

## MIT Open Access Articles

### *Modeling Passenger Travel and Delays in the National Air Transportation System*

The MIT Faculty has made this article openly available. *Please share*  
how this access benefits you. Your story matters.

**Citation:** Barnhart, Cynthia, Douglas Fearing, and Vikrant Vaze. "Modeling Passenger Travel and Delays in the National Air Transportation System." *Operations Research* 62, no. 3 (June 2014): 580–601.

**As Published:** <http://dx.doi.org/10.1287/opre.2014.1268>

**Publisher:** Institute for Operations Research and the Management Sciences (INFORMS)

**Persistent URL:** <http://hdl.handle.net/1721.1/89059>

**Version:** Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

**Terms of use:** Creative Commons Attribution-Noncommercial-Share Alike



# Modeling passenger travel and delays in the National Air Transportation System

Cynthia Barnhart

Department of Civil and Environmental Engineering, Massachusetts Institute of Technology

Douglas Fearing

Operations Research Center, Massachusetts Institute of Technology,

[dfearing@mit.edu](mailto:dfearing@mit.edu)

Vikrant Vaze

Department of Civil and Environmental Engineering, Massachusetts Institute of Technology

---

**Abstract:** Many of the existing methods for evaluating an airline's on-time performance are based on flight-centric measures of delay. However, recent research has demonstrated that passenger delays depend on many factors in addition to flight delays. For instance, significant passenger delays result from flight cancellations and missed connections, which themselves depend on a significant number of factors. Unfortunately, lack of publicly available passenger travel data has made it difficult for researchers to explore the nature of these relationships. In this paper, we develop methodologies to model historical travel and delays for U.S. domestic passengers. We develop a discrete choice model for estimating historical passenger travel and extend a previously-developed greedy re-accommodation heuristic for estimating the resulting passenger delays. We report and analyze the estimated passenger delays for calendar year 2007, developing insights into factors that affect the performance of the National Air Transportation System in the United States.

---

Draft completed August 2<sup>nd</sup>, 2010.

## 1 Introduction

Over the past two years, flight and passenger delays have been on the decline due to reduced demand for air travel as a result of the recent economic crisis. As the economy rebounds, demand for air travel in the United States is also expected to recover (Tomer & Puentes, 2009). Thus, after a brief reprieve, the U.S. will once again face a looming transportation crisis due to air traffic congestion. In calendar year 2007, the last year of peak air travel demand before the economic downturn, flight delays were estimated to have cost airlines \$19 billion (U.S. Congress Joint Economic Committee, 2008) compared to profits of just \$5 billion (Air Transport Association, 2008). In 2007, passengers were also severely impacted, with the economic costs of time lost due to delays estimated at \$12 billion according to the Joint Economic Committee report. A similar analysis performed by the Air Transport Association estimated the

economic costs of passenger delays at approximately \$5 billion for 2007. While there are differences in methodologies, the huge discrepancy between these estimates suggests the need for a more transparent and rigorous approach to measuring passenger delays. Accurately estimating passenger delays is important not only as a means to understand system performance, but also to motivate policy and investment decisions for the National Air Transportation System.

Another important consideration is that neither of the passenger delay cost estimates listed above includes the delays associated with itinerary disruptions, such as missed connections or cancellations. Analysis performed by Bratu and Barnhart (2005) suggests that itinerary disruptions and the associated delays represent a significant component of passenger delays. Their analysis was performed using one month of proprietary passenger booking data from a legacy carrier. The challenge in extending this analysis system-wide is that publicly available data sources do not contain passenger itinerary flows. For example, on a given day, there is no way to determine how many passengers planned to take the 7:05am American Airlines flight from Boston Logan (BOS) to Chicago O'Hare (ORD) followed by the 11:15am flight from Chicago O'Hare (ORD) to Los Angeles (LAX), or even the number of non-stop passengers on each of these flights. Instead, the passenger flow data that is publicly available is aggregated over time, either monthly or quarterly, and reports flows based only on the origin, connection, and destination airports. The methodologies we develop in this work are precisely to address these limitations.

Beyond the analysis of historical passenger delays, we expect our approach to be valuable in extending passenger analyses to other contexts where previously only flight information has been available. For example, much of the research on traffic congestion considers only flight delays, due to both the lack of passenger data and the complexities associated with passenger-centric objectives. Thus, to encourage further passenger-centric research, we have made estimated passenger itinerary flows for 2007 publicly available<sup>1</sup>.

## **1.1 Literature Review**

As mentioned above, our work is largely motivated by the findings of Bratu and Barnhart (2005). Using one month of booking data from a major U.S. carrier, their research demonstrated that itinerary disruptions in the form of flight cancellations and missed connections contribute significantly to overall passenger delays. To generate this result, the authors use a passenger delay calculator to estimate passenger delays by greedily re-accommodating passengers traveling on disrupted itineraries.

---

<sup>1</sup> For further information, please visit <http://web.mit.edu/nsfnats/README.html>, which provides detailed instructions for accessing the data.

The primary challenge we address in our work is estimating disaggregate passenger itinerary flows from publicly available aggregate flow data using a small set of proprietary booking data. In her Master's thesis, Zhu (2009) attempted to address this problem using an allocation approach based on linear programming. One challenge with this approach is the inability to incorporate secondary factors, such as connection time, which play an important role in passenger delays. The nature of the extreme point optimal solutions to the linear programming model also creates challenges, because a much larger proportion of flights end up being either empty or full as compared to the proprietary data. These limitations have led us to apply instead a discrete choice modeling approach. In a related context, Coldren, Koppelman and others have applied discrete choice models to estimate airline itinerary shares from booking data (Coldren, Koppelman, Kasurirangan, & Mukherjee, 2003 and Coldren & Koppelman, 2005). In the airline itinerary shares estimation problem, the goal is to predict the share of passenger demand for a market (i.e., all air travel from an origin to destination) that will utilize each of a set of available itinerary choices. Thus, the itinerary shares problem is more general in that all routes between the origin and destination are considered simultaneously. In our problem, due to the manner in which publicly available passenger flow data is aggregated, we are interested in estimating the share of passenger demand for a single carrier and route combination across different itineraries. Nonetheless, the success of the Coldren and Koppelman models suggest that application of a discrete choice model is reasonable in this area.

Other researchers have performed passenger delay analyses without first disaggregating passenger itinerary flows, but these approaches tend to require rather substantial assumptions. Sherry, Wang, and Donohue (2007) estimate passenger delays by treating all passengers as non-stop and assuming that all flights on an origin-destination segment operate at the monthly average load factor. Tien, Ball, and Subramanian (2008) develop a structural model of passenger delays, but in order to use the model are forced to make unverifiable assumptions regarding key parameter values (e.g., the delay thresholds for missed connections). Each of these approaches would benefit from access to estimated passenger itinerary data from which to enhance or validate the model.

Additional studies on air transportation passenger choice have helped us determine which features to include in our model. Theis, Adler, Clarke, & Ben-Akiva (2006) demonstrate that passengers travelling on one-stop itineraries are sensitive to connection times, specifically exhibiting a disutility associated with both short and long connection times. The referenced study by Coldren & Koppelman (2005) suggests that passengers prefer travelling on larger aircraft. Last, recent work has shown that flight cancellation decisions are affected by flight load factors – the fraction of seats filled on each flight (Tien, Churchill, & Ball, 2009). This suggests flight cancellations are an important factor to consider, because we would

expect fewer passengers to have been booked on canceled flights. That is, though we do not expect passengers to predict cancellations, in hindsight, cancellations provide valuable information regarding the *historical* distribution of passengers across itineraries.

## 1.2 Contributions

The contributions of our research fall broadly into two categories: i) an approach for disaggregating publicly available passenger demand data, and ii) an analysis of historical passenger delays using these disaggregate passenger itinerary flows. The outline of the paper follows this structure.

In Section 2, we describe the components of the passenger itinerary allocation process. First, we join passenger and flight data from multiple sources into a large Oracle database. Next, we process the data to establish the necessary inputs for passenger allocation, such as potential itineraries and flight seating capacities. Last, we develop a discrete choice model for passenger itinerary allocation, training and validating the results using a small set of proprietary booking data.

In Section 3, we utilize the disaggregated passenger itinerary flows to analyze domestic passenger delays for 2007. First, we extend the passenger delay calculator developed by Bratu and Barnhart (2005) to support a multi-day, multi-carrier rebooking process. Next, we analyze the sensitivity of our approach and validate the calculated delays against those estimated from the proprietary booking data. Last, we analyze passenger delays from 2007 to develop further insights into the relationship between flight delays and passenger delays and develop a simplified regression-based approach for estimating passenger delays directly.

We conclude the paper with a discussion of other problems to which this data is either already being applied or could be applied in the future.

## 2 Passenger Itinerary Allocation

In this section, we describe the process of allocating passengers to individual itineraries. We consider an itinerary to be a sequence of connecting flights that represents a one-way trip, including scheduled departure, connection (if any), and arrival times. Thus, round-trip travel would be represented by two one-way itineraries. To describe this process, we first define the following terminology.

- *carrier-segment*: the combination of an operating carrier, origin, and destination, where the operating carrier provides non-stop flight access between the origin and destination; and

- *carrier-route*: a sequence of *carrier-segments* that represents the flight path a passenger could travel from the origin of the first *carrier-segment* to the destination of the last *carrier-segment*.

With these definitions in mind, we can describe passenger itinerary allocation as the effort to combine *carrier-segment* demand data that is aggregated monthly with *carrier-route* demand data that is aggregated quarterly to allocate passengers to plausible itineraries. For example, a plausible one-stop itinerary would be taking the 7:55am American Airlines flight from Boston Logan (BOS) to Chicago O’Hare (ORD) followed by the 11:15am flight from Chicago O’Hare (ORD) to Los Angeles (LAX) on Thursday, August 9<sup>th</sup>. The *carrier-segment* data would tell us how many passengers traveled on American Airlines flight legs from BOS to ORD and ORD to LAX in August, whereas the *carrier-route* data would tell us how many passengers traveled on American Airlines from BOS to LAX connecting in ORD in the 3<sup>rd</sup> quarter of 2007. Note that when we discuss itineraries in this paper, we are including both the specific dates and times of travel in our definition of an itinerary. In Section 2.1, we describe each of the data sources in detail, followed by a description of the data processing in Section 2.2. In Section 2.3, we describe the methodological core of our paper – the discrete choice model used to allocate passengers to itineraries. Last, in Section 2.4, we validate the discrete choice allocations against a small set of proprietary booking data.

## 2.1 Data Sources

The U.S. Bureau of Transportation Statistics (BTS) provides a wealth of data related to airline travel (Bureau of Transportation Statistics). The Airline Service Quality Performance (ASQP) database provides planned and realized flight schedules for many airlines. Reporting is mandatory for all airlines that carry at least 1% of U.S. domestic passengers. For calendar year 2007, the database contains information for 20 airlines, ranging from Aloha Airlines with 46,360 flights to Southwest Airlines with 1,168,871 flights. BTS also maintains the Schedule B-43 Aircraft Inventory which provides annual lists of aircraft in inventory for most airlines. Most importantly for our purposes, the Schedule B-43 provides the seating capacity for each aircraft, matching approximately 75% of the flights in ASQP by tail number. We cannot match 100% of flights this way, because tail number information is sometimes inaccurate or non-existent in both ASQP and Schedule B-43.

The Federal Aviation Administration (FAA) maintains the Enhanced Traffic Management System (ETMS) database, which includes schedule information for all flights tracked by air traffic control. This database is not publicly available, due to the presence of sensitive military flight information, but a filtered version was made accessible for the purposes of this research. The benefit of this database over ASQP is that, in addition to the planned and realized flight schedules, it contains the International Civil

Aviation Organization (ICAO) aircraft equipment code for each flight. Using the ETMS database, we are able to determine the ICAO aircraft code and seating capacity for many of the 25% of ASQP flights that could not be matched through Schedule B-43.

There are two BTS datasets that we depend on for passenger demand information. The first is the T-100 Domestic Segment (T-100) database, which contains passenger and seat counts for each *carrier-segment* and equipment type aggregated monthly. For example, from this data we can see that in September 2007, American Airlines performed 79 departures from BOS to ORD using Boeing 757-200s with 14,852 seats available and 11,215 passengers. T-100 is a particularly useful database in that it contains information on both passenger counts and aircraft types. If the variation in seating capacity is sufficiently low for a *carrier-segment*, we estimate the seating capacity of each matching flight by dividing the number of seats available by the number of departures performed. We say that the variation in seating capacity is sufficiently low if the coefficient of variation (the standard deviation divided by the mean) is less than 2.5%. By combining T-100 with the data from Schedule B-43 and ETMS described above, we are able to estimate accurate seating capacities for approximately 98.5% of the ASQP flights. For the remaining 1.5% of ASQP flights, because the variation in seating capacity is high, we use the T-100 data to estimate a seating capacity that is slightly higher than average. For these flights, the seating capacity we use equals the average seating capacity across the matching T-100 rows plus one standard deviation to account for variation across aircraft types. The second passenger demand database we depend on is the Airline Origin and Destination Survey (DB1B), which provides a 10% sample of domestic passenger tickets from reporting carriers, including all of the carriers in ASQP, aggregated quarterly by removing information on flight times. For example, in the 3rd quarter of 2007, 128 passenger tickets were sampled that included a one-way trip on American Airlines from BOS to ORD to LAX. We use this data to determine the approximate number of monthly passengers travelling on each non-stop or one-stop *carrier-route*.

The last data set we use contains proprietary booking data from a large network carrier for the 4<sup>th</sup> quarter of 2007. This data is used for training our passenger flow estimation model and for validating our results. All of the data sets, including the proprietary booking data, are joined using a large Oracle database.

## **2.2 Data Processing**

There are two data processing steps that are performed prior to the discrete choice allocation of passengers to itineraries. The first step is estimating the set of potential itineraries on which passengers might have travelled. These itineraries will be used to create the choice sets in the discrete choice allocation model. The second step is estimating the number of passengers travelling on each *carrier-*

*route* for each month. These passengers will be allocated to matching itineraries using the discrete choice allocation model described in the next section.

We generate the set of potential itineraries for the year based on the flights in ASQP and the *carrier-routes* represented in DB1B. For the purpose of our analysis, we include only non-stop and one-stop itineraries. Itineraries with more than one stop account for only 2.5% of the one-way trips in DB1B. A non-stop itinerary is generated for each flight in ASQP, whereas a one-stop itinerary is generated only for valid flight pairs. The rules we use for determining valid flight pairs are as follows:

1. The *carrier-route* represented by the flight pair exists within DB1B. This filters out nonsensical routes, such as BOS – IAH – PVD (Boston to Houston to Providence), but allows for multi-carrier and code-share itineraries as long as at least one DB1B passenger utilized the corresponding multi-carrier *carrier-route*.
2. The planned connection time (the difference between the planned departure time of the second flight and the planned arrival time of the first flight) is at least 30 minutes and not more than 5 hours.
3. For a given first flight and matching *carrier-route*, we generate at most 2 connections. This ensures that passengers are not assigned to a much longer connection when multiple shorter connection times are available. Passenger utility associated with connection time is also explicitly considered within our discrete choice allocation model.

Using the 2007 ASQP and DB1B data sets, this procedure leads to 273,473,424 itineraries, of which 7,455,428 are non-stop. These itineraries are stored in our Oracle database for ease of querying during passenger flow estimation.

We estimate the number of passengers travelling each month on each *carrier-route* as follows.

1. First, for each *carrier-segment*,  $s$ , we calculate a monthly scaling factor,  $\alpha_s$ . The scaling factor is calculated as the ratio between the monthly *carrier-segment* demands specified by T-100 and the 10%-sample quarterly *carrier-segment* demands calculated from DB1B. For all itineraries represented in the DB1B 10% ticket sample, we aggregate the number of passengers on each *carrier-segment*. If DB1B sampled across *carrier-segments* uniformly, including international itineraries, and there were no monthly variations across *carrier-segments* in T-100, we would expect this ratio to equal 3.33 for all *carrier-segments*. That is, we would expect 3.33 times the DB1B sampled, quarterly *carrier-route* demand to be a good estimate of the total, monthly *carrier-route* demand (i.e., multiplying by 10 to account for the sampling and then dividing by 3,



the number of months in a quarter). Instead, because DB1B samples only domestic itineraries and there are no guarantees that the sampling is exactly 10% for each *carrier-segment*, we find that the calculated ratio varies around a mean of approximately 4.1.

2. For each one-stop *carrier-route* represented in DB1B, we estimate the monthly passenger demand by first scaling the 10%-sample quarterly DB1B passenger counts by the minimum  $\alpha_s$  across the corresponding sequence of monthly *carrier-segments*. Because we use the minimum scaling factor across all *carrier-segments* in the *carrier-route*, this approach underestimates the number of one-stop passengers. We prefer this approach, because other approaches cause the scaled then aggregated *carrier-route* demands to exceed the original *carrier-segment* demands (i.e., they allocate too many passengers on the *carrier-segment*). To resolve the undercounting of one-stop passengers, we subsequently apply a uniform scaling to all *carrier-route* demands to ensure that the percentage of one-stop passengers is consistent with DB1B.
3. For each non-stop *carrier-route* represented in DB1B, we calculate the monthly passenger demand by subtracting the passenger allocated on all matching *carrier-routes* based on the estimated one-stop demands in step 2 from the total *carrier-segment* demands provided by T-100 the one-stop demands estimated in step 2 for all *carrier-routes* that contain the corresponding *carrier-segment*. Note that because some one-stop *carrier-route* passengers are not able to be allocated to matching itineraries (e.g., due to lack of available seats), we wait to calculate the non-stop *carrier-route* demands until the one-stop allocation is complete. This ensures that, when aggregated by *carrier-segment*, our estimated *carrier-route* demands match the original T-100 data set.

In the next section, we describe the discrete choice allocation model we use for allocating the monthly *carrier-route* passengers to the generated itineraries.

### 2.3 Discrete Choice Allocation

As described in the previous section, for each month and *carrier-route*, we estimate passenger demands and generate a set of potential passenger itineraries. Next, we estimate the number of passengers corresponding to each itinerary. For each itinerary choice,  $i$ , we assign a passenger utility,  $u(x_i)$ , based on features of the corresponding itinerary. Then, for each passenger, we randomly select an itinerary choice from the ones matching the passenger's *carrier-route* according to the proportions,  $P(i)$ , described by the discrete choice multinomial logit function in Equation 2.1.

$$P(i) = \frac{e^{u(x_i)}}{\sum_j e^{u(x_j)}}, \quad \forall \text{ itineraries } i. \quad (\text{Equation 2.1})$$

The utility function,  $u(\cdot)$ , we use for our discrete choice model includes parameters for the interaction of the local time of departure and day of week, parameters for a piecewise linear function of connection time (to model the disutility associated with short and long connection times), as well as parameters for flight cancellations and seating capacities. To describe the utility function, we first define the following notation:

- $x_i^{\text{day}}$  = the day of week for itinerary  $i$  with Sunday = 1 and Saturday = 7;
- $x_i^{\text{time}}$  = the local time of departure for itinerary  $i$ ;
- $x_i^{\text{connect}}$  = the connection time for itinerary  $i$  with  $x_i^{\text{connect}} = 0$  for non-stop itineraries;
- $x_i^{\text{cancel}}$  = 1 if any flight in itinerary  $i$  was marked as canceled in ASQP, 0 otherwise;
- $x_i^{\text{seats}}$  = the minimum seating capacity for the flights in itinerary  $i$ ;
- $\mathcal{T}_n$  = the  $n^{\text{th}}$  four-hour daily block of departures, with  $\mathcal{T}_1 = 1:00 - 4:59\text{am}$  and  $\mathcal{T}_6 = 9:00\text{pm} - 12:59\text{am}$  the following day;
- $c_m$  = the  $m^{\text{th}}$  threshold for the piecewise linear utility for connection time in minutes with  $c_0 = 0$ ,  $c_1 = 45$ ,  $c_2 = 60$ ,  $c_3 = \infty$ ; and
- $\mathcal{I}(\cdot)$  = the indicator function for the expression argument.

Using this notation, the mathematical formulation of the itinerary choice utility function is provided in Equation 2.2. In the equation, the second sum represents the piecewise linear utility associated with connection times, with  $(\cdot)^+$  indicating the positive part of the inner expression, i.e.  $\max\{0, \cdot\}$ .

$$u(x_i) = \sum_{d=1}^7 \sum_{n=1}^6 \beta_{dn}^{\text{day-time}} \mathcal{I}(x_i^{\text{day}} = d) \mathcal{I}(x_i^{\text{time}} \in \mathcal{T}_n) + \sum_{m=1}^3 \beta_m^{\text{connect}} \left( \min\{x_i^{\text{connect}}, c_m\} - c_{m-1} \right)^+ + \beta^{\text{cancel}} x_i^{\text{cancel}} + \beta^{\text{seats}} x_i^{\text{seats}} . \quad (\text{Equation 2.2})$$

We include flight cancellations in our model, because all else being equal, a carrier is more likely to cancel a flight with fewer passengers than one with more. This is intuitively reasonable because there is a significant cost associated with rebooking the disrupted passengers. Thus, in hindsight, we would expect fewer passengers to have been scheduled to travel on itineraries that include a canceled flight. This decision is additionally supported by the passenger delay validation we perform in Section 3.2. Additionally, we include the minimum seating capacity on the itinerary as a measure of aircraft sizes.

The discrete choice model represented by Equation 2.2 is estimated with BIOGEME (Bierlaire, 2003) using one quarter of the proprietary booking data from a single major carrier extended to include unselected itineraries from our set of generated itineraries. Adding in the unselected itineraries is

important to accurately assess the disutility associated with very long connections. Without these itineraries it would appear that the very long connections are more strongly preferred, because unselected alternatives would not appear in the choice set. When extending the booking data, we consider only the generated itineraries with connection times of one hour or longer to eliminate choice set issues due to airport-specific minimum connection times. Additionally, this approach ensures that the distribution of connection times in our allocation aligns with the distribution of connection times in the booking data. Because there are often hundreds of choices for each month and *carrier-route*, we use sampling of alternatives to limit the size of the choice set to 10 alternatives for each observation, where each passenger in the booking data represents a single observation. Sampling of alternatives limits the computational effort required to train the model while still ensuring consistent parameter estimates. There is substantial literature on sampling of alternatives ranging from general applications (Ben-Akiva & Lerman, 1985) to specific considerations in a route choice context (Frejinger, Bierlaire, & Ben-Akiva, 2009). The estimated parameter values and statistics from this model are listed in Table A1.1 in Appendix 1. All of the parameter estimates are significant at the 0.99 confidence level using a classic Student's *t*-test, except for  $\beta_{5,3}^{\text{day-time}}$ , which is extremely close to zero. The overall model is also extremely significant, with a likelihood-ratio test value of 1563127.059 highly unlikely to occur under a  $\chi^2$  distribution with 46 degrees of freedom (corresponding to a p-value of less than  $10^{-30}$ ). Additionally, we feel that the parameter estimates are subjectively reasonable suggesting the highest utility for travel on Sunday, Thursday afternoon, and Friday, and the lowest utility for late night and pre-dawn travel.

Using the estimated parameters of this model, we calculate the utility associated with each of the generated itineraries and then, for each passenger, we sample a [0, 1] uniform random variable to select an itinerary allocation based on the proportions calculated using Equation 2.1. When a flight becomes full, we remove all corresponding itineraries from the choice set and update the expected proportions for the remaining itineraries. Because of this step, the order in which *carrier-routes* are processed is an important issue, as *carrier-routes* processed first are more likely to find seats available on all flights. Thus, to maintain the aggregate connecting percentage in the allocation, we process one-stop passengers before non-stop passengers. Within each group (i.e., one-stop or non-stop), we process a single passenger at a time. Passengers are processed in a random order, which reduces order-based biases in the results (e.g., having no seats available for *carrier-routes* that are processed last). To determine the random order, we sample a [0, 1] uniform random variable to set a priority for each individual passenger and then sort the passengers according to these priorities.

## 2.4 Validating Passenger Itinerary Flows

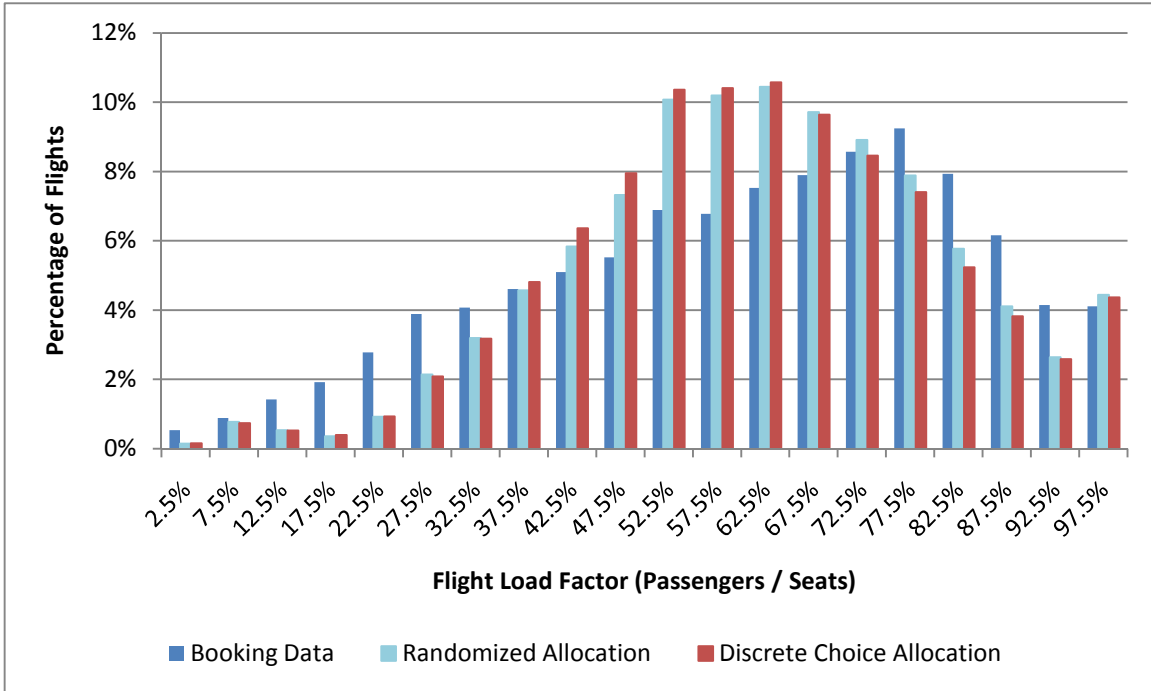
Unlike T-100, which includes passengers travelling on both domestic and international itineraries, our one quarter of proprietary booking data includes only domestic itineraries. Thus, the aggregate passenger counts for each *carrier-segment* and month are significantly lower than the T-100 data on average. Direct validation between the allocation described above and the proprietary data would lead to results that are heavily biased by this discrepancy. Instead, we perform a validation allocation where we scale the DB1B data (as described in Section 2.2) by the monthly *carrier-segment* passenger counts from the booking data (instead of T-100). Using this approach, the total number of validation passengers allocated is approximately the same as the number of passengers represented in the booking data.

For validation purposes, we are primarily concerned with distributional properties of our allocation approach. That is, because we have no way of determining the actual itinerary for each passenger, we instead focus on ensuring that our allocation is reasonable in an aggregate sense. We do so by comparing aggregate distributions of our validation allocation against the aggregate distributions of the proprietary booking data. The distributions we consider are:

1. Distribution of flight load factor,
2. Distribution of daily average load factor,
3. Distribution of percentage of connecting passengers, and
4. Distribution of connection time for one-stop passengers.

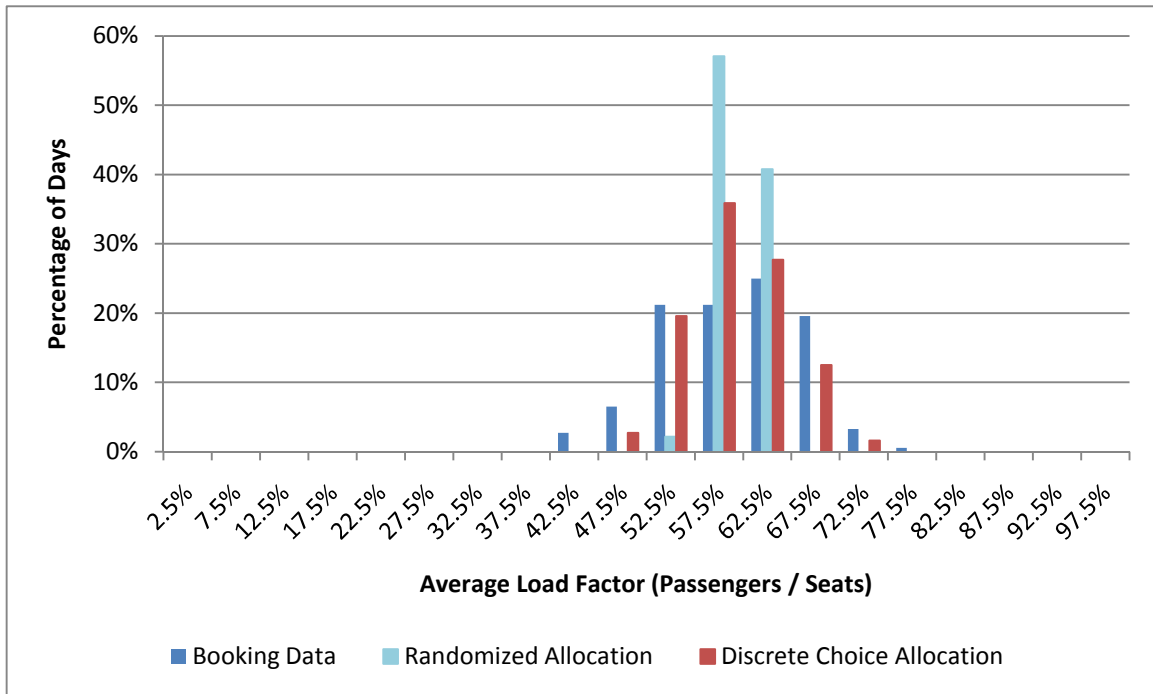
For each of these distributions, we compare our validation allocation to the booking data and to a randomized allocation (in which we assume all itinerary utilities,  $u(x_i)$ , are equal). The randomized allocation allows us to test the sensitivity of our approach to the individual parameter values of our discrete choice model.

In Figure 2.1, we consider the distribution of flight load factor. In the plot, we bucket flights by load factor in increments of 5%, with the percentage on the  $x$ -axis representing the mid-point of the bucket. The  $y$ -axis lists the percentage of flights falling into each bucket. With regards to load factors, the discrete choice allocation performs similarly to the randomized allocation. Each of these approaches under-estimates the number of flights with load factors between 0% and 35% and between 75% and 95%, and over-estimates the number of flights with load factors between 40% and 70%. We believe these discrepancies are due to the impacts of revenue management and our inability to model price as a dependent feature in our allocation model. Nonetheless, each of these approaches appears to perform quite well.



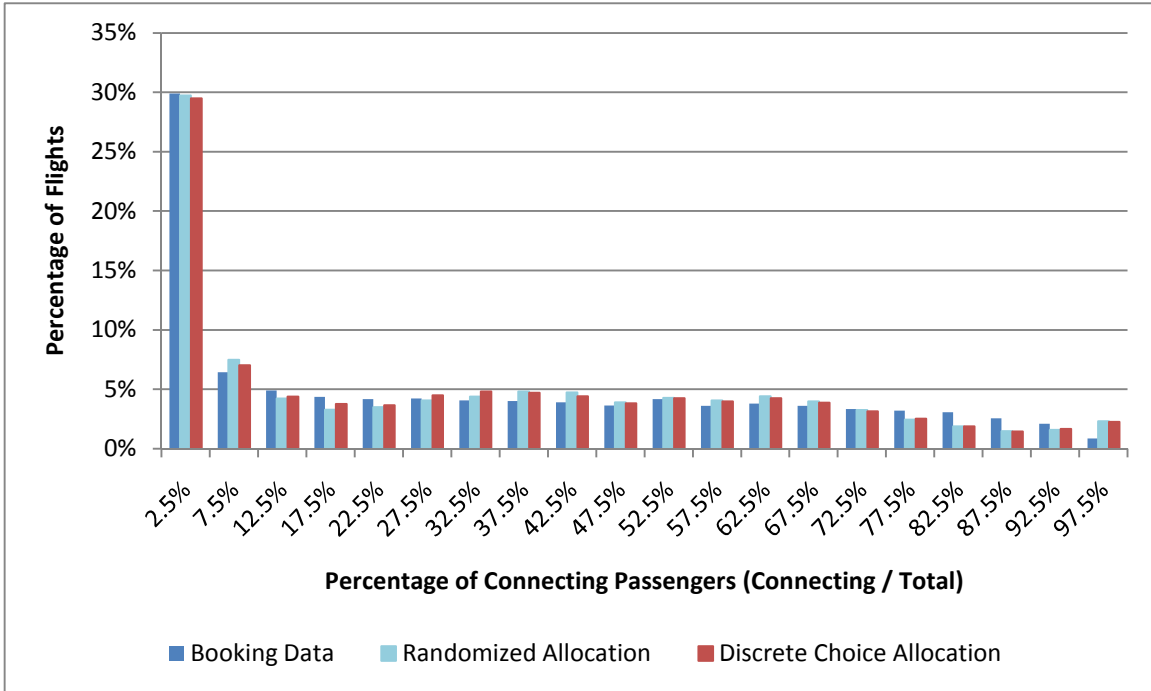
**Figure 2.1: Distribution of flight load factor for booking data, randomized allocation, and discrete choice allocation.**

In Figure 2.2, we consider the distribution of daily average load factor – the fraction of seats filled each day. As with flight load factors, daily average load factors are grouped into 5% buckets. In this plot, we see a significant improvement using the discrete choice model as compared to a randomized allocation. Although the discrete choice model increases the spread of average daily load factors, the proprietary booking data suggests even further variability. The proprietary booking data we use is from the fourth quarter of 2007, thus we attribute this additional variability to the impact of holiday travel. Because each holiday impacts travel differently, and because we have access to only one quarter of booking data, we do not attempt to model the impact of holiday travel directly.



**Figure 2.2: Distribution of daily average load factor for booking data, randomized allocation, and discrete choice allocation.**

In Figure 2.3, we consider the distribution of connecting passengers. That is, for each flight we determine the percentage of passengers on the flight who are on the first or second leg of a one-stop itinerary. The flight connecting passenger percentages are subsequently bucketed as with load factors above. Other than the 95% to 100% bucket, both the randomized allocation and the discrete choice allocation match the booking data well. The higher percentage of flights filled with connecting passengers is most likely due to our decision to allocate one-stop passengers prior to non-stop passengers (as described in Section 2.3). If we were to consider a larger bucket from 85% to 100%, all three data sets correspond quite well (5.5% for booking data, 5.4% for randomized allocation, and 5.4% for discrete choice allocation).



**Figure 2.3: Distribution of connecting passengers for booking data, randomized allocation, and discrete choice allocation.**

In Figure 2.4, we consider the distribution of connection times for one-stop passengers. In this plot, we consider only one-stop passengers, and bucket these passengers by the connection time between the two flights. In this plot, we utilize 10-minute buckets, with each point on the  $x$ -axis specifying the mid-point of the bucket. The  $y$ -axis lists the percentage of one-stop passengers falling into each bucket. This plot demonstrates the power of the discrete choice modeling approach. Using this approach we are able to very accurately match the distribution of connection times that exists in the proprietary booking data. The randomized allocation exhibits no preference towards connection times, so all variation is due to availability of connections, with connection times of an hour to an hour and a half being the most frequent.

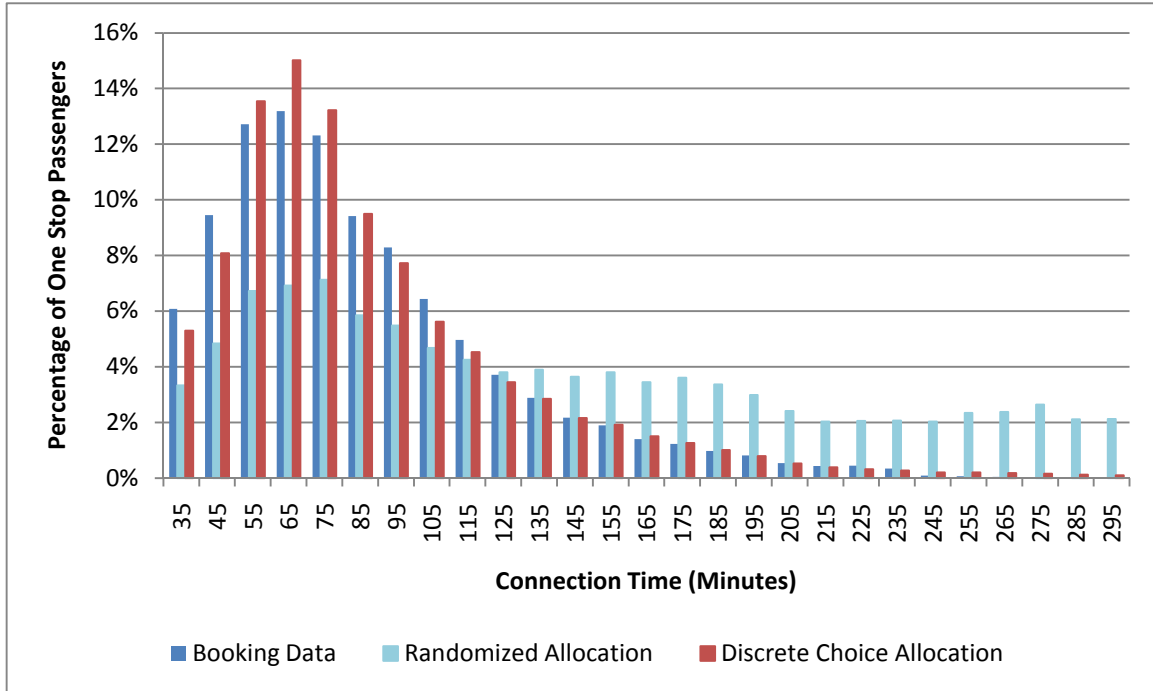


Figure 2.4: Distribution of connection times for booking data, randomized allocation, and discrete choice allocation.

In Section 3.2, we extend our validation to the analysis of passenger delays. In the context of estimating passenger delays, the number of passengers traveling on itineraries with flight cancellations is critical due to the immense impact of itinerary disruptions. Our results show that the inclusion of flight cancellations in the model allows us to pick up on an important factor, namely the tendency of airlines to cancel flights with fewer passengers, further justifying our discrete choice approach.

Overall, it appears that the discrete choice model and randomized allocation both do quite well according to system-wide metrics such as the distribution of load factors and connecting passengers. On the other hand, for more disaggregate measures, such as connection times or daily load factors, the discrete choice model appears to dominate the randomized allocation.

### 3 Analyzing Passenger Delays

In this section, we turn our focus from estimating passenger travel to the analysis of passenger delays. First, in Section 3.1, we describe the procedure we use to calculate passenger delays based on the estimated passenger itinerary flows developed in Section 2. Next, we use the estimated passenger delays to further validate our allocation approach. In Section 3.3, we highlight a few key findings from our passenger delay analysis, such as how passenger delays break down annually, by carrier, by airport, by time of year, by day of week, and by time of day. Last, in Section 3.4, we develop a linear regression



model to estimate passenger delays directly, allowing us to bypass the passenger allocation and delay calculation procedures for certain estimation tasks.

### **3.1 Passenger Delay Calculator**

The procedure we use for calculating passenger delays is an extension of the passenger delay calculator developed by Bratu and Barnhart (2005). In order to calculate the passenger delays associated with the estimated passenger itinerary flows, the authors used the realized flight schedules in ASQP, which provide information about flight delays, flight cancellations and diversions. The original passenger delay calculator was applied to only a single carrier, and thus assumed a default passenger delay value for passengers accommodated on a different carrier. For the purpose of our study, involving all of the 20 ASQP-reporting domestic carriers, we extend the algorithm to estimate the delays for passengers rebooked on a carrier different than planned.

The first step in passenger delay calculation is to determine which passengers have had their itinerary disrupted and therefore need to be re-accommodated to their final destinations. A non-stop itinerary is disrupted only if the corresponding flight is canceled or diverted. A one-stop itinerary is disrupted if one or both of the two flights is canceled or diverted, or if the first flight is delayed to such an extent that the corresponding passengers are unable to make their connection to the second flight (we assume this is the case if the available connection time is less than 15 minutes). For non-disrupted itineraries and the corresponding passengers, the passenger delay is simply equal to the (non-negative) flight delay associated with the last flight in the itinerary. Note that for passengers on non-disrupted one-stop itineraries this means that delay on the first leg is absorbed into the planned connection time. If an itinerary is disrupted, each of the passengers on the itinerary must be re-accommodated from the point of disruption to the final destination of the itinerary. We assume that each of these passengers is re-accommodated on the best available alternative itinerary, where best is defined as the alternative scheduled to arrive the earliest. The passenger delay for these passengers is then the time they reach their final destination minus the planned arrival time, ignoring negative values. Thus, the primary work of the passenger delay calculator is the re-accommodation of passengers whose itineraries have been disrupted.

Disrupted passengers are re-accommodated in an order based on the itinerary's time of disruption. For canceled or diverted flights, we use the planned departure time as the time of disruption. For missed connections, we use the actual arrival time of the first flight as the time of disruption. Note that rather than ignoring the diversions, we treat them the same as cancellations. This is not due to any limitations with the algorithm, but to the fact that the destination to which the flight is diverted is not provided in ASQP. The number of flight diversions is equal to about 10% of the number of flight cancellations (or

about 0.23% of total flights), so we do not expect the method of handling diversions to impact the final results significantly, as long as they are not ignored entirely. Under this assumption, the point of disruption for canceled or diverted flights is the origin of the flight, whereas for missed connections, the point of disruption is the planned connecting airport.

One challenge in calculating passenger delays is that the ASQP database does not include all possible flight options, such as those of non-reporting carriers. Therefore, in order to be conservative in our estimates, we put a limit on the re-accommodation delay for each disrupted passenger based on the time of disruption. For passengers disrupted during daytime hours, between 5:00am and 5:00pm, we limit the re-accommodation delay to 8 hours. For passenger disrupted during evening or pre-dawn hours, between 5:00pm and 5:00am, we set the limit to 16 hours to allow for re-accommodation the following day. Then, in order to re-accommodate each passenger, we check if there are any valid recovery itineraries amongst the 273,473,424 itineraries generated in Section 2.2. A recovery itinerary is valid if it departs from the point of disruption at least 45 minutes after the time of disruption (to allow time for rebooking and transfer), has available seat(s), and is scheduled to arrive at the passenger's final destination in time to satisfy the re-accommodation delay limit. We first search for itineraries that use airlines matching the original itinerary (e.g., the two carriers on a multi-carrier one-stop itinerary), along with any sub-contracted or parent airlines. For example, when a Continental itinerary is disrupted, we look for recovery itineraries on Continental or ExpressJet (or any combination of the two). If we are unable to find a valid itinerary using these airlines, we attempt to re-accommodate the passengers using any valid itinerary across all carriers. If we are unable to find an alternative at this point, we assign the passenger a delay equal to the re-accommodation delay limit, assuming that he or she will be re-accommodated in some other fashion. For passengers who are recovered on a new itinerary, the passenger delay is calculated based on the actual arrival time of the recovery itinerary. Note that we allow disruption chaining, that is we allow for the possibility that the recovery itinerary to which a passenger is assigned may in turn get disrupted and the passenger may be required to be rebooked again. Although we allow such disruption chaining in our passenger delay calculator, we maintain the re-accommodation delay limit throughout. Thus, passengers are often defaulted to the re-accommodation delay limit after a second disruption. This ensures that our disruption chains do not become overly long, because in many cases airlines have knowledge of future disruptions at the time of re-accommodation (e.g., a severe weather event that is projected to last throughout the day).

### **3.2 Validating Passenger Delays**

In this section, we validate against three potential sources of error in our end-to-end approach for estimating passenger delays. First, we consider the sensitivity of our passenger delay estimates to the re-

accommodation delay limits described in the previous section. Second, we validate our passenger delay estimates against those estimated from the proprietary bookings data. The purpose of this validation is to ensure that there are not important factors missing from our discrete choice allocation. Last, we measure the impact of our discrete choice sampling variation on our aggregate passenger delay estimates to ensure that a single allocation is sufficient for our subsequent analyses. That is, we want to ensure that at the levels of aggregation we are interested in, the variance between samples is low.

As described in the previous section, each time a passenger is disrupted, the passenger delay calculator attempts to re-accommodate the passenger on an alternative itinerary for which the planned arrival time satisfies the re-accommodation delay limit. Because these limits have been introduced with the intent to be conservative in our estimates, the results are sensitive to the limits chosen. In Table 3.1, we compare delay estimates for 2007 utilizing different re-accommodation delay limits for daytime (5:00am to 5:00pm) and evening (5:00pm to 5:00am) disruptions. In each column of the table, we list the corresponding daytime and evening delay limits separated by a “/”, so “8 hour / 16 hour” would refer to the limits described in Section 3.2. Even when these limits are increased to 24 hours for both daytime and evening disruptions, we are still unable to find alternative itineraries for just over 8% of the disrupted passengers, because there are either no flights or no seats available. This explains why the choice of re-accommodation delay limit has such a significant impact on the estimated delay for disrupted passengers.

	<b>6 hour / 12 hour Delay Limits</b>	<b>8 hour / 16 hour Delay Limits</b>	<b>12 hour / 24 hour Delay Limits</b>	<b>24 hour / 24 hour Delay Limits</b>
Average Passenger Delay (min)	27.78	30.15	33.09	37.47
Average Disrupted Passenger Delay (min)	376.33	448.63	538.41	672.08
Average Daytime Disrupted Passenger Delay (min)	259.37	308.17	385.92	589.89
Average Evening Disrupted Passenger Delay (min)	554.36	662.42	770.52	797.19
% of Passengers Disrupted	3.3%	3.3%	3.3%	3.3%
% of Disrupted Passengers Receiving Default Delay	33.5%	20.7%	15.5%	8.0%

**Table 3.1: Comparison of delay estimates using different daytime / evening delay limits in the passenger delay calculator.**

Of the daytime / evening limit combinations that we tested, we chose to use the 8 hour / 16 hour delay limits as the basis for reporting our results for the following reasons. First, the 8 hour daytime limit ensures that passengers who are both scheduled to arrive and subsequently disrupted between the hours of 5:00am and 5:00pm are rescheduled (either on an alternative or by default) to reach their destination

before the following morning. This would not hold if the daytime delay limit were larger than 12 hours. The 16 hour evening limit ensures that, for passengers who are disrupted during the evening, the passenger delay calculator will consider at least a few hours' worth of rebooking alternatives the following morning. This would not hold for evening limits of 12 hours or less. Last, based on the fact that 20% of the disrupted passengers are not re-accommodated within the system, we feel that these are sufficiently conservative values.

Next, in order to further validate our allocation approach, we consider our estimated passenger delays alongside those estimated from the proprietary booking data. Because the proprietary booking data does not contain information on passenger re-accommodations, we cannot compare our passenger delay estimates directly to actual delays. Instead, we first estimate the passenger delays associated with the proprietary bookings data by applying the passenger delay calculation procedure from the previous section. The intent of this validation is to ensure that any itinerary-level differences in passenger counts are washed out by the aggregate passenger delay calculations. Thus, in Table 3.2, we compare the passenger delay results based on our estimated passenger flows to those based on proprietary booking data for one major carrier for the 4<sup>th</sup> quarter of 2007. For each of these two data sets, we list the number of passengers impacted by each type of disruption (or lack thereof) as well as the total number of hours of delay accumulated. For the case of disruption chaining, we categorize disrupted passengers and their delays based on the cause of the first itinerary disruption (i.e., cancellation / diversion or missed connection). By construction, as described in Section 2.4, the total number of passengers is very close (within 0.4%), but for all other categories, the error ranges from 2.0% to 4.4%, which we believe to be quite good. One interesting result is that our estimated passenger delays appear to be consistently biased high across all categories. This suggests that airlines do (slightly) more to mitigate passenger delays than we are able to pick up in our allocation approach. For example, much like with cancellations, airlines may choose to push delays to flights with fewer passengers.

Cause	Passenger Counts			Delays (Hours)		
	Booking Data	Estimated Flows	Percentage Difference	Booking Data	Estimated Flows	Percentage Difference
<b>Flight Delays</b>	7,113,553	7,141,404	0.39%	1,968,253	2,007,925	2.02%
<b>Flight Cancellations</b>	114,654	119,174	3.94%	933,486	962,681	3.13%
<b>Missed Connections</b>	80,439	77,082	4.17%	558,722	583,296	4.40%

<b>Total</b>	7,308,646	7,337,660	0.40%	3,460,460	3,553,903	2.70%
--------------	-----------	-----------	-------	-----------	-----------	-------

**Table 3.2: Passenger delays by cause**

Last, because the passenger itinerary allocation methodology is based on a probabilistic allocation, the estimated passenger flows and hence the calculated passenger delays are not deterministic but rather are subject to sampling errors. For instance, the validation results presented in Table 3.2 are for a single allocation. Thus, it is critical that we also identify the extent of errors due to sampling. To do this, we perform two allocations using the entire year’s worth of data and calculate the passenger delays associated with each of these allocations. To determine the sampling error, we aggregate the passenger delays on a daily, monthly, and annual basis. From these aggregated delays, we calculate a percentage error equal to the absolute value of the difference divided by the smaller of the two. In Table 3.3, we summarize these errors by presenting the minimum, maximum, average and median percentage error for each aggregation level. Note that for the annual delays, there is just one aggregated delay value for each sample, as opposed to 365 for daily and 12 for monthly. The table demonstrates that the sampling errors are very low, even when calculated on a daily basis. Moreover, as we would expect, the range of sampling errors decreases significantly as the level of aggregation increases. This suggests that sampling error is not a significant source of concern, especially for the levels of aggregation we consider in the following section (e.g., annual, by carrier, or by month).

<b>Aggregation Level</b>	<b>Minimum</b>	<b>Maximum</b>	<b>Average</b>	<b>Median</b>
<b>Daily</b>	0.0034%	2.0780%	0.3948%	0.3309%
<b>Monthly</b>	0.0149%	0.1611%	0.0729%	0.0599%
<b>Annual</b>	0.0472%	0.0472%	0.0472%	0.0472%

**Table 3.3: Summary of sampling error in passenger delay estimates**

### **3.3 Passenger Delay Results**

In this section, we first summarize the results of our passenger delay calculations. Next, using the results of our allocation and delay calculation process, we discuss several key findings from our analysis. These findings serve two purposes: 1) to further the understanding of the complex performance characteristics of the National Air Transportation System, and 2) to demonstrate the breadth of analytical possibilities based on the methodologies we have developed.

In Table 3.4, we present various characteristics of our passenger delay estimates by carrier (and in aggregate) for calendar year 2007. The 20 carriers listed are those that are represented in the ASQP data set. Each passenger is identified based on the carrier operating the first flight in the itinerary. Passengers traveling on multi-carrier itineraries account for approximately 8.9% of all passengers allocated. As one

might expect, the average delay for non-disrupted passengers is typically quite close to the average flight delay. However, including disrupted passengers, an average passenger experiences approximately twice as much delay as an average flight, due to the impact of flight cancellations and missed connections. We estimate that flight cancellations are responsible for approximately 30% of passenger delays, whereas missed connections account for approximately 18%.

Carrier	Flights	Avg Flight Delay (min)	Passengers	Avg Passenger Delay (min)	Percent of Total Delay due to Cancellations	Percent of Total Delay due to Missed Connections	Avg Disrupted Passenger Delay (min)
Pinnacle Airlines (9E)	258,851	13.81	6,563,926	41.33	38.45%	26.76%	465.86
American Airlines (AA)	633,857	19.15	59,548,442	40.28	35.18%	16.39%	446.98
Aloha Airlines (AQ)	46,360	4.11	3,727,087	6.59	25.81%	6.06%	253.36
Alaska Airlines (AS)	160,185	14.92	14,600,205	22.87	26.37%	7.87%	395.33
JetBlue Airways (B6)	191,450	20.37	19,019,504	33.50	33.79%	6.76%	562.65
Continental Airlines (CO)	323,151	16.38	30,566,023	26.98	22.57%	15.17%	474.54
Delta Airlines (DL)	475,889	12.63	44,150,213	28.21	26.63%	22.40%	438.14
Atlantic Southeast Airlines (EV)	286,234	21.86	7,108,944	53.20	31.31%	37.62%	440.87
Frontier Airlines (F9)	97,760	11.52	7,780,342	17.85	12.12%	20.36%	473.93
AirTran Airways (FL)	263,159	13.41	18,522,194	25.34	23.78%	20.57%	477.73
Hawaiian Airlines (HA)	56,175	4.25	5,860,062	9.29	13.95%	13.26%	300.29
American Eagle Airlines (MQ)	540,494	17.81	12,393,301	49.78	44.90%	20.14%	447.81
Northwest Airlines (NW)	414,526	16.30	32,507,200	34.60	29.79%	23.00%	481.22
Comair (OH)	236,032	17.81	6,934,004	46.09	42.66%	19.53%	509.79

Carrier	Flights	Avg Flight Delay (min)	Passengers	Avg Passenger Delay (min)	Percent of Total Delay due to Cancellations	Percent of Total Delay due to Missed Connections	Avg Disrupted Passenger Delay (min)
Skywest Airlines (OO)	597,880	13.43	14,542,785	37.28	34.35%	26.58%	461.24
United Airlines (UA)	490,002	18.26	44,245,649	36.77	30.87%	20.18%	436.66
US Airways (US)	485,447	16.28	40,894,218	31.83	27.09%	18.55%	428.94
Southwest Airlines (WN)	1,168,871	10.47	95,579,734	15.62	19.45%	10.11%	368.66
Expressjet Airlines (XE)	434,773	16.68	12,531,100	37.83	39.42%	17.58%	516.93
Mesa Airlines (YV)	294,362	16.44	9,457,714	41.45	40.07%	21.67%	447.20
<b>Total</b>	<b>7,455,458</b>	<b>15.32</b>	<b>486,532,647</b>	<b>30.15</b>	<b>30.36%</b>	<b>18.36%</b>	<b>448.63</b>

**Table 3.4: 2007 passenger delays by carrier**

Using the aggregated results in Table 3.4 combined with the disaggregated results derived from our approach, we highlight nine key findings regarding the breakdown and causes of passenger delays. In each case, we begin by stating the finding, and then providing further details, including any definitions or assumptions, as well as further discussion of the result.

**Key Finding #1:** *The ratio of average passenger delay to average flight delay is maximum for regional carriers and minimum for low-cost carriers, owing primarily to the cancellation rates and the connecting passenger percentages.*

As above, a passenger is identified based on the carrier operating the first flight in the itinerary. We categorize American Airlines (AA), Continental Airlines (CO), Delta Airlines (DL), Northwest Airlines (NW), United Airlines (UA), and US Airways (US) as the *legacy network carriers*; JetBlue Airways (B6), Frontier Airlines (F9), AirTran Airways (FL), and Southwest Airlines (WN) as the *low cost carriers*; and Pinnacle Airlines (9E), Atlantic Southeast Airlines (EV), American Eagle Airlines (MQ), Comair (OH), Skywest Airlines (OO), Expressjet Airlines (XE), and Mesa Airlines (YV) as the *regional carriers*.

Across all carriers in 2007, the ratio of average passenger delay to average flight delay is 1.97. For individual carriers, it ranges between 1.49 for Southwest Airlines (WN) and 2.99 for Pinnacle Airlines (9E). For the legacy network carriers, this ratio ranges from 1.65 to 2.23, with an average value of 2.03.

For regional carriers, it ranges from 2.27 to 2.99 with an average value of 2.61. Last, for low cost carriers, it ranges from 1.49 to 1.89 with an average value of 1.61.

The reasons for such disparity become clear when we look at the cancellation percentages and the percentages of connecting passengers. In the year 2007, the overall percentage of canceled flights was 2.4% and the percentage of connecting passengers was 27.2%. The regional carriers had both the greatest percentage of cancellations (3.4%) as well as the greatest percentage of connecting passengers (39.6%). Low-cost carriers had the lowest percentage of cancellations (1.2%) and the lowest fraction of connecting passengers (17.0%). Legacy network carriers fell between these two extremes, both for the percentage of cancellations (2.2%) and the percentage of connecting passengers (31.0%). As we show later in Section 3.4, the percentage of canceled flights and the percentage of connecting passengers are highly correlated with the ratio of average passenger delay to average flight delay.

**Key Finding #2:** *Passengers scheduled to transfer in one of 6 airports: Newark (EWR), Chicago O'Hare (ORD), New York La Guardia (LGA), Washington Dulles (IAD), New York Kennedy (JFK) or Philadelphia (PHL), were exposed to the longest average connecting passenger delays. For each of these airports, over 10% of scheduled connecting passengers had their itineraries disrupted. These 6 airports were also among the worst airports with respect to both average delays for departing flights and departure cancellations.*

We restrict this analysis to only the connecting passengers and consider data from the top 50 transfer airports in the U.S. These airports account for nearly 98.7% of all domestic connecting passengers in the U.S. On average, 12.2% of the passengers scheduled to connect through the 6 airports listed had their itineraries disrupted compared to just 6.9% at the remaining 44 airports. These airports were the worst transfer airports in terms of average connecting passenger delay. Across these 6 airports, the average delay per connecting passenger of 78.5 minutes was 32.9 minutes more than that at the remaining 44 airports (45.6 minutes). These 6 airports are among the 9 worst transfer airports in terms of departure cancellation rates and the 7 worst transfer airports in terms of average delays for departing flights. The worst transfer airports based on departure cancellation rates also includes Reagan (DCA), Boston (BOS), and Dallas / Fort Worth (DFW). DFW is also on the list of transfer airports with the worst average delays for departing flights, rounding out that list.

**Key Finding #3:** *Domestic passenger connections are highly concentrated at the top three transfer airports: Atlanta (ATL), Chicago O'Hare (ORD), and Dallas / Fort Worth (DFW), representing approximately 43.2% of planned passenger connections. As such, ATL, ORD, and DFW are responsible*



*for more than 40% of domestic missed connections, and contribute to more than 43% of all delays to connecting passengers.*

As above, we restrict this analysis to only connecting passengers and consider data from the top 50 transfer airports in the U.S. Approximately 43.2% of these connecting passengers connect either at ATL, ORD, or DFW. Consequently, the largest numbers of connecting passengers either miss their connections or are alternatively disrupted at one of these three airports, representing 44.5% of all disrupted connecting passengers and 40.5% of all misconnecting passengers. In comparison, only 15.3% of all connecting passengers connect at the next three largest transfer airports: Denver (DEN), Phoenix (PHX) and Houston (IAH), which contribute to 16.8% of the missed connections and 15.3% of the delays to connecting passengers. Among all transfer airports, ATL contributes the most to total connecting passenger delays (15.8%), because it has the highest number of connecting passengers (17.9% of all connecting passengers), even though its average connecting passenger delay is below average (43.9 minutes vs. 49.9 minutes on average). DFW contributes 13.0% of the total connecting passenger delays while servicing just 11.0% of scheduled connecting passengers due to a higher than average connecting passenger delay (59.07 minutes). For ORD, the discrepancy is even more extreme, as it services just 9.1% of all connecting passengers, but corresponds to 14.4% of the total connecting passenger delays. This substantial discrepancy is due to ORD having the second highest average connecting passenger delays (78.4 minutes behind EWR at 93.1 minutes).

**Key Finding #4:** *Average evening passenger delay is 86.8% greater than the average morning passenger delay. One important reason for this difference is the 89.4% greater average evening flight delay compared to the average morning flight delay. The other important reason is the greater ease of rebooking for the passengers disrupted in the morning compared to those disrupted in the evening, as reflected by 66.3% higher average disrupted passenger delay to evening passengers compared to that for the morning passengers.*

For this analysis, all passengers and flights are categorized as morning or evening depending on the planned departure time from their origin airport. Any passenger (or flight) with planned local departure time between midnight and 11:59 am is denoted as a morning passenger (or morning flight) while any passenger (or flight) with planned local departure time between noon and 11:59 pm is categorized as an evening passenger (or evening flight). Note that one-stop passengers are categorized depending on the planned departure time of the first flight in the itinerary. According to this definition, 41.0% of the flights were classified as morning flights and 43.8% of the passengers were classified as morning passengers.

The contribution of non-disrupted passenger delay to the total passenger delay depends mainly on the flight delays, while the contribution from the disrupted passengers depends on the percentage of disrupted passengers and average delay to disrupted passengers. For calendar year 2007, average delay for morning passengers was 20.3 minutes compared to 37.8 minutes for evening passengers. A large part of this difference can be attributed to the higher average delay to evening flights (18.5 minutes), compared to morning flights (9.8 minutes). In fact, 73.2% of overall flight delays are due to delays to evening flights.

The greater ease of rebooking is suggested by the fact that the average disrupted passenger delay to morning passengers is 320.3 minutes while that for the evening passengers is 532.6 minutes. This difference is, in part, due to the different maximum delay values used for morning and evening disruptions, though it is also heavily influenced by the near-term availability of rebooking alternatives. Evening passengers are much more likely to be disrupted at times where the next available rebooking alternative requires an overnight stay-over. Though the delay to disrupted passengers differs dramatically, the percentage of disrupted passengers does not differ much between morning (2.96%) and evening passengers (3.52%), which is due in part to the smaller difference between the percentage of canceled (or diverted) flights in the morning (2.1%) and evening (2.6%). As a result, the relative disparity between delays to morning and evening flights is greater than the disparity between delays to the morning and evening passengers. In other words, the ratio of average passenger delay to average flight delay in the morning (2.07) is slightly higher than that in the evening (2.04).

**Key Finding #5:** *The average passenger delay for the three months of summer and the three months of winter was 56% higher than for the remaining six months, with June being the worst month for air travel in terms of both total as well as average passenger delays.*

For this analysis we consider June, July, and August as the *summer months*; and December, January, and February as the *winter months*. Average passenger delay in the summer months was 37.4 minutes while that in winter months was 36.0 minutes. In the remaining 6 months, however, the average passenger delay was only 23.5 minutes. June and February were the only two months with average passenger delays greater than 40 minutes. On the other end of the spectrum, September and November were the only two months with average passenger delays of less than 20 minutes. In terms of total passengers, the summer and winter months fall on opposite extremes, with average passengers per month being 9.9% above the annual average in the summer and 10.3% below the annual average in the winter. The end result is that total monthly passenger delays are 36.5% higher during the summer as compared to an average month, whereas total monthly passenger delays are only 7.0% higher during the winter. If

load factors were to increase during the winter, it is possible that the winter months would become the worst months for travel based on passenger delays.

**Key Finding #6:** *Delay to the non-stop disrupted passengers depends on the ease of rebooking and therefore is lower for origin-destination pairs with higher daily frequency. Average delay to disrupted non-stop passengers on routes with at least 10 daily flights per carrier is 31.4% lower than the overall average for disrupted passengers, and on routes with at most 3 daily flights per carrier, it is 15.3% higher than the overall average.*

Disruptions to non-stop passenger itineraries occur due only to flight cancellations. The average delays for these disrupted passengers are dependent on the ease of rebooking which, in part, depends on the number of direct flights offered by the carrier for the corresponding origin-destination pair. The overall average delay to disrupted non-stop passengers is 443.6 minutes. When the carrier has a daily frequency of at least 10 flights for the origin-destination pair, this average decreases to 304.1 minutes. On the other hand, when the carrier has at most 3 flights per day, it is more difficult to obtain a suitable recovery itinerary, increasing the average delay of disrupted non-stop passengers to 511.5 minutes.

**Key Finding #7:** *The relative benefits of flight frequency in terms of the ease of rebooking depend significantly on load factors. On carrier-segments that have less than a 75% average load factor, average delays to disrupted non-stop passengers are 385.9 minutes as compared to 216.5 minutes when considering only those carrier-segments with 10 or more flights per day, representing a 43.9% improvement due to increased frequency. On carrier-segments with at least a 75% average load factor, a frequency of 10 or more flights per day leads to a relative reduction in disrupted non-stop passenger delays of only 16.9% (455.5 minutes vs. 378.5 minutes).*

For this analysis, we consider only those combinations of carriers and segments which have at least 2 flights per day. Average delays to disrupted nonstop passengers on carrier-segments with at least 75% average load factor (455.5 minutes) is 18% higher than on carrier-segments with less than 75% average load factor (385.9 minutes). On average, nonstop passengers disrupted on low load factor (less than 75% full), high frequency (10+ flights per day) carrier-segments experience 55.0% less delay than their low load factor, low frequency (2 - 6 flights per day) counterparts, whereas for high load factor (at least 75% full) carrier-segments, the increasing flight frequency only reduces average delays by 23.6%. That is, though increasing flight frequency is beneficial for all disrupted nonstop passengers, the impact is largest when there are ample seats available for rebooking.

**Key Finding #8:** *Monday and Saturday have, by far, the lowest ratio of average passenger delay to average flight delay and these are the only two days when the ratio is lower than the overall average value for the week. One part of the reason is the lower percentage of canceled flights and another is the significantly higher percentage of morning passengers on these two days.*

The ratio of average passenger delay to average flight delay on Monday is 1.75 and on Saturday it is 1.88. For the remaining five days of the week, this ratio ranges between 2.00 and 2.03 compared to an overall average of 1.97 for the week. One reason for this difference is the lower percentage of canceled (and diverted) flights on these two days; 2.2% for Mondays and 1.9% for Saturday, compared to the average of 2.5% for the remaining 5 days of the week. Another important reason for this difference is the higher percentage of morning travelers on these two days. On Monday and Saturday, 47.6% and 48.3% of passengers respectively are morning passengers while only 42.3% are morning passengers for the remaining 5 days. As discussed in Finding #4, average delays for morning passengers are significantly lower than for evening passengers, due to shorter flight delays and easier rebooking earlier in the day. It is interesting to note that on Monday, average flight delays (10.3 minutes) are almost equivalent to the average flight delays throughout the week (10.2 minutes), which implies that the difference in ratio is entirely due to reduced passenger delays.

**Key Finding #9:** *Southwest Airlines has the lowest average passenger delay, nearly 55% lower than its competitors, even though its average flight delay is only 36.3% lower than other airlines. The primary reason for the smaller magnitude of passenger delays is the relative infrequency of disruptions to passenger itineraries; both in terms of a much lower number of flight cancellations and much lower percentage of missed connections.*

Over the last 5 calendar years (2005 - 2009), Southwest Airline has been ranked as the airline with highest overall on-time flight arrival performance (BTS, 2010), among all the airlines that predominantly serve the continental United States. This excludes Aloha Airlines (AQ) and Hawaiian Airlines (HA), among all the ASQP reporting carriers. Southwest has also had the overall best on-time arrival performance for 4 out of these 5 years, including 2007, which is the year of our analysis. The following analysis is performed for all the ASQP reporting carriers, excluding Aloha Airline and Hawaiian Airlines.

Southwest Airlines has an average passenger delay of 15.6 minutes, less than half of the 33.7 minute average value for the other airlines. The primary driver of this difference is not the difference between the average flight delays of Southwest Airlines (10.5 minutes) and that of the other airlines (16.2 minutes), nor the difference in average delay to non-disrupted passengers for Southwest (11.1 minutes) and that of other airlines (17.2 minutes). The major driver of Southwest Airline's passenger on-time

performance is the relative infrequency of itinerary disruptions. For instance, only 1.0% of Southwest flights are canceled as compared to 2.8% of flights for other carriers. Consequently, the percentage of passengers on canceled flights for Southwest (0.9%) is nearly a third of that for other airlines (2.4%). In addition, the percentage of passengers missing a connection (out of all passengers) on Southwest (0.4%) is nearly one fourth of that for the other carriers (1.4%). This is due to the smaller percentage of one-stop passengers (15.5% for Southwest compared to 30.0% for the other carriers) and to the propensity for longer connection times (41.9% of passenger connections are longer than 1.5 hours for Southwest, compared to 36.1% for the other carriers). Thus, for Southwest airlines, only 29.6% of all passenger delay is due to itinerary disruptions, while for other airlines, delays due to itinerary disruptions are responsible for 50.9% of all passenger delays.

### 3.4 Regression Model for Passenger Delay Estimation

In this section, we develop a linear regression model to 1) identify critical characteristics of airline networks, schedules, and passenger itineraries that affect passenger delays; and 2) estimate passenger delays directly given public data, thus bypassing the process of passenger allocation and re-accommodation. This simplified process is schematically depicted along the right path in Figure 3.1.

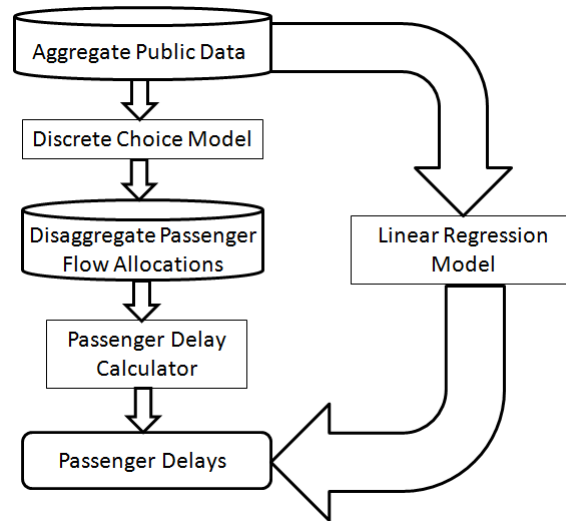


Figure 3.1: Linear regression model to bypass the passenger allocation and re-accommodation process.

Flight delays influence passenger delays in the most direct way. In the absence of cancellations and missed connections, and assuming equal numbers of passengers per flight, average passenger delay will equal average flight delay. Other factors such as cancellations and missed connections tend to increase average passenger delays beyond the average flight delay. In our regression models, we use the average passenger delay as the dependent variable,  $y_i$ .

We restrict the independent variables in our models to those which are available in public data sets such as ASQP, DB1B, etc. The purpose of this model is to see how well we can predict passenger delays based on the publicly available data without utilizing the complex allocation and delay calculation processes described in Sections 2.3 and 3.1 respectively. For model estimation, we use the results of our passenger delay calculations. Each observation corresponds to a single day and airline combination. Thus, the 20 airlines in our data and the 365 days in 2007 correspond to the availability of 7300 observations for estimating the model. To describe the model, we utilize the following notation, where each observation  $i$  corresponds to a carrier-day combination:

- $d_i^p$  = average passenger delay corresponding to observation  $i$ ;
- $d_i^f$  = average flight delay corresponding to observation  $i$ ;
- $LF_i$  = load factor for the carrier and month corresponding to observation  $i$ ;
- $f_i^{canc}$  = fraction of canceled flights corresponding to observation  $i$ ;
- $f_i^{conn}$  = fraction of connecting passengers for the carrier and quarter corresponding to observation  $i$ ;
- $f_i^{60}$  = fraction of flights with at least 60 minutes of delay corresponding to observation  $i$ ;
- $\mathcal{I}(\cdot)$  = the indicator function for the expression argument;
- $\beta_0$  = intercept; and
- $\beta_j$  = coefficient of independent variable  $x_j$ .

Average daily flight delays, the daily fraction of canceled flights, and the daily fraction of flights with at least 60 minutes of delay can be obtained from the Airline Service Quality Performance (ASQP) data for each carrier for each day. Average monthly load factors can be obtained from T-100 segment data for each carrier for each month. Average connecting passenger percentages can be obtained from DB1B data for each carrier for each quarter. Various model specifications were tested and we present here the model specification that was found to be the most suitable. The independent variables for this model are calculated as shown in Equations 3.1 through 3.5.

$$x_{i1} = d_i^f, \quad \forall i \quad \text{(Equation 3.1)}$$

$$x_{i2} = f_i^{canc}, \quad \forall i \quad \text{(Equation 3.2)}$$

$$x_{i3} = f_i^{canc} * \mathcal{I}(LF_i > 0.8), \quad \forall i \quad \text{(Equation 3.3)}$$

$$x_{i4} = f_i^{conn}, \quad \forall i \quad \text{(Equation 3.4)}$$

$$x_{i5} = f_i^{60} * f_i^{conn}, \quad \forall i \quad \text{(Equation 3.5)}$$

Using these definitions, the linear regression model is given by Equation 3.6.

$$\sum y_i = \beta_0 + \sum_{j=1}^5 \beta_j x_{ij}, \quad \text{for all observations } i. \quad (\text{Equation 3.6})$$

The estimation results for this model are as shown in the Table 3.5. The model has a very good fit with an  $R^2$  value of 0.9506. All of the coefficients are statistically highly significant, with each different from zero with at least a 99.99% confidence level.

Parameter Description	Parameter	Estimate	Std Error	p-value
Intercept	$\beta_0$	-1.34	0.24	0.00
Average flight delay	$\beta_1$	1.00	0.01	0.00
Fraction of canceled flights	$\beta_2$	458.77	2.92	0.00
Fraction of canceled flights $\times$ High load factor dummy	$\beta_3$	96.79	4.62	0.00
Fraction of connecting passengers	$\beta_4$	10.14	0.50	0.00
Fraction of connecting passengers $\times$ Fraction of flights with at least 60 minutes of delay	$\beta_5$	139.14	4.53	0.00

**Table 3.5: Regression parameter estimates**

The estimate of  $\beta_1$  is 1.00, suggesting that all else being equal, change in average flight delay results in equal change in average passenger delays. The magnitude of the estimate for  $\beta_2$  demonstrates that the fraction of canceled flights has a very strong impact on passenger delays. The greater the fraction of canceled flights, the greater the average passenger delay, because passengers on canceled flights must be re-accommodated on later itineraries. Once a passenger is disrupted, rebooking requires seat availability on alternate itineraries, which means that higher average load factors reduce the probability that the passengers will be quickly re-accommodated. Thus, we would expect cancellations to have a greater impact on passenger delays when load factors are high. This effect is demonstrated in our model results through the significant positive estimated value of  $\beta_3$ , which parameterizes the interaction of cancellations and load factors. In the publicly available data, there is no information on connection times, thus the best proxy for missed connections is the percentage of passengers with connections. In this context, the positive estimated value of  $\beta_4$  is reasonable because the higher the ratio of connecting passengers to total passengers, the higher the proportion of missed connections, and hence the higher the average passenger delay. Out of all the connecting passengers, the fraction missing their connection depends on the fraction of flights that have large delay. We would expect the fraction of connecting passengers to have a greater

impact on passenger delays when the fraction of flights with large delays is high. This effect is demonstrated through the significant positive value of estimate for  $\beta_5$

Next, we assess the error in passenger delay estimates obtained from this simple regression model using publicly available data. Using the estimated parameter values reported in Table 3.5, we calculate the average passenger delay for each carrier-day combination. Average passenger delays for each carrier for each month are then calculated using simple averaging of the daily values. T-100 segments database provides the total segment passengers for each month for each carrier, while the DB1B database provides the fraction of connecting passengers for each carrier for each quarter. Combining the two, we estimate the number of monthly passengers for each carrier. Multiplying the average passenger delays by the number of passengers for each carrier-month combination provides an estimate of the total passenger delay for each carrier-month. Aggregating across all carriers for the entire year, the estimate comes out to be 247,602,145 hours compared to the 244,468,965 hours estimated through the passenger allocation and delay calculation processes. That is, using our simplified regression approach, we are able to estimate annual delays within 1.28% of the totals listed in Table 3.4.

<b>Aggregation Level</b>	<b>Passenger Allocation and Delay Calculation</b>	<b>Regression-based Delay Estimation</b>
By Carrier-Day	11.1%	15.1%
Daily	10.3%	12.4%
Monthly	3.3%	8.0%
Quarterly	2.7%	8.0%

**Table 3.6: Summary of error in passenger delay estimates at different aggregation levels**

To further validate our regression approach, we compare the passenger delays obtained by the regression analysis with those obtained by applying the passenger delay calculator directly to the proprietary booking data. The percentage errors in estimates at different levels of aggregation are presented in Table 3.6. The second column lists the error in passenger delay estimates obtained by applying passenger allocation followed by delay calculation using our estimated passenger itinerary flows. The third column lists the error in passenger delays obtained from our simplified regression-based delay estimation approach. In both cases, the error is with respect to delay estimates based on the proprietary booking data. Table 3.6 demonstrates that the error decreases with increased level of aggregation for both approaches. At all aggregation levels, the errors are higher for the regression-based approach as compared to the passenger allocation and delay calculation approach. For the entire quarter, the regression model estimates are within 8.0% of the estimates based on the proprietary booking data. This



suggests that the simplified regression approach provides a good alternative for estimating total delays or the total cost of delays if a more thorough analysis is not required.

As an example of the potential applications of this simplified approach to passenger delay estimation, we applied the model to estimate the passenger delays for the year 2008. The model inputs such as flight delays, flight cancellations rates, connecting passenger percentages and load factors were obtained from public data for 2008. Passenger delays for the entire 2008 year were estimated using the regression parameter estimates listed in Table 3.5. Table 3.7 compares aggregate statistics on flight schedules and passenger itineraries for the years 2007 and 2008 and summarizes total passenger delays. For 2008, the estimated average passenger delays were 6.7% less than those for 2007, mainly due to 8.8% lower average flight delays and a 7.6% lower cancellation rate. However, because of a 6.0% reduction in the number of passengers, the total passenger delay for 2008 was estimated to be 12.2% lower than that for 2007.

	2007	2008	Difference
# Flights	7,455,458	7,009,726	-6.0%
Avg Flight Delay (min)	15.3	14.0	-8.8%
% of Flights Canceled	2.4%	2.2%	-7.6%
Avg Load Factor	76.6%	76.1%	-0.6%
% of Connecting Passengers	31.7%	32.5%	2.4%
% of Flights with at least 60 Minutes of Delay	7.2%	6.6%	-8.7%
Passengers	474,003,389	445,704,815	-6.0%
Total Passenger Delay (hours)	247,602,145	217,310,671	-12.2%
Avg Passenger Delay (min)	31.3	29.3	-6.7%
Ratio of Avg Passenger Delay to Avg Flight Delay	2.05	2.09	2.3%

Table 3.7: Delay estimates using the regression-based approach for 2007 and 2008

## 4 Further Applications

In this paper, we have developed methodologies for modeling and estimating historical passenger travel and delays. We have applied these approaches to analyze and develop insights into some of the key factors affecting passenger air travel performance. By developing and sharing this estimated passenger data, our goal is to enable and encourage further passenger-centric air transportation research, and we

believe this data has many applications. For example, beyond the analyses provided herein, we are applying the passenger data developed in this paper in multiple ways. First, in collaboration with George Mason University, we have used these approaches to estimate the overall costs of passenger delays as one component of NEXTOR’s Total Delay Impact Study (Ball, et al., 2010). This report was commissioned by the FAA and has been developed through collaboration between all of the universities affiliated with the National Center of Excellence for Aviation Operations Research (NEXTOR). The goal of the report is to provide a complete and rigorous assessment of the total costs of delays in order to inform aviation policy decision-making. Second, we are jointly analyzing flight and passenger data to better understand how airline network structures and scheduling decisions impact passengers. Last, we are using the estimated passenger itineraries to analyze how potential changes in Traffic Flow Management allocation policies propagate through to both airlines and passengers. In each of these areas, access to estimated passenger travel and delay data has enabled research opportunities that would not exist otherwise.

## Appendix 1

Parameter Description <sup>2</sup>	Parameter	Estimate	Std Error	p-value
Sunday, 1:00am – 4:59am	$\beta_{1,1}^{\text{day-time}}$	-0.183	0.07030	0.01
Sunday, 5:00am – 8:59am	$\beta_{1,2}^{\text{day-time}}$	-0.132	0.00730	0.00
Sunday, 9:00am – 12:59pm	$\beta_{1,3}^{\text{day-time}}$	0.103	0.00738	0.00
Sunday, 1:00pm – 4:59pm	$\beta_{1,4}^{\text{day-time}}$	0.148	0.00736	0.00
Sunday, 5:00pm – 8:59pm	$\beta_{1,5}^{\text{day-time}}$	0.165	0.00864	0.00
Sunday, 9:00pm – 12:59am*	$\beta_{1,6}^{\text{day-time}}$	-0.366	0.01620	0.00
Monday, 1:00am – 4:59am	$\beta_{2,1}^{\text{day-time}}$	-0.333	0.06660	0.00
Monday, 5:00am – 8:59am	$\beta_{2,2}^{\text{day-time}}$	<b>0.000</b>	<b>fixed</b>	
Monday, 9:00am – 12:59pm	$\beta_{2,3}^{\text{day-time}}$	0.066	0.00723	0.00
Monday, 1:00pm – 4:59pm	$\beta_{2,4}^{\text{day-time}}$	-0.062	0.00764	0.00
Monday, 5:00pm – 8:59pm	$\beta_{2,5}^{\text{day-time}}$	-0.202	0.00983	0.00

<sup>2</sup> The starred time intervals start at 9:00pm on the specified day and include all flights departing up until 12:59am on the following day.

Parameter Description <sup>2</sup>	Parameter	Estimate	Std Error	p-value
Monday, 9:00pm – 12:59am*	$\beta_{2,6}^{\text{day-time}}$	-0.348	0.01660	0.00
Tuesday, 1:00am – 4:59am	$\beta_{3,1}^{\text{day-time}}$	-0.567	0.07930	0.00
Tuesday, 5:00am – 8:59am	$\beta_{3,2}^{\text{day-time}}$	-0.273	0.00711	0.00
Tuesday, 9:00am – 12:59pm	$\beta_{3,3}^{\text{day-time}}$	-0.109	0.00774	0.00
Tuesday, 1:00pm – 4:59pm	$\beta_{3,4}^{\text{day-time}}$	-0.130	0.00790	0.00
Tuesday, 5:00pm – 8:59pm	$\beta_{3,5}^{\text{day-time}}$	-0.134	0.00953	0.00
Tuesday, 9:00pm – 12:59am*	$\beta_{3,6}^{\text{day-time}}$	-0.289	0.01570	0.00
Wednesday, 1:00am – 4:59am	$\beta_{4,1}^{\text{day-time}}$	-0.596	0.07900	0.00
Wednesday, 5:00am – 8:59am	$\beta_{4,2}^{\text{day-time}}$	-0.223	0.00693	0.00
Wednesday, 9:00am – 12:59pm	$\beta_{4,3}^{\text{day-time}}$	-0.042	0.00757	0.00
Wednesday, 1:00pm – 4:59pm	$\beta_{4,4}^{\text{day-time}}$	-0.075	0.00770	0.00
Wednesday, 5:00pm – 8:59pm	$\beta_{4,5}^{\text{day-time}}$	-0.062	0.00919	0.00
Wednesday, 9:00pm – 12:59am*	$\beta_{4,6}^{\text{day-time}}$	-0.174	0.01500	0.00
Thursday, 1:00am – 4:59am	$\beta_{5,1}^{\text{day-time}}$	-0.520	0.07150	0.00
Thursday, 5:00am – 8:59am	$\beta_{5,2}^{\text{day-time}}$	-0.149	0.00671	0.00
Thursday, 9:00am – 12:59pm	$\beta_{5,3}^{\text{day-time}}$	0.001	0.00752	0.91
Thursday, 1:00pm – 4:59pm	$\beta_{5,4}^{\text{day-time}}$	-0.021	0.00770	0.01
Thursday, 5:00pm – 8:59pm	$\beta_{5,5}^{\text{day-time}}$	0.126	0.00892	0.00
Thursday, 9:00pm – 12:59am*	$\beta_{5,6}^{\text{day-time}}$	-0.192	0.01660	0.00
Friday, 1:00am – 4:59am	$\beta_{6,1}^{\text{day-time}}$	-0.404	0.07160	0.00
Friday, 5:00am – 8:59am	$\beta_{6,2}^{\text{day-time}}$	-0.114	0.00674	0.00
Friday, 9:00am – 12:59pm	$\beta_{6,3}^{\text{day-time}}$	0.062	0.00739	0.00
Friday, 1:00pm – 4:59pm	$\beta_{6,4}^{\text{day-time}}$	0.062	0.00745	0.00
Friday, 5:00pm – 8:59pm	$\beta_{6,5}^{\text{day-time}}$	0.085	0.00887	0.00

Parameter Description <sup>2</sup>	Parameter	Estimate	Std Error	p-value
Friday, 9:00pm – 12:59am*	$\beta_{6,6}^{\text{day-time}}$	-0.255	0.01920	0.00
Saturday, 1:00am – 4:59am	$\beta_{7,1}^{\text{day-time}}$	-0.307	0.08220	0.00
Saturday, 5:00am – 8:59am	$\beta_{7,2}^{\text{day-time}}$	-0.144	0.00747	0.00
Saturday, 9:00am – 12:59pm	$\beta_{7,3}^{\text{day-time}}$	0.064	0.00835	0.00
Saturday, 1:00pm – 4:59pm	$\beta_{7,4}^{\text{day-time}}$	-0.118	0.00910	0.00
Saturday, 5:00pm – 8:59pm	$\beta_{7,5}^{\text{day-time}}$	-0.233	0.01370	0.00
Saturday, 9:00pm – 12:59am *	$\beta_{7,6}^{\text{day-time}}$	-0.178	0.01550	0.00
Connection time (minutes) $\leq$ 45	$\beta_1^{\text{connect}}$	0.007	0.00013	0.00
Connection time (minutes) $>$ 45 and $\leq$ 60	$\beta_2^{\text{connect}}$	0.028	0.00055	0.00
Connection time (minutes) $>$ 60	$\beta_3^{\text{connect}}$	-0.018	0.00004	0.00
Flight cancellation	$\beta^{\text{cancel}}$	-0.143	0.00956	0.00
Minimum seating capacity	$\beta^{\text{seats}}$	0.005	0.00010	0.00

**Table A1.1: Estimated itinerary choice utility function parameters and standard errors, with p-values listed based on a classic Student's *t*-test.**

## References

- Air Transport Association. (2008). *2008 Economic Report*. Washington, D.C.
- Ball, M., Barnhart, C., Dresner, M., Hansen, M., Neels, K., Odoni, A., et al. (2010). *Total delay impact study: a comprehensive assessment of the costs and impacts of flight delay in the United States*. Washington, D.C.: FAA.
- Ben-Akiva, M., & Lerman, S. (1985). *Discrete choice analysis: theory and application to travel demand*. Cambridge: The MIT Press.
- Bierlair, M., & Frejinger, E. (2008). Route choice modeling with network-free data. *Transportation Research Part C*, 187-198.
- Bierlaire, M. (2003). BIOGEME: a free package for the estimation of discrete choice models. *Proceedings of the 3rd Swiss Transportation Research Conference*. Ascona, Switzerland.
- Bogers, E., Bierlaire, M., & Hoogendoorn, S. (2007). Modeling learning in route choice. *Transportation Research Record*, 1-8.
- Bratu, S., & Barnhart, C. (2005). An analysis of passenger delays using flight operations and passenger booking data. *Air Traffic Control Quarterly*, 13 (1), 1-27.

- Bureau of Transportation Statistics. (n.d.). *TranStats*. Retrieved September 3, 2009, from On-Time Performance: [http://transtats.bts.gov/TableInfo.asp?Table\\_ID=236&DB\\_Short\\_Name=On-Time&Info\\_Only=0](http://transtats.bts.gov/TableInfo.asp?Table_ID=236&DB_Short_Name=On-Time&Info_Only=0)
- Coldren, G., & Koppelman, F. (2005). Modeling the competition among air-travel itinerary shares: GEV model development. *Transportation Research Part A* , 39, 345-365.
- Coldren, G., Koppelman, F., Kasurirangan, K., & Mukherjee, A. (2003). Modeling aggregate air-travel itinerary shares: logit model development at a major U.S. airline. *Journal of Air Transport Management* , 9, 361-369.
- Frejinger, E., Bierlaire, M., & Ben-Akiva, M. (2009). Sampling of alternatives for route choice modeling. *Transportation Research Part B* , 43 (10), 948-994.
- Sherry, L., Wang, D., & Donohue, G. (2007). Air travel consumer protection: metric for passenger on-time performance. *Transportation Research Record* , 22-27.
- Theis, G., Adler, T., Clarke, J., & Ben-Akiva, M. (2006). Risk aversion to short connections in airline itinerary choice. *Transportation Research Record* , 28-36.
- Tien, S., Ball, M., & Subramanian, B. (2008). Constructing a passenger trip delay metric: an aggregate level approach. *NEXTOR Working Paper* .
- Tien, S., Churchill, A., & Ball, M. (2009). Quantifying the relationship between airline load factors and flight cancellation trends. *Transportation Research Record* , 2106, 39-46.
- Tomer, A., & Puentes, R. (2009). *Expect delays: an analysis of air travel trends in the United States*. Washington, D.C.: Brookings Metropolitan Infrastructure Initiative Series.
- U.S. Congress Joint Economic Committee. (2008). *Your flight has been delayed again: flight delays cost passenger, airlines, and the U.S. economy billions*. Washington, D.C.
- Zhu, Y. (2009). *Evaluating airline delays: the role of airline networks, schedules, and passenger demands*. Massachusetts Institute of Technology.