



**TARMAC DELAY POLICIES:
A PASSENGER-CENTRIC ANALYSIS**

Chiwei Yan

Operations Research Center, Massachusetts Institute of Technology, USA

Vikrant Vaze

Thayer School of Engineering, Dartmouth College, USA

Allison Vanderboll

*Department of Civil and Environmental Engineering,
Massachusetts Institute of Technology, USA*

Cynthia Barnhart

*International Center for Air Transportation,
Massachusetts Institute of Technology, USA*

Report No. ICAT-2016-01

January 2016

MIT International Center for Air Transportation (ICAT)
Department of Aeronautics & Astronautics
Massachusetts Institute of Technology
Cambridge, MA 02139 USA

Tarmac Delay Policies: A Passenger-Centric Analysis

Chiwei Yan^{a,1}, Vikrant Vaze^b, Allison Vanderboll^c and Cynthia Barnhart^a

^aOperations Research Center, Massachusetts Institute of Technology, USA

^bThayer School of Engineering, Dartmouth College, USA

^cDepartment of Civil and Environmental Engineering, Massachusetts Institute of Technology, USA

Abstract: In this paper, we analyze the effectiveness of the 2010 Tarmac Delay Rule from a passenger-centric point of view. The Tarmac Delay 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 of the departure airport; and that an opportunity for passengers to deplane must occur no later than three hours after the touchdown at the arrival airport. The Tarmac Delay Rule aims to protect enplaned passengers on commercial aircraft from excessively long delays on the tarmac upon taxi-out or taxi-in, and monetarily penalizes airlines that violate the stipulated three-hour tarmac time limit. Comparing the actual flight schedule and delay data after the Tarmac Delay Rule was in effect with that before, we find that the Rule has been highly effective in reducing the frequency of occurrence of long tarmac times. However, another significant effect of the rule has been the rise in flight cancellation rates. Cancellations result in passengers requiring rebooking, and often lead to extensive delay in reaching their final destinations. Using an algorithm to estimate 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 Tarmac Delay Rule in effect for that same year. Our delay estimates are calculated using U.S. Department of Transportation data from 2007. Through our results and several sensitivity analyses, we show that the overall impact of the current Tarmac Delay Rule is a significant increase in passenger delays, especially for passengers scheduled to travel on the flights which are at risk of long tarmac delays. We evaluate the impacts on passengers of a number of rule variations, including changes to the maximum time on the tarmac, and variations in that maximum by time-of-day. Through extensive scenario analyses, we conclude that a better balance between the conflicting objectives of reducing the frequency of long tarmac times and reducing total passenger delays can be achieved through a modified version of the existing rule. This modified version involves increasing the tarmac time limit to 3.5 hours and only applying the rule to flights with planned departure times before 5pm. Finally, in order to implement the Rule more effectively, we suggest the tarmac time limit to be defined in terms of the time when the aircraft begin returning to the gate instead of being defined in terms of the time when passengers are allowed to deplane.

Keywords: Aviation; Tarmac delay rule; Passenger disruption and delay

¹ Corresponding author
77 Massachusetts Avenue, E40-130, Cambridge, MA 02139, USA
Email address: chiwei@mit.edu (Chiwei Yan)

1. Introduction

In 2007, flight delay levels in the U.S. were very high in general. But on February 14, 2007, in the midst of what came to be known as the “Valentine’s Day Blizzard”, passengers on flights originating at New York City’s John F. Kennedy International Airport (JFK) suffered extremely long delays. Some of these passengers endured as much as seven hours of delay on their aircraft, often without access to food. Boarded and pushed back 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 times, defined as those lasting more than three hours, were fairly common in 2007. That year, there were 1,654 instances of three hour or longer *taxi-out times*, defined as the period of time between cabin door closure and aircraft lift-off. In this paper, we will use the terms *tarmac time* and *taxi-out time* interchangeably. Moreover, the actual number of instances with taxi-out times greater than or equal to three hours was much higher, as the 1,654 count does not include the flights that pushed back from their gates, joined the departure queue, later were cancelled, and then taxied back to a gate to deplane. Additionally, if we include flights with intermediate taxi-out times, that is, those between one and three hours, the number increases dramatically. As shown in Table 1, using data from the Bureau of Transportation Statistics (BTS) (2007), the number of flights with taxi-out times between one and three hours was approximately 50 times the number of flights with taxi-out times of three hours or longer. Note that, for reasons explained later in this section, we will focus our analysis on taxi-out times (rather than taxi-in times).

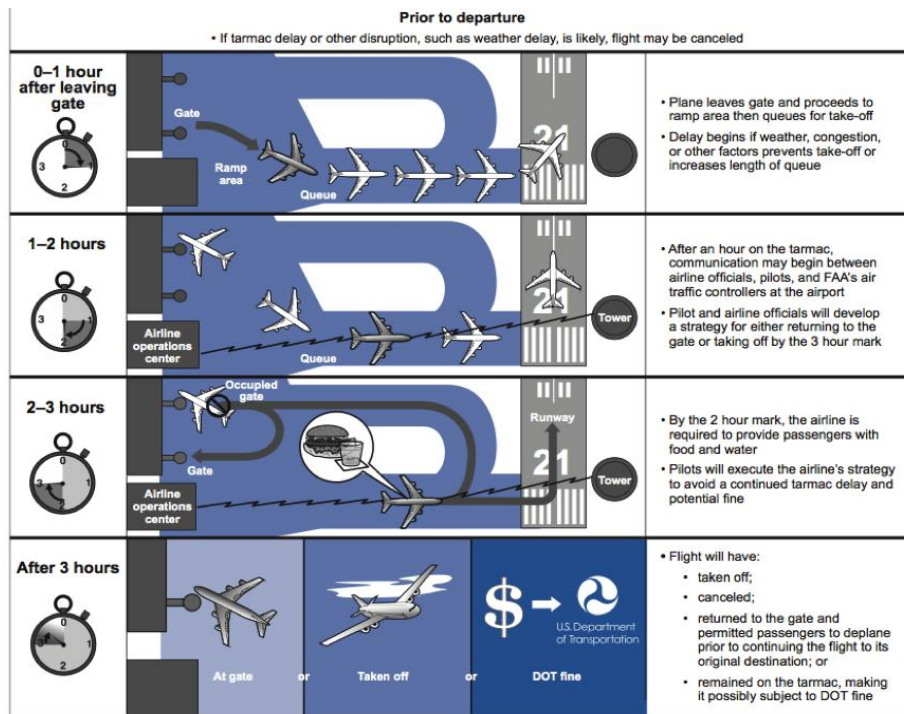
Length of taxi-out times (minutes)	Number of occurrences
60 to 119	75,833
120 to 179	7,507
180 to 239	1,370
240 to 299	239
300 to 359	36
360 or greater	9

Table 1: Non-cancelled flights (including diversions) that experienced lengthy tarmac times during taxi-out in 2007, as reported by BTS

1.1 The Tarmac Delay Rule and Airline Response

Following these events, amid pressure from consumer advocacy groups, the U.S. Department of Transportation announced a policy known as the Tarmac Delay Rule (the “Rule”) on December 21, 2009, which 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 at the departure airport; and that an opportunity for passengers to deplane must occur no later than three hours after touchdown at the arrival airport. There are two exemptions: 1) if the pilot determines that moving from the departure queue or deplaning passengers would constitute a safety or security risk; and 2) if local air traffic control decides that airport operations would be significantly disrupted by the

delayed aircraft returning to a gate or deplaning area. 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 Rule requires that carriers and individual airports develop a plan that is mutually agreeable for deplanement in case a violation is imminent. In case of flights delayed at the departure airport, 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 Rule; that is, the aircraft cannot begin to head back to a gate or other deplanement area at the end of the three-hour period. Instead, passengers wanting to be deplaned 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 flights (operated by any carrier), originating or landing at U.S. airports (in this latter case the limit on time on the tarmac is four hours). Flights operated by aircraft with less than 30 seats are exempt. The Rule's penalty to the airlines for non-compliance is a fine of up to \$27,500 per passenger. In reality, the fine level varies from case to case. As of Jan 15th 2015, the Department of Transportation had issued 17 orders assessing \$5.24 million dollars in total for violations of the Rule (U.S. Department of Transportation, 2015). The largest penalty was on January 2nd into January 3rd, 2014, when the Department of Transportation fined Southwest \$1.6 million dollars for 16 flights violating the rule. Shown in Figure 1, taken from the U.S. Government Accountability Office (GAO) report (2011), are the various points in the taxi-out process when decisions must be made.



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

Figure 1: Schematic of airline decision-making when faced with a long taxi-out delay (GAO, 2011)

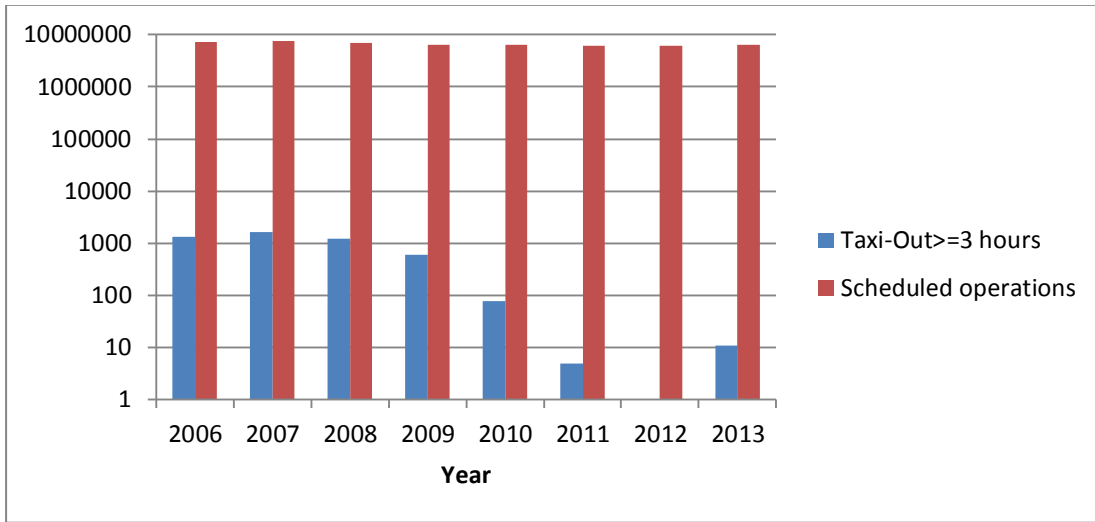


Figure 2: 2006-2013 Number of operated flights with taxi-out time exceeding three hours, and total number of scheduled operations

Since the announcement and implementation of the Rule, frequency of taxi-out times of three hours or longer has significantly decreased, as depicted in Figure 2, using data from BTS (2006-2013). We compare the annual average number of operated flights with tarmac time of three hours or longer, and the annual average number of scheduled operations, from 2006 to 2008, the three years just prior to the announcement of the Rule, with the same numbers for 2011 to 2013, the first three years after the implementation of the Rule. The annual average number of operated flights with three hours or longer tarmac time decreased by 99.6% from the pre-Rule period of 2006-2008 (1408.3 flights) to the post-Rule period of 2011-2013 (5.7 flights). The annual average number of scheduled operations, however, decreased only by 14.1% (from 7.2 million to 6.2 million flights). This data suggests that the Rule has been highly effective in keeping passengers off the tarmac for lengthy periods of time during the taxiing-out operation.

In order to control for the difference in the number of scheduled operations across this time period (and thus to indirectly control for airport congestion), we compare the 2013 numbers with the 2009 numbers. The Rule did not get implemented until 2010 and was not announced until the last 10 days of 2009, while by the start of year 2013, it had been over two years since the implementation of the Rule. The total number of scheduled operations was almost the same (only 1.25% different) for 2009 and 2013. From 2009 to 2013, the capacities of all major airports in the U.S. remained virtually unchanged and average flight delays actually increased by about 10% from 2009 (11.6 min) to 2013 (12.7 min) (BTS 2009, 2013). Yet, the number of operated flights with three hours or longer of tarmac time decreased by 98.2% from 2009 (604 flights) to 2013 (11 flights).

While the Rule seems effective in keeping passengers from experiencing lengthy delays on the tarmac during the taxiing-out operation, we aim to explore other consequences of the Rule in this paper. The GAO study (2011) interviewed airline officials who stated that airlines changed their cancellation criteria in response to the Rule. In order to test this qualitative finding, the authors of the aforementioned GAO (2011) report used available data on tarmac delays before and after the implementation of the Rule, and developed two regression models to evaluate whether cancellation rates increased after the Rule went into effect. The regression models controlled for other factors that are related to cancellations. These other factors included level of airport congestion, origin/destination weather conditions, ground delay programs, airport on-time performance, size of airline, airport status as a hub, passengers per flight, route distance, day-of-week, and scheduled departure hour. Their results suggested that after the implementation of the Rule, flights experiencing any level of taxi-out time were more likely to be cancelled than before the Rule implementation. In Table 2, we present how the likelihood of cancellation rapidly increases as the duration of taxi-out time increases. Other studies (e.g., the U.S. Department of Transportation (2014) report, Marks and Jenkins (2010), and the U.S. Department of Transportation (2009) report) analyzing the effects of the Rule have also concluded that the Rule has increased, to various degrees, the cancellation probability for flights with long taxi-out times.

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

Table 2: U.S. GAO-reported change in likelihood of flight cancellation, by taxi-out time (GAO, 2011)

1.2 The Rule’s Impact on Passenger Delay

Motivated by the observation that the Rule has led to an increase in the likelihood of flight cancellations and in consequent passenger disruptions, in this paper we quantify the impact of the Rule on passenger delays for those aboard tarmac-delayed flights in the U.S. National Airspace System (NAS). Passenger delay is defined as the actual arrival time of the passenger’s actual itinerary at the passenger’s final destination minus the scheduled arrival time of the passenger’s scheduled itinerary at the passenger’s final destination. Passenger delay is differentiated from flight delay as the former also accounts for passenger disruptions, resulting from flight cancellations, diversions, and passenger *misconnections* (a passenger is assumed to misconnect if his/her first flight arrives less than 15 minutes before the actual departure of the 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 time misses his/her connecting flight leg, and has to wait, say three more hours, for the next flight with an available seat to his/her final destination. This situation results in a passenger 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 or diverted flights) is rebooked in order to reach his/her final destination. Note that some passengers may choose not to get rebooked and thus to abandon their air travel plans due to a flight disruption. However, due to lack of data on the percentage of such passengers, we don't explicitly incorporate this effect in our analysis.

A simple comparison of the passenger delay in a year before the Rule was implemented to the passenger delay in a year after does not represent a valid assessment of the Rule's impact, since such direct comparison would fail to properly control for a number of factors, including year-to-year variations in airline schedules, congestion levels, passenger demand fluctuations, capacity changes, and weather differences. In fact, passenger delay calculation itself presents a challenge primarily due to lack of available data. We describe in Section 2 the approach we used for calculating passenger delay. To understand the impacts of the Rule on passengers, we experiment with a simulation using pre-Rule operations as follows. First we identify flights from year 2007 with significant taxi-out times; next, we create a number of scenarios in which some or all of these flights are cancelled; and finally, we 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, etc. However, given the lack of granularity in our data about individual passengers and their value of time, we focus on one metric – passenger delays – which we can estimate with some degree of certainty.

In selecting the set of flights for our analysis, we focus on those with taxi-out delays, instead of taxi-in delays, because airlines have a higher degree of control over the operational actions taken when a taxi-out delay occurs. For example, when a taxi-out delay occurs, a decision can be made to return to a gate, especially when the lift-off is likely to get substantially delayed. For a taxi-in delay, however, the aircraft has only the option to wait for a gate; it can't take off again and return to its origin airport. Moreover, the number of flights with taxi-in times of three hours or longer is far fewer than the number of flights with taxi-out times of three hours or longer (see Table 3, with data from BTS 2006-2010).

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

Table 3: Lengthy taxi-out and taxi-in incidents, 2006-2010

Finally, we selected year 2007 as our representative pre-Rule operational scenario, because it had the highest number of lengthy (three hours or longer) taxi-out incidents of any year from 2006-2010. Additionally, that year featured several notable lengthy tarmac delays, such as the Valentine's Day Blizzard described previously, that prompted consumer protection groups to lobby Congress for regulations that led to the Tarmac Delay Rule.

1.3 Contributions and Outline

In this paper we quantify the delays to passengers due to cancellations that could result from the Tarmac Delay Rule. We apply an existing methodology, the Passenger Delay Calculator (Barnhart et al., 2014), 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 restrictions defining the Rule. Ours is the first research study that analyzes the effectiveness of the Tarmac Delay Rule from the perspective of the airline passengers, the very group of stakeholders whose interests the Rule is supposed to protect to begin with. A major contribution of this research is the quantification of the extent to which the Rule is effective, and the ways in which it is costly to passengers, those on tarmac-delayed flights and those on flights elsewhere in the NAS. Furthermore, the general framework of that our study lays out can be used for analyzing, from a passenger-centric perspective, other important policy questions which are directly or indirectly related to passenger delays.

Our results provide policy-makers with insights to inform future revisions of the Rule. Our main result is that, while the three-hour tarmac delay rule (in its current form) effectively decreases tarmac delays, especially the extremely long tarmac delays, each passenger-minute of tarmac time saving is achieved at the cost of an increase of approximately three passenger-minutes in total passenger delays. Our methodology and results have been found to be robust under a variety of sensitivity analyses. However, we find that by judiciously imposing certain modified versions of the rule, passengers can enjoy the benefits of reduction in lengthy and inconvenient waiting times on the airport tarmacs, with the total passenger delay increase being less than half the total amount of time saved on the tarmac. Additionally, in order to implement the Rule more effectively, we also suggest that the tarmac time limit should be defined in terms of the time when the aircraft should start to return to the gate instead of being defined in terms of the time when passengers are allowed to deplane.

An outline of the rest of the paper is as follows. In Section 2, we describe the procedure used 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 Section 3, we estimate the passenger delays that would have resulted had the Rule been in effect in 2007, and compare this estimated delay to the delay experienced by passengers in the absence of the Rule. We treat the latter as our pre-Rule baseline. We also perform sensitivity analyses to understand the impact of our modeling assumptions and simplifications on our delay estimates. In Section 4, we identify the characteristics of flights that are most severely impacted and likely to have the greatest increase in passenger delays as a

result of the Rule. We use this information to explore some revised tarmac delay rule policies and compare the resultant delays. We conclude in Section 5 by providing a summary of the findings of this research and detail future research topics that might be explored as more data becomes available.

2. Data and Methodology

The goal of this research is to quantify the impacts on passengers as a result of the Tarmac Delay Rule. We do so by identifying flights that incurred lengthy taxi-out times in 2007, and use them to perform a variety of scenario analyses. The metrics that we obtain as our results are based on passenger delays and tarmac times.

2.1 Literature Review

The methodology used in this paper builds primarily upon the work of Bratu and Barnhart (2005), and Barnhart et al. (2014). We use the Passenger Delay Calculator (PDC) algorithm, originally proposed by Bratu and Barnhart (2005), which calculates passenger delay given inputs of flight schedules (planned and actual), planned itineraries of passengers, and aircraft seating capacity data. Sherry et al. (2007) also calculate passenger delays, but treat all passenger itineraries as non-stops. This approach is not applicable for our purposes, because we wish to explicitly incorporate delay due to missed flight connections into our calculation of passenger delay. Tien et al. (2008) also provide an algorithm for calculating passenger delay, but their approach provides only an aggregate measure of passenger delay based on an aggregate cancellation rate and on aggregate itinerary characteristics, and as such, is unsuitable for our task at hand. Sherry et al. (2010) develop an algorithm to allocate passengers onto itineraries based on publicly available aggregate data to get disaggregate passenger itinerary flows. However, their approach doesn't incorporate passenger preferences, that is, it doesn't incorporate the fact that certain itineraries might be more attractive to passengers than others due to better departure / arrival time and / or day-of-week, and / or more reasonable connection time.

Barnhart et al. (2014) estimated disaggregate passenger itinerary flows, using publicly available aggregate data by training their model on one quarter of booking information from a major U.S. carrier in 2007. They used a multinomial logit modeling approach to disaggregate the itinerary flows by accounting for passenger preferences for time-of-day, day-of-week, connection times, etc. Their work also includes a model for estimating seating capacities of aircraft whose tail numbers are not listed in Schedule B-43 (see Sub-section 2.2). In this paper, we use the following three aspects of their study: 1) estimated disaggregate passenger itinerary flows; 2) the aircraft seating capacity values; and 3) their extended version of the Passenger Delay Calculator (PDC) algorithm. Figure 3 depicts a step-by-step schematic of the PDC algorithm (Barnhart et al., 2014). In Step 1, inputs to the algorithm include planned passenger itineraries and flight schedules, cancellations, and flight delay data. Given this, all passengers are assigned a binary identifier of disrupted or non-disrupted.

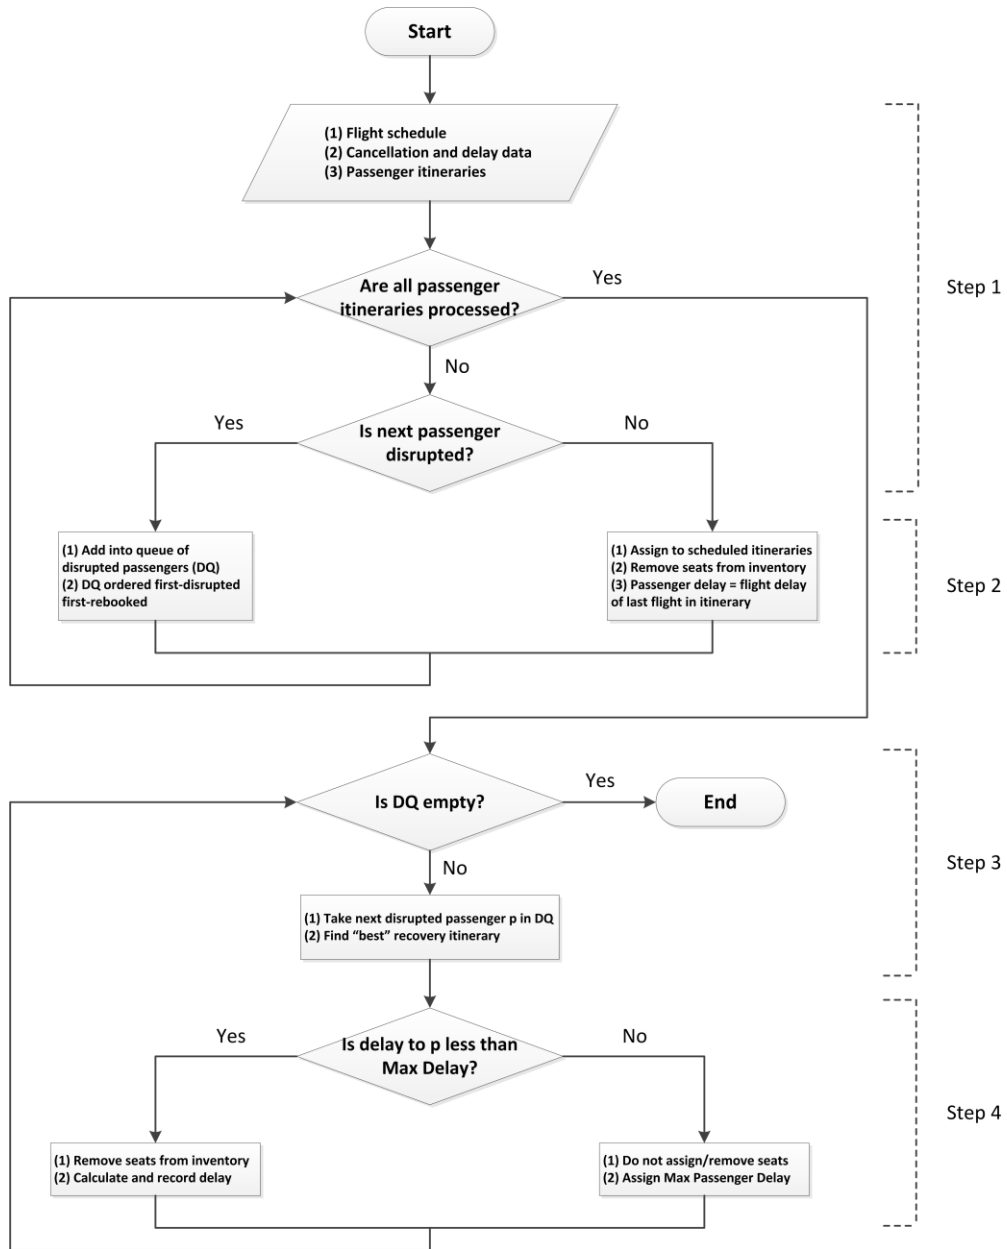


Figure 3: Passenger Delay Calculation Flowchart (Bratu and Barnhart, 2005; Barnhart et al., 2014)

In Step 2, each passenger who is not disrupted is assigned to his or 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 his/her flight, while a non-disrupted passenger on a connecting itinerary is assigned passenger delay equal to the flight delay of the last flight in his/her itinerary. In the case that a passenger arrives at his/her final destination before that passenger's scheduled arrival time, the

passenger delay for that passenger is set to zero. This can occur when a flight flies faster than scheduled, or due to slack in block time, or if the passenger is rebooked onto an itinerary that arrives earlier than that passenger's planned itinerary. Disrupted passengers are placed into the *Disruption Queue* (DQ) and the queue is processed in a first-disrupted, first-rebooked fashion. We choose this policy because we do not have access to detailed information about passengers' airline frequent flier status or fare class, which could allow us to follow 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 the lack of detailed information about airline frequent flier status, fare classes, 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 carriers to the ones 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 (e.g., American Eagle). If no recovery itinerary for p is found on the same or related carriers, 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 or her final destination with delay not exceeding eight hours (for passengers disrupted between 5:00am and 4:59pm), or 16 hours (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 value equal to the difference between the scheduled arrival time of the last flight on p 's planned itinerary and the actual arrival time of the last flight on p 's recovery itinerary. 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). These differences in maximum delay values reflect the difficulty in rebooking later in the day, often 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 the next 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 is constructed and passenger delays are estimated. The bulk of this data is publicly available from BTS. The data inputs to the PDC include:

1. **Airline On-Time Performance (AOTP) database:** This database includes for each flight, scheduled and actual flight departure and arrival locations and times, taxi-out and taxi-in times, wheels-off and

wheels-on times, operating carrier, and flight number. This information is reported monthly by air carriers in the United States that correspond to more than one percent of domestic scheduled passenger revenues. This data is available for all flights operated by these carriers at airports in the 48 U.S. contiguous states. In 2007, this included 20 unique carriers². The 2007 data does not report the airport to which flights were diverted, nor does it include the taxi-out time for any flight that may have departed the gate but was subsequently cancelled prior to take-off.

2. **T-100 Domestic Segment (T-100) database:** 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 flown on each carrier, for each non-stop flight segment and for each aircraft type, aggregated monthly. Thus, a passenger flying OAK-IAD-BOS (Oakland-Washington DC-Boston) on a given carrier and aircraft type(s) is added to the count of both the OAK-IAD and IAD-BOS flight segments. The passenger counts are used as inputs to the multinomial logit passenger itinerary flow model presented in Barnhart et al. (2014).
3. **Form 41 Schedule B-43 Inventory database, and Enhanced Traffic Management System (ETMS) database:** The Form 41 Schedule B-43 database includes aircraft seating capacities (specified by tail number), which are matched to the AOTP database using tail numbers. This allows us to estimate the available seats for each flight reported in the AOTP. About 75% of the flights in AOTP can be matched to an entry in Schedule B-43 dataset. The remaining seating capacities 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 AOTP. The remainder is obtained through an algorithm presented in Barnhart et al. (2014) using the T-100 Domestic Segment database.
4. **Airline Origin and Destination Survey (DB1B) database:** This dataset, aggregated quarterly, is a 10% sample of ticketed passengers on carriers reporting to the AOTP database. 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 to BOS, and connecting in IAD, is reported in DB1B as a connecting passenger with origin of OAK, connection in IAD, and destination of BOS, rather than attributed separately to the two non-stop flight segments. This data is used as input to the multinomial logit passenger itinerary flow model presented in Barnhart et al. (2014).
5. **Booking data:** A proprietary booking (passenger itinerary) dataset from a major U.S. carrier for the fourth quarter of 2007 was used by Barnhart et al. (2014) to train their multinomial logit model for estimating passenger itinerary flows, and to validate results.

² Including Pinnacle Airlines, American Airlines, Aloha Airlines, Alaska Airlines, JetBlue Airways, Continental Airlines, Delta Airlines, Atlantic Southeast Airlines, Frontier Airlines, AirTran Airways, Hawaiian Airlines, American Eagle Airlines, Northwest Airlines, Midwest Airlines, SkyWest Airlines, United Airlines, US Airways, Southwest Airlines, ExpressJet Airlines, Mesa Airlines.

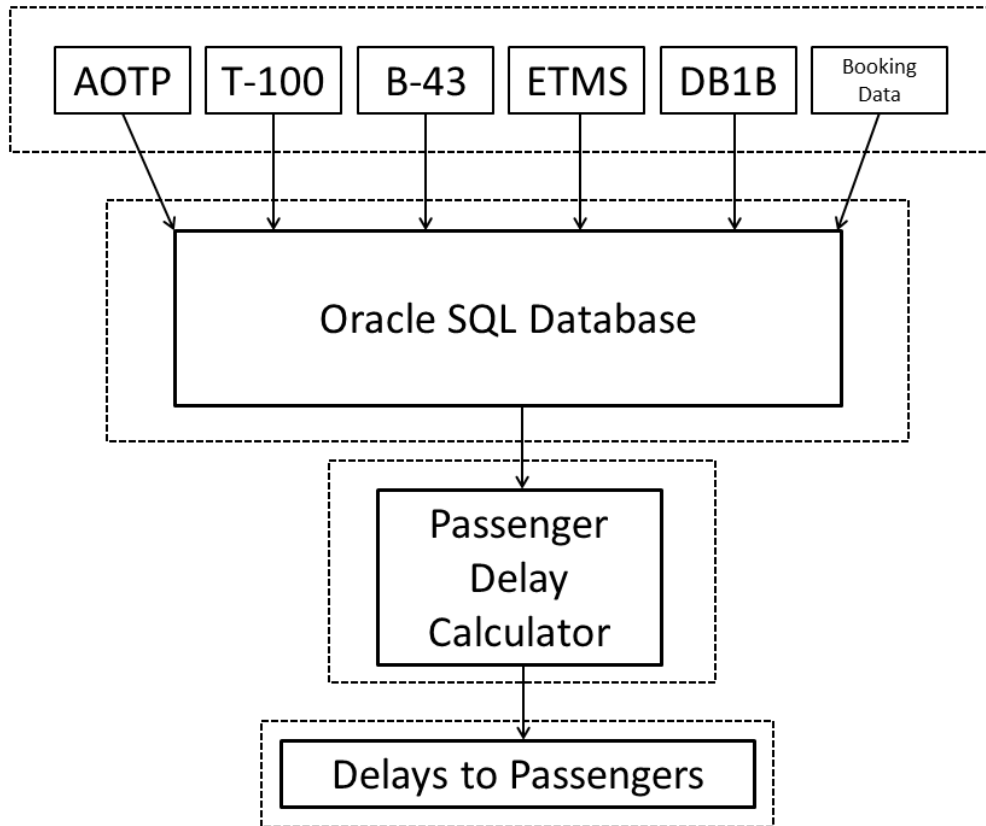


Figure 4: Data inputs and outputs of Passenger Delay Calculator

These six individual datasets are joined in an Oracle SQL database that provides input to the Passenger Delay Calculator (Figure 4). The Passenger Delay Calculator (PDC) is coded in the Java programming language, and connected to the Oracle SQL database. Outputs of the PDC include the delay and the number of passengers associated with each itinerary. The PDC output allows us to estimate actual passenger delay in 2007, which we call as the 2007 pre-Rule baseline delay. Throughout this paper, we compare the 2007 pre-Rule baseline delay to the delay estimated (using PDC) for various hypothetical scenarios that we create. For each scenario, we manipulate the input databases to represent our hypothetical scenario. For example, when we wish to analyze a policy of cancelling flights that taxied-out for three hours or longer in 2007, we change the cancellation flags of selected flights in AOTP. The passengers on these now-cancelled flights are added to the disruption queue, along with other passengers who were actually disrupted in the year 2007, and the PDC algorithm is used to compute the resulting passenger delays.

3. Passenger-Centric Analysis of the Tarmac Delay Rule

In this section, we quantify the impacts of the Tarmac Delay Rule on passengers traveling on flights with three hours or longer tarmac time, as well as on passengers traveling on all other flights in the NAS. In

the analysis that follows, we will use the results of the Passenger Delay Calculator to compare the estimates of the passenger delay experienced in 2007 in the absence of the Rule with the estimated delay that the same passengers would have experienced if the Rule had been in effect in 2007. We begin our analysis in Sub-section 3.1 by creating an operated flight schedule for 2007 for the hypothetical scenario assuming that the Rule was in effect. We do this by manipulating the operational data through cancelling selected flights. We then calculate the resultant passenger delays for this scenario in Sub-section 3.2. We provide an estimate of total passenger delay with and without the Rule, using the same set of assumptions and simplifications for both cases. Our assumptions and simplifications generally result in an underestimate of passenger delay for the post-Rule scenario (as detailed in Sub-section 3.3).

Throughout Section 3 we refer to the *pre-Rule baseline scenario*. This is what occurred operationally in 2007. This is the scenario that resulted in delay equal to the *pre-Rule baseline delay* defined earlier. Additionally, we refer to a set of *affected flights*, denoted by F_{AF} . A flight $f_i \in F_{AF}$ is a flight that was operated (i.e., not cancelled or diverted) in 2007 and experienced a taxi-out time greater than or equal to 180 minutes. We refer to the passengers on flights in this set F_{AF} as *affected passengers* and they are denoted by set P_{AF} . Similarly, we define all the other flights that were scheduled in 2007 but were not in set F_{AF} as *non-affected flights*, denoted by set F_{OF} . The passengers that were not on the set of affected flights F_{AF} are referred to as *non-affected passengers* and are denoted by set P_{OF} . The definition of affected and non-affected flights and passengers will change in Sub-section 3.3.1 where we change airline's cancellation policy to test the sensitivity of our analyses, and in Section 4 where we propose several different tarmac delay policies as potential candidates for improvement over the existing version of the policy. Note that the use of the term *non-affected passengers* is only for notational convenience. As we will see later, these passengers are also indirectly affected by the Rule.

3.1 Hypothetical Flight Schedule Generation under the Tarmac Delay Rule

We create a hypothetical flight schedule by first cancelling the non-cancelled and non-diverted flights that incurred three hours or longer tarmac time. The cancellation of an affected flight can allow other flights with later scheduled departure times and the same departure airport to be assigned earlier wheels-off times (that is, the time at which the aircraft becomes airborne). The assignment process is iterative, beginning by ordering all departing flights for the given airport and day by wheels-off time. We order by wheels-off time rather than by planned or actual gate departure time in order to control for differences between physical distances from individual gates to runways, and for the differences between departure queue lengths. Let flight f_i be the first flight in F_{AF} (ordered by wheels-off time in a non-decreasing manner), and assume that it has a pre-Rule baseline wheels-off time denoted by $WOT(f_i)$. As mentioned earlier, the pre-Rule baseline case represents the actual schedule in 2007. We first identify and cancel flight f_i creating a wheels-off slot, denoted by S , available for use by a subsequent flight in the departure queue. We identify f_{i+1} , the flight with a wheels-off time immediately following that of f_i . For this illustration, assume that flight f_{i+1} is a member of the set F_{OF} . We then check to ascertain if f_{i+1} is able to use the free wheels-off time slot, S . In this step, we test if

the planned gate departure time (PDT) of f_{i+1} plus the actual taxi-out time of f_{i+1} is no later than the wheels-off time of f_i . If this condition is met, f_{i+1} is moved up to time slot S . This in turn opens up the possibility of using the original wheels-off time slot of flight f_{i+1} by the subsequent flight f_{i+2} . Note that the procedure above assumes that the actual taxi-out durations and the time difference between the actual wheels-off and wheels-on time (that is, the time at which the aircraft lands) for each non-affected flight remain unchanged. If, however, the aforementioned criterion is not met, the algorithm keeps flight f_{i+1} in its original wheels-off time slot, slot S remains empty, and the algorithm moves down the wheels-off time list. We continue this iterative process moving up non-cancelled flights into available departure time slots, using a “first departed, first moved-up” flight processing order based on actual wheels-off times. We summarize the entire iterative process in Algorithm 1 below. Additionally, we define ADT and AAT as the actual departure and arrival times, respectively, of a flight.

One may argue that in practice, after cancelling an affected flight, in order to maintain fleet balance, the airline may need to cancel and/or delay one or more other flights. However, due to the fact that these additional cancellations and delays are usually determined by sophisticated recovery algorithms, which vary across airlines, we don’t incorporate these additional schedule revisions into the main body of our analysis. We do, however, relax this assumption in Sub-section 3.3 by providing a simple heuristic for cancelling other flights to maintain fleet balance, and then estimate the passenger delay under this revised operational plan.

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 AOTP. Thus, we do not include in the flight set F_{AF} the flights that taxied-out three hours or longer, then took off and then were diverted. Similarly, we do not include the passengers on such flights in passenger set P_{AF} . This assumption is expected to have a relatively insignificant effect on our results, because out of the total number of non-cancelled flights taxiing out three hours or longer in 2007, only 24 (1.45%) were diverted. Note, however, that the delays to these passengers are included in our calculations of overall passenger delays as well as non-affected passenger delays.

3.2 Post-Rule Baseline Results

We now present results for the hypothetical scenario in which all flights with three hours or longer taxi-out times are cancelled. We refer to this hypothetical scenario as the *post-Rule baseline*. We compare that with the pre-Rule baseline scenario. In this and the subsequent sub-sections, we will present passenger delay results separately for the following two categories of passengers: 1) Passengers P_{AF} who were on the affected flights F_{AF} , and 2) Passengers P_{OF} who were not on the affected flights F_{AF} . In the year 2007, there were a total of 156,470 passengers in set P_{AF} and a total of 486,376,064 passengers in set P_{OF} , as per the estimated passenger flow data.

Algorithm 1: Departure Compression

Order all departing flights by their Wheels-Off Times (WOT) in a non-decreasing order.

Denote this ordered flight set as F .

INITIALIZE

$i = 1, Slot\ List = \emptyset$

WHILE $i \leq size(F)$

IF flight $F(i)$ is an affected flight

 Cancel flight $F(i)$

 Add $WOT(F(i))$ to the end of *Slot List*

ELSE

FOR $j = 1$ to $size(Slot\ List)$

IF $PDT(F(i)) + Taxi\ Out(F(i)) \leq Slot\ List(j)$

$WOT(F(i)) = Slot\ List(j)$

$OFFSET = (WOT(F(i)) - Taxi\ Out(F(i))) - PDT(F(i))$

$ADT(F(i)) = WOT(F(i)) - Taxi\ Out(F(i))$

$AAT(F(i)) = PAT(F(i)) + OFFSET$

 Remove $Slot\ List(k) \forall k \leq j$ from *Slot List*

EXIT FOR LOOP

END IF

END FOR

END IF

$i = i + 1$

END WHILE

Table 4a provides the average passenger delay (in minutes) and total passenger delay (in passenger-minutes) for passengers P_{AF} , passengers P_{OF} , and for all passengers. The columns represent the pre-

Rule baseline, post-Rule baseline, the change from pre-Rule to post-Rule baseline, and the change expressed as percentage of the pre-Rule baseline value. The percentage is calculated as the difference between the pre-Rule baseline value and the post-Rule baseline value divided by the pre-Rule baseline value. Note that the percentage change is the same for total and average passenger delays. We also estimate the potential tarmac time saving with the Tarmac Delay Rule in effect. Because we don't know the exact time required by each affected flight to go back to the gate and deplane passengers, we use the following three methods to estimate tarmac time savings.

1. Minimum Tarmac Time Savings Method (MinTTS Method): Assume that the affected flights arrive back at the gate exactly at the three hour time limit if the Rule is in effect (i.e., each such flight incurs exactly three-hours of tarmac time).
2. Maximum Tarmac Time Savings Method (MaxTTS Method): Assume that the affected flights are cancelled immediately (i.e., each such flight incurs exactly zero tarmac time).
3. Average Tarmac Time Savings Method (AvgTTS Method): Assume that the affected flights arrive back at the gate at 1.5 hour (half of tarmac time threshold) after leaving the gate if the Rule is in effect (i.e., each such flight incurs exactly 1.5 hours of tarmac time).

The first three columns of Table 4b provide the tarmac time savings (in passenger-minutes), and the next three columns provide the ratio of the increase in total passenger delay (in passenger-minutes) to the reduction in total tarmac time (in passenger-minutes) under the three distinct methods. As shown in the last column of Table 4b, under the AvgTTS method, for every minute decrease in tarmac time, the Rule results in approximately 3 minutes of additional passenger delay. Based on the results of the other two (extreme) methods, this number ranges between 1.7 and 11. Note that most (91%) of the passenger delay increase is borne by the passengers P_{AF} who are on the affected flights F_{AF} .

Metric	Pre-Rule Baseline	Post-Rule Baseline	Change	% Change
Avg Delay to Passengers P_{AF} (min)	282.943	616.552	333.609	117.9%
Total Delay to Passengers P_{AF} (min)	44,272,099	96,471,835	52,199,736	
Avg Delay to Passengers P_{OF} (min)	30.963	30.971	0.008	0.0%
Total Delay to Passengers P_{OF} (min)	15,059,986,265	15,065,061,646	5,075,381	
Avg Delay to All Passengers (min)	31.045	31.162	0.117	0.4%
Total Delay to All Passengers (min)	15,104,258,364	15,161,533,481	57,275,117	

Table 4a: Pre-Rule baseline and post-Rule baseline passenger delay comparison

Tarmac Time Saving (min)			Total Delay Increase/Tarmac Time Saving		
MinTTS Method	MaxTTS Method	AvgTTS Method	MinTTS Method	MaxTTS Method	AvgTTS Method
5,181,040	33,345,640	19,263,340	11.055	1.718	2.973

Table 4b: Pre-Rule baseline and post-Rule baseline tarmac time comparison

3.3 Sensitivity Analyses

The results in Sub-section 3.2 were calculated assuming 1) that all the non-cancelled and non-diverted flights with three hours or longer tarmac times were cancelled; 2) that no other additional flights were cancelled; 3) that the passengers on the additional cancelled flights were available for rebooking onto any itinerary whose first flight has a planned departure from their disruption airport at any time that is at least 45 minutes later than the planned departure time of the cancelled flight; and 4) that the PDC algorithm rebooks passengers according to the actual aircraft seating capacity constraints. In this sub-section, we look at each of these assumptions one by one, and evaluate the effects of relaxing or modifying the assumptions.

3.3.1 Impact of Cancelling a Subset of Affected Flights

In Sub-section 3.2, we analyzed the passenger delay effects of cancelling all the non-cancelled and non-diverted flights with three hours or longer tarmac times (i.e., affected flights) and of cancelling no other additional flights. In this sub-section, we analyze the sensitivity of those results to varying assumptions about the percentage of affected flights to be cancelled. The motivation for having this sub-section is as follows. It is difficult to accurately model different airlines' risk management decisions toward the Rule. In other words, it is difficult to estimate the exact trade-offs that the individual airlines make between cancelling flights proactively on one hand, and running the risk of getting fined because of lengthy tarmac delays on the other hand. This decision is especially complicated because of the large variations in fine levels, as discussed in Sub-section 1.1, that have been imposed so far by the Department of Transportation. Furthermore, flight cancellation decisions are related closely to airlines' other operational decisions (such as those related to crew and aircraft), which are also difficult to model accurately due to lack of data. Instead, we use a different approach where we test the effects of cancelling various subsets of these affected flights.

Specifically, we test four scenarios by randomly cancelling 20%, 40%, 60%, and 80% of the affected flights. For each scenario, we conduct 10 simulation runs of PDC and report summary statistics in Table 5 and Figures 6 through 11. Table 5 reports the average values over 10 simulation runs. Figures 6-11 are boxplots³ describing other summary statistics for the 10 simulation runs. Note that we change the definitions of affected and non-affected passengers (P_{AF} and P_{OF}) in this sub-section: P_{AF} is the set of passengers on the *affected flights which are cancelled*, while set P_{OF} denotes the remaining passengers in the NAS. The first row in Table 5 reports the average number of passengers in P_{AF} under each cancellation percentage. The remaining rows present the total passenger delay in P_{AF} , average passenger delay and its percentage increase compared to pre-Rule baseline for passengers in P_{AF} , P_{OF}

³ The red central mark represents mean value, the edges of the box are the 25th (q_1) and 75th (q_3) percentiles, the whiskers extend to largest data point smaller than or equal to $q_3 + 1.5(q_3 - q_1)$ and the smallest data point larger than or equal to $q_1 - 1.5(q_3 - q_1)$. Points larger than $q_3 + 1.5(q_3 - q_1)$ or smaller than $q_1 - 1.5(q_3 - q_1)$ are outliers marked as positive signs.

and all passengers, and tarmac time savings and ratios of increase in total passenger delays to the reductions in total tarmac time under each of the three tarmac time savings calculation methods. All percentages are obtained by subtracting the value for the pre-Rule baseline scenario from that for the scenario mentioned in the header row, and then dividing by the value for the pre-Rule baseline scenario.

From Table 5 and Figures 6 through 11, we find that, the percentage increase in affected passenger delay and the ratio of passenger delay increase to tarmac time savings (under each of the three methods) is very stable across the various cancellation percentages. In each of the five cases with different cancellation percentages, the delay to affected passengers increases by between 113.3% and 117.9% and is monotonically increasing with cancellation percentages. Also, the ratio of passenger delay increase to tarmac time savings for the AvgTTS method varies in a narrow band between 2.849 and 2.973. Thus, for the rest of the analysis in this section (Section 3) we will assume the cancellation percentage to be 100%.

The fact that the percent increase in average passenger delay for the affected passengers is very stable across different cancellation percentages has other interesting implications. As mentioned earlier, our analysis is carried out for year 2007 which had a higher percentage of flights with long tarmac times compared with other pre-Rule years such as 2006 or 2008. The stability of percent increases in the average passenger delay across different cancellation percentages suggests that the second order (interaction) effects between the cancellations of multiple tarmac-delayed flights are relatively insignificant and therefore data from other years with fewer tarmac-delayed flights is also expected to demonstrate similar results in terms of the percent increase in the average passenger delays.

Metric		Cancel 20%	Cancel 40%	Cancel 60%	Cancel 80%	Cancel 100% (Post-Rule Baseline)
Number of Passengers in P_{AF}		30,913	62,885	93,560	124,132	156,470
Total Delay to P_{AF} (min)		18,666,124	38,198,717	57,066,753	75,853,623	96,471,835
Avg Delay to P_{AF} (min)		603.577	607.469	609.862	611.066	616.552
Avg Delay to P_{OF} (min)		31.031	31.016	31.001	30.987	30.971
Avg Delay to All Passengers (min)		31.067	31.090	31.112	31.135	31.162
Increase in Avg Delay to P_{AF} (%)		113.3%	114.7%	115.5%	116.0%	117.9%
Increase in Avg Delay to P_{OF} (%)		0.2%	0.2%	0.1%	0.1%	0.0%
Increase in Avg Delay to All Passengers (%)		0.1%	0.1%	0.2%	0.3%	0.4%
Tarmac Time Saving (min)	MinTTS Method	990,641	2,079,903	3,096,041	4,087,184	5,181,040
	MaxTTS Method	6,555,053	13,399,113	19,936,715	26,430,890	33,345,640
	AvgTTS Method	3,772,847	7,739,508	11,516,378	15,259,037	19,263,340
Total Delay Increase /	MinTTS	11.104	10.606	10.652	10.757	11.055

Tarmac Time Saving	Method					
	MaxTTS Method	1.674	1.645	1.654	1.663	1.718
	AvgTTS Method	2.908	2.849	2.863	2.881	2.973

Table 5: Sensitivity of passenger delays and tarmac times to different cancellation percentages

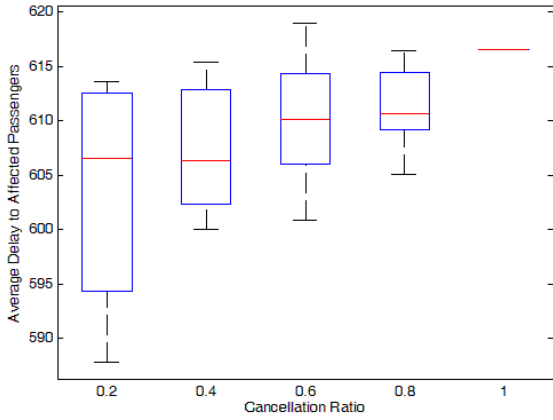


Figure 6: Average Delay to P_{AF} for Different Cancellation Percentages

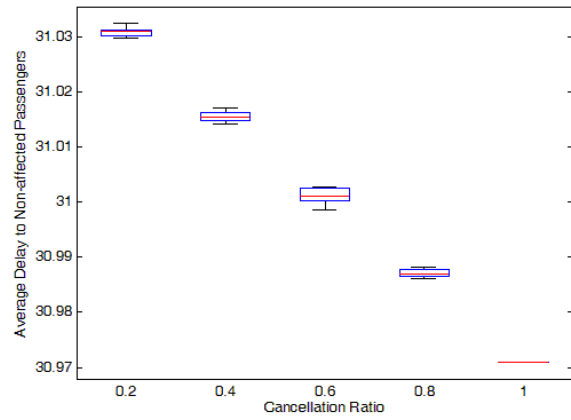


Figure 7: Average Delay to P_{OF} for Different Cancellation Percentages

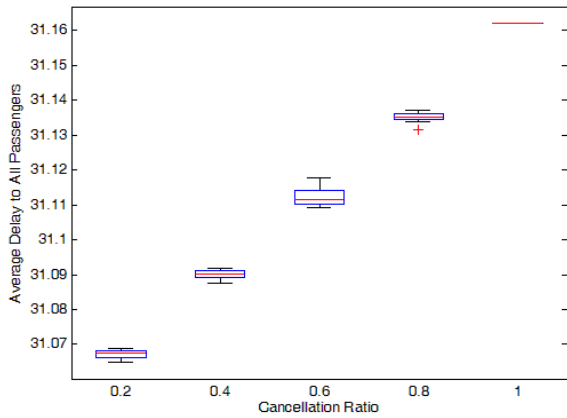


Figure 8: Average Delay to All Passengers for Different Cancellation Percentages

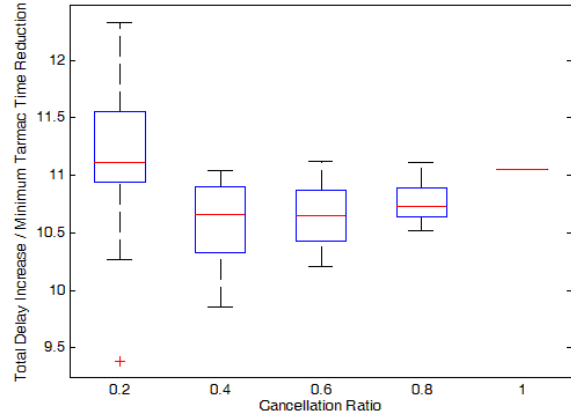


Figure 9: Ratio of Total Passenger Delay Increase to Tarmac Time Saving for Different Cancellation Percentages under MinTTS Method

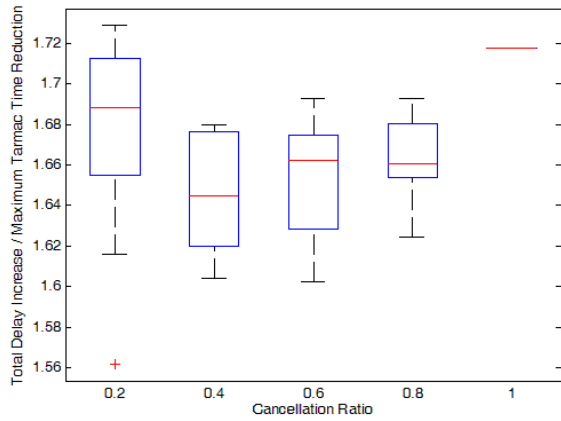


Figure 10: Ratio of Total Passenger Delay Increase to Tarmac Time Saving for Different Cancellation Percentages under MaxTTS Method

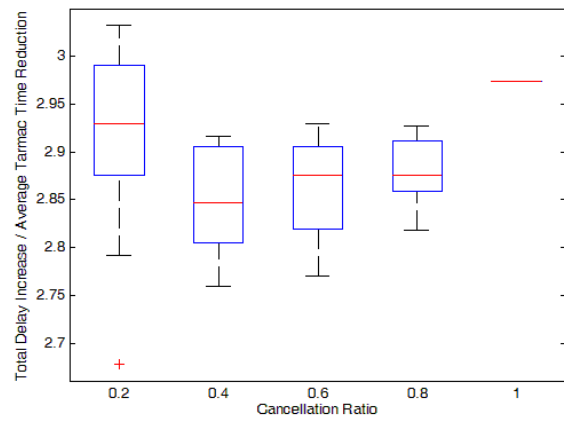


Figure 11: Ratio of Total Passenger Delay Increase to Tarmac Time Saving for Different Cancellation Percentages under AvgTTS Method

3.3.2 Impact of Cancelling Return Flights

When cancelling the affected flights, in order to maintain fleet balance, airlines may need to cancel and/or delay one or more other flights. However, due to the fact that these additional cancellations and delays usually result from sophisticated recovery algorithms, which vary across airlines, it is difficult to identify the exact set of flights that will get cancelled to ensure fleet balance. We instead look at a simple heuristic for cancelling one other flight per affected flight (thus creating pairs of cancelled flights) to maintain fleet balance and operational feasibility, and then estimate the passenger delay under this revised operational plan.

Suppose an aircraft is scheduled to fly from BOS to SFO (Boston to San Francisco), SFO to LAX (San Francisco to Los Angeles), and LAX to BOS (Los Angeles to Boston). 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 repositioned to SFO in order to operate the subsequent flights in the original route of the aircraft. Rosenberger et al. (2004) show that airlines usually choose to cancel additional flight legs to create a cancellation cycle, which preserves aircraft balance, rather than reposition an aircraft. We therefore design a simple decision rule to generate cancellation cycles that include exactly two flights; for example, one from A to B and the other from B to A.

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

1. Flight f_i^r departs no earlier than the planned arrival time plus minimum turn time of flight f_i ;
2. Flight f_i^r is operated by the same carrier as flight f_i ;

3. Flight f_i^r is operated by an aircraft of the same type as that of f_i ; and
4. Flight f_i^r is the flight departing the earliest, among all flights satisfying conditions 1, 2 and 3.

If no flight is found to match these four criteria, we do not cancel a return flight f_i^r . We denote by F_{CPS} the set of return flights, (usually) one for each flight in F_{AF} . The results in this sub-section are obtained based on the assumption that we cancel all flights in F_{CPS} and in F_{AF} .

Metric		Post-Rule Baseline	Pair Canc.	Thr=105	Thr=165	Thr=225
Number of Passengers in P_{AF}		156,470	156,470	156,470	156,470	156,470
Total Delay to P_{AF} (min)		96,471,891	96,271,610	98,254,867	100,399,914	101,678,431
Avg Delay to P_{AF} (min)		616.552	615.272	627.947	641.656	649.827
Avg Delay to P_{OF} (min)		30.971	31.076	30.956	30.956	30.955
Avg Delay to All Passengers (min)		31.162	31.264	31.148	31.152	31.154
Increase in Avg Delay to P_{AF} (%)		117.9%	117.5%	121.9%	126.8%	129.7%
Increase in Avg Delay to P_{OF} (%)		0.0%	0.4%	0.0%	0.0%	0.0%
Increase in Avg Delay to All Passengers (%)		0.4%	0.7%	0.3%	0.3%	0.4%
Tarmac Time Saving (min)	MinTTS Method	5,181,040	5,181,040	5,181,040	5,181,040	5,181,040
	MaxTTS Method	33,345,640	33,345,640	33,345,640	33,345,640	33,345,640
	AvgTTS Method	19,263,340	19,263,340	19,263,340	19,263,340	19,263,340
Total Delay Increase / Tarmac Time Saving	MinTTS Method	11.055	20.566	9.708	10.083	10.287
	MaxTTS Method	1.718	3.195	1.508	1.567	1.598
	AvgTTS Method	2.973	5.531	2.611	2.712	2.767

Table 6: Sensitivity of passenger delays and tarmac times to flight pair cancellations and to rebooking time thresholds

Same as Table 5, Table 6 presents results under different scenarios in different columns. The row structure in Table 6 is identical to that in Table 5. The second column is a repeat of the last column of Table 5. The third column refers to the scenario with flight pair cancellations. The last three columns refer to the scenarios explained in Sub-section 3.3.3. Comparing the second and third columns, we find that the percentage change in average delay to passengers P_{AF} remains almost unchanged with and without pairs cancellation, while that for the passengers P_{OF} and for all passengers increases considerably. This can be clearly seen by looking at the last three rows of Table 6. These rows refer to the ratio of total passenger delay increase to the total tarmac time savings, under each of the tarmac time savings calculation methods. The ratios under the flight pairs cancellation scenario for each

calculation method are almost twice those for the post-Rule baseline scenario. Note that the total tarmac time savings remain approximately the same because: 1) the tarmac time saving for the affected passengers remains unchanged regardless of whether the other flight in the pair is cancelled or not; and 2) the tarmac time saving for the cancelled return flights is negligible because they typically have a small tarmac delay value.

While the flight pairs cancellation scenario is not necessarily the most accurate depiction of airline recovery strategies, it shows that the fleet balancing necessitated by the cancellation of affected flights further adds to passenger delays. Thus our post-Rule baseline scenario underestimates passenger delay increases resulting from the Rule.

3.3.3 Impact of Passenger Rebooking Time

Results in Sub-section 3.2 were obtained assuming that the passengers on cancelled flights were available for rebooking onto any itinerary departing at least 45 minutes later than the planned departure time of the cancelled flight. However, this assumption can be operationally unrealistic because decision to cancel a flight is often not made until considerably after the flight's planned departure time; and, rebooking of passengers can be time consuming, especially on a day with large NAS delays and significant disruptions.

In this sensitivity analysis, we consider three different scenarios, assuming that passengers are available for rebooking one, two, and three hours, respectively, after the planned departure time of their original flight. Hence, affected passengers can be re-accommodated on any itinerary for which the origin of the first flight is the disruption airport and the planned departure time is at least 105 minutes, 165 minutes and 225 minutes, respectively, after the planned departure time of their original flight. We refer to these three scenarios as those with rebooking thresholds of 105 minutes, 165 minutes and 225 minutes respectively. We compare the passenger delays with those for our original rebooking threshold of 45 minutes. The results are summarized in the last three columns of Table 6.

As the rebooking threshold increases from 45 minutes to 225 minutes in steps of 60 minutes, the total and average passenger delays to passengers P_{AF} increase slightly and steadily. However, the delay to passengers P_{OF} and to all passengers does not change significantly. Similarly, the ratio of total delay increase to tarmac time savings under any of the three calculation methods does not vary much. Hence, our passenger delay results are robust to the rebooking time assumptions.

3.3.4 Impact of Load Factors

Results in Tables 4a and 4b indicate that the passenger delays are considerably higher, especially for the passengers in set P_{AF} , for the post-Rule baseline scenario as compared with the pre-Rule baseline scenario. Obviously, some of the delay increase is due to the difficulty in rebooking the passengers on affected flights onto other itineraries, which in turn partly depends on the flight load factors on other

flights. In this sub-section, we analyze the extent to which the increase in passenger delays is caused by load factors. Therefore, we estimate passenger delays for a variety of load factor scenarios under the presence of the Rule. We create four hypothetical scenarios by multiplying the number of seats on all flights in the NAS by factors of 1.5, 2, 2.5 and infinity (represented by a very large number), respectively. We compare these hypothetical scenarios with the post-Rule baseline scenario results in Table 7. The row structure of Table 7 is the same as that of Tables 5 and 6. Unlike Tables 5 and 6, however, each cell in rows 2 through 5 of Table 7 lists two numbers (instead of one), one for the pre-Rule scenario and one for the post-Rule scenario, separated by a '/'. This is so because, when we multiply the seating capacity for the post-Rule scenario by a certain factor, we need to do the same for the corresponding pre-Rule scenario, to ensure a fair comparison. Note that the column headings for the last five columns are the numbers by which we multiply the seating capacities, thus effectively reducing the load factors. They are not the load factors themselves.

Metric		Seating Capacity Multiplier				
		1 (Post-Rule Baseline)	1.5	2	2.5	∞
Number of Passengers in P_{AF}		156,470	156,470	156,470	156,470	156,470
Total Delay to P_{AF} (min) (Pre-Rule / Post-Rule)		44,272,091	42,845,398	42,720,065	42,701,132	42,685,485
		/	/	/	/	/
Avg Delay to P_{AF} (min) (Pre-Rule / Post-Rule)		96,471,891	82,951,319	78,896,086	77,775,604	76,610,998
		/	/	/	/	/
Avg Delay to P_{OF} (min) (Pre-Rule / Post-Rule)		282.943	273.825	273.024	272.903	272.803
		/	/	/	/	/
Avg Delay to All Passengers (min) (Pre-Rule / Post-Rule)		616.552	530.142	504.225	497.064	489.621
		/	/	/	/	/
Increase in Avg Delay to P_{AF} (%)		30.963	28.440	28.141	28.077	28.039
Increase in Avg Delay to P_{OF} (%)		/	/	/	/	/
Increase in Avg Delay to All Passengers (%)		30.971	28.428	28.127	28.063	28.023
Increase in Avg Delay to P_{AF} (%)		31.045	28.518	28.219	28.156	28.117
Increase in Avg Delay to P_{OF} (%)		/	/	/	/	/
Increase in Avg Delay to All Passengers (%)		31.162	28.589	28.280	28.213	28.172
Tarmac Time Saving (min)		117.9%	93.6%	84.7%	82.1%	79.5%
Total Delay Increase / Tarmac Time Saving		0.0%	0.0%	0.0%	-0.1%	-0.1%
Total Delay Increase / Tarmac Time Saving		0.4%	0.2%	0.2%	0.2%	0.2%
Tarmac Time Saving (min)	MinTTS Method	5,181,040	5,181,040	5,181,040	5,181,040	5,181,040
	MaxTTS Method	33,345,640	33,345,640	33,345,640	33,345,640	33,345,640
	AvgTTS Method	19,263,340	19,263,340	19,263,340	19,263,340	19,263,340
Total Delay Increase / Tarmac Time Saving	MinTTS Method	11.055	6.657	5.691	5.385	5.093
	MaxTTS	1.718	1.034	0.884	0.837	0.791

	Method					
	AvgTTS Method	2.973	1.790	1.531	1.448	1.370

Table 7: Sensitivity of passenger delays and tarmac times to load factors

The percentage change listed in rows 6 through 8 of Table 7 is calculated as the delay under each hypothetical scenario described by the header row minus the delay under the corresponding pre-Rule baseline scenario, expressed as a percentage of the corresponding pre-Rule baseline. Since the pre-Rule baseline here is different for each seating capacity multiplier, the pre-Rule baseline delays are computed separately for each column of Table 7. The last three rows of Table 7 list the ratios of the total delay increase to the tarmac time savings for each seating capacity multiplier scenario and for each tarmac time savings calculation method. These last three rows show that even in the absence of any seating capacity limitations, passenger delays are greater with the Rule than without (since all numbers are positive).

As expected, the increase in delays to passengers P_{AF} decreases with an increase in the seating capacity multiplier, from 117.9% to 79.5%. Additionally, delays to passengers P_{AF} decrease by only about 3.6% for the pre-Rule baseline scenarios when the seating capacity multiplier changes from 1 to infinity, while those for the post-Rule scenarios decrease by 20.6%. This indicates that seating capacity affects delays to passengers P_{AF} much more significantly under the post-Rule scenario than under the pre-Rule baseline scenario. This is expected because under the post-Rule scenario, most of the passengers in P_{AF} are disrupted and hence seating capacity on other flights affects the ease of their rebooking and hence their delays. On the other hand, under the pre-Rule baseline scenario, several of the passengers in P_{AF} are not disrupted and hence their delays are relatively less affected by the changes in flight seating capacities. The percentage change in delays to passengers P_{OF} remains mostly unaffected or decreases slightly with increases in seating capacity multipliers. This is so because most of these passengers are not disrupted in both pre-Rule and post-Rule scenarios, and hence the delays to this group of passengers are not too sensitive to changes in flight seating capacities.

Combining the aforementioned effects on passengers P_{AF} and passengers P_{OF} , we find that, when the seating capacity multiplier changes from 1 to infinity, the passenger delay to all passengers as a whole decreases more under the post-Rule scenario than under the pre-Rule baseline scenario. Stated differently, under the infinite seating capacity scenario, the percentage change due to the Rule in delays to all passengers is reduced to half (0.2%) of what it is for the scenario with seating capacity multiplier equal to one (0.4%). In summary, when comparing the infinite seating capacity scenario to the baseline scenario, the passenger delay increase due to the Rule reduced to about two-third (from 117.9% to 79.5%) when considering only the passengers P_{AF} and reduced to about half (from 0.4% to 0.2%) when considering all passengers.

Thus, load factors are responsible for about one-third of the delay increase to passengers P_{AF} and about half of the delay increase to all passengers. The remaining delay increases are due to what we can call as

schedule effect, that is, the phenomenon that, irrespective of the seat availability, the disrupted passengers need to wait for at least the first available itinerary to reach their destinations.

4. Possible Revisions to the Current Policy

In Sub-section 3.2, we showed that while the current tarmac delay rule is effective in reducing the time spent by the passengers on the tarmac, the Rule significantly increases passenger delays, especially for affected passengers. Ideally, it is desirable to have a situation that can result simultaneously in a short tarmac time and low passenger delays. While such an ideal situation is difficult to achieve, in this section, we explore variants of the Rule that might result in a better tradeoff between the two objectives of short tarmac times and low passenger delays. In particular, in Sub-section 4.1 we explore the possibility of varying the Rule's tarmac time threshold, which is currently set at three hours. In Sub-section 4.2, we explore the effects of applying the Rule selectively depending on the planned departure time of the flights. Finally, in Sub-section 4.3 we combine the insights from Sub-sections 4.1 and 4.2, and investigate the effects of a combined Rule which incorporates ideas from both Sub-sections 4.1 and 4.2.

4.1 Different Tarmac Time Limits

The analyses of the ratio of total passenger delay increase to tarmac time saving conducted in previous sub-sections show that setting the tarmac time threshold involves a critical trade-off between reducing total time spent on tarmac and reducing the total passenger delay. In this sub-section, we evaluate the effects of varying the tarmac time threshold in order to find a better balance between these two objectives. The results in this sub-section are obtained for different values of tarmac time thresholds, namely, 2.0, 2.5, 3.0, 3.5, and 4.0 hours. For each tarmac time threshold, the set of flights that get cancelled due to the Rule and the set of passengers on these flights vary. The results are summarized in Table 8. The total number of passengers in the system remains unchanged at 486,532,534 across all scenarios. Note that we re-define F_{AF} in this sub-section as the set of flights that were operated (i.e., not cancelled or diverted) in 2007 and experienced a taxi-out time greater than or equal to the corresponding tarmac time threshold. Also, we define P_{AF} as the passengers on the flights in set F_{AF} . As before, F_{OF} is the set of all flights not in F_{AF} and P_{OF} is the set of all passengers not in P_{AF} .

Table 8 has a row structure same as Tables 5, 6, and 7. The first two rows after the header row in Table 8 list the total number of affected passengers P_{AF} , and the total delay to these passengers. The next three rows list the average delay to affected passengers P_{AF} , the average delay to non-affected passengers P_{OF} , and the average delay to all passengers. The three rows after that list the percentage change in average delay to affected passengers P_{AF} , percentage change in average delay to non-affected passengers P_{OF} , and percentage change in average delay to all passengers. In each of these three cases, the change in passenger delay is calculated with respect to the average delay to the relevant set of passengers in the pre-Rule baseline scenario and the percentage is calculated by dividing by the average delay to the relevant set of passengers in the pre-Rule baseline scenario. The next three rows of Table 8

list the reduction in tarmac time under the three different tarmac time savings calculation methods, namely, MinTTS, MaxTTS and AvgTTS. As described before, the three methods assume that the tarmac time under the tarmac delay rule equals the corresponding tarmac time threshold, zero, and half of the corresponding tarmac time threshold, respectively. As a result, these tarmac time savings numbers are different for each different value of tarmac time threshold. Finally, the last three rows of Table 8 list the ratio of total passenger delay increase to total tarmac time saving under each combination of tarmac time threshold and tarmac time savings calculation method.

Metric		Tarmac Time Threshold (hours)				
		2	2.5	3 (Post-Rule Baseline)	3.5	4
Number of Passengers in P_{AF}		831,023	356,629	156,470	63,183	26,591
Total Delay to Passengers P_{AF} (min)		484,727,755	212,585,244	96,471,835	39,097,912	16,132,740
Avg Delay to Passengers P_{AF} (min)		583.290	596.096	616.552	618.804	606.699
Avg Delay to Passengers P_{OF} (min)		30.700	30.901	30.971	31.002	31.026
Avg Delay to All Passengers (min)		31.644	31.315	31.162	31.078	31.057
Increase in Avg Delay to P_{AF} (%)		106.2%	110.7%	117.9%	118.7%	114.4%
Increase in Avg Delay to P_{OF} (%)		-0.8%	-0.2%	0.0%	0.1%	0.2%
Increase in Avg Delay to All Passengers (%)		1.9%	0.9%	0.4%	0.1%	0.0%
Tarmac Time Saving (min)	MinTTS Method	28,801,407	12,344,277	5,181,040	2,098,961	824,022
	MaxTTS Method	128,524,167	64,118,727	33,345,640	14,918,201	7,002,822
	AvgTTS Method	77,070,927	38,231,502	19,263,340	6,409,620	2,317,050
Total Delay Increase / Tarmac Time Saving	MinTTS Method	10.116	10.651	11.055	7.776	7.239
	MaxTTS Method	2.267	2.051	1.718	1.094	0.852
	AvgTTS Method	3.780	3.439	2.973	2.546	2.575

Table 8: Passenger delays and tarmac times for different tarmac time thresholds

As shown in the first and second row of Table 8, as the tarmac time threshold increases, the number of affected passengers P_{AF} and the total delay to the affected passengers P_{AF} , both decrease approximately as geometric sequences. In fact, for each half an hour increase in tarmac time threshold, the number of affected passengers P_{AF} decreases by a factor between 40% and 44%, and the total delay to affected passengers P_{AF} decreases by a factor between 40% and 46%. With increasing tarmac time threshold, the average delay to all passengers decreases monotonically. Finally, the percentage increase in delay to all passengers, with respect to the pre-Rule baseline, also decreases significantly and monotonically, with the largest percentage reduction (71%) occurring from 3 hour to 3.5 hour threshold.

The ratios of total delay increase to tarmac time saving for the three calculation methods show that the ratios generally (but not always) have a decreasing trend with increasing threshold value. The lowest value of the ratio for the AvgTTS method is at the 3.5 hour threshold and for the MinTTS and MaxTTS methods, by far the largest drop in the ratio occurs when going from 3 to 3.5 hour threshold. Note that, all else being equal, a lower value of the ratio is considered better, because these ratios represent the price in the form of additional total passenger delay paid for a reduction in the long tarmac times. While the results in Table 8 describe a Pareto frontier between the contradictory objectives of minimizing long tarmac times and minimizing additional passenger delays, they suggest that a good tradeoff between these objectives can be obtained at a tarmac time threshold of 3.5 hours.

In addition, we also want to point out a caveat here. The current version of the Rule penalizes all flights that do not allow passengers to deplane by the three hour mark. As described in Section 1 and illustrated in Figure 1, the actual process of decision-making about whether or not to turn a flight back to a gate or other deplaning area starts taking place much before the three hour mark is reached. Depending on the specific airline and airport, the airline needs to come to a decision about whether or not to turn back from the departure queue to a gate or other deplaning area somewhere around the two hour mark or so. At that moment, the airline is unlikely to have an exact estimate of whether or not a particular flight will be able to take off before the three hour mark and the time it will take to return to a gate or other deplaning area. Under such uncertainty, and given the fact that the airline runs a risk of losing millions of dollars in fines if it errs on the side of being too optimistic about the remaining tarmac times, an airline is expected to be slightly conservative and turn back to the gate even if the projected total tarmac time is somewhat less than the three hour mark. Therefore, the actual tarmac time limit for an airline is likely to be somewhat lower than that stated in the Rule. In order to implement the Rule more effectively, we suggest the tarmac time limit to be defined in terms of the time when the aircraft begin returning to the gate instead of being defined in terms of the time when passengers are allowed to deplane. This reduces the uncertainty and heterogeneity in risk preferences among airlines when responding to the Rule. In this case, the tarmac time threshold on the time of beginning to return to the gate could be set to be a bit lower (by 30 minutes or so) compared with those listed in Table 8. This difference represents an estimate of the time it takes to leave the departure queue and reach back at the gate at a specific airport.

4.2 Selective Tarmac Delay Rule Based on Flight Departure Time

In this sub-section, we explore the effects of applying the Rule selectively based on the planned departure time of the flights. We define the *flight delay multiplier* for each flight $f \in F_{AF}$ as the ratio of the total passenger delay to all passengers on that flight under the Rule to that without the Rule. Thus if the multiplier is greater than 1, passengers on that flight incur more passenger delay if the Rule is in effect than if it is not. In Figure 12, we plot as boxplots the flight delay multipliers for all flights in set F_{AF} against the flight's planned departure hour (in local time). From the figure, we observe that for flights later in the day (e.g., after 5 pm or so), the flight delay multipliers tend to be higher. This is because

passengers on later flights have fewer same-day rebooking opportunities. This effect is exaggerated, however, by the different maximum passenger delay threshold values (eight hours for passengers disrupted between 5:00am and 4:59pm, and 16 hours for passengers disrupted between 5:00pm and 4:59am) used by the passenger delay calculator. To correct for this effect, we again plotted in Figure 13 the flight delay multipliers, this time assuming the same maximum passenger delay threshold value of eight hours for passengers disrupted at any time of the day. Similar to Figure 12, Figure 13 also shows, but to a lesser extent, an increase in the flight delay multipliers for all flights with planned departure times after 5 pm, compared to those before. These observations from Figures 12 and 13 indicate that it becomes progressively more difficult to find recovery itineraries for passengers disrupted later in the day, indicating that there is value in exploring the effectiveness of a tarmac delay rule that stipulates a maximum tarmac time limit only for flights with planned departure times before a certain time of the day.

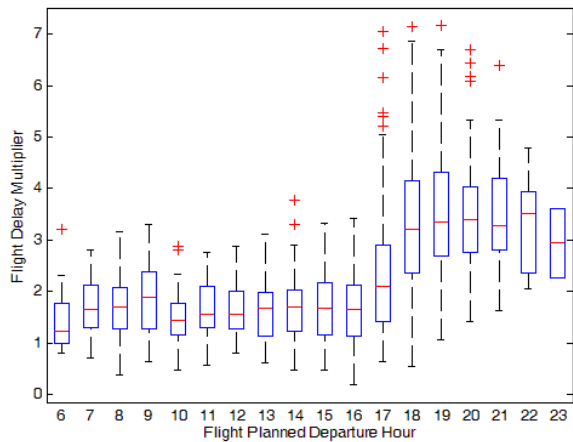


Figure 12: Flight Delay Multipliers for Different Values of Planned Flight Departure Times

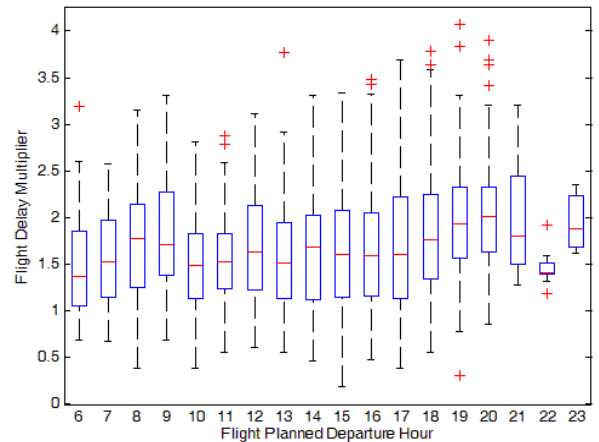


Figure 13: Flight Delay Multipliers for Different Values of Planned Flight Departure Times with Maximum Passenger Delay Threshold of 8 hours throughout the Day

We test 1pm, 3pm, 5pm, and 7pm as four candidate time-of-day thresholds. Note that, for each different time-of-day threshold, flights belonging to the set of affected flights F_{AF} are different. The results are listed in Table 9. Similar to Sub-section 4.1, we re-define F_{AF} in this sub-section as the set of flights that were operated (i.e., not cancelled or diverted) in 2007 and experienced a taxi-out time greater than or equal to 180 minutes and had a planned departure time before the corresponding time-of-day threshold. P_{AF} is still the set of passengers on the flights in set F_{AF} . As before, F_{OF} is the set of all flights not in F_{AF} and P_{OF} is the set of all passengers not in P_{AF} . The row structure of Table 9 is the same as that of Tables 5 through 8. The first eight rows list the number of passengers in set P_{AF} , total delay to passengers P_{AF} , average delay to passengers P_{AF} , average delay to passengers P_{OF} , average delay to all passengers, percent increase in average delay to P_{AF} , percent increase in average delay to P_{OF} and percent increase in average delay to all passengers. The percent increase in average passenger delays is

calculated by subtracting the pre-Rule baseline delays to the corresponding set of passengers and then dividing by the same quantity. The next three rows list the tarmac time savings for each of the three different tarmac time savings calculation methods, and the last three rows list the ratio of the passenger delay increase to tarmac time saving for each of these three tarmac time savings calculation methods. Each column of Table 9 corresponds to a different time-of-day threshold.

As expected, the number of passengers in the set P_{AF} , the total and average delays to passengers in the set P_{AF} , and the percent increase in delays to passengers P_{AF} when compared with the pre-Rule baseline scenario, all increase monotonically from left to right as the time-of-day threshold increases. None of the other passenger delay metrics show monotonic behavior due to various second order effects. The ratio of the passenger delay increase to tarmac time saving for a time-of-day threshold of 5pm is found to be the lowest across the five columns for each of the three tarmac time savings calculation methods. In fact, under the AvgTTS method, Table 9 shows that 5:00pm is the only time-of-day threshold value, among those tested, for which the ratio of passenger delay increase to tarmac time saving is found to be less than 1.0, which implies that for each minute of tarmac time saving calculated using the AvgTTS method, the increase in passenger delay is less than 1 minute. Thus, from Table 9, 5:00pm is identified as the optimal value of the time-of-day threshold, among those tested, in terms of providing the best tradeoff between the contrasting objectives of achieving maximum tarmac time savings and minimum increases in passenger delays.

Metric	Planned Flight Departure Time					
	1:00pm	3:00pm	5:00pm	7:00pm	Post-Rule Baseline	
Number of Passengers in P_{AF}	30,091	55,179	85,933	129,556	156,470	
Tot Delay to Passengers P_{AF} (min)	12,850,419	23,676,044	37,076,912	73,212,831	96,471,835	
Avg Delay to Passengers P_{AF} (min)	427.052	429.077	431.463	565.106	616.552	
Avg Delay to Passengers P_{OF} (min)	31.029	31.014	30.995	31.070	30.971	
Avg Delay to All Passengers (min)	31.053	31.059	31.065	31.212	31.162	
Increase in Avg Delay to P_{AF} (%)	50.9%	51.6%	52.5%	99.7%	117.9%	
Increase in Avg Delay to P_{OF} (%)	0.2%	0.2%	0.0%	0.3%	0.0%	
Increase in Avg Delay to All Passengers (%)	0.0%	0.0%	0.1%	0.5%	0.38%	
Tarmac Time Saving (min)	MinTTS Method	1,084,241	1,808,328	2,837,282	4,235,673	5,181,040
	MaxTTS Method	5,238,279	11,318,628	17,733,002	26,844,393	33,345,640
	AvgTTS Method	3,792,431	6,563,478	10,285,142	15,540,033	19,263,340
Total Delay Increase / Tarmac Time Saving	MinTTS Method	3.852	3.804	3.529	7.734	11.055
	MaxTTS Method	0.642	0.608	0.564	1.220	1.718
	AvgTTS Method	1.101	1.048	0.973	2.108	2.973

Table 9: Passenger delays and tarmac times for different flight departure time thresholds

4.3 Combination of the Two Revisions

The main takeaway from Sub-section 4.1 was that across the various tarmac time threshold policies, the best tradeoff between passenger delay increases and tarmac time savings was obtained at 3.5 hours

threshold, among the tested values. The main takeaway from Sub-section 4.2 was that across the different time-of-day thresholds, the best tradeoff between passenger delay increases and tarmac time savings was obtained at 5:00pm, among the tested values. In this sub-section, we explore a combined policy imposing the tarmac time threshold of 3.5 hours for all flights whose planned departure time is before 5:00pm local time at the origin airport. Table 10 lists the results for the combined policy. The row structure is the same as that of Tables 5 through 9. The second column lists the post-Rule baseline results for comparison with the results of the combined policy, which are listed in the last column. The number of affected passengers P_{AF} under the combined policy scenario is roughly about 22% of those under the post-Rule baseline scenario, while the average delay to the affected passengers P_{AF} is approximately 28% lower under the combined policy compared with the post-Rule baseline scenario. The increase in average delay to all passengers is found to be significantly smaller for the combined policy. In fact, the combined policy incurs negligible additional average passenger delay beyond that of the pre-Rule baseline case. Under each of the three methods of calculating the tarmac time savings, the ratio of passenger delay increase to tarmac time saving for the combined policy is found to be approximately one-sixth of that under the post-Rule baseline. Under the AvgTTS method, the ratio of passenger delay increase to tarmac time saving is approximately 0.481 implying that each minute of tarmac time saving can be achieved with less than 0.5 minute increase in passenger delay. Thus, at least when measured on this metric, this policy significantly outperforms all the other policies analyzed in Sections 3 and 4. It is noteworthy that under this combined policy, on average, the additional passenger delay incurred due to the Rule is less than half of the tarmac time savings. This result suggests that by judiciously imposing the rule, passengers can enjoy the benefits of reduction in lengthy and inconvenient waiting times on the airport tarmacs, while incurring a total delay increase that is less than half of the total amount of time saved on the tarmac.

Metric		Post-Rule Baseline	Combined Policy
Number of Passengers in P_{AF}		156,470	34,201
Tot Delay to Passengers P_{AF} (min)		96,471,835	15,056,147
Avg Delay to Passengers P_{AF} (min)		616.552	440.225
Avg Delay to Passengers P_{OF} (min)		30.971	31.021
Avg Delay to All Passengers (min)		31.162	31.049
Increase in Avg Delay to P_{AF} (%)		117.9%	55.6%
Increase in Avg Delay to P_{OF} (%)		0.0%	0.2%
Increase in Avg Delay to All Passengers (%)		0.4%	0.0%
Tarmac Time Saving (min)	MinTTS Method	5,181,040	1,162,917
	MaxTTS Method	33,345,640	8,026,767
	AvgTTS Method	19,263,340	4,594,842
Total Delay Increase / Tarmac Time Saving	MinTTS Method	11.055	1.899
	MaxTTS Method	1.718	0.275
	AvgTTS Method	2.973	0.481

Table 10: Passenger delays and tarmac times under a combined policy of a 3.5 hour tarmac time threshold imposed on flights with planned departure times before 5:00pm

5. Conclusions and Future Research Directions

In this paper, we analyze the impact of the 2010 Tarmac Delay Rule on passenger delays. 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 hypothetical scenarios with the Rule in effect for that same year. Our main result is that, while the three-hour tarmac delay rule effectively decreases tarmac delays, especially extremely long tarmac delays, each passenger-minute of tarmac time saving is achieved at the cost of an increase of approximately three passenger-minutes in total passenger delays. This is due primarily to increases in flight cancellations.

We performed a number of sensitivity analyses to assess the robustness of our results. We considered various hypothetical scenarios by varying the set of cancelled flights, and the passenger rebooking time threshold. We found that our results are robust to variations in these parameters and in most cases, underestimate the increase in passenger delay due to the Rule relative to the tarmac time saving benefits achieved by the Rule. We also analyzed scenarios with reduced load factors relative to their actual values and found that roughly one-third of the delay increase to passengers on tarmac-delayed flights, and roughly half of the delay increase to all passengers, is caused by the load factors experienced in 2007. Remaining delays can be attributed to what we call as the flight schedule effect.

We then analyzed scenarios involving variants of the tarmac delay policy. We varied the tarmac time threshold and the latest planned flight departure time for which the tarmac delay rule is imposed. We found that across the various tarmac time threshold policies that we tested, the best tradeoff between passenger delay increases and tarmac time savings was obtained for a 3.5-hour threshold. We found that across the different time-of-day thresholds tested, the best tradeoff between passenger delay increases and tarmac time savings was obtained for a time-of-day threshold of 5:00pm. By combining those two improvements, we obtained a policy in which a 3.5 hour tarmac time threshold is imposed on flights with planned departure times before 5:00pm. For this combined policy scenario, we were able to obtain a ratio of increase in passenger delay to tarmac time saving of less than 0.5, implying that each passenger-minute of tarmac time savings can be achieved with less than 0.5 passenger-minutes increase in passenger delays. This result suggests that by judiciously imposing the rule, passengers can enjoy the benefits of reduction in lengthy and inconvenient waiting times on the airport tarmacs, with the total passenger delay increase being less than half the total amount of time saved on the tarmac. Additionally, in order to implement the Rule more effectively, we also suggest that the tarmac time limit should be defined in terms of the time when the aircraft begins returning to the gate instead of the time when passengers are allowed deplane.

In summary, our study has identified and quantified the positive and negative effects of the current tarmac delay policy, and our methodology has been found to be robust under sensitivity analyses. We also explored and assessed different variants of the current policy and identified especially promising alternatives to guide tarmac delay policy decision-making. A future research direction is to a study the

Rule's impact on airlines with different network structures, such as, international carriers, regional carriers, legacy carriers, and low-cost carriers. It has been shown that smaller aircraft on short-haul flights are assigned disproportionately more ground delay time (Vossen et al., 2003). Thus, it would be interesting to examine if regional airlines incur disproportionately more delay as a result of the Rule, thereby introducing inequitable impacts. Another direction for future research is a comparative longitudinal analysis of tarmac time savings and passenger delay increases due to the Rule across different years, using the methodology described in this study. Finally, this study focused on operational, rather than strategic, aspects of the tarmac delay rule. It is important to investigate the impacts of the tarmac delay rule on airline schedule decision-making, and on capacity allocation over airline networks.

Acknowledgements

This research was supported by the Federal Aviation Administration National Center of Excellence for Aviation Operations Research (NEXTOR II).

References

1. Barnhart, C., Fearing, D., & Vaze, V. (2014). Modeling passenger travel and delays in the National Air Transportation System. *Operations Research*, Issue 3, Volume 62.
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. Bureau of Transportation Statistics (2006-2013). Airline On-Time Performance Data. Accessed January 15, 2015.
http://www.transtats.bts.gov/Tables.asp?DB_ID=120&DB_Name=Airline%20On-Time%20Performance%20Data&DB_Short_Name=On-Time
4. 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
5. 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/
6. 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
7. 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
8. 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
9. Coldren, G., & Koppelman, F. (2005). Modeling the competition among air-travel itinerary shares: GEV model development. *Transportation Research Part A*, Number 4, Volume 39.

10. Marks, J. and Jenkins, D. (2010). Impact of Three-Hour Tarmac Delay Rules and Fines on Passenger Travel Time and Welfare.
11. Rosenberger, J., Johnson, E., & Nemhauser, G. (2004). A Robust Fleet Assignment Model with Hub Isolation and Short Cycles. *Transportation Science*, Number 3, Volume 38.
12. Sherry, L., Wang, D., & Donohue, G. (2007). Air travel consumer protection: metric for passenger on-time performance. *Transportation Research Record*, Issue 1.
13. Sherry, L., Samant, A., Calderon-Meza, G. (2010). Trends in Airline Passenger Trip Delays (2007): A Multi-Segment Itinerary Analysis. *10th AIAA Aviation Technology, Integration and Operations (ATIO) Conference*. Fort Worth, Texas.
14. Tien, S., Ball, M., & Subramanian, B. (2008). Constructing a passenger trip delay metric: an aggregate level approach. *NEXTOR Working Paper*.
15. U.S. 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.
16. U.S. Department of Transportation (2014). Independent Review and Analysis of the Impact of the Three-Hour Tarmac Delay Rule.
17. U.S. Department of Transportation (2009). Final Regulatory Impact Analysis of Rulemaking on Enhanced Airline Passenger Protections.
18. U.S. Department of Transportation (2015). U.S. Department of Transportation Fines Southwest \$1.6 Million for Violating Tarmac Delay Rule. Accessed June 7, 2015.
<http://www.transportation.gov/briefing-room/us-department-transportation-fines-southwest-16-million-violating-tarmac-delay-rule>
19. 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.