


Tarmac Delay Policies: A Passenger-Centric Analysis

by

Allison “Sunny” Elizabeth Vanderboll

B.S.C.E., Stanford University (2011)

ARCHIVES



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

Master of Science in Transportation

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2013

© Massachusetts Institute of Technology 2013. All rights reserved.

Author
Department of Civil and Environmental Engineering
May 6, 2013

Certified by
Cynthia Barnhart
Ford Professor of Civil & Environmental Engineering
Associate Dean, School of Engineering
Thesis Supervisor

Accepted by
Heidi M. Nepf
Chair, Departmental Committee for Graduate Students

Tarmac Delay Policies: A Passenger-Centric Analysis

by

Allison “Sunny” Elizabeth Vanderboll

Submitted to the Department of Civil and Environmental Engineering
on May 6, 2013, in partial fulfillment of the
requirements for the degree of
Master of Science in Transportation

Abstract

In this work, we analyze the effectiveness of the 2010 Tarmac Delay Rule from a passenger-centric point of view. The Tarmac Delay Rule aims to protect enplaned passengers on commercial aircraft from excessively long delays upon taxi-out or taxi-in, and monetarily penalizes airlines that violate the stipulated three-hour time limit. Using an algorithm to calculate passenger delay, we quantify delays to passengers in 2007, before the Tarmac Delay Rule was enacted, and compare these delays to those estimated for hypothetical scenarios with the rule in effect for that same year. Our delay estimates are achieved using U.S. Department of Transportation data from 2007, and one quarter of booking data purchased from a large legacy carrier to validate our results. The results suggest that the rule has been a highly effective deterrent for airlines to keep tarmac times under three hours. This benefit is offset, however, because coincident with shortened tarmac delays are flight cancellations. Cancellations result in passengers requiring rebooking, and extensive delays. Through our analysis, we show that the overall impact of the Tarmac Delay Rule is a significant increase in passenger delays. We evaluate the impacts of variations to the rule, including changing the rule to apply to flights that are delayed for both less and more than the three hours stipulated in the rule, and identifying other variants of the rule that might better meet the objective of benefiting the flying public. Through extensive scenario analysis, we determine that the rule should be applied selectively, depending on flight departure times and specific network characteristics.

Thesis Supervisor: Cynthia Barnhart

Title: Ford Professor of Civil & Environmental Engineering

Associate Dean, School of Engineering

0.1 Acknowledgments

I would like to thank Professor Cynthia Barnhart for her guidance throughout my research. In addition, my deepest gratitude goes out to Dr. Vikrant Vaze for his support throughout the project, and to Professor Douglas Fearing for answering my endless questions on the Passenger Delay Calculator. I would like to extend my gratitude to Dan Murphy and the FAA for guidance, suggestions, and financial support. My thanks goes out to Professor Amedeo Odoni for his involvement with the project as well.

My two years at MIT would not have been the same without my friends in the Transportation program and the MIT graduate community. Thank you JP, Simmy, Yi-Hsin, Andrés, Fabián, Sumit, and so many more who made my MIT experience what it was. I would also like to acknowledge the many people I worked closely with on extra-curricular events, especially Gery and Wei on the Orientation Women's Welcome Lunch, and Alex on the Transportation Showcase. It was my pleasure to work with you all!

I am grateful for the support of my family, and of my friends, especially Philip, for always showing up.

Contents

0.1	Acknowledgments	5
1	Introduction	13
1.1	Lengthy tarmac delays in 2007	13
1.2	Contribution and Outline	18
2	Methodology and Data	21
2.1	Literature Review	22
2.2	Data inputs	25
3	Passenger delay stemming from the Tarmac Delay Rule	29
3.0.1	Notation	30
3.0.2	Initial assumptions	30
3.1	Passenger delays as a function of cancellation of flights in F_{AF}	30
3.2	Understanding the impacts of load factor on passenger delays	35
3.3	Effects to passengers not travelling on tarmac-delayed flights	38
4	Multipliers as a method to quantify impact to passengers on tarmac-delayed flights	43
4.1	Initial analysis and motivation for multipliers	43
4.2	Characteristics of flights with high and low multipliers	45
4.2.1	Characteristics of flights with multipliers less than one	47
4.2.2	Characteristics of flights with high multipliers	48
4.3	Validation of findings	48

4.4	Selective cancellations to minimize passenger delay on F_{AF}	50
5	Sensitivity Analysis: Implication of model assumptions on previous findings	53
5.1	Cancellation-time threshold	53
5.2	Passenger rebooking time	57
5.3	Cancelling return flight pairs	58
6	Conclusion and Future Directions	61
6.1	Review of methodologies and findings	61
6.2	Directions for future work	64
A	Tables	67
B	Figures	69

List of Figures

1-1	Schematic of airline decision-making when faced with a long taxi-out delay	15
1-2	2006-2010 non-cancelled flights taxiing-out more than three hours, and total scheduled operations	16
2-1	Passenger Delay Calculator algorithm schematic	23
2-2	Data inputs and outputs to Passenger Delay Calculator	26
3-1	P_{AF} delay for July 27 as a function of percentage of F_{AF} flights cancelled	31
3-2	P_{AF} delay for June 21 as a function of percentage of F_{AF} flights cancelled	32
3-3	P_{AF} delay for January 21 as a function of percentage of F_{AF} flights cancelled	34
3-4	P_{AF} delay for July 27, comparing actual and unlimited aircraft seating capacities	36
3-5	P_{AF} delay for June 21, comparing actual and unlimited aircraft seating capacities	37
3-6	P_{AF} delay for January 21, comparing actual and unlimited aircraft seating capacities	38
3-7	Schematic of algorithm used to move flights up into open slots created by cancellation of F_{AF}	40
4-1	Years of delay under different scenarios to passengers P_{AF} for all days in 2007 where there were nine or more flights in F_{AF}	44
4-2	One-flight multipliers for flights in set F_{AF} on July 27	45

4-3	One-flight multipliers for flights in set F_{AF} on June 26	46
4-4	One-flight multipliers for flights in set F_{AF} on February 14	46
4-5	One-flight multiplier for flights in F_{AF} , from airports and days in Section 3.3, ordered by push-back time	49
4-6	Selective cancellations for July 27	51
4-7	Selective cancellations for June 26	52
4-8	Selective cancellations for February 14	52
5-1	One-day multiplier for July 27 for all flights in F_{AF} , with varying cancellation-time thresholds	55
5-2	One-day multiplier for June 26 for all flights in F_{AF} , with varying cancellation-time thresholds	56
5-3	One-day multiplier for February 14 for all flights in F_{AF} , with varying cancellation-time thresholds	56
B-1	One-flight multipliers for flights in F_{AF} on January 17, departing DFW	69
B-2	One-flight multipliers for flights in F_{AF} on June 26, departing ORD .	70
B-3	One-flight multipliers for flights in F_{AF} on May 31, departing EWR .	70
B-4	One-flight multipliers for flights in F_{AF} on July 27, departing JFK . .	71
B-5	One-flight multipliers for flights in F_{AF} on July 10, departing ATL . .	71
B-6	One-flight multipliers for flights in F_{AF} on June 19, departing PHL .	72

List of Tables

1.1	Non-cancelled flights (<i>including</i> diversions) that experienced lengthy tarmac delays during taxi-out in 2007, as reported to BTS	14
1.2	U.S. G.A.O.-reported likelihood of flight cancellation change, by tarmac time	17
1.3	Lengthy taxi-out and taxi-in incidents, 2006-2010	18
3.1	Key findings of the impacts of cancellations	33
3.2	Days and locations selected for analysis of impact of F_{AF} flights on the wider NAS system	39
3.3	Key findings of earlier wheels-off time analysis	41
5.1	Comparing delay multipliers resulting from two-hour rebooking-delay	58
5.2	Percent increase in delay when return flight pairs are cancelled	60
A.1	IATA airport codes and names referenced in this work	67
A.2	2007 ASQP-reporting carriers, and associated codes	68
A.3	Information on days in 2007 selected for analysis	68

Chapter 1

Introduction

1.1 Lengthy tarmac delays in 2007

On February 14, 2007, in the midst of what came to be known as the “Valentines Day Blizzard”, passengers on flights originating at New York City’s John F. Kennedy International Airport (JFK)¹ suffered extremely long delays. In 2007, delays were not uncommon. However, some of these passengers originating at JFK endured as much as 7 hours of delay on their aircraft, often with only peanuts to eat. Boarded and pushed away from the gates, the aircraft were unable to return to a gate to allow passengers to deplane in the deteriorating weather conditions. The media learned about the situation of the trapped passengers, and outrage ensued. Lengthy tarmac delays, defined as those lasting more than three hours, were fairly common in 2007. That year, there were 1,654 instances of an aircraft taxiing out for three hours or longer. This figure could be much higher, as it does not count any aircraft that pushed back and joined the departure queue but later cancelled and taxied back to a gate to deplane. As we show in Table 1.1, using data from the U.S. Bureau of Transportation Statistics (BTS), the number of aircraft that taxied out between one and three hours was more than 75 times greater than the number of flights with taxi-out times of three hours or more.

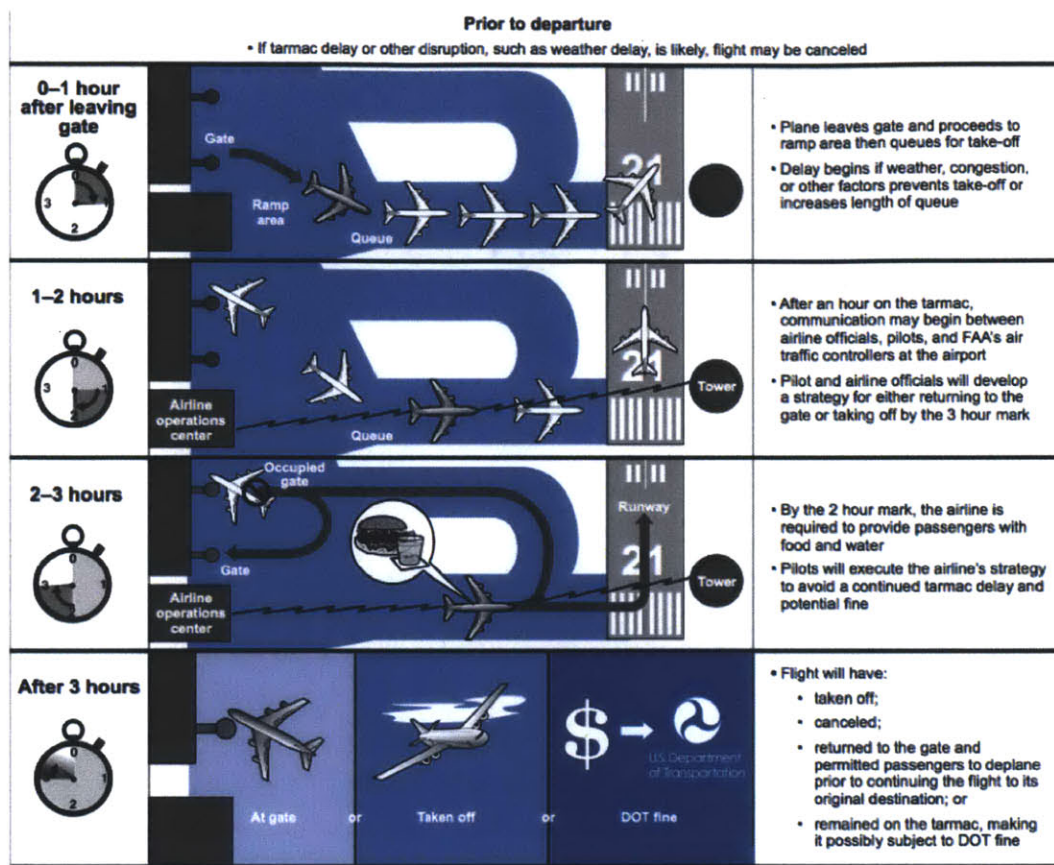
Amid heavy consumer advocacy group pressure, the U.S. Department of Trans-

¹See Table A.1 for list of International Air Transport Association (IATA) airport codes.

Length of taxi-out delay	Number of flights affected
1 hr to 1:59	75,833
2 hr to 2:59	7,507
3 hr to 3:59	1,370
4 hr to 4:59	239
5 hr to 5:59	36

Table 1.1: Non-cancelled flights (*including* diversions) that experienced lengthy tarmac delays during taxi-out in 2007, as reported to BTS

portation announced a policy known as the Tarmac Delay Rule (the “Rule”) on December 21, 2009; it went into effect on April 29, 2010. The Rule stipulates that aircraft lift-off, or an opportunity for passengers to deplane, must occur no later than three hours after the cabin door closure at the gate. There are two exemptions: if the pilot determines that moving from the departure queue or deplaning passengers would constitute a safety or security risk, or if local air traffic control decides that airport operations would be significantly disrupted by the delayed aircraft returning to a gate or deplaning area. In the Rule, it is suggested that carriers and individual airports develop a plan that is mutually agreeable for deplanement in case a violation is imminent. It should be noted that significant latitude for local decision-making is written into the Rule allowing local air traffic control to decide what constitutes a significant disruption to operations. The pilot must request clearance to leave the departure queue to taxi to a gate or other deplanement area in sufficient time to comply with the three-hour rule; that is, the aircraft cannot begin to taxi-in at the *end* of the three-hour period. Instead, passengers must be fully deplaned at the three-hour limit. Additionally, food and water must be made available no later than two hours from push-back (for departing aircraft) and from touchdown (for arriving aircraft). Operable lavatory facilities must be available as well. The Rule currently applies to U.S. flag carriers operating domestic flights, and to international carriers, originating or landing at U.S. airports (in this case the limit is four hours). Flights operated by aircraft with under 30 seats are exempt. The Rule’s penalty for non-compliance is a fine of up to \$27,500 per passenger. In Figure 1-1, from the U.S. G.A.O. Report (2011), we see at what point in the taxi-out process decisions must be made.



Sources: GAO analysis of tarmac delay rule and aviation stakeholder interviews.

Figure 1-1: Schematic of airline decision-making when faced with a long taxi-out delay

Since the announcement and implementation of the Rule, taxi-out incidents of three hours or more have significantly decreased, as depicted in Figure 1-2, using data from BTS.

A decrease of 13.5% in scheduled operations between 2007 and 2010 was accompanied by a 95% reduction in non-cancelled flights with taxi-out delay of three hours or more (1,654 in 2007 to 78 in 2010). This data would suggest that the Rule has been highly effective in keeping passengers off the tarmac for lengthy periods of time. We benchmark against the reduction of operations and tarmac delays in 2000 and 2002. In 2000, we observed 5.68 million scheduled operations, with 1,587 taxi-outs lasting longer than three hours (BTS data). In 2002, operations decreased 7.24% over 2000




Figure 1-2: 2006-2010 non-cancelled flights taxiing-out more than three hours, and total scheduled operations

(e.g. 5.27 million), but there were still 930 flights with taxi-out times lasting longer than three hours, comprising a 41% reduction over 2000. In 2010, however, with more scheduled operations (6.45 million), only 78 such incidents occurred. While the Rule seems effective in keeping passengers from experiencing lengthy delays on the tarmac, we explore other consequences of the Rule in this work. The September 2011 U.S. Government Accountability Office report uses available data on tarmac delays before and after the implementation of the Rule, and develops two regression models to evaluate whether cancellation rates increased after the Rule went into effect. The authors interview airline officials who state that airlines changed their decision-making about cancellations in response to the Rule. In order to test this qualitative finding, the authors develop two regression models to control for other factors that are related to cancellations. These other factors considered include level of airport congestion, origin/destination weather conditions, ground delay programs, airport on-time performance, size of airline, status as a hub, passengers per flight, route distance, day of week, and scheduled departure hour. Their results suggest that flights experiencing

any level of taxi-out delay were more likely to be cancelled after, rather than before the implementation of the Rule. In Table 1.2, we observe how the likelihood of cancellation rapidly increases as the duration of taxi-out delay rises.

Taxi-out time	Increased likelihood of cancellation in 2010 versus 2009
Before taxi out (at gate)	24% more likely
1-60 minutes	31% more likely
61-120 minutes	214% more likely
121-180 minutes	359% more likely

Table 1.2: U.S. G.A.O.-reported likelihood of flight cancellation change, by tarmac time

In this work we attempt to quantify the impact of the Tarmac Delay Rule with respect to passengers on tarmac-delayed flights in the U.S. National Aviation System (NAS). We cannot simply compare the passenger delay in a year before the Rule was implemented to the passenger delay in a year after. This is due to a variety of factors, which include changes in airline schedules year-to-year, differences in congestion levels, demand fluctuations, capacity changes, and weather differences. Additionally, passenger delay calculation itself presents a challenge due to lack of available data. We describe in Chapter 2 the approach we adopted and adapted for calculating passenger delay. To understand the impacts to passengers resulting from the Rule, we experiment with a simulation using pre-Rule operations in which we identify 2007 flights with significant (more than three hours) taxi-out delay, create a number of scenarios in which some or all of these flights are cancelled, and calculate the resultant delay to the passengers on these flights.

There are many ways to measure the impact of a flight cancellation on a passenger, including quantifying monetary loss and logistical hassles, or the loss of a day at a conference, meeting, or vacation. However, given the lack of granularity in our data about individual passengers and their value of time, we focus on one metric we can measure with some degree of certainty. We quantify the difference between a passenger’s planned arrival time at their final destination and their actual arrival time.

In selecting the set of flights on which to perform our analysis, we focus on flights

with taxi-out delays, instead of taxi-in delays. This is because airlines have a higher degree of control over the operational actions of an aircraft with a taxi-out delay. For example, a decision can be made to return to a gate, though of course with no guarantee of efficiency in this process. In a taxi-in delay, however, the aircraft has essentially only the option to wait for a gate; it does not take off again and return to its origin airport. In addition, the number of aircraft taxiing in longer than three hours is far fewer than the number of taxi-out events greater than three hours (see Table 1.3, with data from BTS).

Year	Taxi-outs more than 3 hours	Taxi-ins more than 3 hours
2006	1,341	61
2007	1,654	43
2008	1,231	19
2009	606	2
2010	79	4

Table 1.3: Lengthy taxi-out and taxi-in incidents, 2006-2010

Finally, we chose to select 2007 as our representative pre-Rule operational scenario, as 2007 had the highest number of lengthy (greater than three hours) taxi-in incidents of any year from 2006-2010. Additionally, that year featured several notable lengthy tarmac delays that prompted consumer protection groups to lobby Congress for regulations that led to the “Tarmac Delay Rule”, such as the Valentines Day Blizzard described previously.

1.2 Contribution and Outline

In this work we develop a strategy for quantifying the delay impact to passengers of cancellations resulting from the Tarmac Delay Rule. There is currently no data that exists to describe this impact. We apply an existing methodology (the Passenger Delay Calculator) to flight schedule and operational data for a year *before* the Rule was implemented, and analyze the impacts of varying levels of cancellation rates and alternative wordings of the Rule. We aim to discern how the policy is effective, and in

which ways it is costly to passengers, both on tarmac-delayed flights, and elsewhere in the NAS; and to provide policy makers with insights to inform future policy.

In Chapter 2, we describe the procedure we use to calculate passenger delays, an overview of other methods of passenger delay calculation, and a brief discussion of why we chose this particular method for our research. In Chapter 3, we estimate the delays that might have resulted to passengers had the Rule been in effect in 2007, and compare this estimated delay to the delay experienced by passengers in our 2007 base-case analysis. In Chapter 4, we identify the characteristics of flights that are the most impacted and likely to have the greatest increase in delays as a result of the Tarmac Delay Rule. We use this information to test various cancellation policies and compare the resultant delay. In Chapter 5, we perform sensitivity analyses to understand the impact of our modeling assumptions and simplifications on our delay estimates. In Chapter 6, we provide a summary of the findings of this research and detail future research topics that might be explored as more data becomes available.

Chapter 2

Methodology and Data

The goal of this work is to quantify the impacts to passengers as a result of the Tarmac Delay Rule. We do so by identifying flights that incurred lengthy taxi-out delays in 2007, and use them to perform a variety of scenario analyses. The metric we obtain in our results is passenger delay; we will define this metric and the method by which it is calculated next.

Passenger delay is defined as the difference between the scheduled itinerary arrival time at the passenger's final destination, and their actual arrival time. Passenger delay is differentiated from flight delay as it also considers passenger disruptions, resulting from flight cancellations, diversions, and passenger *misconnections* (a passenger misconnects if their first flight arrives less than 15 minutes before the actual departure of their second flight). Flight delay alone can considerably underrepresent the delay to passengers. For example, as a result of a two-hour flight delay, a passenger on this delayed flight with a one-hour connection misses his or her connecting flight leg, and has to wait, say three hours, for the next flight with an available seat to his/her final destination. This results in an itinerary delay of four hours, double the two-hour flight delay. As observed from this example, passenger delay depends on the itinerary of the passenger, and thus is greatly impacted by the flight schedule and number of available seats. A *recovery itinerary* is a flight or sequence of flights on which a *disrupted passenger* (one who misconnects or whose itinerary has one or more cancelled flights) is rebooked in order to reach his/her final itinerary destination.

2.1 Literature Review

This thesis builds primarily upon the work of Bratu and Barnhart (2005), and Barnhart, Fearing and Vaze (2013), and relies heavily on the Passenger Delay Calculator (PDC), an algorithm formalized in the work of Bratu and Barnhart, which calculates passenger delay given inputs of flight schedules (planned and actual), itineraries, and aircraft capacity data. Sherry, Wang and Donohue (2007) also calculate passenger delays but treat all passenger itineraries as non-stops. This approach is not applicable for our purposes, as we wish to incorporate missed flight connections into our calculation of passenger delay resulting from the Tarmac Delay Rule. Tien, Ball, and Subramanian (2008) also provide an algorithm for calculating passenger delay, but their approach uses fixed parameter values for the percentage of flights cancelled. Thus, we select the Bratu and Barnhart PDC as our calculation method. The PDC that Bratu and Barnhart developed is a greedy algorithm accommodating disrupted passengers in the order in which they are disrupted. When a flight is cancelled (and thus all the passengers on that flight have the same disruption time), the passengers are randomly (though any specified ordering is possible) placed into the disruption queue for rebooking onto a recovery itinerary. Passengers on cancelled flights are assumed available for rebooking at the planned time of departure of the cancelled flight, and passengers who misconnect are available for rebooking at the actual arrival time of the first flight of their itinerary. These disrupted passengers can be rebooked onto flights departing 45 minutes or more after their earliest rebooking time. The Bratu and Barnhart PDC is validated using one month of booking data from a major U.S. carrier from August 2000. Due to data availability, only passengers with domestic itineraries containing at most one connection were considered.

In this work, we also adopt the Barnhart, Fearing and Vaze (2013) work in which they estimate disaggregate passenger itinerary flows, using publicly available aggregate data and train their model on one quarter of booking information from a major U.S. carrier. They use a multinomial logit modeling approach to disaggregate the itinerary flows, as this method proved successful in Coldren and Koppelman (2005).

Barnhart, Fearing and Vaze also build upon the work of Bratu and Barnhart to extend the PDC to consider recovery itinerary options from within all 20 carriers from which we have operational data in 2007¹. Their work also includes a model for estimating seating capacities of aircraft whose tail numbers are not listed in the Schedule B-43 (see Section 2.2).

We present in Figure 2-1 a step-by-step schematic of the PDC algorithm (Bratu and Barnhart 2005), with updates from Barnhart, Fearing and Vaze (2013) as utilized in this work. In Step 1, inputs to the algorithm include passenger itineraries and

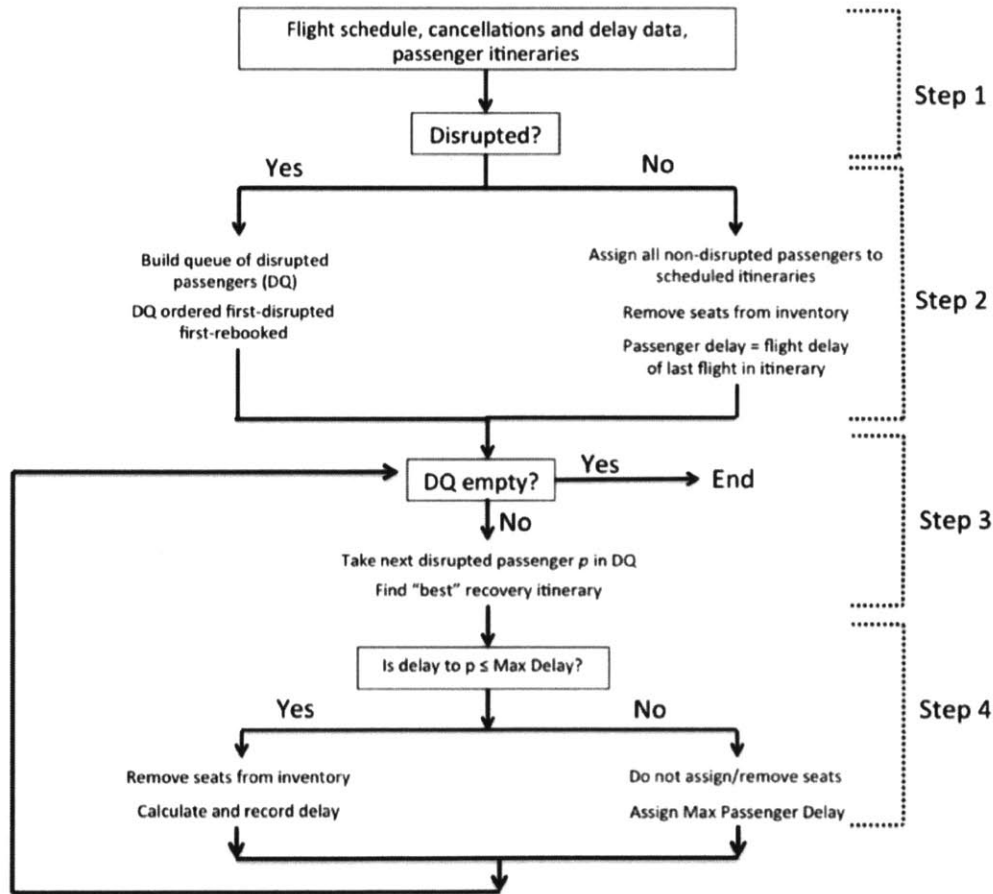


Figure 2-1: Passenger Delay Calculator algorithm schematic

flight schedules, cancellations, and delay data. Given this, all passengers are assigned

¹These carriers report to Airline On-Time Performance (ASQP) data, defined in Section 2.2

a binary identifier of disrupted or non-disrupted.

In Step 2, each passenger not disrupted is assigned to his/her planned itinerary and the pool of available seats is accordingly reduced on the flight legs in the planned itinerary. Passenger delay, if any, is recorded. A non-disrupted passenger on a nonstop itinerary is assigned passenger delay equal to the flight delay of their flight; a non-disrupted passenger on a connecting itinerary has passenger delay equal to the flight delay of the last flight in the itinerary. In the rare case that a passenger arrived *before* their scheduled itinerary arrival time and thus had “negative” passenger delay, their delay is set to zero. This occurs when their flight flew faster than scheduled, or if they were rebooked onto an itinerary that arrived earlier than their scheduled itinerary. Disrupted passengers are placed into a *Disruption Queue* (DQ) and are ranked in a first-disrupted, first-rebooked fashion in DQ. This ordering policy is chosen because we do not have access to detailed information about passengers’ airline frequent flier status or fare class, which could allow us to utilize other rebooking priority schemes. Passengers who have the same disruption time (for example, passengers on the same cancelled flight) are randomly ordered in the queue. This is also due to lack of detailed information about passenger fares, cabin status, etc.

In Step 3, if DQ is empty, the algorithm ends. If DQ is not empty, the next disrupted passenger, p , is selected. The algorithm searches first for a recovery itinerary for p on the same or related carrier operating any of the flights in the planned itinerary of passenger p . Related carriers are the parent carrier (e.g., American Airlines) or the subcontracting/regional carrier (American Eagle). If no recovery itinerary for p is found on the same or related carrier, all other carriers are considered.

Once a recovery itinerary is identified in Step 3, the algorithm moves to Step 4, where the recovery itinerary is checked against the maximum passenger delay time. If the passenger is scheduled to arrive at his/her final destination with no more than eight hours (for passengers disrupted between 5:00am and 4:59pm), or 16 hours of delay (for passengers disrupted between 5:00pm and 4:59am), passenger p is assigned to the itinerary, the seat(s) are removed from the flight(s) comprising the recovery itinerary, and p is assigned a delay time value equal to the difference

between p 's scheduled arrival time of his/her planned last flight, and the actual arrival time of the last flight in the recovery itinerary; this delay is then recorded. As discussed previously, delay is always non-negative; delay is set to zero if the passenger arrived *earlier* than their scheduled itinerary arrival time. If passenger p cannot be accommodated on any carrier without incurring more than the maximum passenger delay, no itinerary is selected, no seats are removed from the inventory, and passenger p is instead assigned a maximum value of delay (eight hours for passengers disrupted between 5:00am and 4:59pm and 16 hours for passengers disrupted between 5:00pm and 4:59am); this value is then recorded. The differences in maximum delay values depending on the time of disruption reflect the difficulty in rebooking later in the day, due to reduced frequency of flights during the night.

After delay is recorded for passenger p at the end of Step 4, the algorithm returns to Step 3 to check for another passenger in DQ. If DQ is not empty, Steps 3 and 4 repeat. If DQ is empty, the algorithm ends.

2.2 Data inputs

Next, we describe the data inputs to the PDC from which the disruption queue (DQ) is constructed and passenger delays are estimated. The bulk of this data is publicly available from the U.S. Bureau of Transportation Statistics (BTS), which makes available databases pertaining to commercial air transportation in the United States. The data inputs to the PDC include:

1. **Airline On-Time Performance Data (ASQP)**: This is a database that includes for each flight, scheduled and realized flight departure and arrival locations and times for flights operating at airports in the 48 U.S. contiguous states; taxi-out and taxi-in times; wheels-off and wheels-on times; operating carrier; and flight number, reported monthly by air carriers in the United States that serve more than one percent of domestic scheduled passenger revenues. In 2007, this included 20 unique carriers (see Table A.2). The 2007 data does not report the airport to which flights were diverted, nor does it include the taxi-out time

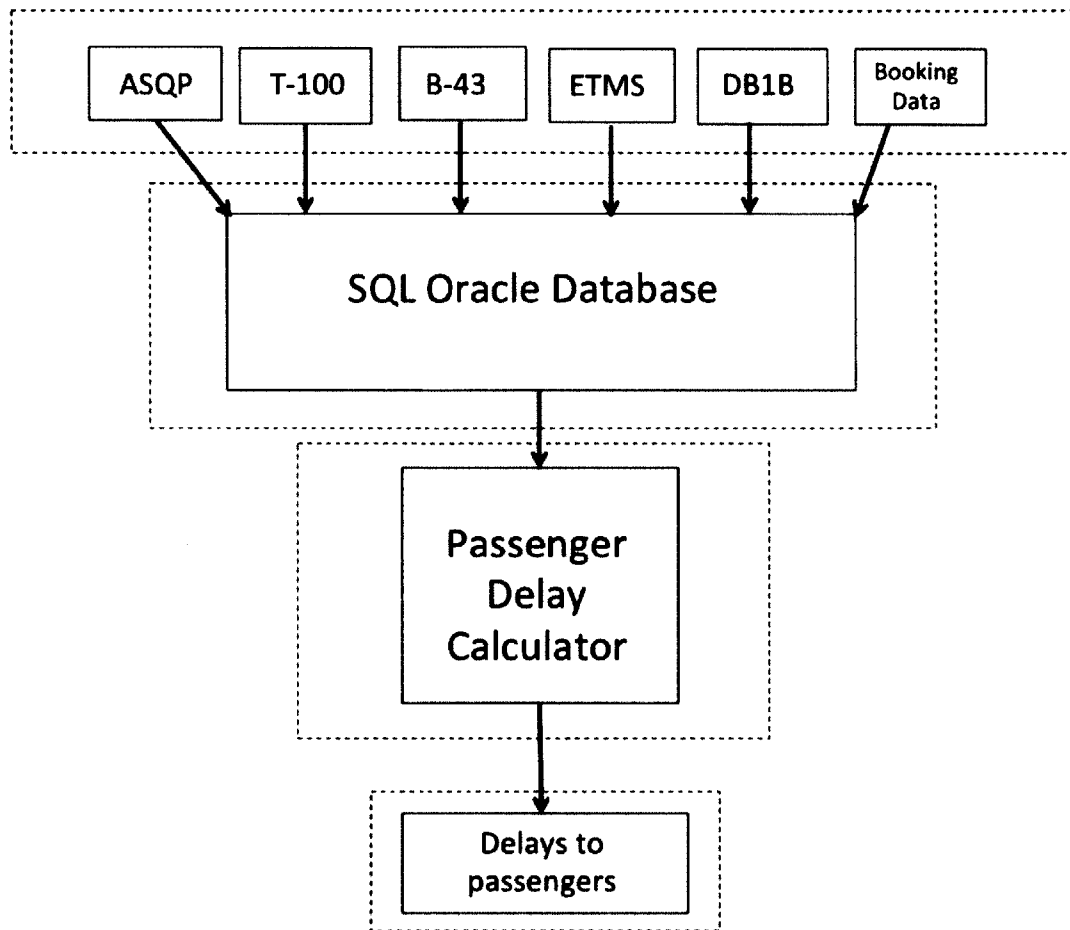


Figure 2-2: Data inputs and outputs to Passenger Delay Calculator

for a flight that may have departed the gate but was subsequently cancelled prior to lift-off.

2. **T-100 Domestic Segment database (T-100):** This dataset allows us to estimate *load factors* (the ratio of total passengers flown to total seats flown) by providing us with the number of seats flown and passenger flows on each flight leg, carrier segment and aircraft type, aggregated monthly. Thus, a passenger flying OAK-IAD-BOS on a given carrier and aircraft type(s) is added to the count of both the OAK-IAD and IAD-BOS flight segments, for all such flight segments operated that month by that carrier of the given aircraft type(s). The passenger counts are used to feed the multinomial logit passenger itinerary flow

model presented in Barnhart, Fearing, and Vaze (2013).

3. **Form 41 Schedule B-43 Aircraft Inventory Data, and Enhanced Traffic Management System (ETMS):** The Schedule B-43 database includes aircraft seating capacities (specified by tail number), which we match to the ASQP database using tail numbers. This allows us to estimate the available seat count for each flight reported in the ASQP. About 75% of the flights in ASQP can be matched to an entry in this dataset; the remaining are obtained by using the FAA's ETMS database (not publicly available). Together, Schedule B-43 and ETMS provide us with the seating capacities of 98.5% of the flights in ASQP. The remainder is obtained through an algorithm presented in Barnhart, Fearing, and Vaze (2013) using the T-100 Domestic Segment database.
4. **Airline Origin and Destination Survey (DB1B):** This dataset, aggregated quarterly, is a 10% sample of ticketed passengers on ASQP-reporting carriers. Each carrier reports all ticket-coupons ending in '0' (thus the carrier would report the information on ticket number XYZ10, XYZ20, and so on, assuming the last two digits increase sequentially as 10, 11, , 19, etc.). This results in a randomized sample of reported passenger itineraries. DB1B differs from T-100 data in that the same passenger, flying from OAK-BOS, and connecting in IAD, is reported in DB1B as a connecting passenger with an origin of OAK, and a destination of BOS. This data is used in the calculation of origin-destination passenger itinerary flows, as presented by Barnhart, Fearing, and Vaze (2013).
5. **Booking data:** The fourth quarter of 2007 proprietary booking (passenger itinerary) data of a legacy carrier is used by Barnhart, Fearing and Vaze (2013) to train the passenger itinerary flows multinomial logit model, and to validate results.

These six individual SQL databases are joined in an Oracle SQL database that provides input to the Passenger Delay Calculator (Figure 2-2). The Passenger Delay Calculator (PDC) is coded in Java programming language, and connected to the

Oracle SQL database. Outputs of the PDC include an itinerary-by-itinerary comma-separated value (CSV) data file in which each passenger itinerary is associated with a delay time value, and with the number of passengers on that itinerary.² This output file provides us with the means to calculate actual 2007 passenger delay, which we use as our baseline delay. Throughout this thesis, we compare the 2007 baseline “as-flown” delay value to other hypothetical scenarios that we create and analyze. For each scenario, we manipulate the input databases to represent our hypothetical situation. For example, when we wish to introduce a policy of cancelling flights that taxi-out longer than three hours in 2007, we use the SQL language to change the cancellation flags of selected flights in ASQP. The passengers on the now-cancelled flights are added to the Disruption Queue, along with other passenger who were disrupted in the baseline scenario, and the PDC algorithm computes the resulting passenger delays.

In addition to manipulating the database inputs to the PDC, we also systematically exclude diverted flights from our analysis because diversion airports are not reported in ASQP. Thus, we do *not* include flights that taxied-out longer than three hours and were diverted in the set of flights subject to cancellation as a result of the Tarmac Rule, and we do not count passengers on these flights as people affected by tarmac delays. Hence, we effectively exclude diverted flights from being subject to the Tarmac Delay Rule, and assume that these passengers’ disruptions statuses remain unchanged in our hypothetical scenarios compared to the baseline scenario. This assumption should have an insignificant effect on our results: the total number of non-cancelled flights taxiing out three hours or more in 2007 was 1,654, and only 24 (1.45%) of these flights were diverted. However, we note that these passengers’ delays *are* included in the calculation of overall delay in the NAS.

²Multiple passengers can have the same itinerary and disruption; for example, a family travelling together.

Chapter 3

Passenger delay stemming from the Tarmac Delay Rule

In this section, we quantify the impacts of the Tarmac Delay Rule on passengers on tarmac-delayed flights, as well as for passengers on flights in the same departure queue as the tarmac-delayed flights. We investigate the impact of cancellation of flights that would have violated the Tarmac Delay Rule had the Rule been in place in 2007. In the analysis that follows, we will use the results of the Passenger Delay Calculator to compare passenger delay experienced in 2007 with the delay these same passengers might have experienced if the Tarmac Rule had been in effect. We begin our analysis by cancelling some or all flights that had a taxi-out delay of three hours or more in 2007. We then calculate resultant passenger delays after we have manipulated the operational data by cancelling the selected flights. Of note, we do not aim to quantify the exact amount of delay that would have been experienced by passengers had their flights been cancelled; our model makes a variety of assumptions and simplifications which we believe to result in an underestimate of delay. Some of these simplifying assumptions are analyzed in more detail in Chapter 5, their implications on our delay calculation results analyzed, and our hypothesis of model output underestimation tested.

3.0.1 Notation

Throughout this thesis we refer to the *baseline scenario*. This is the 2007 as-flown case, and is what occurred operationally in 2007, including all cancellations and lengthy tarmac delays. We refer to *baseline delay* as the delay that passengers experienced in 2007 due to cancellations and tarmac delays.

Additionally, we refer to a set of *affected flights*, denoted F_{AF} . A flight $f_i \in F_{AF}$ is a flight f_i that was operated (not cancelled or diverted) and experienced a taxi-out time in 2007 greater than or equal to 180 minutes. We refer to the passengers on this set of flights F_{AF} as passengers P_{AF} .

3.0.2 Initial assumptions

For each flight $f_i \in F_{AF}$ that we decide to cancel, we assume that passengers on f_i are available for rebooking at the planned departure time of f_i . We allow each passenger to be rebooked only on flights that depart at least 45 minutes after his/her originally scheduled departure time. This 45-minute window allows for passenger transfer between terminals, if necessary. For the purposes of our initial experiment, but later relaxed, we also assume that a flight leg can be cancelled in isolation, without cancelling a subset of flight legs to maintain balance in the schedule. In Chapter 5, we investigate the implications of these assumptions on our passenger delay estimation.

3.1 Passenger delays as a function of cancellation of flights in F_{AF}

Here, we examine how resultant delays to passengers P_{AF} are related to the percentage of cancelled flights in F_{AF} . For this analysis we consider three different days (days 1, 4 and 6 of Table A.3) and randomly cancel 0%, 10%, 20%,...,90%, 100% of flights in F_{AF} for each.

For each cancellation, of $x\%$ of flights in F_{AF} , we run the Passenger Delay Calculator model 10 times, each time cancelling a random sample of $x\%$ of the flights in

F_{AF} , while allowing the other flights not cancelled in F_{AF} to remain as-flown in 2007. We average the delays to passengers P_{AF} calculated for the 10 samples to obtain the expected delay for each cancellation rate. Flights not cancelled in each random sample, accrued the same amount of tarmac delay as in 2007, and departed at their 2007 wheels-off time. For this analysis we therefore generate one value for the total delay for all passengers P_{AF} given a 0% cancellation rate (this represents the baseline scenario), one value for a 100% cancellation rate (the *all-cancel scenario*), and ten separate values, all measuring total expected delays for each cancellation rate of 10%, 20%,...,90%. To quantify the range of potential passenger delays, we obtain the average, the minimum and the maximum total passenger delay for each of the 10%, 20%,...,90% cancellation rate scenarios, and report one value each for the baseline and all-cancel scenarios.

In Figures 3-1 to 3-3, we see that the relationship between passenger delay and percentage of flights cancelled is linear for the three days analyzed. Trendlines fitted to the average expected value of each cancellation rate are observed to be linear, with R^2 values for each day above 0.99 (see Table 3.1). Additionally, upper bound and lower bound delays are observed also to be linearly related.




Figure 3-1: P_{AF} delay for July 27 as a function of percentage of F_{AF} flights cancelled




Figure 3-2: P_{AF} delay for June 21 as a function of percentage of F_{AF} flights cancelled

For these three days, we depict in Table 3.1 the values of average delay per passenger and total passenger delay for the as-flown case and for the 100% cancellation scenario; the slope of the trend line fitted to average passenger delays; the ratio of total baseline delay to the delay corresponding to the 100% cancellation scenario; and the correlation coefficient of the average delay values to the trend line.

Based on this analysis, we conclude that cancelling F_{AF} flights results in more passenger delay than allowing the flights to remain more than three hours on the tarmac before taking off. This finding is significant: passengers P_{AF} are worse-off (in terms of final destination arrival time delay) when their flights are cancelled than when allowed to taxi-out more than three hours and then depart. We use this result to motivate much of our analysis in subsequent chapters. We note that the slope of the trend line, or rate of increase of passenger delay per percentage cancellation, is much higher for July 27 and June 21 (0.8 and 1.0 hours/passenger, respectively), than for January 21 (0.3 hours/passenger). This difference may be partially explained by varying load factors. July and June 2007 have high summer load factors of 81.5% and 82% (T-100, BTS) respectively, while January 2007, a winter month, has an average load factor of 67.4% (T-100, BTS). In the next section we investigate the impacts of

Metric	July 27	June 21	January 21
Average delay (hours) per passenger $p \in P_{AF}$, baseline scenario	4.76	4.13	5.91
Average delay (hours) per passenger $p \in P_{AF}$, all-cancel scenario	12.87	14.28	8.51
Additional delay per $p \in P_{AF}$ for each 10% increase in the cancellation rate of F_{AF}	0.8 hours	1.0 hours	0.3 hours
Additional delay (years) to P_{AF} for each 10% increase in cancellation rate of F_{AF} , baseline scenario	4.74	2.495	0.85
Delay to passengers P_{AF} , 100% cancellation scenario (years)	12.83	8.63	1.23
Ratio of delay to passengers P_{AF} for 100% cancellation to baseline scenario	2.7	3.5	1.45
R^2 value	0.9986	0.995	0.99

Table 3.1: Key findings of the impacts of cancellations




Figure 3-3: P_{AF} delay for January 21 as a function of percentage of F_{AF} flights cancelled

load factor on passenger delays resulting from F_{AF} cancellations.

3.2 Understanding the impacts of load factor on passenger delays

Next, we examine the sensitivity of our findings in Section 3.1 with respect to load factors. In this section, we separate out the effects of two main drivers of passenger delay: schedule (that is, next flight availability) and seat availability (or load factor). By quantifying the contributions of each of these drivers to delays experienced by passengers affected by lengthy taxi-out times, we discern why passengers P_{AF} in the all-cancel scenario from Section 3.1 experienced more delay than they did under the baseline scenario, when their flights incurred long tarmac delays but eventually took off.

To isolate the effects of schedule, we eliminate load factor as a constraint by setting all aircraft seating capacities to 999, effectively accommodating all passengers on the next scheduled flight to their respective destinations. We then perform the delay analysis described in Section 3.1 for days 1, 4 and 6 of Table A.3. In Figures 3-4 through 3-6, we plot the resulting delays and compare them to the case with actual aircraft seating capacities. For each seating capacity scenario, we report the average expected value of delay to passengers P_{AF} for each cancellation rate 10%, 20%,...,90%.

Under the all-cancel scenario for July 27, seat unavailability during rebooking accounted for 41% of the resulting delay to passengers P_{AF} . The remaining 59% of the delay is attributable to the flight schedule, that is, the inherent delay caused by waiting for the next scheduled flight. We conclude from this result that flight load factors play an important but not dominant role in passenger delay. We also observe that passenger delay impacts from load factors and schedules today would be greater now than in 2007. Load factors in July 2012 averaged 82.95%, up from 81.47% (T-100 data) in July 2007, and there has been a 16% reduction in scheduled flights reported in ASQP from July 2007 to July 2012. Similar results are achieved for June 21 and January 21, shown in Figures 3-5 and 3-6.

An all-cancel scenario for June 21 with actual seating capacities results in an increase in total passenger delay of 3.5 times the baseline delay (specifically, 8.63




Figure 3-4: P_{AF} delay for July 27, comparing actual and unlimited aircraft seating capacities

versus 2.50 years of passenger delay, respectively). However, when we remove capacity constraints, this reduces to an increase of 2.5 times that of the baseline (6.20 versus 2.50 years of passenger delay). Overall, 60% of the total delay for the all-cancel scenario is attributed to schedule and 40% to load factor.

For January 21, 2007, with lower load factors during that month, 69% of passenger delay for the all-cancel scenario is driven by lack of available scheduled flights, while only 31% of the delay is attributable to lack of available seats for passengers P_{AF} requiring rebooking.

Our results suggest that the negative impact to passengers P_{AF} results from lack of scheduled flights, and the resulting absence of recovery itineraries for disrupted passengers. Because dynamic adjustment of flight schedules to create additional capacity for passengers whose flights have been cancelled due to the Tarmac Delay Rule is very difficult to achieve, disrupted passengers often experience a lengthy delay in




Figure 3-5: P_{AF} delay for June 21, comparing actual and unlimited aircraft seating capacities

their arrival times at their final destinations (see Section 3.1).




Figure 3-6: P_{AF} delay for January 21, comparing actual and unlimited aircraft seating capacities

3.3 Effects to passengers not travelling on tarmac-delayed flights

We extend our analysis of the impact of cancellations of all flights in F_{AF} , turning our attention to *other flights*, F_{OF} , that is, flights not in the set F_{AF} , and the passengers travelling on these flights, P_{OF} , on the days in question. The goal is to provide a NAS-wide perspective on the potential impacts of the Tarmac Delay Rule, specifically investigating how passengers on ‘affected’ and on ‘other’ flights are impacted by the Rule. We consider six different airports on six different days during which the number of F_{AF} flights at the given airport is 10 or more. In our experiment, we cancel all flights in F_{AF} , and allow departing flights to utilize the departure slots that become available as a result. We then measure the reduction in passenger delays for these earlier departing flights. The days and airports we consider in this analysis are detailed in Table 3.2.

For each day at each of the given airports, F_{OF} represents the set of non-cancelled

Impact	Date	Airport	Flights in F_{AF} at given airport
High	January 17	DFW	19
High	June 26	ORD	17
Medium	May 31	EWR	13
Medium	July 27	JFK	12
Low	July 10	ATL	10
Low	June 19	PHL	10

Table 3.2: Days and locations selected for analysis of impact of F_{AF} flights on the wider NAS system

flights incurring less than three hours of taxi-out time in the 2007 baseline scenario. Note that the complete set of flights is partitioned into the two sets F_{AF} and F_{OF} .

We assume the cancellation of a flight in F_{AF} allows other flights in F_{OF} with later scheduled departure times to be assigned earlier *wheels-off slot times* (the time at which the aircraft becomes airborne). For a given day and airport, we begin this iterative process by cancelling the first flight $f_i \in F_{AF}$. This creates a free wheels-off time slot in the departure schedule. We illustrate the process of assigning wheels-off time slots to other flights in the departure queues using the example depicted in Figure 3-7.

We begin by ordering all departing flights for the given airport and day by wheels-off time. We choose to order by wheels-off time rather than planned or actual gate departure time in order to control for physical distance between particular gates and terminals, and the departure queue.

Flight f_i , the first flight in F_{AF} (ordered by wheels-off time), has a baseline wheels-off time, denoted $WOT(f_i)$. The baseline case, again, represents the as-flown schedule in 2007. For this illustration, assume flights f_{i+1} and f_{i+2} are members of the set F_{OF} . In Step 1, we identify and cancel flight f_i , creating a wheels-off slot, denoted $S(f_i)$, available for use by a subsequent aircraft in the departure queue. In Step 2, we identify f_{i+1} , the flight with a wheels-off time immediately following that of f_i . In Step 3, we check to ascertain if f_{i+1} is able to use the free wheels-off time slot, $S(f_i)$. In this step, we test if the planned gate departure time (PDT) of f_{i+1} is not later than the wheels-off time of f_i . If this condition is met, f_{i+1} is moved up to time slot $S(f_i)$.

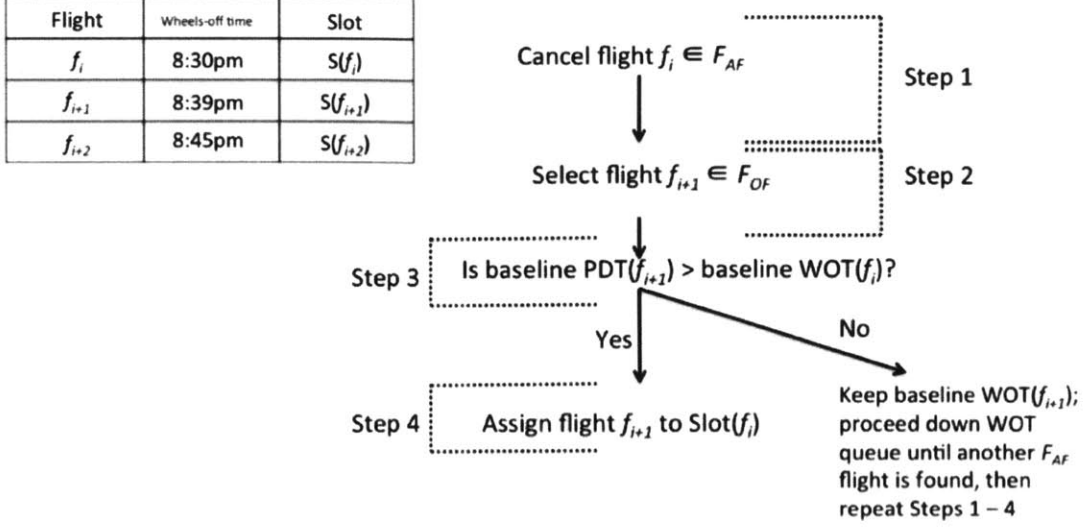


Figure 3-7: Schematic of algorithm used to move flights up into open slots created by cancellation of F_{AF}

If, however, this criterion is not met, the algorithm keeps flight f_{i+1} in slot $S(f_{i+1})$; slot $S(f_i)$ remains empty; and the algorithm continues down the wheels-off time list until the next flight $f_i \in F_{AF}$ is found, at which point it repeats the process illustrated in in Figure 3-7. We continue this iterative process moving up non-cancelled flights into available departure time slots, for each cancelled flight, using a “first-departed, first-moved-up” flight processing order based on actual wheels-off time. We denote the set of moved-up flights as $F_{OF,Moved}$ and record their corresponding earlier wheels-off times.

We then calculate total delay for passengers P_{AF} and P_{OF} for the following scenarios:

1. The baseline 2007 scenario.
2. The scenario with all flights in F_{AF} cancelled and updated wheels-off times for flights in $F_{OF,Moved}$.

We then calculate the net impact to passengers P_{AF} and P_{OF} by subtracting the sum of delay to passengers P_{AF} and P_{OF} in the second scenario above from the delay

Date	Airport	Flights in F_{AF}	Flights in $F_{OF,Moved}$	Aggregate impact (passenger-days)	Push-back time of first F_{AF} flight
January 17	DFW	19	190	19.9	7:17am
June 26	ORD	17	247	17.3	1:08pm
May 31	EWB	13	68	-530.2	5:07pm
July 27	JFK	12	53	-262.2	4:16pm
July 10	ATL	10	358	114.1	2:31pm
June 19	PHL	10	53	-292.2	4:31pm

Table 3.3: Key findings of earlier wheels-off time analysis

to those same passengers in the baseline scenario. In our results, shown in Table 3.3, we observe an expected net reduction in total delay resulting from the Rule at the investigated airports which had departures of the first F_{AF} flight in the morning or afternoon (DFW, ORD, and ATL). For the days with a net increase in expected delays, the airports investigated all had the first F_{AF} flights pushing back late in the afternoon or evening (JFK, PHL and EWB). Thus, passengers on flights in F_{AF} who were disrupted earlier in the day experienced relatively less delay. This likely results because these passengers have a greater opportunity to be rebooked that same day. (The time of disruption of passengers on F_{AF} is further explored in Sections 4.2 and 4.3.) We observe that the days for which the Rule produces the largest delay reductions have the largest number of flights in set $F_{OF,Moved}$, a consequence of the early push-back times of flights F_{AF} . On these days, when the first flight in F_{AF} pushes back earlier in the day, and wheels-off time slots are made available earlier, more flights have the opportunity to move up, and reduce their delay. Thus, we see again that the impact of cancellation resulting from a Rule violation is highly correlated with the time of day of the cancellations. This result further suggests that it is unwise to apply the Rule indiscriminately across all flights. To minimize overall passenger delay, time of day and other characteristics must be considered when deciding whether cancellations or extended tarmac delay allowances are most effective. In Chapter 4 we explore these characteristics further.

Chapter 4

Multipliers as a method to quantify impact to passengers on tarmac-delayed flights

In this section, we hone our understanding of the impacts of cancellations on passengers travelling on flights that are affected by lengthy tarmac delays, P_{AF} . We present a method to compare the impacts of cancellations of flights in the set F_{AF} for various days. Using this method, we identify characteristics of flights in F_{AF} that have higher cancellation impacts than other flights in the same set. We use this insight to develop a targeted cancellation strategy, and then measure the impacts to passengers P_{AF} .

4.1 Initial analysis and motivation for multipliers

In this section, we provide motivation for devising a metric for comparing the impact of cancellations. Using our insights gained, we then develop a method for impact comparison, which we define as a *multiplier*.

We examine all days, denoted D , in which the number of flights in set F_{AF} NAS-wide was greater than nine for the day. We then use the Passenger Delay Calculator to calculate two values for each day: baseline 2007 delays to passengers P_{AF} , and delay to passengers P_{AF} if all F_{AF} flights been cancelled that day. We graph the

results in Figure 4-1. The paired columns represent one day. The first column in the pairing shows the passenger-years of delay to passengers P_{AF} if all flights in set F_{AF} were cancelled. The second shows the delay to passengers P_{AF} in the 2007 baseline case. We observe that for all 39 days, the all-cancel scenario results in higher delay for the affected passengers P_{AF} for that day.




Figure 4-1: Years of delay under different scenarios to passengers P_{AF} for all days in 2007 where there were nine or more flights in F_{AF}

As shown in Figure 4-1, these 39 days have vastly different values for passenger delay, making comparison across days difficult. To facilitate comparison of passenger delay for the all-cancel and the baseline scenario, we define a *multiplier* as the ratio of the all-cancel delay to passengers P_{AF} to the 2007 baseline passenger delays to passengers P_{AF} . The *one-day multiplier* equals the sum of delay to passengers P_{AF} for a given day with all flights in F_{AF} cancelled, divided by the sum of the baseline 2007 delay for those same passengers P_{AF} . The *one-flight multiplier* is defined as the ratio of passenger delay on a particular flight $f_i \in F_{AF}$ when flight f_i is cancelled and the baseline 2007 passenger delay on flight f_i .

4.2 Characteristics of flights with high and low multipliers

In this section, we use the ‘one-flight’ definition of the multiplier to examine individual flights and their impacts on delay. First we calculate the multiplier for individual flights on specifically selected days. We then examine the characteristics of flights with high and low multipliers to discern the differences between flights for which passengers incurred very little or very large amounts of incremental delay as a result of cancellation.

We select days 1, 2 and 3 from Table A.3 on which to perform this analysis.

For each of these three days, we calculate and graph the individual flight multipliers in Figures 4-2 through 4-4. Each column represents a multiplier for a particular flight in F_{AF} . We note that any multiplier higher than one represents a flight for which the passengers on the given flight incurred more delay when the flight was cancelled than their baseline 2007 delay.




Figure 4-2: One-flight multipliers for flights in set F_{AF} on July 27

Only five of the 76 flights in F_{AF} for July 27 have a multiplier less than one. 21 flights have a multiplier of less than two for this day. Thus, in the all-cancel scenario,

most passengers P_{AF} on July 27 would have incurred more than double their baseline delay had all F_{AF} been cancelled.




Figure 4-3: One-flight multipliers for flights in set F_{AF} on June 26

Only four of the 71 flights in set F_{AF} for June 26 have multipliers less than one. 56 flights had a multiplier of less than two this day; thus most passengers P_{AF} would have incurred between one and two times their baseline delay for June 26.




Figure 4-4: One-flight multipliers for flights in set F_{AF} on February 14

There are 64 flights in set F_{AF} on February 14; only three have a multiplier less than one. 40 flights had a multiplier of less than two this day; most flights fall in the

one-to-two multiplier range.

In this analysis we find that when cancelled, very few flights in F_{AF} have multipliers less than one. That is, there are few flights where passenger delay on $f_i \in F_{AF}$ would be less if f_i had been cancelled and the passengers rebooked, compared to allowing f_i to remain on the tarmac and eventually take off. Many flights have multipliers between one and two; although we observe some flights with very high multipliers between four and six.

These results raise questions about the causes of the differences in multipliers within individual days. What are the common characteristics of flights with low or high multipliers, and how can that information be incorporated into decision-making to minimize passenger delay? To begin to answer this question, we analyze ASQP data related to the individual flights with low and high multipliers on the three days of analysis presented above.

4.2.1 Characteristics of flights with multipliers less than one

Passengers on flights whose multiplier is less than one experienced less delay when the flight was cancelled than their baseline 2007 delay. These flights had the following similar characteristics:

- The flight was already delayed by the time the passengers boarded the aircraft.
- The flights incurred extremely long tarmac delays; that is, much greater than three hours of delay on the tarmac.
- The flights tended to be earlier in the day.
- The carrier had more flights on the same route throughout the day.

Cancelling flights with these characteristics enables passengers to be rebooked more easily onto later flights that same day to the same destination on the same carrier. In addition, these passengers already have incurred delay before even enplaning due to prolonged flight delay, so cancelling the flight at its planned time of gate departure,

as is done in the Passenger Delay Calculator, allows passengers the opportunity to begin the rebooking process without incurring the delay.

4.2.2 Characteristics of flights with high multipliers

Next we examine flights with multipliers higher than average for the given day. We characterize flights with very high multipliers as follows:

- The flights tend to be later in the day or an overnight flight, often the carrier's last flight of the day to that destination.
- The flight has a taxi-out time of only slightly more than 180 minutes.

Excluding flights with these characteristics from the Tarmac Delay Rule prevents passengers on these flights from waiting overnight to fly on recovery itineraries the next day, thus incurring far less delay. Additionally, as we observe in Section 3.2, lack of schedule density accounts for most of the delay experienced by disrupted passengers. It is not uncommon for passengers on flights with just over three hours of taxi-out time to wait more than three hours for a recovery itinerary. Thus, cancelling these flights as a result of the Rule yields particularly high delay for passengers on the flights.

As observed in Figures 4-2 through 4-4, there exist some very high multipliers (more than 5.33). Passenger delay is capped at eight hours by the PDC algorithm for passengers disrupted between 5:00am and 4:59pm, and at 16 hours for passengers disrupted between 5:00pm and 4:59am. From this, any multiplier should be at most 5.33 (that is, 16 hours delay \div 3 hours taxi-out time). However, these flights, once airborne, sometime fly faster than their scheduled block times.

4.3 Validation of findings

In the previous section, we observe that flights with low multipliers tend, among other characteristics, to be earlier in the day, while those with high multipliers tend to be later in the evening. In this section, we evaluate if flights with later push-back times

have higher multipliers than those with push-back times earlier in the day, keeping all other factors equal. We plot in Figure 4-5 the one-flight multipliers for flights in F_{AF} departing from the six airports on the six days presented in Section 3.3. We order the one-flight multipliers in increasing actual gate departure (push-back) time and observe trends in Figure 4-5. This analysis confirms our observations in Section 4.2 that passengers on flights with tarmac delays occurring later in the day incur particularly long delays if the flight is cancelled, suggesting that cancellations from the Tarmac Delay Rule particularly penalize passengers on flights departing later in the day.




Figure 4-5: One-flight multiplier for flights in F_{AF} , from airports and days in Section 3.3, ordered by push-back time

We observe a trend toward higher multipliers associated with later push-back times, just as we observed in Section 4.2. For graphs of multipliers for the six different origin airports and days, please see Figures B-1 through B-6.

4.4 Selective cancellations to minimize passenger delay on F_{AF}

In this section, we use our insights from the previous sections to develop selective cancellation strategies in which we only cancel flights with specific characteristics. We then quantify and analyze the changes in delays for passengers P_{AF} under each strategy and identify the most optimal one.

We determined previously that passengers on flights with multipliers less than one incurred less delay when the flight was cancelled, as compared to the baseline 2007 delay. We hypothesize, given our previous findings, that cancelling flights with low multipliers while allowing flights with high multipliers take off, will result in lower overall delay for passengers P_{AF} than for the all-cancel scenario. In this analysis, we use days 1, 2 and 3 from Table A.3, and run the Passenger Delay Calculator on four different scenarios to calculate delay to passengers P_{AF} for each day:

1. Baseline 2007 case: no flights in set F_{AF} cancelled, all flights incur three-hours or more taxi-out delay, as experienced in 2007;
2. All flights $f_i \in F_{AF}$ with multipliers less than one cancelled;
3. All flights $f_i \in F_{AF}$ with multipliers less than two cancelled; and
4. All flights in set F_{AF} cancelled.

We present the results graphically in Figures 4-6 through 4-8.

There were 76 flights in set F_{AF} on July 27; five flights with a multiplier less than one, and 21 flights with a multiplier less than two. Delay to passengers P_{AF} for the case where multipliers less than one are cancelled is 4.876 years, while the baseline delay is 4.925 years.

There were 71 flights in set F_{AF} on June 26; four flights with a multiplier less than one, and 56 flights with a multiplier less than two. Delay to passengers P_{AF} for the case in which flights are cancelled only if their multipliers are less than one is 3.605 years, while baseline delay is 3.756 years.




Figure 4-6: Selective cancellations for July 27

There were 64 flights in set F_{AF} on February 14; three flights with a multiplier less than one, and 40 flights with a multiplier less than two. Delay for the case where flights with multipliers less than one are cancelled is 3.968 years, while the baseline delay is 4.215 years.

From these results, although few flights in F_{AF} have multipliers less than one, cancelling only these flights and allowing the others to remain on the tarmac and eventually take off resulted in slightly less total passenger delay for passengers P_{AF} . These results suggest that treating all tarmac-delayed flights the same, and approaching the problem by cancelling all flights in F_{AF} , does not minimize passenger delay. Instead, characteristics of the flight and the possibility for passengers to reach their final itinerary destination in a reasonable time must be considered in order to minimize passenger delay and tarmac time.




Figure 4-7: Selective cancellations for June 26




Figure 4-8: Selective cancellations for February 14

Chapter 5

Sensitivity Analysis: Implication of model assumptions on previous findings

In Chapters 3 and 4, we used the Passenger Delay Calculator to investigate the impact of cancellation on passengers P_{AF} , as well as the benefits to other flights of reduced congestion and delays due to F_{AF} cancellations. Our reliance on this delay calculation methodology motivates us to investigate the underlying assumptions and inherent simplifications in the model, and their impact on our delay calculation results. In this chapter we examine the implications of three of these, namely: cancellation-time threshold, passenger rebooking time, and non-cycle flight cancellations.

In the following sections, we explore these three modeling choices and quantify their impacts on delay estimation.

5.1 Cancellation-time threshold

In our previous analysis, we considered flights that taxied-out for three hours or more as members of the set of flights F_{AF} . However, given the wording of the Tarmac Delay Rule, this criterion is operationally infeasible. The Rule stipulates that aircraft must return to the gate and be available for deplanement, or must have taken-off, by the

end of the three-hour period. The three-hour countdown begins from the time the aircraft cabin door closes at the gate. In order to comply with the Rule, an aircraft must make the decision to leave the departure queue *before* three hours in order to return the aircraft to the gate and deplane passengers by the end of the three-hour mark. In some airports, communication about leaving the departure queue begins after one hour of taxi-out time (U.S. G.A.O. Report). Thus, in operational reality, more flights would be members of the set F_{AF} , and the number of passengers P_{AF} affected would likewise be higher. Therefore, we consider cancellation time thresholds that are shorter than three hours to obtain a more accurate representation of the impact of the Rule.

We select days 1, 2 and 3 from Table A.3 on which to perform our analysis. We choose one-and-a-half hours, two hours, two-and-a-half hours, three hours, four hours, and five hours¹ as a representative range of cancellation-time thresholds. We then define a flight to be in the set F_{AF} if it taxied out longer than X hours and was not cancelled or diverted in 2007. We generate six different sets of flights F_{AF} for each day, each corresponding to one of the X different cancellation-time thresholds. For each cancellation-time threshold and the corresponding set of flights F_{AF} and associated passengers P_{AF} , we calculate delays to passengers P_{AF} for the all-cancel and the baseline 2007 scenarios. We then calculate the ratio of the total delay experienced by passengers P_{AF} in the all-cancel scenario and the baseline delay experienced by the same passengers. We calculate this one-day multiplier for each time threshold, for each day. The results are depicted in Figures 5-1 through 5-3.

In all days analyzed, we observe that as the cancellation time threshold is decreased (and consequently more flights are cancelled), the associated multiplier increases. Of note, only June 26, with a taxi-out minimum time threshold of four hours, yielded a one-day multiplier less than one. We conclude from this that our analysis in Chapters 3 and 4 represents an underestimation of the impact of the Tarmac Delay Rule, as we only consider flights which taxied-out at least three hours as candidates

¹We calculate the one-day multiplier up to four hours in the case of June 26, when the longest taxi-out delay was 4.35 hours.




Figure 5-1: One-day multiplier for July 27 for all flights in F_{AF} , with varying cancellation-time thresholds

for cancellation. We project that, operationally, delays to passengers will be even higher than the delays we have calculated in this analysis due to the wording of the Tarmac Delay Rule, which requires that aircraft be ready to deplane at the *end* of the three-hour period. Given that the process of leaving the departure queue and returning to a gate takes longer and can be highly variable at a congested airport, flights at large and congested airports will be forced to have a lower cancellation-time threshold. In order to make the Rule equitable for flights across all types of airports (small and large, congested and non-congested), we suggest the Rule wording be changed to require aircraft to decide to return to the gate and leave the departure queue within three hours of cabin door closure, rather than requiring the flight be *ready to deplane* at the end of three hours.




Figure 5-2: One-day multiplier for June 26 for all flights in F_{AF} , with varying cancellation-time thresholds




Figure 5-3: One-day multiplier for February 14 for all flights in F_{AF} , with varying cancellation-time thresholds

5.2 Passenger rebooking time

We now examine the impacts of passenger rebooking times. In the Passenger Delay Calculator, we have assumed passengers on cancelled flights are available for rebooking at the time of their original flight's planned departure time, and can be reaccommodated on flights departing 45 minutes or more after their original cancellation time. This can be operationally unrealistic for two primary reasons:

1. The information needed to decide which flights to cancel is not available until the flights have already taxied-out and are awaiting take-off clearance, which will be after the flight's planned departure time.
2. Access to perfect information and instantaneous rebooking of passengers is not realistic, especially on a day with a large amount of NAS delay and disruptions.

It is more realistic to assume time is needed between when the flight is cancelled and when the passengers are made available for rebooking. In this analysis, we therefore assume passengers are available for rebooking two hours after the planned departure time of their original flight, and we compare the impacts of this assumption with our original assumption that passengers are available for rebooking at the planned departure time of their flight. We refer to this latter scenario as the *instantaneous rebooking case*. To perform this analysis, we select days 1, 5 and 7 from Table A.3.

In the following scenario for each of the three days, we cancel all flights in F_{AF} , and assume passengers P_{AF} are available for rebooking two hours after the planned departure time of their cancelled flight. We use the Passenger Delay Calculator to calculate the resulting total passenger delay, for each day, and for each of the following scenarios:

1. Delay to passengers P_{AF} with all F_{AF} cancelled, and instantaneous rebooking.
2. Delay to passengers P_{AF} with all F_{AF} cancelled and passengers P_{AF} available for rebooking two hours *after* the planned departure time of the flight in set F_{AF} .

3. Baseline, as-flown 2007 delay to passengers P_{AF} .

We then calculate two ratios for each day. The first is the ratio of total passenger delay with two-hour rebooking delay (item two, above), to total passenger delay with instantaneous rebooking (item one, above, computed by setting the rebooking time equal to the scheduled flight departure time, as analyzed in Chapters 3 and 4). The second ratio is the total P_{AF} delay with an assumed available rebooking time of two hours after the scheduled flight departure time (item two, above) to the passengers' P_{AF} baseline 2007 delay (item three, above). Column 2 of Table 5.1 contains Ratio 1 for each day, while Column 3 of Table 5.1 contains Ratio 2 values. By calculating these multipliers, we quantify by how much a more realistic rebooking delay affects our total passenger delay results.

Date	Multiplier effect, from all-cancelled with instantaneous rebooking	Multiplier effect, from baseline 2007 delays)
July 27	1.0365	2.8893
January 17	1.0988	1.8966
August 23	1.0183	1.6357

Table 5.1: Comparing delay multipliers resulting from two-hour rebooking-delay

As shown in Column 2 of Table 5.1, delaying rebooking for disrupted passengers by two hours results in a small increase in passenger delay, ranging between 1% - 9% above that resulting from instantaneous rebooking. The results of Column 3 then suggest that the delays to passengers P_{AF} arise primarily from flight cancellations and not from rebooking delays. Again, this result shows that passenger delays resulting from cancellations is due primarily to schedule effects (see Section 3.2), rather than rebooking, which is more dynamic and controllable by airlines.

5.3 Cancelling return flight pairs

In this section, we consider the impact of cancelling pairs of flights to maintain aircraft balance and operational feasibility, rather than cancelling only single flight legs in F_{AF} .

If a flight from airport A to airport B is cancelled, other flights downstream must also be cancelled to maintain aircraft flow balance. Suppose an aircraft is scheduled to fly from BOS to SFO, SFO to LAX, and LAX to BOS. If the flight from BOS to SFO is cancelled (and thus the aircraft stays in Boston), either the flights from SFO to LAX and LAX to BOS must be cancelled, or an aircraft must be *ferried* (flown with no passengers) to SFO in order to complete the subsequent flights. Rosenberger, Johnson and Nemhauser (2004) note that airlines usually choose to cancel additional flight legs to create a cancellation cycle, which preserves aircraft balance, rather than ferry an aircraft. We therefore will consider cancellation cycles including two flights; for example, from A to B and from B to A.

We select days 1, 5 and 7 from Table A.3 for analysis. We develop a Cancelled Pairs Scenario (CPS), defined as follows.

For each flight $f_i \in F_{AF}$ from an airport A to an B, we define a “return” flight as a flight from airport B to airport A, denoted f_{BA} , such that:

1. Flight f_{BA} departs no earlier than the planned arrival time of flight f_i ;
2. Flight f_{BA} is the flight departing the earliest after the scheduled arrival time of flight f_i ;
3. Flight f_{BA} is operated by an aircraft of the same type (or close to the same type) as f_i , implying that it has a seating capacity within five seats of that operating flight f_i ; and
4. Flight f_{BA} is operated by the same carrier as flight f_i .

We denote as F_{CPS} the set of return flights that correspond to the flights in F_{AF} . Then, in our Cancelled Pairs Scenario, we cancel all flights in F_{CPS} and in F_{AF} .

In this analysis, we calculate the years of delay for passengers on the flights F_{CPS} , and calculate the percent increase in delay, equal to delay for passengers on F_{CPS} plus the delay to passengers P_{AF} with all flights in F_{CPS} and F_{AF} cancelled, divided by delay for passengers P_{AF} with all flights in F_{AF} cancelled. We present our results in Table 5.2. We note that the number of flights in F_{CPS} may not be exactly twice

that of the number of flights in F_{AF} because in some cases, a return flight f_{BA} was not found within 24 hours of the scheduled arrival time of a flight $f_i \in F_{AF}$. In these cases, we did not match a flight in F_{AF} with a return flight in F_{CPS} .

Date	Number of flights in F_{AF}	Number of flights in F_{CPS}	% increase in passenger delay
July 27	76	75	64%
January 17	32	30	32%
August 23	11	10	89%

Table 5.2: Percent increase in delay when return flight pairs are cancelled

We see that for all days considered, delays increase when return pair flights are cancelled, as expected. To relate this finding back to the impact of cancellations resulting from the Tarmac Delay Rule, we revisit the five flights in set F_{AF} on July 27 which had a multiplier of less than one (see Section 4.2 and Figure 4-2). We consider the return flights in F_{CPS} corresponding to each of these flights, and pair each of the five flights in F_{AF} with their corresponding return flight in F_{CPS} , thus resulting in five flight pairs. We calculate baseline 2007 delay for passengers on each pair of flights, and the delay for the same passengers when both of the flights in the pair are cancelled. We then calculate the ratio of these two numbers, resulting in a *one-pair multiplier* for each of the five pairs. We found that every one of the one-pair multipliers for the five pairs was higher than one, suggesting that when operationally feasible cancellations are implemented, while passengers P_{AF} on flights in F_{AF} with multipliers less than one incurred less than their baseline delay, the cancellation of paired flights F_{CPS} results additional delays that cause the time-savings benefit to be negated.

Chapter 6

Conclusion and Future Directions

This work focuses on capturing the effects to passengers of the 2010 three-hour Tarmac Delay Rule. In our approach, we quantify passenger delays resulting from the Rule by applying the Passenger Delay Calculator (Bratu and Barnhart, 2005) to a set of hypothetical scenarios. We use data from 2007, a year in which the Rule was not in place, identify operated flights with taxi-out delays of at least three hours, and consider them candidates for cancellation. This is based on the U.S. G.A.O. Report (2011), which linked the implementation of the Rule to a higher likelihood of cancellation for tarmac-delayed flights. We compare the delays realized in 2007 to those we estimate assuming that all or selected flights are cancelled due to the Rule. This chapter is divided into two sections: first, we review our major methodologies and findings from this thesis, and then we discuss directions for future work.

6.1 Review of methodologies and findings

In our work, we first assume that all operated flights with more than three hours of taxi-out delay in 2007 are cancelled, and passengers are rebooked onto their final destinations on the available options that arrive first at their destination, given carrier constraints. We do not consider flights that are diverted, nor do we consider flights that incur a lengthy tarmac delay and are subsequently cancelled and returned to the gate.

In Chapter 2, we describe the procedure we use to calculate passenger delays. Passenger itineraries, flight schedules (realized and planned), and aircraft seating capacity data are joined in an Oracle SQL database, and passenger delays are calculated using the Passenger Delay Calculator, an algorithm formalized by Bratu and Barnhart (with extensions from Barnhart, Fearing, and Vaze), and written into Java code.

In Chapter 3, we simulate various cancellation scenarios in which some or all of the 2007 flights that taxied-out for at least three hours are cancelled, and calculate resultant passenger delays. In our first set of analyses, we examine the correlations between percentage of flights cancelled and passenger delay. We find that passenger delay increases linearly with the percentage of lengthy tarmac-delayed flights that are cancelled. In our next analysis, we compare the impacts of load factor and schedule on passenger delays by measuring the difference in delay between the *actual* load factor case and the *unlimited* load factor case. We find that in all days studied, the impact of schedule is greater than that of load factor, accounting for 60-70% of additional delays for the days we analyzed. Finally, we analyze the impacts that flight cancellations have on passengers on other (non-cancelled) flights in the departure queue. We allow these other flights to move up to the earlier wheels-off time slots that are made available by the cancelling of flights. We find that depending on the time of day of disruptions, the benefits of allowing non-cancelled flights to take-off earlier can be outweighed by the increased delays to passengers on cancelled tarmac-delayed flights.

In Chapter 4, we introduce a method of comparing the delay impacts of cancellations on different days and different flights. We define a multiplier for each flight leg (and each day) by dividing the delay incurred by passengers on a cancelled flight (or set of cancelled flights) by the baseline delay to the same flight(s), as operated in 2007. We then use the multiplier to identify flights, and find that flights with high multipliers: 1) tend to be later in the day, usually an airline's last flight of the day to its destination; and 2) often incur just over three hours of tarmac delay on taxi-out. The flights with multipliers below one (for which the passengers experience less delay if the flight is cancelled) tend to be flights: 1) with large amounts of pre-enplanement

delay; 2) with taxi-out delays of much more than three hours; and 3) that are scheduled earlier in the day. Using this information about flight characteristics with high and low multipliers, we selectively cancel flights and find that cancelling only flights with low multipliers (allowing flights with high multipliers to operate) results in the least amount of expected passenger delay, less than the baseline 2007 delay, and significantly less than the delay incurred if *all* flights with long tarmac delays are cancelled.

In Chapter 5, we explore the impacts of our modeling assumptions and simplifications. We begin by changing the threshold of maximum taxi-out time before cancellation, and measure passenger delay at these lower and higher thresholds. We find that the multiplier is highest when we decrease the cancellation threshold to one hour of taxi-out time, and the multiplier decreases as the cancellation threshold time increases. This suggests that operationally, with the Rule requiring aircraft to return to the gate *within* three hours, passenger delay will be even greater than we estimate using a three-hour tarmac delay threshold. Next, we explore our assumption that passengers are available for rebooking, following a cancellation, immediately after the planned departure time of their cancelled flights. We delay the rebooking time for passengers to reflect a more realistic operational scenario, and measure the increase in their delay as a result of the rebooking delay. As expected, we show that delay increases as rebooking time increases, further suggesting that our estimates of passenger delays underestimate the potential delays resulting from the Rule. Finally, we examine our simplifying assumption that a single flight leg is cancelled, violating conservation of flow. We instead cancel the tarmac-delayed flight and a *return* flight to ensure conservation of flow by aircraft type. Our results show an increase in passenger delay when pairs of flights are cancelled, even for pairs of flights in which the tarmac-delayed flight had a multiplier of less than one when it was cancelled individually.

We find consistently that passengers incur longer delays when flights are cancelled to avoid tarmac-delay fines than when they are allowed to depart after lengthy tarmac delays. Through our analysis in this work we have observed that applying the

Tarmac Delay Rule uniformly across all flights taxiing-out longer than three hours does not result in minimizing passenger delay. Instead, this approach results in a few passengers incurring less delay than if their flights departed after at least three hours of tarmac delay, and most passengers incurring more delay when long tarmac-delayed flights are not operated. Passengers on some flights experience very long delays as a result of cancellations, as much as six or seven times their baseline delay in some cases. These results suggest that the Rule should be applied selectively if the goal is to minimize overall passenger delay on flights with long taxi-out delays. When we consider passengers *not* on flights with long taxi-out delays, the cancellations of flights resulting from the Rule has a positive impact, as congestion is decreased and flights are allowed to depart earlier. However, this is an incomplete picture of the Rule impacts. Cancellations as a result of the Rule impact not only the passengers on tarmac-delayed flights, but also downstream passengers whose flights are cancelled to maintain aircraft flow balance given cancellations of tarmac-delayed flights.

We hope that the findings of this work provide policy makers with insights that can be used for future policy decisions regarding the Tarmac Delay Rule. In the next section, we conclude with a look forward, describing possible future extensions of this work.

6.2 Directions for future work

This work has focused strictly on the effects to passengers as a result of the Rule. A further course of inquiry is an analysis of airline operational changes that result from the Rule. Specifically, an interesting question is whether or not the Rule has led to an increase in the number of flight cancellations, and if so, to what extent. Answering this question would build upon the work of the U.S. G.A.O. Report (2011), and would require a longitudinal investigation of weather patterns and cancellations, as varying weather conditions over the years strongly influence cancellation rates. Finally, a study of the Rule's impact on airlines with different network structures, for example, regional versus legacy carriers, could prove insightful. It is known that smaller aircraft

on short-haul flights are assigned disproportionately more ground delay time (Vossen et al., 2003). Thus, it would be worthwhile to consider if regional airlines incur proportionally more delay as a result of the Rule, thereby introducing inequitable impacts to various airlines.

Appendix A

Tables

IATA Code	Airport Name
ATL	HartsfieldJackson Atlanta International
BOS	Boston Logan International
DFW	Dallas/Fort Worth International Airport
EWB	Newark Liberty International
IAD	Washington Dulles International
JFK	John F Kennedy International
LAX	Los Angeles International
OAK	Oakland International
ORD	Chicago O'Hare International
PHL	Philadelphia International
SFO	San Francisco International

Table A.1: IATA airport codes and names referenced in this work

IATA Code	Carrier Name
9E	Pinnacle Airlines
AA	American Airlines
AQ	Aloha Airlines
AS	Alaska Airlines
B6	JetBlue Airways
CO	Continental Airlines
DL	Delta Air Lines
EV	Atlantic Southeast Airlines
F9	Frontier Airlines
FL	AirTran Airways
HA	Hawaiian Airlines
MQ	American Eagle Airlines
NW	Northwest Airlines
OH	Midwest Airlines
OO	SkyWest Airlines
UA	United Airlines
US	US Airways
WN	Southwest Airlines
XE	ExpressJet Airlines
YV	Mesa Airlines

Table A.2: 2007 ASQP-reporting carriers, and associated codes

Identifier	Day	Number of flights in F_{AF}	Number of pas- sengers P_{AF}	Average monthly load factor (BTS)
1	July 27	76	8,730	81.5%
2	June 26	71	7,010	82%
3	February 14	64	6,881	71.7%
4	June 21	43	5,295	82%
5	January 17	32	1,963	67.4%
6	January 21	20	1,267	67.4%
7	August 23	11	1,006	79.9%

Table A.3: Information on days in 2007 selected for analysis

Appendix B

Figures




Figure B-1: One-flight multipliers for flights in F_{AF} on January 17, departing DFW




Figure B-2: One-flight multipliers for flights in F_{AF} on June 26, departing ORD




Figure B-3: One-flight multipliers for flights in F_{AF} on May 31, departing EWR




Figure B-4: One-flight multipliers for flights in F_{AF} on July 27, departing JFK




Figure B-5: One-flight multipliers for flights in F_{AF} on July 10, departing ATL




Figure B-6: One-flight multipliers for flights in F_{AF} on June 19, departing PHL

Bibliography

- [1] Barnhart, C., Fearing, D., & Vaze, V. (2013). Modeling passenger travel and delays in the National Air Transportation System. *Operations Research*, Under Review.
- [2] Bratu, S., & Barnhart, C. (2005). An analysis of passenger delays using flight operations and passenger booking data. *Air Traffic Control Quarterly*, Number 1, Volume 13.
- [3] Sherry, L., Wang, D., & Donohue, G. (2007). Air travel consumer protection: metric for passenger on-time performance. *Transportation Research Record*, Issue 1.
- [4] Tien, S., Ball, M., & Subramanian, B. (2008). Constructing a passenger trip delay metric: an aggregate level approach. *NEXTOR Working Paper*.
- [5] Coldren, G., & Koppelman, F. (2005). Modeling the competition among air-travel itinerary shares: GEV model development. *Transportation Research Part A*, Number 4, Volume 39.
- [6] Rosenberger, J., Johnson, E., & Nemhauser, G. (2004). A Robust Fleet Assignment Model with Hub Isolation and Short Cycles. *Transportation Science*, Number 3, Volume 38.
- [7] United States Government Accountability Office (2011). Airline Passenger Protections: More Data and Analysis Needed to Understand Effects of Flight Delays. *Report to Congressional Requesters*, Report Number GAO-11-733.

- [8] Vossen, T., Ball, M.O., Hoffman, R. & Wambsganss, M. (2003). A General Approach to Equity in Traffic Flow Management and its Application to Mitigating Exemption Bias in Ground Delay Programs. *Air Traffic Control Quarterly*, Issue 4, Volume 11.
- [9] Bureau of Transportation Statistics (2006-2010). Airline On-Time Performance Data. Accessed April 22, 2013.
http://www.transtats.bts.gov/Tables.asp?DB_ID=120&DB_Name=Airline%20On-Time%20Performance%20Data&DB_Short_Name=On-Time
- [10] Bureau of Transportation Statistics (2007, 2012). Form 41 Traffic T-100 Domestic Segment Data. Accessed April 22, 2013.
http://www.transtats.bts.gov/Fields.asp?Table_ID=259
- [11] Bureau of Transportation Statistics (2007). Form 41, Schedule B-43. Accessed April 22, 2013.
http://2bts.rita.dot.gov/programs/airline_information/schedule_b43/2007/
- [12] Bureau of Transportation Statistics (2007). Airline Origin and Destination Survey (DB1BCoupon). Accessed April 22, 2013.
http://www.transtats.bts.gov/Fields.asp?Table_ID=289
- [13] Bureau of Transportation Statistics (2007). Airline Origin and Destination Survey (DB1BMarket). Accessed April 22, 2013.
http://www.transtats.bts.gov/Fields.asp?Table_ID=247
- [14] Bureau of Transportation Statistics (2007). Airline Origin and Destination Survey (DB1BTicket). Accessed April 22, 2013.
http://www.transtats.bts.gov/Fields.asp?Table_ID=272