held constant at their sample means. The range of *average fare* is from \$24 to \$158, the 0.05 and 0.95 sample percentiles respectively. Figure 4.5 shows that the chance of cancellation increases with the *average fare*. However, the magnitude of this variable's effect on cancellation probability is much smaller compared to *GID* and *aircraft size*.

Figure 4.5 Impact of Average Fare on the Cancellation Decisions

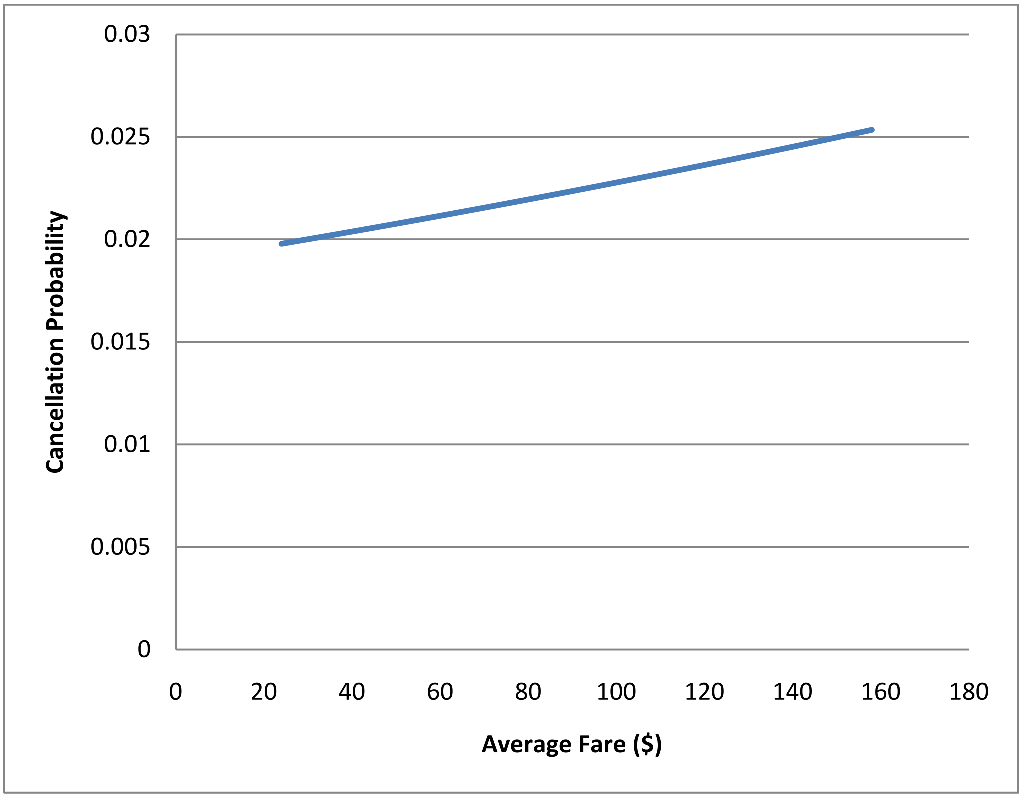

4.4.2 Piece-wise Linear Model C-2

Model C-2 is developed to further test the hypothesis that the impact of *GID* is piece-wise linear in the utility. Model C-2 uses the Model C-1 specification except that the *GID* variable is broken down into four *GID* variables-- *GID1* through *GID4*.

$$\begin{aligned} &GID1=\min(GID,15) \\ &GID2=\min(\max(0,GID-15),30) \\ &GID3=\min(\max(0,GID-45),45) \\ &GID4=\max(0, GID-90) \end{aligned} \quad (4.4)$$

The deterministic part of the cancellation utility is specified as follows.

$$\beta_0 + \beta_1 * \ GID1 + \beta_2 * \ GID2 + \beta_3 * \ GID3 + \beta_4 * \ GID4 + \beta_5 * Delay\ saving + \beta_6 * \ Distance + \beta_7 * Hub\ destination\ dummy + \ \beta_8 * Major\ airlines\ dummy \ + \beta_9 * Internal\ delay \ + \beta_{10} * Frequency + \beta_{11} * Seats + \beta_{12} * Average\ fare + \beta_{13} * \ Load\ factor + \ \beta_{14} * Market\ fare \quad (4.5)$$

The estimation results of Model C-2, along with Model C-1, are summarized in Table 4.8.

Table 4.8 Estimation Results of Model C-2

	Model C-2		Model C-1	
	Estimate	Std Error	Estimate	Std Error
Intercept	**-4.2868**	0.1157	**-4.3372**	0.1012
GID (in minutes)			**0.0100**	0.0002
GID1 (in minutes)	-0.0021	0.0054		
GID2 (in minutes)	**0.0081**	0.0022		
GID3 (in minutes)	**0.0214**	0.0011		
GID4 (in minutes)	**0.0075**	0.0003		
Delay savings (in minutes)	**0.0018**	0.0001	**0.0017**	0.0001
Distance (in miles)	**-0.0005**	0.0001	**-0.0005**	0.0001
Hub destination dummy	**0.1977**	0.0333	**0.2058**	0.0335
Major airlines dummy	**0.1690**	0.0676	**0.1491**	0.0678
Internal delay (in minutes)	**0.0074**	0.0005	**0.0073**	0.0004
Frequency	**0.0223**	0.0048	**0.0272**	0.0048
Number of seats (in seats)	**-0.0052**	0.0007	**-0.0052**	0.0007
Average fare (in $)	**0.0018**	0.0005	**0.0019**	0.0005
Load factor	**-0.0040**	0.0011	**-0.0042**	0.0011
Market fare (in $)	**0.0018**	0.0004	**0.0018**	0.0004

Note: Bolded estimates are statistically significant at the 0.05 level

The results of Model C-2 are very similar to those of Model C-1 for all variables excluding *GID*. The result of the piecewise linear model of *GID* indicates that when delay is less than 15 minutes, the impact of a cancellation decision is insignificant--delays less than 15 minutes are not of concern to airlines. This could be due to the fact that in the current system, flights with delays of less than 15 minutes are considered to be on-time for purposes of measuring airline on-time performance. For the other ranges of *GID*, the coefficients are all positive and significant but of different magnitudes, with the largest coefficient for the 45-90 minute range (*GID3*). This model confirms the hypothesis that the impact of *GID* on cancellation utility is non-linear; cancellation probability is impacted at an increasing rate until about 90 minutes, and then at a decreasing rate for delays above 90 minutes.

Figure 4.6 plots the impact of *GID* on cancellation probability, comparing the Model C-2 with Model C-1. The comparison shows that the piecewise linear model predicts a higher cancellation probability when the GID exceeds 63 minutes, and the difference becomes quite large fairly quickly.

Figure 4.6 Impact of GID on the Cancellation Decisions --Comparison of Model C-1 and Model C-2

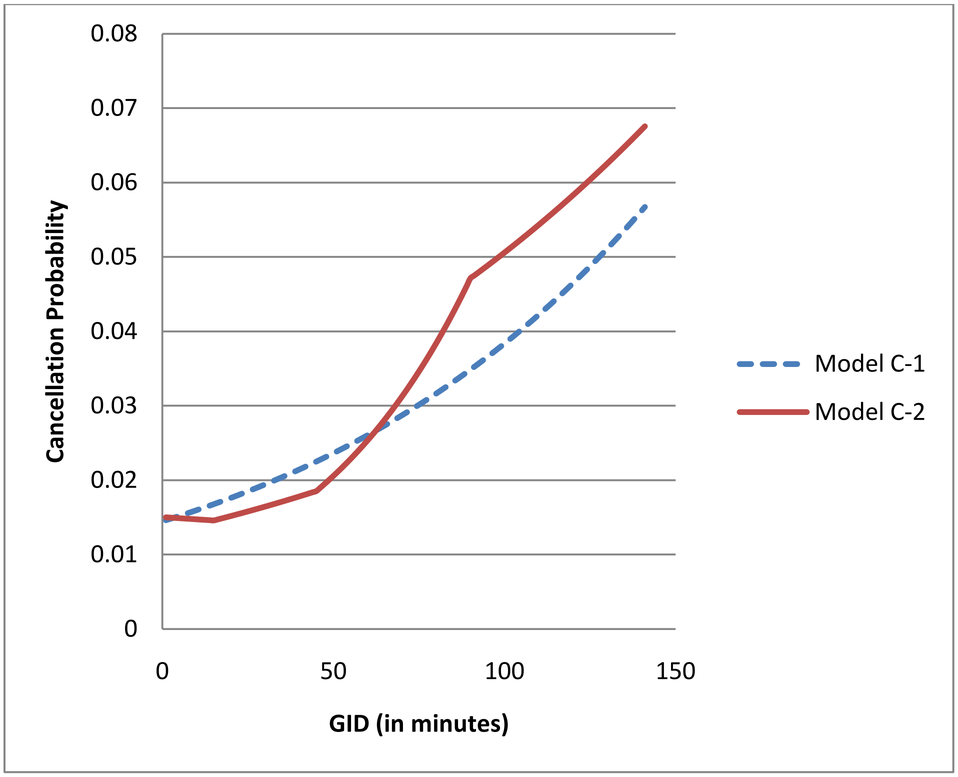

4.4.3 Model C-3 with Quadratic/Interactive Terms

The attributes entering into the cancellation utility of Model C-1 are in first-order forms. It is possible that some of these attributes influence cancellation utility in a non-linear manner; for example *GID*, as confirmed by its piece-wise linear representation in Model C-2. While small *GID* values may be of little consequence, as *GID* increases it may exert a much stronger system-wide impact and thus have a greater impact on the cancellation utility. If this were true, a quadratic *GID* term would be expected to have a positive coefficient in the cancellation utility function. Model C-3 is developed to consider the non-linear effects of variables like *GID*.

Delay savings (from a hypothetical cancellation) on five flights of five minutes each are expected to have a different impact on the system than a *delay savings* on one flight of 25 minutes. While the total delay savings are 25 minutes in both cases, the sum of squared delay savings is different: $5 \cdot 5^2 = 125 \min^2$ in the first case versus $1 \cdot 25^2 = 625 \min^2$ in the second. It is hypothesized that airlines prefer large *delay savings* on a small number of flights (second case) to small *delay savings* on a large number of flights (first case). Therefore a quadratic term for delay savings, called *delay saving squared*, is included and expected to be positive. It is calculated by summarizing the square of *delay savings* to individual flights using the algorithm depicted in Table 4.1, over all flights experiencing delay savings.

Furthermore, the impact of *GID* on cancellation decisions may depend on certain flight characteristics, such as *aircraft size*, and *average fare* paid by the passengers. Two interaction terms are introduced to the model: *GID multiplied by aircraft size* and *GID multiplied by average fare*.

Similarly, flight heterogeneity can be considered in *delay saving* metrics. *Delay saving weighted by seats* is developed by multiplying minutes of *delay saving* to individual flights by the *aircraft size* of the individual flights, and then summarizing over all flights experiencing delay savings. It should be noted that this may not be the maximum seat-minute savings that could be obtained—it is simply the savings that are obtained from applying the algorithm depicted in Table 4.1. The hypothesis is that airlines derive a greater utility from a given amount of *delay savings* if the flights for which delay is being reduced have more seats.

Therefore five newly developed quadratic and interactive terms are added to the baseline Model C-1 to capture non-linear effects. The deterministic utility function of Model C-3 is:

$$\beta_0 + \beta_1 * GID + \beta_2 * (GID * GID) + \beta_3 * (GID * Seats) + \beta_4 * (GID * Average\ fare) + \beta_5 * Delay\ saving + \beta_6 * Delay\ saving\ squared + \beta_7 * Delay\ saving\ weighted\ by\ seats + \beta_8 * Distance + \beta_9 * Hub\ destination\ dummy + \beta_{10} * Major\ airlines\ dummy + \beta_{11} * Internal\ delay + \beta_{12} * Frequency + \beta_{13} * Seats + \beta_{14} * Average\ fare + \beta_{15} * Load\ factor + \beta_{16} * Market\ fare \tag{4.6}$$

The estimation results for Model C-3 are summarized in Table 4.9. All coefficients are statistically significant.

The coefficients on the first order variables in Model C-3 are similar to those of Model C-1. The overall impact of *GID* on cancellation utility in this model is a little complicated. If one takes the derivative of cancellation utility with respect to *GID*, the marginal effect of *GID* on cancellation utility becomes a function of seats and fare

$$\beta_1 + 2\beta_2 + \beta_3 * seats + \beta_4 * fare \tag{4.7}$$

$where$ $\beta_1, \beta_2, \beta_3, \beta_4$ are estimates from the Model C-3.

Table 4.9 Estimation Results of Model C-3

	Estimate	Std Error
Intercept	**-4.557200**	0.110500
GID (in minutes)	**0.015500**	0.000490
GID*GID (in minutes squared)	**-0.000005**	0.000000
GID*seats (in seat-minutes)	**-0.000030**	0.000003
GID*fare (in $-minutes)	**-0.000006**	0.000003
Delay savings (in minutes)	**0.000454**	0.000159
Delay saving squared (in minutes squared)	**0.000002**	0.000001
Delay saving weighted by seats (in seat-minutes)	**0.000012**	0.000001
Distance (in miles)	**-0.000480**	0.000056
Hub destination dummy	**0.194800**	0.034500
Major airlines dummy	**0.184400**	0.068000
Internal delay (in minutes)	**0.007860**	0.000437
Frequency	**0.028900**	0.004830
Number of seats (seats)	**-0.004290**	0.000793
Average fare (in $)	**0.002600**	0.000513
Load factor	**-0.005420**	0.001150
Market fare (in $)	**0.001510**	0.000412

Note: Bolded estimates are statistically significant at the 0.05 level

Given the coefficient estimates and sample means of *aircraft size* (110 seats) and *average fare* ($126.6), the marginal effect of *GID* is 0.0114, slightly larger than the 0.0100 of the baseline Model C-1. Figure 4.7 shows a comparison of the effects of *GID* on cancellation decisions according to Model C-1 and Model C-3. It again confirms that the marginal effect of GID is very similar in these two models for an average flight.

Moreover, Model C-3 confirms all hypotheses about *delay savings* metrics: Given a total delay savings, airlines prefer it is from large delay savings on a small number of flights; they also prefer delay savings on flights operated with larger aircraft.

Figure 4.7 Impact of GID on the Cancellation Decision -- Comparison of Model C-1 and Model C-3

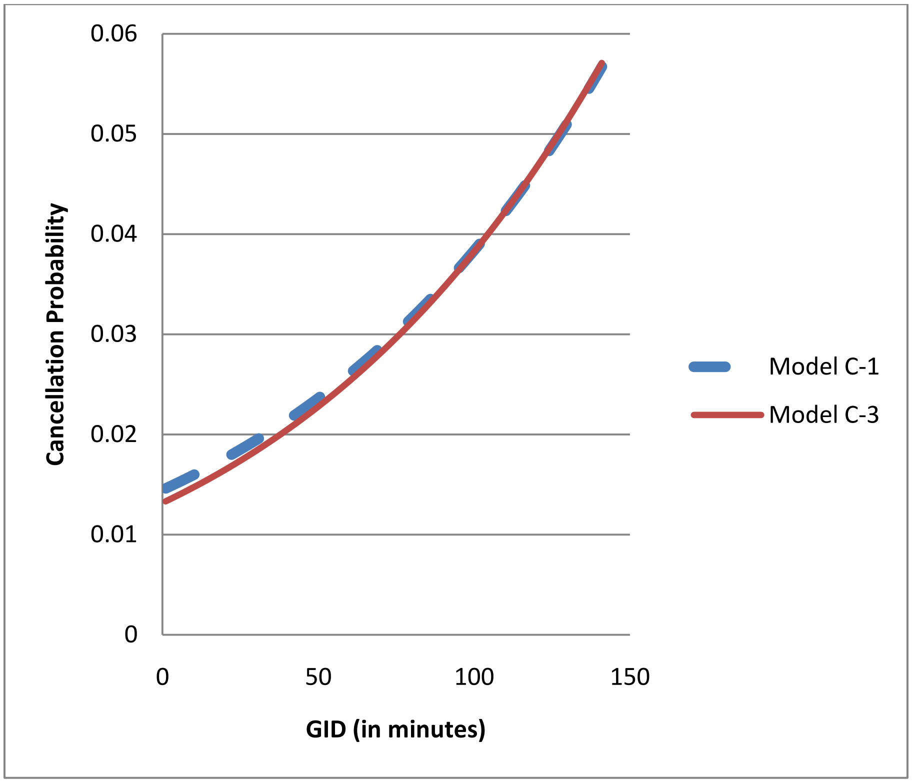

4.4.4 Mixed Logit Model C-4

The models presented in the previous sub-sections assume that the same utility function governs all flight cancellation decisions. However it is more reasonable to relax this assumption and allow some variation among the flights. For example, cancellation utilities for some flights may be more sensitive to *GID* than those of other flights. The *GID* coefficients for such flights would be higher.

The mixed logit model (also called the random coefficient model) is able to capture the "taste variation" among the flights by assuming that the coefficients of the cancellation utility are random variables instead of deterministic values. The SAS MDC procedure provides a convenient way to estimate a mixed logit model, when the number of coefficients whose values are assumed to be random is small. Normal distributions are assumed for the coefficients of the four variables that are of most interest: *GID*, *delay savings*, *aircraft size* and *average fare*. The coefficients on the remaining variables are assumed to be deterministic, as before.

The results are presented in Table 4.10, which also shows the standard logit model results for comparison. The standard deviation of the random coefficients—*GID*, *aircraft size* and *average fare*—are statistically significant, indicating that the impacts of these three variables indeed vary within the flights. However the standard deviation of the random coefficient on *delay savings* is not statistically significant, indicating that the effect of *delay savings* is fairly consistent among the flights.

Table 4.10 Estimation Results of Model C-4

		Mixed Logit Model		Standard Logit Model	
Variables	Parameter	Estimate	Std Error	Estimate	Std Error
Intercept	Mean coefficient	**-4.28240**	0.10870	**-4.33720**	0.10120
GID (in minutes)	Mean coefficient	**0.01050**	0.00026	**0.01000**	0.00019
	Std. dev. of coefficient	**0.00391**	0.00029		
Delay savings (in minutes)	Mean coefficient	**0.00174**	0.00007	**0.00168**	0.00006
	Std. dev. of coefficient	0.00003	0.00109		
Number of seats (in seats)	Mean coefficient	**-0.01010**	0.00168	**-0.00516**	0.00073
	Std. dev. of coefficient	**0.00612**	0.00099		
Average fare (in $)	Mean coefficient	**0.00196**	0.00050	**0.00189**	0.00046
	Std. dev. of coefficient	**0.00035**	0.00213		
Distance (in miles)	Mean coefficient	**-0.00054**	0.00006	**-0.00051**	0.00006
Hub destination dummy	Mean coefficient	0.18300	0.03480	**0.20580**	0.03350
Major airlines dummy	Mean coefficient	**0.31470**	0.08580	**0.14910**	0.06780
Internal delay (in minutes)	Mean coefficient	**0.00794**	0.00035	**0.00733**	0.00043
Frequency	Mean coefficient	**0.02470**	0.00520	**0.02720**	0.00478
Load factor	Mean coefficient	**-0.00345**	0.00120	**-0.00415**	0.00114
market fare (in $)	Mean coefficient	**0.00198**	0.00040	**0.00181**	0.00041

Note: Bolded estimates are statistically significant at the 0.05 level

Consider first the *GID* estimates. The distribution of the coefficient of *GID* has an estimated mean of 0.01050 and estimated standard deviation of 0.00391, such that 99.6% of the distribution is above zero and 0.4% is below. Thus *GID* is a positive inducement for cancellation for 99.6% flights. While the variation in the *GID* coefficient is significant, the mean value is very close to that obtained from the standard logit model.

In the case of *aircraft size*, the estimated mean is –0.01010 and estimated standard deviation is 0.00612. Thus *aircraft size* is a negative factor for 95.07% of the flights and a positive factor for the remaining 4.93%. In addition, there is a substantial difference between the estimate for the mean value of the coefficient and that obtained in the standard logit model. The effect of *aircraft size* on cancellation utility is almost twice as large in the mixed logit model as what was estimated from the standard logit model.

The distribution of the *average fare* coefficient has an estimated mean of 0.00196 and estimated standard deviation of 0.00035, such that *average fare* is a positive factor for all flights. While the variation in the average fare coefficient is significant, the mean value is very close to that obtained from the standard logit model.

To summarize the findings on random coefficients, the standard deviations are relatively small so the flight-to-flight variations in utility are not great. The fixed coefficients on the other variables are similar to those of the standard logit model.

4.4.5 Individual Airline Model C-5

Individual airline models can be developed by applying the model structure in the Model C-1 to individual airline data. Five out of eleven airlines are selected for this study: Airline A, Airline B, Airline C, Airline D, and Airline L. The first four are legacy airlines and operate conventional hub-spoke networks. Airline L is a low cost airline and operates a point-to-point network. Moreover, the aircraft fleets of the legacy airlines include a variety of sizes while low cost airline only operates one type of aircraft. Consequently, *aircraft size* will not affect Airline L's flights cancellation decisions and will not be included in the model. Furthermore, Airline L has no commuter affiliates or hub airport, eliminating the need for the *major airline dummy* and *hub destination dummy* in the model. The model results are summarized in Table 4.11.

Table 4.11 Estimation Results of Model C-5

	Airline A		Airline B		Airline C		Airline D		Airline L	
	Estimate	Std Error	Estimate	Std Error	Estimate	Std Error	Estimate	Std Error	Estimate	Std Error
Intercept	**-3.3800**	0.2697	**-3.4421**	0.2553	**-3.2948**	0.3315	**-3.3711**	0.2454	**-4.8916**	2.2463
GID (in minutes)	**0.0090**	0.0003	**0.0070**	0.0006	**0.0086**	0.0004	**0.0102**	0.0005	**0.0126**	0.0021
Delay savings (in minutes)	**0.0024**	0.0001	**0.0017**	0.0001	**0.0022**	0.0001	**0.0012**	0.0002	**0.0024**	0.0006
Distance (in miles)	**-0.0005**	0.0001	**-0.0004**	0.0002	**0.0003**	0.0001	**-0.0014**	0.0003	0.0003	0.0013
Hub destination dummy	**0.0591**	0.0710	0.0758	0.0800	0.1256	0.1031	-0.1406	0.1129		
Major airlines dummy	**0.8324**	0.2059	**1.2621**	0.1803	**0.9404**	0.1230	**0.6269**	0.2276		
Internal delay (in minutes)	**0.0058**	0.0006	**0.0148**	0.0011	**0.0055**	0.0011	**0.0054**	0.0016	-0.0298	0.0848
Frequency	0.0141	0.0083	-0.0088	0.0112	**0.0402**	0.0117	-0.0143	0.0140	**0.2520**	0.1138
Number of seats (in seats)	**-0.0075**	0.0023	**-0.0134**	0.0018	**-0.0085**	0.0013	-0.0035	0.0025		
Average fare (in $)	0.0006	0.0012	**0.0066**	0.0016	**-0.0035**	0.0013	-0.0001	0.0016	-0.0013	0.0266
Load factor	**-0.0158**	0.0025	0.0039	0.0028	**-0.0247**	0.0033	**-0.0114**	0.0032	-0.0120	0.0211
Market fare (in $)	**0.0023**	0.0010	**-0.0062**	0.0010	**0.0041**	0.0011	-0.0004	0.0012	-0.0139	0.0302

Note: Bolded estimates are statistically significant at the 0.05 level

As expected, each model's intercept is negative and statistically significant. Airline L's intercept has the highest value while the four legacy airlines' intercepts have similar values. This indicates that in general, low cost airline is much less willing to cancel flights compared to the other legacy airlines.

A higher *GID* increases the cancellation utility for all airlines. Following the same procedure used to calculate the *GID* impact in Model C-1, the impact of *GID* on each individual airline is developed. Figure 4.8 compares the *GID* impacts on cancellation probability among airlines. Airline L is always the least likely to cancel flights regardless of the quantity of *GID*, followed by Airline D and Airline C. Airline A and Airline B have similar results.

Figure 4.8 Impact of GID on the Cancellation Decisions --Comparison over Airlines

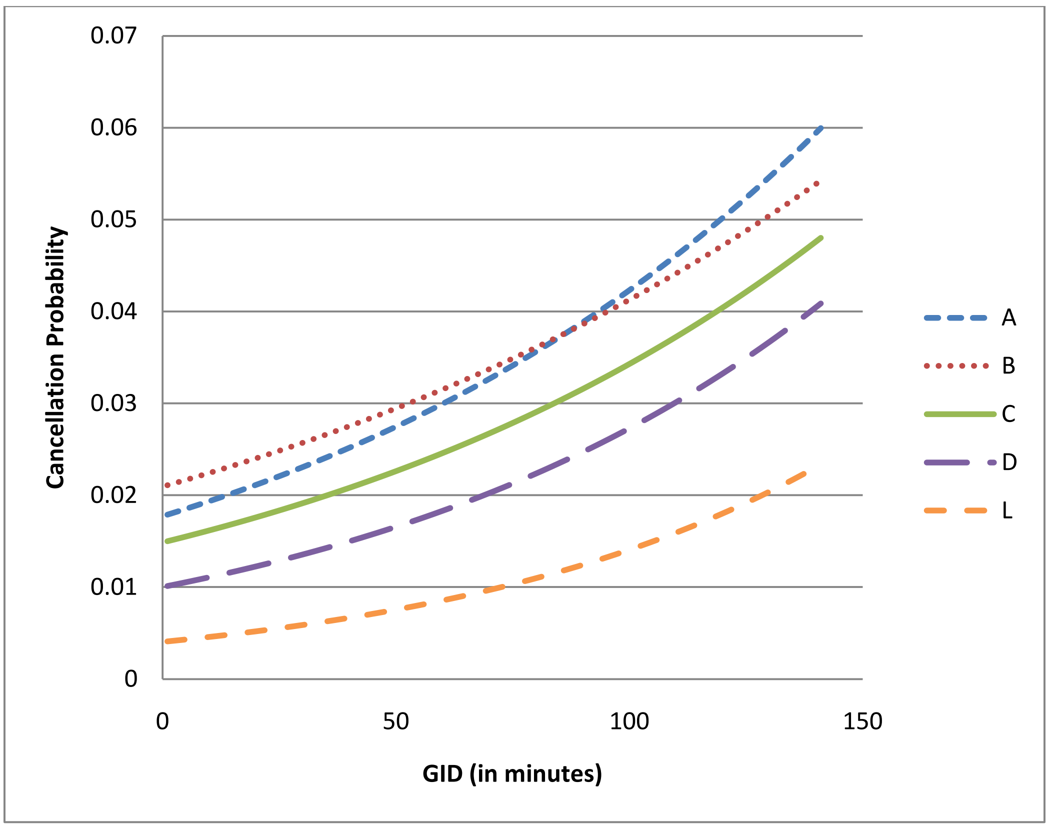

Aircraft size has a negative and statistically significant coefficient in the cancellation utility functions for Airline A, Airline B and Airline C. It is negative but insignificant for Airline D. When comparing the four legacy airlines with regard to *aircraft size* impact, as plotted in Figure 4.9, the results show that Airline B is most likely to cancel their small aircraft, while Airline D is least likely to do so. When dealing with large size aircraft, such as a flight with more than 150 seats, airlines behave more consistently with one another.

Figure 4.9 Impact of the Aircraft Size on the Cancellation Decisions --Comparison over Airlines

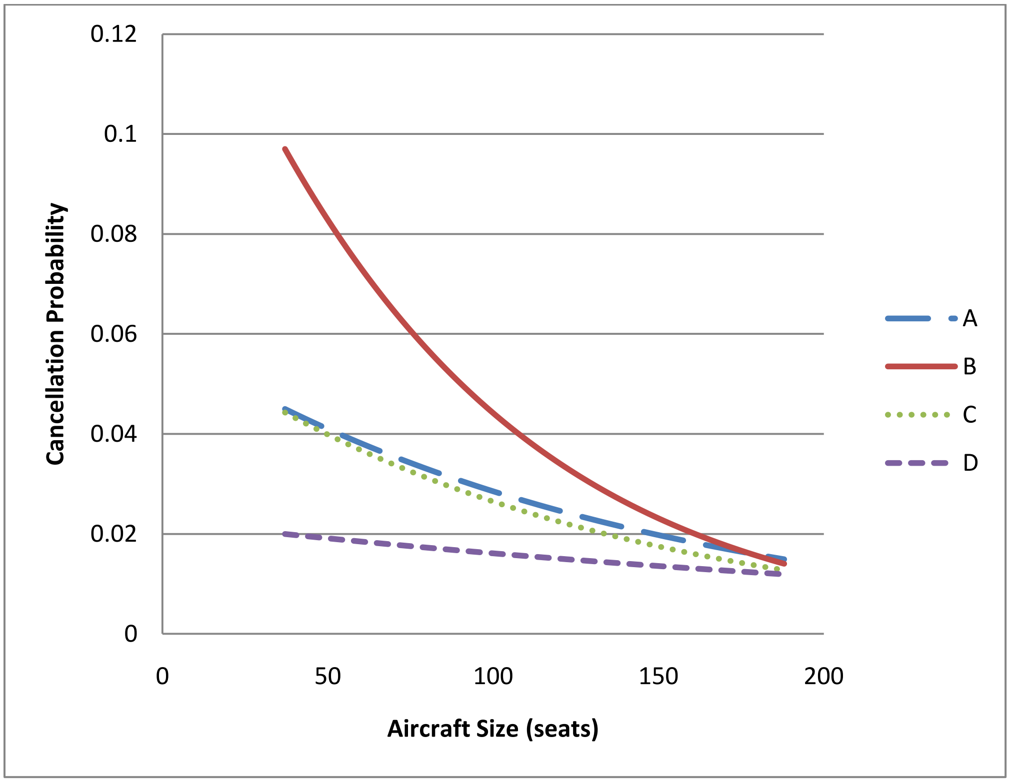

At last, the coefficient of *average fare* is significant only for the Airline C and Airline B. Figure 4.10 compares the impact of *average fare* on the cancellation probability. The *average fare* increases cancellation utility for the Airline B's flights, while it decreases the utility for the Airline C. Thus Airline C is the only carrier whose behavior matches author's expectation insofar as they are more likely to cancel flights with high value passengers onboard.

Figure 4.10 Impact of Average Fare on the Cancellation Decisions— Comparison over Airlines

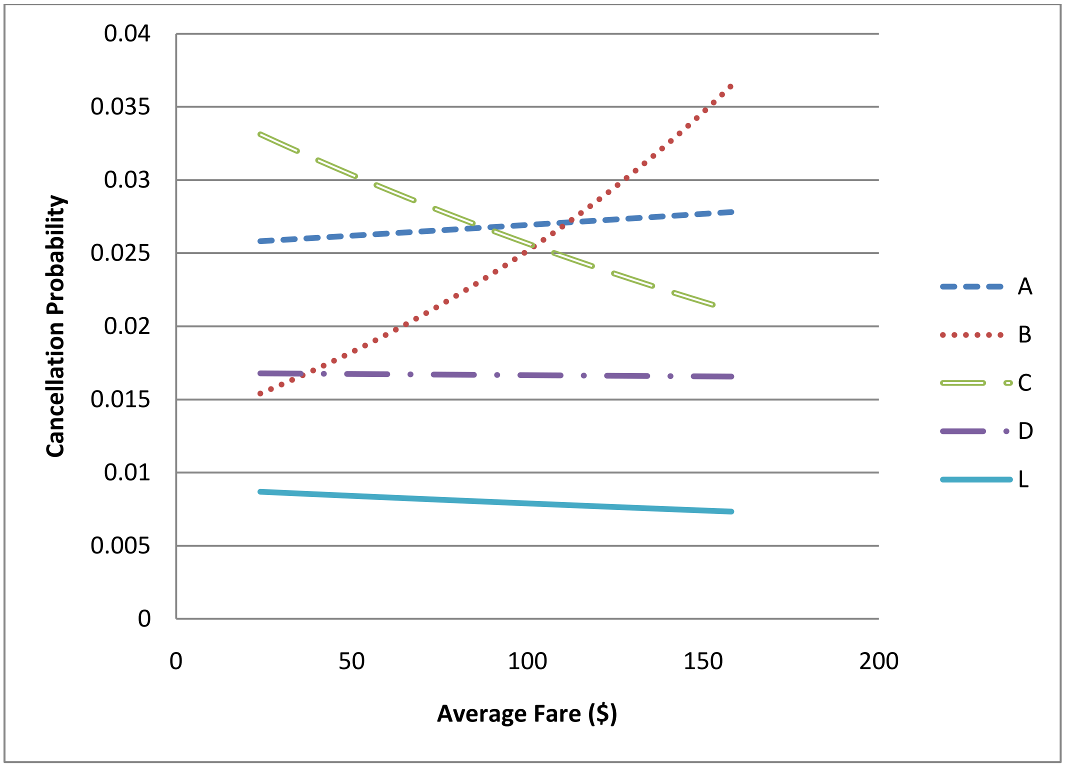

4.5 Cancellation and Delay Tradeoff

When the deterministic part of the cancellation utility is zero, the probability of a flight being cancelled is 0.5. The airline leans toward cancellation if the deterministic part of utility is greater than zero. Based on the estimated function of the deterministic part of utility, for example from baseline Model C-1, the number of minutes of *GID* and *delay savings* required for a flight to have a 0.5 cancellation probability can be calculated. This is done by setting the deterministic part of utility to zero and adopting the sample means of variables other than *GID* and *delay savings.* The analysis is based on the estimated coefficients ($\beta_0, \beta_{1..}\beta_{11}$) for Model C-1 shown in Table 4.7.

A flight's cancellation probability is 0.5 if its deterministic utility is zero:

$$\beta_0 + \beta_1 * \ GID + \beta_2 * Delay\ savings + \beta_3 * \ Distance + \beta_4 * Hub\ destination\ dummy + \ \beta_5 * Major\ airlines\ dummy\ + \beta_6 * Internal\ delay\ + \beta_7 * Frequency + \beta_8 * Seats + \ \beta_9 * Average\ fare{+}\beta_{10} * \ Load\ factor + \ \beta_{11} * Market\ far{=}0 \qquad (4.8)$$

Therefore, if all other variables are set to their sample means, Equation (4.8) becomes

$$0.01 * GID + 0.00168 * Delay\ saving = 4.7073 \tag{4.9}$$

When *GID* is zero, *delay savings* must be as large as 2801 minutes for a flight's cancellation probability to be 0.5. When no *delay savings* can be achieved, 471 minutes of GID will result in 0.5 cancellation probability. These results show that, for the vast majority of flights, delay considerations alone do not strongly incline airlines toward cancelling flights. Most cancellations occur when other factors—including the other variables included in the model as well as unobservable ones such as down-line connectivity—are also at play.

If *aircraft size* and *frequency* are not fixed at their sample means, the relationship can be further developed as a function of *GID*, *delay savings*, *aircraft size* and *frequency.*

$$0.01 * GID + 0.00168 * Delay\ savings - 0.00516 * Seats + 0.0272 * Frequency$$
$$= 4.26439 \tag{4-10}$$

The above equation shows the tradeoff between *GID* and *delay savings* that can result in 0.5 flight cancellation probability, given certain characteristics of the flight (seat capacity and frequency). Figure 4.11 plots such a relationship under four combinations of flight size and frequency.

Figure 4.11 *GID* and *Delay Savings* Required to Make Cancellation Probability 0.5

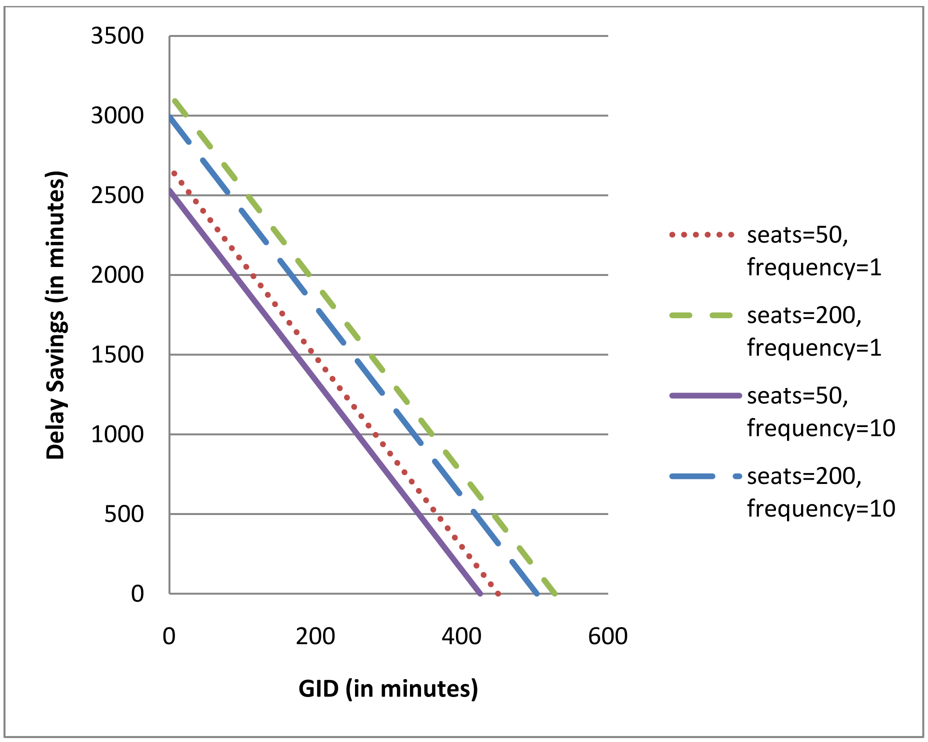

Figure 4.11 shows that at a particular frequency, for larger aircraft, larger *GID* and *delay savings* are required to maintain a flight cancellation probability of 0.5. Similarly, given an aircraft size, at lower frequency, larger *GID* and *delay savings* are

required to maintain a flight cancellation probability of 0.5. However, even in the case when the aircraft is small and the frequency is large, very few if any flights would have *GID* and *delay savings* values that make a cancellation likely.

4.6 Summary

In this chapter, the flight cancellation decision is investigated using a discrete choice model. Factors considered in the model include delay factors, flight characteristics, and flight segment characteristics. Five models are presented in this study. The first three models assume all flights have the same utility function. The fourth model relaxes such assumption and allows for variation among flights. The last model compares individual airline's cancellation utility functions.

The baseline Model C-1 suggests that airlines are more likely to cancel a flight if this flight receives a large delay from GDP (*GID*), or if its cancellation will result in much *delay savings* on other flights in the airlines' fleet. Airlines are also more likely to cancel a flight if is a short haul flight, if it is easier to rebook passengers due to more frequent service, and when the airline has an internal problem with the flight. Cancellation probability also increases slightly with *average fare*, but the magnitude of this effect is small. Cancellation utility decreases with *aircraft size* and *load factor* on the segment. Finally, airlines do not appear to cancel flights operated by a commuter affiliate to benefit their own fleet; however, they are more likely to cancel flights that are shorter and employ smaller aircraft, which are inherent characteristics of most affiliate flights.

Model C-2 demonstrates that there exists a piece-wise relationship between *GID* and cancelation utility. In particular, cancellation probability is impacted at an increasing rate between 15 minutes and 90 minutes, and then at a decreasing rate for delays above 90 minutes.

The non-linear impact of *GID* and *delay savings* on cancellation utility is further tested by introducing quadratic and interactive terms of these variables in Model C-3. The model confirms that for a given total delay savings, airlines prefer it is from a small number of flights and on larger aircraft.

Model C-4 employs a mixed logit specification and shows that there is significant variation in certain cancellation utility coefficients (*GID*, *aircraft size* and *average fare*) among the observations analyzed. However, for almost all flights, the signs of the random coefficients remain the same. In the case of *delay savings*, the variation is statistically insignificant.

Model C-5 investigates the individual airlines' cancellation utilities. Airline B stands out in that they are most aggressive in cancelling small size aircraft, and they cancel flights with high fare passengers. Low cost Airline L, as expected, is the least likely to cancel flights among all five airlines regardless of the amount of *GID* or *average fare*.

Finally, the tradeoff between cancellation and delay is analyzed using the estimation results from the baseline model. When *GID* is zero, *delay savings* must be as large as 2801 minutes for a flight’s cancellation probability to be 0.5. When no *delay savings* can be achieved, 471 minutes of GID will result in 0.5 cancellation probability. These results show that, for the vast majority of flights, delay considerations alone do not strongly incline airlines toward cancelling flights.

Chapter 5 SCS Model

5.1 Introduction

As briefly discussed in Chapter 2, Slot Credit Substitution (SCS) not only enables inter-airline substitution but also reveals the flight assignments that airlines prefer. Given the same set of flights and arrival slots, other feasible arrangements are then generated and assumed to be less preferable to airlines, compared with the proposed arrangement in the SCS messages. A multinomial logit model is applied to estimate airlines' underlying utility in making the SCS arrangements.

This chapter first describes the detailed SCS mechanism in section 5.2 to provide econometric modeling framework. An algorithm for generating all feasible arrangements is presented in section 5.3 and the resulting correlation issue is discussed in section 5.4. Section 5.5 describes the data. Section 5.6 develops explanatory variables. Two models are estimated and presented in section 5.7. Section 5.8 summarizes the chapter.

5.2 SCS Mechanism

A hypothetical example (shown in Table 5.1) is generated to illustrate how SCS works. Under GDP, three United Airlines' flights are assigned delayed arrival slots by the FAA (also viewed as the arrangement before the SCS message was sent), as seen in the first and second column of Table 5.1. However, this schedule is not favored by the airline and they attempt to reassign slots by sending a conditional message to the FAA: "I would like to cancel Flight 1 if Flight 2 can be moved to a new arrival slot between time 9:30 and time 10:00. And then my Flight 3 can be moved to take Flight 2's 10:00 slot." The FAA identifies a Delta Airlines' flight occupying the 9:40 slot can be moved to 9:00. As a result, the SCS message is fulfilled and the new schedule gets updated (column 3 in Table 5.1). After SCS, United Airlines' Flight 1 is cancelled. Flight 2 is advanced to 9:40 and Flight 3 can utilize the 10:00 slot pre-owned by Flight 2. The inter-airline substitution has been successful with assistance from the FAA.

Table 5.1 A Hypothetical Example of SCS

Flight	Schedule before SCS	Schedule after SCS
UA 1	9:00	10:30(cancelled)
UA 2	10:00	9:40
UA 3	10:30	10:00

The SCS message clearly reflects the airlines' tradeoff behavior—Flight 2 and Flight 3 benefit from the cancellation of Flight 1. It also presents the opportunity to estimate airlines' utility by considering joint decision making among the multiple flights mentioned in a single SCS message (also called a SCS packet). In comparison, the cancellation model discussed in Chapter 4 treats an individual flight as an independent observation-- cancelling a flight is assumed not to influence the cancellation decision on other flights.

It is obvious that airlines prefer the flight arrangement after SCS than schedule before SCS; otherwise they would not send the SCS message. A further assumption is made that the flight arrangement after SCS is the best among all feasible arrangements for the same set of flights and slots. As a result, the SCS message can be modeled using a multinomial choice model—the arrangement after SCS is the chosen alternative and all other feasible arrangements based on the same set of flights and slots are non-chosen alternatives. The chosen alternative is recorded in the database. However, the non-chosen alternatives need to be generated. This procedure is further described in the next section.

5.3 Choice Set Generation

The detailed choice set generation can be described by continuing with the example previously shown in Table 5.1. Based on the flights and slots assignment after SCS (copied to column 1 and 2 in Table 5.2, all possible assignments (including the chosen one) can be generated. Given three slots and three flights, and also the possibility of cancelling any but not all of these flights, the number of possible arrangements is:

$$P(3,3) + C(3,1) * P(3,2) + C(3,2) * P(3,1) \tag{5.1}$$

Where,

$P(3,3)$ is the permutation where no flights are cancelled.

$C(3,1) * P(3,2)$ is the permutation where one flight is cancelled.

$C(3,2) * P(3,1)$ is the permutation where two flights are canceled.

Thus, there are 24 possible arrangements in total. A sample of these is listed in Table 5.2.

Table 5.2 Choices Permutations

Flight	Arrangement After SCS	Alternative 1	…	Alternative10	…	Alternative24
UA 1	10:30 (cancelled)	9:40	…	10:00	…	10:00 (cancelled)
UA 2	9:40	10:00	…	10:30 (cancelled)	…	10:30
UA 3	10:00	10:30	…	9:40	…	9:40 (cancelled)

Not all 24 alternatives are feasible schedules due to *Earliest Estimate Time of Arrival (EETA)* constraints—a flight cannot arrive before its reported *EETA* time. For example, Flight 3 has an *EETA* of 9:50, so it cannot take the slot earlier than 9:50. Therefore Alternative 10 is not feasible and should be excluded from the choice set. However, if Flight 3 is cancelled, as shown in Alternative 24, the *EETA* constraint does not need to be met. Because assigning a slot to a cancelled flight is just to ensure the airline retains the cancelled slot. It is a "conceptual" assignment and there is no further meaning to associate the cancelled slot with the cancelled flight. Therefore, Alternative 24 is valid. After checking with *EETA* constraints, the remaining feasible alternatives are considered the choice set.

The algorithm to develop choice sets was programmed in Matlab. The Figure 5.1 shows the flow of the algorithm.

Figure 5.1 Choice Generation Algorithm

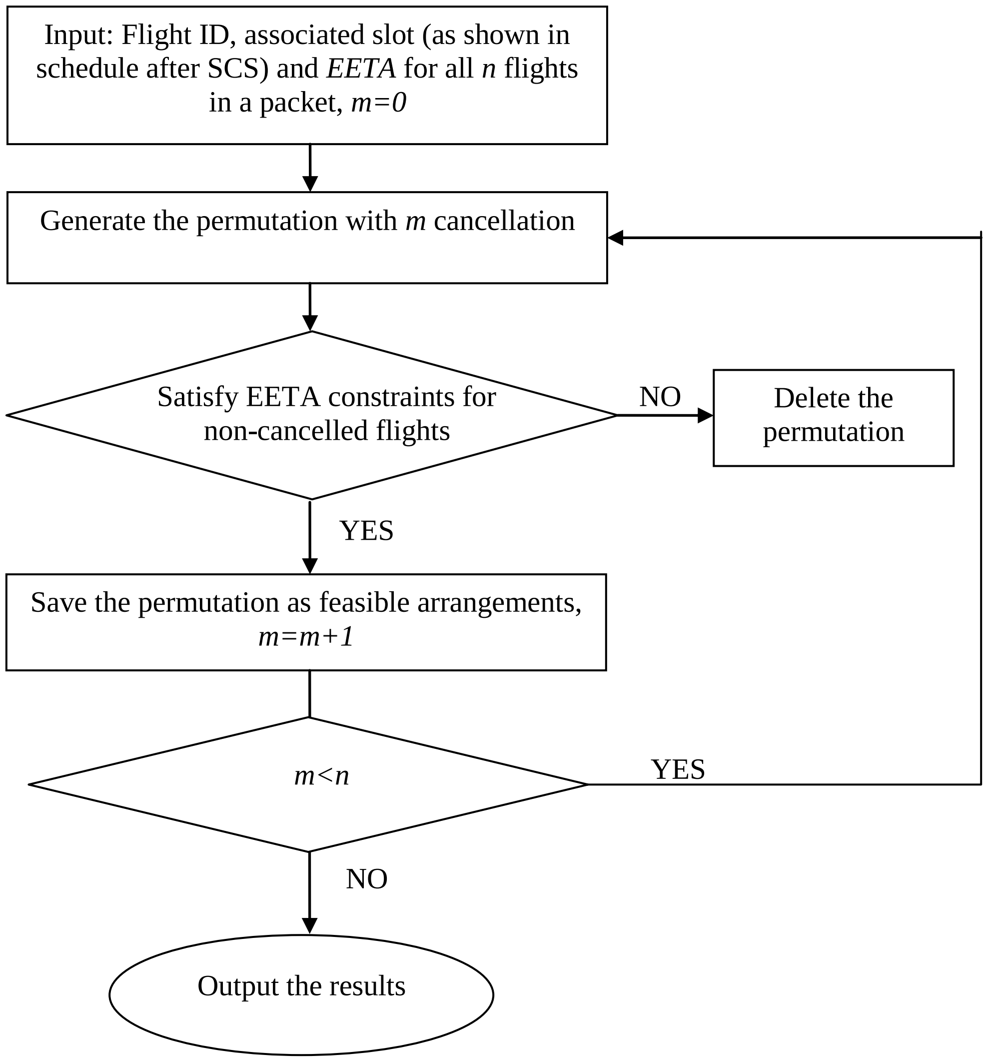

It is expected that the number of choice alternatives grows exponentially with the number of flights in a packet. Therefore, packets with more than five flights are not considered in the model. It is assumed that packets with less than or equal to five flights are representative and are generated under the same utility structure as the packets with more than five flights.

Due to the fact that the alternatives are generated based on one set of flights and slots, there may exist correlations between the alternatives such that the IIA assumption in the standard logit model is violated. The details are discussed in the next section.

5.4 Correlated Alternatives

The alternatives are correlated due to slots overlapping. Table 5.3 demonstrates an example. Alternative 1 is more similar to Alternative 2 than Alternative 3. This is due to

United flight 1 being assigned the same slot 9:40 in both Alternatives 1 and 2, while no flight is assigned to the same slot in Alternatives 1 and 3.

Table 5.3 Example of Generated Alternatives

Flight	Alternative 1	Alternative 2	Alternative 3
UA 1	9:40	9:40	10:30
UA 2	10:00	10:30	9:40
UA 3	10:30	10:00	10:00

This problem shares a similarity with the route choice problem, where generated route alternatives are usually correlated due to overlapping paths. The route choice problem is studied by Ben-Akiva and Bierlaire [36], who summarize several approaches that consider the correlations among alternatives. Rammin [37] extends Ben-Akiva and Bierlaire's study by comparing different approaches using results from empirical studies.

According to the literature, there are several methods of specifying models to account for correlations among alternatives. One would be to define a commonality factor and include it as an explanatory variable. This commonality factor is alternative specific and should capture how the alternative is perceived within a choice set.

A second method to address correlation among alternatives is to make assumptions about the covariance matrix, and then estimate the parameters of the covariance matrix explicitly from data. One reasonable assumption would be that the covariance is proportional to a correlation factor. The correlation factor could be defined based on how identical two alternatives are; for example, the percentage of overlapping slots between two alternatives. This approach is complex and computationally intensive. In fact, the work of this dissertation is not concerned with how exactly alternatives are correlated with each other. And the main purpose of developing the models is not to forecast. In other words, the covariance specification and estimation approach is favored for forecast models, in which the correlation patterns needs to be captured explicitly.

This study employs a third approach, by specifying a mixed logit model to account for the correlations among alternatives. A mixed logit model allows for a flexible error covariance structure. Various correlation patterns can be obtained by choosing appropriate variables to specify the error components. This approach avoids defining the commonality factor or the correlation factor, which are necessary in the previous two approaches.

5.5 Data

All SCS messages sent in 2006 are studied for modeling purposes. However, flight information contained in the messages is very limited, including only Flight IDs and partial schedules after SCS. Thus, flights in SCS messages are first merged with the CDM archival database to obtain a complete view of flight schedules just before the SCS messages are sent, along with the schedules immediately after. This can be done because the SCS message time is recorded, and in the CDM archival database, each flight's schedules are published every five minutes or whenever there is any change--both schedule updates and the update time are reported.

As also done in the Chapter 4 cancellation model, the number of seats and average fare are developed for flights in the SCS messages. Only domestic flights are included in the model. SCS packets with more than five flights are excluded to facilitate choice sets generation. Finally, packets with incomplete information are excluded. As a result, 6956 flights from 2542 SCS packets are left for the model estimation. Table 5.4 shows the distribution of packet sizes in the sample data.

Table 5.4 Distribution of Packet Sizes

Number of Flights in a Packet	Flight Count	Packet Count
2	2898	1449
3	1584	528
4	1404	351
5	1070	214
Total	6956	2542

The "after SCS schedules" and the *EETAs* of the sample packets are input to Matlab to generate all feasible alternative schedules. The upper bound of the number of alternatives for each packet is given by Equation (5.2):

$$P(n,n) + C(n,1) * P(n,n-1) + C(n,2) * P(n,n-2) + C(n,3) * P(n,n-3) + C(n,4) * P(n,n-4) \tag{5.2}$$

Where n is the number of flights in a packet. n is less or equal to 5.

Considering the constraints of *EETA*, the actual numbers of alternatives are much fewer than given by Equation (5.2). Table 5.5 summarizes the actual number of alternatives and corresponding packet frequency. The maximum number of alternatives is 1545 and 57% of packets have 6 alternatives or less. The cumulative distribution of packets over the number of alternatives is shown in Figure 5-2.

Table 5.5 Number of Alternatives and Corresponding Packet Frequency

Number of Alternatives	Packet Frequency	Number of Alternatives	Packet Frequency	Number of Alternatives	Packet Frequency
2	1	106	17	495	6
3	1	113	24	500	1
4	407	140	44	522	6
6	1040	174	33	602	7
14	59	202	5	636	13
17	2	208	160	709	5
19	68	229	1	777	6
26	106	254	4	918	11
33	293	321	9	952	2
51	10	324	1	1127	16
65	1	388	7	1336	10
66	27	408	3	1545	100
86	35	419	1		

Figure 5.2 Cumulative Distributions of Packets over the Numbers of Alternatives

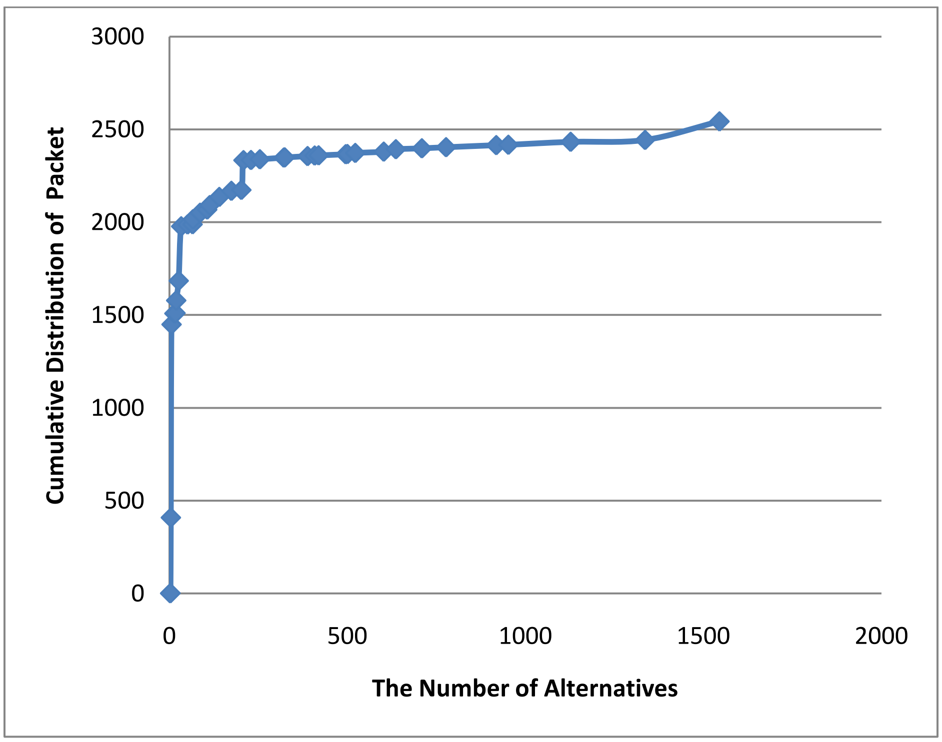

5.6 Explanatory Variables

Each SCS packet is considered an observation. The schedule observed after SCS is the chosen alternative while the rest of generated arrangements are non-chosen alternatives. In regards to the explanatory variables, some are associated with individual flights in a packet, while others are packet-attributes and summarized over flights.

The first 10 explanatory variables are at the flight level. Recall that a packet consists less than or equal to five flights in the sample data. The first set of variables X_1 to X_5 are defined as the *delay reduction* (in minutes) of Flight 1 to 5 in a packet. The delay reduction for non-cancelled flights is calculated using the flights' "before SCS schedule" minus the generated schedule, which is based on "after SCS schedule". For cancelled flights, delay reduction is set to be zero due to the fact that the slot assigned to the cancelled flight is "conceptual"— to hold the cancelled slot for the airline. Thus the delay reduction calculated based on the "conceptual assignment" is inappropriate. Also note that for a flight with a delay increase of 20 minutes, the delay reduction is counted to be -20 minutes. The second set of variables Y_1 to Y_5 are *cancellation dummy* variables for Flight 1 to 5. It is set to be one when a flight is cancelled and zero otherwise. For packets with less than five flights, zeros are also used as placeholders for the extra *X* and *Y* variables. For example, if a packet only has four flights, one set of *Xs* and *Ys* are zero.

In addition, a flight is not explicitly associated with its flight number. A flight can be called Flight 1, or Flight 5. The flight number is used to distinguish among flights in a packet, with no other meaning. However, within the same packet, each flight should be assigned a unique flight number. Furthermore, it is reasonable to believe that the coefficient of a cancellation or a delay reduction is the same, regardless of whether it is for Flight 1 or Flight 5. The coefficients of *delay reduction* are expected to be positive, while the coefficients of *cancellation dummy* are expected to be negative.

The main purpose of generating flight level variables (*Xs* and *Ys*) is to capture the correlations among alternatives by specifying these variables in the error term. The ten variables can represent how alternatives are similar to one another. To be more specific, if two alternatives share a slot, then the corresponding *X* (*delay reduction*) from the two alternatives are the same. Similarly, if the same flight is cancelled in two alternatives, the corresponding *Y* (*cancellation dummy*) is the same. By including these ten variables into the error component, it can represent covariance between any two alternatives.

Delay-reduced seat (in seat-minutes) and *Delay-reduced fare* (in $-minutes) are obtained by multiplying the *delay reduction* by the aircraft seat capacity and average fare respectively. Similarly, *Cancelled seat* (in seats) and *Cancelled fare* (in $s) are generated by multiplying the *cancellation dummy* by the aircraft seat capacity and average fare respectively. Then summarizing these parameters over a packet, another four explanatory variables can be developed as packet attributes. They are *Pkt_reduced_seat, Pkt_reduced_fare, Pkt_cnx_seat, Pkt_cnx_fare* respectively. It is hypothesized that airlines prefer to reduce delays for flights that are operated with large size aircraft and with high fare passengers on board. Conversely, when cancelling flights, airlines are more likely to cancel flights that are operated with small size aircraft and low fare passengers. Therefore, the coefficients of *Pkt_reduced_seat* and *Pkt_reduced_fare* are expected to be positive, while the coefficients of *Pkt_cnx_seat* and *Pkt_cnx_fare* are expected to be negative.

5.7 Model Specification and Results

5.7.1 Standard Logit Model

The deterministic part of the airlines' utility is modeled as a linear function in parameters as follows

$$\sum_{i=1}^{5}(\alpha_i * X_i) + \sum_{j=1}^{5}(\gamma_j * Y_j) + \beta_1 * Pkt\_reduced\_seat + \beta_2 * Pkt\_reduced\_fare + \beta_3 * Pkt\_cnx\_seat + \beta_4 * Pkt\_cnx\_fare \tag{5.3}$$

Where,

X_i	is delay reduction (in minutes) on Flight *i*
Y_j	is the cancellation dummy for Flight *j*
Pkt_reduced_seat	is delay reduced seats (in seat-minutes) over a packet
Pkt_reduced_fare	is delay reduced fare (in $-minutes) over a packet
Pkt_cnx_seat	is cancelled seats (in seats) over a packet
Pkt_cnx_fare	is cancelled fare (in $s) over a packet
$\alpha_i, \gamma_j, \beta_1 \ldots \beta_4$	are the coefficients to be estimated

A multinomial standard logit model is estimated with Procedure MDC in SAS 9.2. The results are shown in Table 5.6.

Table 5.6 Estimation Results of the Standard Logit Model

Variables	Coefficient	Std Error
X_1	**0.01180**	0.00285
X_2	**0.01370**	0.00286
X_3	**0.01260**	0.00286
X_4	**0.01200**	0.00288
X_5	**0.01220**	0.00284
Y_1	**-0.55760**	0.18700
Y_2	**-0.49650**	0.18730
Y_3	**-0.43780**	0.18280
Y_4	**-0.43640**	0.18490
Y_5	**-0.67970**	0.18960
Pkt_reduced_seat	**0.00009**	0.00002
Pkt_reduced_fare	0.00001	0.00002
Pkt_cnx_seat	**-0.02130**	0.00141
Pkt_cnx_fare	**-0.00535**	0.00113
Log-likelihood	-4398	

Note: Bolded coefficients are statistically significant at the 0.05 level.

The coefficients on *delay reduction* (X_1 to X_5) are all positive and statistically significant, which confirms the hypotheses that airlines are favor of reducing delays. The coefficients on *cancellation dummy* (Y_1 to Y_5*)* are all negative and statistically significant, which also confirms that cancellations decrease airlines' utility. Moreover, the magnitudes of the *X* coefficients are very similar. So are those of the *Y* coefficients. Since the utility of an individual flight should be the same regardless of its Flight Number, the mean coefficients values of the *Xs* and *Ys* are calculated to represent the impacts of delay and cancellation, respectively, to all flights.

$$\text{Coefficient of } \bar{X} = \sum_{i=1}^{5} X_i \,/5 = 0.01260 \tag{5.4}$$

$$\text{Coefficient of } \bar{Y} = \sum_{j=1}^{5} Y_j \,/5 = -0.52900 \tag{5.5}$$

The coefficient of *Pkt_reduced_seat* is statistically significant and plays a positive role in forming the packet arrangement favored by the airlines. This confirms the hypothesis that airlines prefer to reduce delay on large aircraft. However the coefficient of *Pkt_reduced_fare* is positive but not statistically significant. The coefficients of *Pkt_cnx_sea* and *Pkt_cnx_fare* are both negative and statistically significant, as expected. Airlines prefer cancelling flights operated with small aircraft and with low-fare passengers onboard.

Given the mean of aircraft size (108 seats) and average fare paid by onboard passengers ($118) from the sample flights, the "willingness to cancel" in terms of delay reduction can be calculated based on the estimated coefficients shown in Table 5.6.

$$Willingness\ to\ Cancel = -\frac{\sum Utilities\ of\ Cancellations}{\sum Utilities\ of\ Delay\ Reductions}$$

$$= -\frac{(-0.52900) + (-0.02130) * 108 + (-0.00535) * 118}{0.01260 + 0.00009 * 108 + 0.00001 * 118}$$

$$=147.64 \tag{5.6}$$

The standard logit model estimates that airlines are willing to cancel a flight if the cancellation can reduce 147.64 minutes of delay for other flights. The delay reduction on the cancelled flight itself is not counted. However, in the real operations, airlines often cancel flights with lower than average seats and fares, which will decrease the numerator, and reduce delay on flights with higher than average seats and fares, which will increase the denominator. Therefore, the willingness to cancel may be lower than 147.64 minutes in the real operations.

5.7.2 Mixed Logit Model

The same specification for the deterministic part of utility is used for the mixed logit model as the standard logit model. Ten variables are selected to enter the stochastic part of utility, or the error component, to account for correlations among alternatives.

The ten error components are all specified at the flight level. The first error component is generated by a normal deviate multiplied by the *delay reduction* on Flight 1 (X_1). The second, third, fourth and fifth error components are generated respectively to X_2, X_3, X_4, and X_5. The sixth error component is generated by a normal deviate multiplied by the *cancellation dummy* on Flight 1 (Y_1) and the rest are based on Y_2, Y_3, Y_4 and Y_5.

Therefore the stochastic part of utility for alternative *i* becomes

$$\sum_{m=1}^{5} \propto_m X_{im} Z_{im} + \sum_{n=1}^{5} \delta_n Y_{in} Z_{in} + \varepsilon_i \tag{5.7}$$

where

$\propto_m, \delta_n$	The coefficients to be estimated
Z_{im}, Z_{in}	Normal deviate
X_{ik}	Variables related to flight reductions
Y_{in}	Variables related to flight cancellations
ε_i	Distributed iid, extreme value

The mixed logit model is also estimated using the MDC procedure in SAS 9.2. Table 5.7 shows the estimated parameters and standard errors for the mixed logit model, along with the results from standard logit model discussed in the previous section.

Table 5.7 Estimation Results of the Mixed Logit Model

	Mixed Logit		Standard Logit	
	Coefficient	Std Error	Coefficient	Std Error
Variables:				
X_1	**0.031400**	0.004030	**0.011800**	0.002852
X_2	**0.032900**	0.003932	**0.013700**	0.002862
X_3	**0.034000**	0.003967	**0.012600**	0.002858
X_4	**0.034500**	0.003836	**0.012000**	0.002883
X_5	**0.033500**	0.003792	**0.012200**	0.002841
Y_1	-0.781000	0.400100	**-0.557600**	0.187000
Y_2	**-1.525000**	0.523900	**-0.496500**	0.187300
Y_3	**-0.973200**	0.432600	**-0.437800**	0.182800
Y_4	-0.421900	0.297800	**-0.436400**	0.184900
Y_5	**-1.536100**	0.593400	**-0.679700**	0.189600
Pkt_reduced_seat	**0.000112**	0.000017	**0.000093**	0.000017
Pkt_reduced_fare	-0.000002	0.000025	0.000005	0.000018
Pkt_cnx_seat	**-0.025100**	0.001529	**-0.021300**	0.001410
Pkt_cnx_fare	**-0.006525**	0.001224	**-0.005354**	0.001133
Error Components:				
X_1	**0.026700**	0.002563		
X_2	**0.019000**	0.002917		
X_3	**0.023500**	0.002762		
X_4	**0.022300**	0.002425		
X_5	**0.024300**	0.002661		
Y_1	0.907700	0.609500		
Y_2	**1.822300**	0.467900		
Y_3	**1.333400**	0.472300		
Y_4	0.437200	0.794700		
Y_5	**1.683600**	0.560200		
Log-likelihood	-4173		-4398	

Note: Bolded coefficients are statistically significant at the 0.05 level.

Among the ten error components, eight of them are statistically significant, which confirms that correlations exist across the alternatives.

It is found that the coefficients of the deterministic part of utility in the mixed logit model are generally larger than those of the standard logit model. This is the same result as obtained by Brownstone and Train [33]. The reason is that the variance of the error term in the standard logit model is larger than that of the mixed logit model. This in turn is due to the fact that in the mixed logit model, some of the variance in the stochastic part of utility is captured by the error components. As a result, when utilities are scaled to have the same variance for the stochastic part of utility, the deterministic part of the standard logit model is scaled down.

Some previous studies have found that although coefficients in the standard logit and mixed logit models are different, the ratios of some coefficients (such as willingness to pay, cost of time etc.) are very similar [38] [39]. Other studies have found a substantial difference in the ratios [40]. The economically meaningful ratio in our study is willingness to cancel for delay reduction to other flights, which is 106 minutes according to the mixed logit model and 148 minutes from the standard logit model.

According to the estimation results obtained from the mixed logit model, the willingness to cancel can be calculated based on several other combinations of aircraft size and average fare, as shown in Figure 5-3. For example, when cancelling a flight with 150 seats and average fare $150 (as shown in the box), the delay savings to other flights should be at least 116 minutes. In general, a flight cancellation should be able to reduce 68 to 127 minutes of delay, depending on the flight characteristics (seat capacity and average fare). Again, this does not include any delay reductions on the cancelled flight itself.

Figure 5.3 Willingness to Cancel

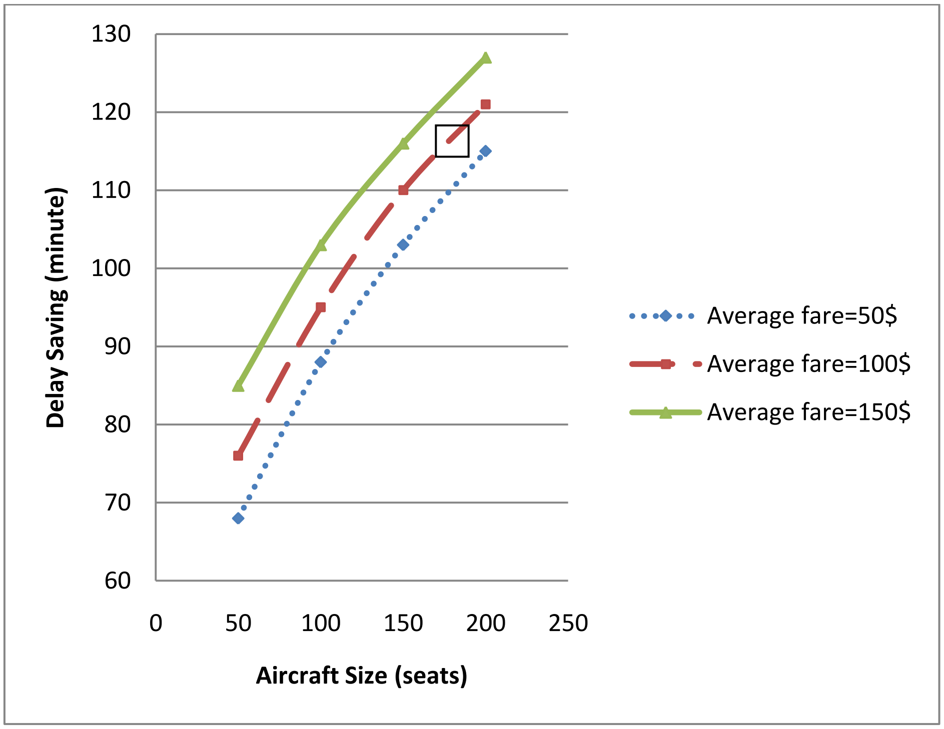

5.8 Summary

The airlines' preference structure is estimated through their SCS behavior. An airline's proposed flight arrangement in the SCS message is considered to be preferred to the other feasible arrangements based on the same set of flight and arrival slots. A multinomial logit model is adopted to estimate the airlines' utilities in making the preferred arrangements. The proposed flight arrangement in the SCS message is the chosen alternative among all feasible arrangements that are generated by the author.

Since the entire sets of alternatives are developed based on the same set of flight and arrival slots, correlation among alternatives exists. The problem is addressed by using a mixed logit model specification. The mixed logit model allows for a flexible error structure and accounts for correlation among alternatives by choosing appropriate variables to enter the error components.

The estimation results of the standard logit model reveal that airlines are generally in favor of reducing delays, particularly on flights operated with large aircraft. On the other hand, airlines utility decreases with flight cancellations. Cancellations are most often made on flights served by small aircraft and with low fare passengers. These are all confirmed hypotheses.

The mixed logit model adopts the same specification of the deterministic part of the utility as the standard logit model. Ten error components are added into the stochastic part of the utility to account for correlations among alternatives. The ten variables are carefully selected and constructed to introduce appropriate correlation patterns. The results show that eight out of ten variables in the error components are statistically significant, and therefore covariance across alternatives does exist.

Finally, the willingness to cancel in terms of delay reduction to other flights, not including delay reduction on the cancelled flight itself, is one of the key findings of this chapter. It is found to be 106 minutes according to the mixed logit model and 148 minutes from the standard logit model, given a flight with the average seat capacity (108) and fare ($118) from the sample.

Chapter 6 Conclusions

This final chapter summarizes the findings from this dissertation and closes with a discussion of the future work.

6.1 Summary of Findings

The airlines' preference structure is investigated through their collaborative responses to the FAA's Traffic Management Initiatives (TMIs), and in particular, the Ground Delay Program (GDP). By observing the actual choices made by airline dispatchers when presented with alternatives, the airlines' utility functions can be inferred through the use of discrete choice models.

However, airlines face a countless number of choice situations during a GDP. This research identifies two scenarios in which a reasonable number of choice situations can be created using data available in the CDM archival database. In one scenario, the airlines' cancellation utility is estimated from their flight cancellation decisions using a binary choice model. Chapter 4 discusses such cancellation models. The other scenario investigates airlines' Slot Credit Substitution (SCS) messages, which consider both cancellations and flight reassignments simultaneously. The submitted flight arrangement in the SCS messages is considered superior to other feasible arrangements for the same set of flights, allowing the utility of the arrangements identified in SCS messages to be estimated using multinomial choice models. The findings of the SCS models are presented in Chapter 5.

The cancellation model captures how airlines value delays on flight itself and potential delay savings on other flights in making a flight cancellation decision. Aircraft size, along with segment frequency and load factor, are all significant factors in cancellation decisions; larger, fuller, and less frequent flights are less likely to be cancelled. Somewhat surprisingly, a higher average fare is found to increase cancellation probability. Hub-bound flights are found more likely to be cancelled than spoke-bound flights. The model also confirms airlines' hedging behavior by preferentially cancelling short-haul flights. In addition, a piece wise linear specification of the utility function confirms that the delay impact is non-linear. A random coefficient model finds that the flight-to-flight variations in utility are not great.

The SCS model captures airlines' tradeoff behavior in dealing with flight cancellations and delays. It confirms that cancelling flights decreases airlines' utility while reducing delays increases the utility. Moreover, airlines are sensitive to the aircraft size and average fare of flights in performing these actions.

The cancellation and delay tradeoff is studied in both models. The cancellation model shows that a cancellation can remove the delay imposed on the flight itself along with the delay savings on other flights in the airline's fleet. However, the model results also suggest that for the vast majority of flights, delay considerations alone do not strongly incline airlines toward cancelling flights. Delay savings and other observable factors affect the likelihood that a flight will be cancelled, just as a variety of factors influence whether a certain car trip results in an accident. In both cases, however, the level of determinism is low, so that the actual occurrences are still unlikely. The SCS model estimates that airlines are willing to cancel a flight if the cancellation can reduce 106 minutes of delays on other flights in their fleet. However, SCS messages are self-selected samples. In other words, the cancellations proposed in the SCS messages are opportunities identified by the airlines and are expected to have greater value from delay reduction than other cancellations.

The differences in the two models, particularly with regard to the tradeoff between cancellation and delay, point to the continued importance of inertia in shaping airline responses to GDPs. In the cancellation model, the choice of whether or not to cancel a flight is also a choice of whether to "stand pat" or undertake a series of actions that involves reassigning slots and rebooking passengers. The results suggest that the predisposition here is toward inaction. In SCS, change is a given, and the question is which of a set of alternative changes is the most desirable. The absence of an inherent "status quo" in the context of an SCS allows a clearer view of an airline's preference structure when its behavior is not overwhelmed by inertia.

An intriguing question is whether the inertia observed in the cancellation model is a reflection of the true interests of the airline, or the more narrow ones of the individuals making decisions on its behalf. Is the workload such that dispatchers must carefully "pick their spots" for making cancellations and substitutions? Would greater pro-action result in a better outcome for the airline? Does this suggest that workload reducing technologies (or employing more dispatchers) would allow airlines to take better advantage of CDM? These are all questions that are raised by the differing results of the two models.

In summary, the three main contributions of this dissertation research are as follows. The first contribution is the identification of opportunities for an airline behavioral study; the second contribution is the estimation of airlines' preference structures through two choices scenarios; the third contribution involves the analysis of the tradeoff between flight cancellations and flight delays, and quantifying the cost of flight cancellations in terms of delay.

6.2 Future Work

There are two directions in which this dissertation research can be continued. The first is to refine the current models and validate the model results. The other is to use the findings to address other aviation issues.

6.2.1 Model Estimation and Validation

Some factors might affect airlines decisions are not included in the model, because these factors were not revealed in published databases. For example, flight crew assignment, aircraft assignment and stand-by resources are treated as exogenous while in fact, such factors might be endogenous. Although random utility theory does account for unobservable factors in the error component, the model could be improved by obtaining relevant data from airlines and incorporating it as additional variables in the deterministic portion of the utility.

In regards to the cancellation model, for a given airline and GDP, all impacted flights are managed by one dispatcher or sometimes several dispatchers in a collaborative manner. In such cases, it is more reasonable to consider that the cancellation decisions for flights are not independent from one another. One expects that cancelling Flight A would decrease the probability of cancelling Flight B, particularly when two flights are serving the same Origin-Destination pair, because Flight B is expected to be used to re-book the passengers from Flight A, if cancelled. The mixed logit model specification can be used to address this issue.

In the case of the SCS model, it would be interesting to apply other modeling approaches, such as the ordered responses model.

This dissertation research focuses on using revealed preferences to estimate airlines' utility, by modeling actual choices made by airline dispatchers' in real-time GDPs. The use of stated preference is an alternative approach. A survey designed with hypothetical questions could be distributed to airline dispatchers. The answers a dispatcher provides would be their stated preference. For example, in the survey, a dispatcher may be presented with a question like: if a flight is experiencing 120 minutes of delays in a GDP, will you cancel it? Hypothetical questions can be designed to accommodate any hypothesis a researcher would like to test. Based on the survey results, a stated preference model could be estimated. The stated preference model could then be used to validate the revealed preference model. In addition, a combined stated and revealed preference model could be estimated to infer airlines' preference structure.

6.2.2 Applications

The estimated airlines' utility function can be used in various situations. Firstly, the findings about flight cancellation cost and heterogeneous delay cost can be used in developing enhanced metrics to evaluate aviation system performance.

Secondly, the airlines' utility function could be used to quantify the benefits of CDM. Before CDM was implemented, airlines lost ownership of slots that held their cancelled flights. By measuring the change in consumer surplus from before and after CDM implementation, the value of CDM can be quantified.

Thirdly, the cancellation model can be used as inputs to many other projects. For example, the model can be integrated into the FAA's simulation tool--National Airspace

System Performance Analysis Capability (NASPAC). NASPAC refers to "an integrated set of computer program modules designed to model the entire National Airspace System, the en-route structure, and traffic flows as a network of inter-related components, reflecting the effects of weather conditions, air-traffic control procedures, and air-carrier operating practices" [41]. Flight cancellations are not considered in the current NASPAC. As a result, the cancellation model developed in this dissertation can be adopted as a component of NASPAC, and can be used to predict how many flights would be cancelled under different airport capacity/weather scenarios. Similarly, the cancellation model can be used as an input to the Future ATM Concepts Evaluation Tool (FACET), a real-time planning applications developed by the NASA research group.

Finally, individual airlines' utility can be further explored, and the results could be used to build a decision support tool to facilitate the airlines' disruption management.

Reference

1. FAA Aerospace Forecasts FY 2008-2025. Available at:
http://www.faa.gov/data_research/aviation/aerospace_forecasts/2008-2025/

2. U.S. Senate Joint Economic Committee Majority Staff, "Your Flight has Been Delayed Again: Flight Delays Cost Passengers, Airlines, and the US Economy Billions", 2008. Available at:
http://jec.senate.gov/index.cfm?FuseAction=Reports.Reports&ContentRecord_id=11116dd7-973c-61e2-4874-a6a18790a81b&Region_id=&Issue_id=

3. FAA CDM Workgroup, Collaborative Decision Making (CDM) 2005 Structure Report, May 20, 2005

4. Clarke, M., "Irregular Airline Operations: A Review of the State-of-the-practice in Airline Operations Control Centers", Journal of Air Transport Management, Volume 4, Issue 2, April 1998, Pages 67-76

5. Ball, M. *et al.*, Air transportation: Irregular Operations and Control, Handbook in OR&MS, Vol. 14. Elsevier, Amsterdam

6. Air Transport Association, "Cost of Delays: Direct Cost", 2008. Available at: http://www.airlines.org/Economics/DataAnalysis/Pages/CostofDelays.aspx

7. m2p Consulting, "Optimized Flight Operations Delay Management", AGIFORS Operations Conference Denver, CO, May 20th –24th, 2007

8. Willemain, T., Contingencies and Cancellations in Ground Delay Programs, NEXTOR Report WP-01-3, May 30, 2001

9. Mukherjee, A., M. Hansen and S. Grabbe, "Ground Delay Program Planning Under Uncertainty in Airport Capacity", AIAA Guidance, Navigation, and Control Conference 10 - 13 August 2009, Chicago, Illinois

10. Ball, M. *et al.*, Total Delay Impact Study: A Comprehensive Assessment of the Costs and Impacts of Flight Delay in the United States, NEXTOR Report, May 2010. (not published yet)

11. Cook, A. *et al.*, Evaluating the True Cost to Airlines of One Minute of Airborne or Ground delay, Final Report. Eurocontrol, Brussels, Belgium, 2004

12. m2p Consulting, "Optimized Flight Operations Delay Management", Presented at AGIFORS Operation conference, 2007

13. Roger, B., R. Hsu, L. Berry and J. Rome, “Preliminary Evaluation of Flight Delay Propagation through an Airline Schedule,” Air Traffic Control Quarterly, Volume 7, November 1999

14. Air Transport Association, “Cost of Delays: Additional Cost”, 2008. Available at: http://www.airlines.org/Economics/DataAnalysis/Pages/CostofDelays.aspx

15. GRA, Inc., Economic values for FAA investment and regulatory decisions-a guide, Report prepared for FAA Office of Aviation Policy and Plans, Washington D.C.

16. Morrison, S. and C. Winston, “Another look at airport congestion pricing”, American Economic Review, Vol. 97, No. 5, pp. 1970-1977

17. Forbes, S., “The Effect of Air Traffic Delays on Airline Prices”, International journal of industrial organization, 26, 1218-1232

18. Alder, T., C. Falzarano, S. Spitz, “Modeling Service Trade-offs in Air Itinerary Choices”, Paper presented in the 84th Annual Meeting of Transportation Research Board, 2005

19. Cynthia, B. and S. Bratu, “Airline Passenger Delays”, Sloan Industry Centers Annual Conference, Atlanta, GA, 2004

20. Bratu, S. and B. Cynthia, “An Analysis of Passenger Delays Using Flight Operations and Passenger Booking Data”, Air Traffic Control Quarterly, Vol. 13 (1), pp. 1-27, 2005

21. Wang, D. *et al.*, “Statistical Comparison of Passenger Trip Delay and Flight Delay Metrics”, Transportation Research Record: Journal of the Transportation Research Board Issue Number: 2052, 2009

22. Sherry, L. and G. Donohue, “U.S. Airline Passenger Trip Delay Report”, Center for Air Transportation Systems Research, George Mason University, April 2009. Available at: http://catsr.ite.gmu.edu

23. Zhu, Y., Evaluating Airline Delays: The Role of Airline Networks, Schedules, and Passenger Demands, MIT Master of Science in Transportation Thesis, 2009

24. Hansen, M., D. Gillen and Djafarian-Tehrani, “Aviation Infrastructure Performance and Airline Operation Cost: A Statistical Cost Estimation Approach”, Transportation Research E, 37, 2001, pp. 1-23

25. Irrgang, M., D. Jenkins, and C. Ray, Airline Irregular Operations, Handbook of Airline Economics (1st edition), McGrawHill Aviation Week Group: New York, pp. 349-365, 1995

26. Sridar, B. “Relationship between Weather, Traffic and Delay Based on Empirical Methods”, Presented at NEXTOR Asilamor annual conference, 2007

27. Metron Project Report (not published)

28. American Airlines, “The Cost of Flight Diversion and Cancellation”, AGIFORS Meeting, Colorado, 2007

29. Rupp, N., “Flight Delays and Cancellations”, East Carolina University, JEL Classification: L13, L93, 2005

30. Rupp, N. and G. M. Holmes, “An Investigation into the Determinants of Flight Cancellations”, Economica, Vol. 73, pp. 749-783, 2006

31. Train, K., Discrete Choice Methods with Simulation, Cambridge University Press, 2003

32. Revelt, D. and K. Train, “Mixed Logit with Repeated Choices: Households’ Choices of Appliance Efficiency Level”, Review of Economics and Statistics, No. 4, pp. 647-657, 1998

33. Brownstone, D. and K. Train, “Forecasting New Product Penetration with Flexible Substitution Patterns”, Journal of Econometrics, Vol. 89, pp. 109-129, 1998

34. Ball, M., G. Dahl, and T. Vossen, Matchings in Connection with Ground Delay Program Planning, Robert H. Smith School Research Paper No. RHS 06-040, 2006. Available at: http://ssrn.com/abstract=939567

35. Eisenhardt, K., “Agency theory: An Assessment and Review”, Academy of Management Review, 14 (1): 57-74, 1989

36. Ben-Akiva, M. and M. Bierlaire, Discrete Choice Models with Applications to Departure Time and Route Choice, Book Series International Series in Operations Research & Management Science, Vol 56, 2006

37. Ramming M., Network Knowledge and Route Choice, Phd Dissertation, Massachusetts Institute of Technology, February 2002

38. Bhat, C., “Accommodating Variations in Responsiveness to Level-of-Service Measures in Travel Model Choice Modeling”, working paper, Department of Civil Engineering, University of Massachusetts at Amherst, 1996

39. Train, K., “Unobserved Taste Variation in Recreation Demand Models”, working paper, Department of Economics, University of California, Berkeley, 1996

40. Bhat, C., "Incorporating Observed and Unobserved Heterogeneity in Urban Work Travel Mode Choice Modeling", working paper, Department of Civil Engineering, University of Massachusetts at Amherst, 1996

41. NASPAC webpage. Available at: http://www.tc.faa.gov/acb300/330_naspac.asp

42. BTS webpage. Available at http://www.bts.gov/programs/statistical_policy_and_research/source_and_accuracy_compendium/form41_schedule.html

43. Data Base Product. Available at http://www.airlinedata.com/

44. BTS webpage. Available at http://www.transtats.bts.gov/Tables.asp?DB_ID=125&DB_Name=Airline%20Origin%20and%20Destination%20Survey%20%28DB1B%29

Appendix 1 Aircraft Size and Load Factor Development

Aircraft size and load factor are developed from T100 databases, which contain data collected by the Bureau of Transportation Statistics (BTS). Domestic air carriers are required to submit information about their operations on a monthly basis by completing Form 41, which BTS inputs into the databases. The database reports segment data on a monthly basis, including total number of departures, origin and destination (OD), aircraft type, operating carrier, total number of onboard passengers, total number of available seats, and average load factor [42]. Data Base Product [43] cleans the data in T100 database and makes corrections to make the T100 data more accurate. The T100 data used for this analysis are provided by the Data Base Product.

T100 data is monthly based, and thus does not include data for individual flights. Monthly average load factor and average number of available seats per flight, developed from the T100 database, are merged with the GDP flights by month, OD, aircraft type, and operating carrier.

Appendix 2 Average Fare Development

The average fare of a GDP impacted flight is developed from the Airline Origin and Destination Survey (DB1B) database [44]. Airlines are required to report 10% of their air passengers' tickets to the BTS and these tickets are stored in the DB1B database. Data Base Product cleans the data in DB1B database and makes corrections to make the DB1B data more accurate. The data used for this analysis are provided by the Data Base Product under "Hub" database.

"Hub" data is itinerary based. A record in the Hub database, which is a one-way trip, consists of year and quarter, operating carrier, OD airports, number of stops and the airports where stops are made, distances, along with total number of passengers travelled on this itinerary in the quarter and their average fare. Table A-1 provides an example of a record in the "Hub" database.

This record shows that in the first quarter of 2006, 26 passengers, whose tickets were stored in DB1B database, travelled from LAX to SEA and made two stops along the way: one at LAS and another at PDX. Thus, there are three segments of flight. The first segment is operated by Alaska Airlines from LAX to LAS, and the distance traveled is 236 miles. The second segment is operated by Alaska Airlines from LAS to PDX, and the distance traveled is 762 miles. The last segment is operated by Alaska Airlines from PDX to SEA, and the distance traveled is 129 miles. The average fare of these 26 passengers is $184.13. For the purpose of this analysis, the itinerary based fare is decomposed to the segment fare, which is based on segment distance Table A-2 shows the decomposition of the itinerary based record shown in Table A-1, into segment data.

To estimate an average fare of one particular segment from the "Hub" database, all records in 2006 are downloaded and transformed from itinerary-based data to the segment-based data, using the method illustrated in Table A-1 and Table A-2. Zero-fare passengers are excluded and so are international passengers. For segments with the same quarter, OD, and operating carrier, the average fare weighted by the number of passengers is calculated. This average fare is then merged with GDP flights by quarter, OD, and operating carrier to estimate the average fare of a GDP flight.

Table A-1 A Sample Record in the DB1B Hub Database

Quarter	Origin	Destination	Stop 1	Stop 2	Carrier 1	Carrier 2	Carrier 3	Distance 1	Distance 2	Distance 3	Passenger counts	Fare
200601	LAX	SEA	LAS	PDX	AS	AS	AS	236	762	129	26	184.13

*SEA: Seattle Tacoma International Airport
*LAX: Los Angeles International Airport
*LAS: Las Vegas McCarran International Airport
*PDX: Portland International Airport
*AS: Alaska Airlines

Table A-2 Decomposition of a Record by Segment

Quarter	Origin	Destination	Carrier	Distance	Passenger counts	Segment Fare
200601	LAX	LAS	AS	236	26	236/(236+762+129)*184.13=38.56
200601	LAS	PDX	AS	762	26	762/(236+762+129)*184.13=124.50
200601	PDX	SEA	AS	129	26	129/(236+762+129)*184.13=21.08

Appendix 3 Market Fare Development

The market fare is developed to represent the market type between origin and destination of GDP impacted flights. It is also developed from the "Hub" database which is described in the Appendix 2. However, the origin and destination of a flight in GDP are now considered the two markets, which means multiple routes and airlines serve these two markets. Figure A-1 shows an example of the market pair, SFO and JFK, which can be served with non-stop flights, one stop flights (over ORD), or two stop flights (over DEN and ORD). On each route, multiple airlines can compete. The market fare is calculated using the average fare from all possible routes and airlines.

Figure A-1 Example of SFO-JFK Market Served by Multiple Routes

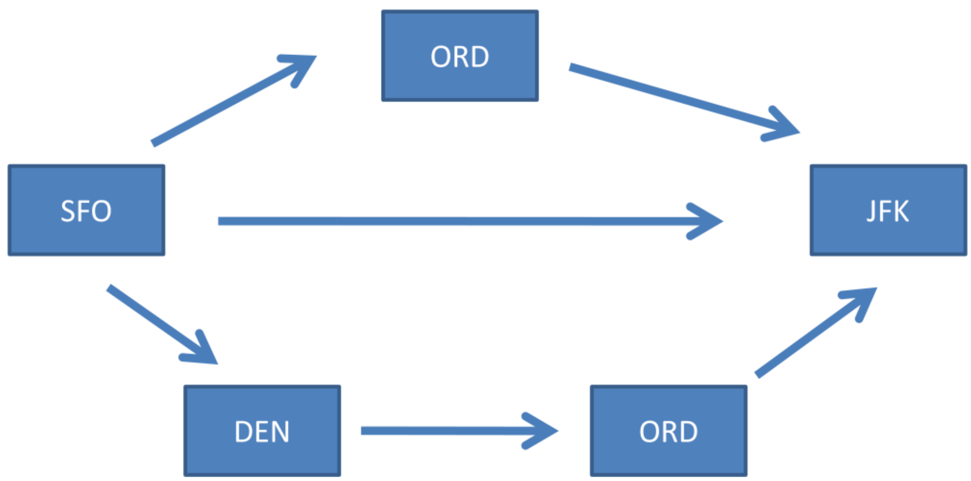

*ORD: Chicago O'Hare International Airport
*DEN: Denver International Airport
*JFK: John F. Kennedy International Airport
*SFO: San Francisco International Airport

Therefore, the itinerary based "Hub" data does not need to be decomposed. For all the records in the "Hub" database with the same OD airports and quarter, the weighted average fare by the number of passengers is calculated. Finally, this weighted average fare is merged with GDP flights by quarter and OD and to obtain the average market fare.