Title
Essays on Price Discrimination and Congestion Externalities in the U.S. Airline Industry

Author
Luttmann, Alexander Phillip

Publication Date
2019

UNIVERSITY OF CALIFORNIA,
IRVINE

Essays on Price Discrimination and Congestion Externalities in the U.S. Airline Industry

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Economics

by

Alexander Luttmann

Dissertation Committee:
Chancellor's Professor Jan Brueckner, Chair
Professor Matthew Freedman
Associate Professor Jiawei Chen

2019

DEDICATION

To my parents, Charles and Lawan Luttmann, for their endless love, support, and encouragement. I would not be in the position I am today if it wasn't for their countless sacrifices.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I am grateful to my committee members, Chancellor's Professor Jan Brueckner, Professor Matthew Freedman, and Associate Professor Jiawei Chen, for their invaluable guidance on my dissertation and assistance throughout my graduate studies. I am also grateful to former faculty member Kevin Roth for helping me with the first chapter of my dissertation. Specifically, I am thankful to Jan Brueckner for his invaluable mentorship and for inspiring my interest in industrial organization and transportation economics. Jan was also instrumental in helping me navigate the referee process for my first two publications. He is dedicated to my academic and professional development, for which I am forever grateful. Jan will always be my role model and I can only hope to have a career half as productive as him!

I am also grateful to my parents (Charles and Lawan Luttmann) and sister (Victoria Luttmann) for encouraging me and ensuring that I am always in good spirits. Thank you to my classmates for the many thoughtful discussions over the past five years. In particular, Kai Yoshioka, Paul Jackson, Max Littlejohn, and Cody Nehiba have always enhanced my graduate experience. Finally, I would like to thank several of my friends outside academia for always providing support and for taking my mind off graduate studies when I needed it most. Clement Christensen, Nick Giovacchini, Mike Henn, Nick Mejia, Keith Osentoski, and Jared Taylor, thank you for all the memories and for always willing to explore new countries with me.

Funding was provided by the Program in Corporate Welfare Studies, Department of Economics, School of Social Sciences, and the University of California, Irvine.

CURRICULUM VITAE

Alexander Luttmann

EDUCATION

Doctor of Philosophy in Economics **2019**
University of California, Irvine *Irvine, CA*

Master of Arts in Economics **2016**
University of California, Irvine *Irvine, CA*

Bachelor of Science in Economics and Statistics **2011**
George Washington University *Washington, DC*

RESEARCH AND TEACHING INTERESTS

Empirical industrial organization, transportation economics, applied microeconomics

RESEARCH EXPERIENCE

Research Analyst **2011–2014**
The Brattle Group *Washington, DC*

Research Assistant **2011–2014**
Bureau of Economic Analysis *Washington, DC*

TEACHING EXPERIENCE

Lecturer **2017–2018**
University of California, Irvine *Irvine, CA*

Teaching Assistant **2014–2018**
University of California, Irvine *Irvine, CA*

REFEREED JOURNAL PUBLICATIONS

Evidence of directional price discrimination in the U.S. airline industry **2019**
International Journal of Industrial Organization

Are passengers compensated for incurring an airport layover? Estimating the value of layover time in the U.S. airline industry **2019**
Economics of Transportation

UNDER REVIEW AND WORKING PAPERS

Capacity constraints and service quality: Do airport slot controls reduce flight delays? **2019**

The effects of hours-of-service regulations on the U.S. airline industry (with Cody Nehiba) **2018**
Revise and Resubmit at the Journal of Policy Analysis and Management

HONORS AND AWARDS

Ken Small Award for Best PhD Student in Transportation and Urban Economics **2018**
University of California, Irvine

Economics Merit Fellowship **2018**
University of California, Irvine

Department of Economics Summer Research Fellowship **2016–2018**
University of California, Irvine

C.Y. Hsieh Prize in Economics for Best Undergraduate Thesis **2011**
George Washington University

Phi Beta Kappa **2011**
George Washington University

CONFERENCE AND SEMINAR PRESENTATIONS

Berkeley Research Group; California State University, Northridge; CNA; International Transportation Economics Association Annual Conference (Paris); Matrix Economics; MITRE; U.S. Treasury	**2019**
California State University, Long Beach; International Transportation Economics Association Annual Conference (Hong Kong); Eastern Economic Association Annual Conference (Boston, MA); Midwest Economics Association Annual Conference (Evanston, IL); UC Irvine Workshop in Industrial Organization and Corporate Welfare; UC Irvine Summer Urban and Transport Seminar	**2018**
International Transportation Economics Association Annual Conference (Barcelona); UC Irvine Transportation, Urban and Environmental Seminar	**2017**
UC Irvine Transportation, Urban and Environmental Seminar	**2016**

REFEREE SERVICE

Advances in Econometrics; International Journal of Industrial Organization; Journal of Economic Behavior and Organization; Regional Science and Urban Economics

PROFESSIONAL AFFILIATIONS

American Economic Association (AEA), European Economic Association (EEA), International Transportation Economics Association (ITEA)

ABSTRACT OF THE DISSERTATION

Essays on Price Discrimination and Congestion Externalities in the U.S. Airline Industry

By

Alexander Luttmann

Doctor of Philosophy in Economics

University of California, Irvine, 2019

Chancellor's Professor Jan Brueckner, Chair

This dissertation contains three chapters and focuses on price discrimination and airport congestion in the U.S. airline industry.

The first chapter explores possible determinants that may affect an airline's decision to charge passengers different roundtrip fares depending on trip origin, a case of directional price discrimination. Such fare differences cannot be the result of differences in cost, as the cost of flying a roundtrip passenger does not significantly differ depending on direction. It is argued that directional fare differences result from airlines recognizing that passenger price elasticities differ between route endpoints. A price discriminating airline will then charge a higher roundtrip fare at the endpoint where the passenger price elasticity of demand is comparatively lower. Evidence is found suggesting that airlines do use differences in income to price discriminate when setting roundtrip fares. Fares are found to be \$0.18-\$0.43 higher on average for each \$1000 difference in average per capita income between origin and destination metro areas. This finding is sensible assuming that higher incomes reduce the price elasticity of demand for air travel, with richer passengers being less sensitive to the cost of travel.

The second chapter investigates the trade-off between providing convenient flight connections for passengers and reducing airport congestion. From the passenger perspective, layovers

are detrimental since the addition to total travel time relative to a nonstop itinerary is a cost incurred by the passenger. An airline is able to reduce a passenger's layover time by narrowing the gap between flights at the connecting airport. However, narrowing this flight gap has the adverse effect of increasing airport congestion. Taking these perspectives into account, it is clear that layover time influences a prospective passenger's purchasing decision and an airline's flight scheduling decision. Using published fare and itinerary data from Google Flights, this chapter provides insight into both decisions by providing empirical estimates on the value of layover time in the U.S. airline industry. Passengers are found to be compensated with a fare that is $42.74-$47.60 cheaper per hour of layover time.

The last chapter evaluates the effectiveness of slot controls (restrictions on the number of departing and arriving flights) as a congestion management policy. Utilizing the introduction of slot controls at John F. Kennedy (JFK) and Newark (EWR) airports in 2008 as a quasi-experiment, no evidence is found of a reduction in flight delays at both airports. In the months after slot controls were introduced, the average arrival delay at EWR actually increased by 7 minutes. Further, the length of Delta's departure banks (high-volume periods of departing flights) decreased by about 2 minutes at JFK while the scheduled time of EWR flights decreased by 1.5-2.2 minutes. These findings are consistent with Ater (2012), who suggested that policies aimed at reducing congestion at highly concentrated airports will only have a limited impact because dominant airlines already internalize congestion. The results highlight the need for policymakers to carefully consider how the allocation of airport slots will impact flight scheduling decisions when implementing similar policies in the future.

Chapter 1

Evidence of Directional Price Discrimination in the U.S. Airline Industry

1.1 Introduction

Travelers are all too familiar with the discriminatory pricing behavior of airlines. Two months before departure, a prospective passenger could pay $200 for a seat in economy. If that same passenger waits until one week before departure, the price of an economy seat may have doubled to $400.[1] While frustrating for consumers, the ability of an airline to dynamically price seats in order to segment price-elastic leisure travelers from price-inelastic business travelers remains an integral part of an airline's yield management strategy.[2]

[1] Using a sample of daily fare quotes for travel on the New York-London route, Bilotkach et al. (2010) confirmed that fares typically increase as the departure date of a flight approaches.

[2] In 1998, American and Delta credited yield management techniques for revenue increases of nearly $500 and $300 million a year respectively (Netessine and Shumsky, 2002).

Previous literature on price discrimination in the airline industry has focused on the various mechanisms utilized by airlines to segment passengers with differing price elasticities of demand. Advance-purchase restrictions attached to discount tickets have been shown to be profit-maximizing for airlines wishing to divert demand from peak to off-peak periods since these restrictions are effective in segmenting customers by their value of time (Gale and Holmes, 1993). These restrictions allow airlines to reduce fares for price-elastic leisure travelers, who typically have a low value of time (Dana, 1998). Other restrictions like Saturday night stay, minimum stay, and non-refundable tickets are designed to discourage price-inelastic consumers from buying cheaper tickets (Stavins, 2001).[3] Additionally, Puller and Taylor (2012) find fares purchased on weekends to be 5% lower, supporting the conjecture that airlines price discriminate when the mix of purchasing passengers makes demand more price-elastic.

At the route level, Borenstein and Rose (1994) find evidence of significant variation in fares. The expected absolute variation in fares an airline charges to two different passengers on the same route was found to be 36% of the airline's average ticket price. Borenstein and Rose also observed an increase in the extent of a carrier's price dispersion as the number of competitors in a market grows, a result consistent with price discrimination in monopolistically competitive markets. The present chapter differs from their study by examining price dispersion on a directional basis. For example, instead of treating Los Angeles (LAX) to New York (JFK) roundtrips as the same market regardless of origin, JFK-LAX is treated separately from LAX-JFK. The chapter then contributes to the literature by being the first to explore possible determinants that may affect an airline's decision to charge passengers different roundtrip fares depending on trip origin.[4] Directional fare differences cannot be the

[3] Business travelers typically prefer to return home before the weekend. Therefore, Saturday night stay and minimum stay restrictions are effective in separating business from leisure travelers. Similarly, business travelers prefer to purchase fully refundable tickets in case of the event that business meetings are canceled.

[4] A January 7th, 2015 Wall Street Journal article by Scott McCartney titled "*Airline Fare Riddle: One Route, Two Prices*" highlights an example where roundtrip flights between Los Angeles and Honolulu were 7.5% more expensive if the flight originated in Los Angeles.

result of differences in cost, as the cost of flying a roundtrip passenger does not significantly differ depending on direction.

This chapter argues that directional fare differences are a result of airlines recognizing that passenger price elasticities of demand differ between route endpoints. Three potential sources impacting the price elasticity are considered. Foremost, if high incomes reduce the price elasticity of demand for air travel, with rich passengers being less sensitive to the cost of travel, airlines may income discriminate by charging passengers beginning roundtrip travel at the higher income endpoint a higher average fare.[5] Differences in endpoint populations could also impact the price elasticity if passengers beginning roundtrip travel at the more populous endpoint benefit from greater flight frequency. That is, more frequent service increases the likelihood that passengers are able to book flights with preferred departure times, a benefit passengers are willing to pay more for. Additionally, if the proportion of business and leisure passengers on a route differs depending on trip origin, airlines may set fares directionally, with roundtrip fares being cheaper in the direction where the mix of passengers is more likely to contain a higher proportion of price-elastic leisure travelers.

Given the substantial level of competition in the U.S. airline industry and assuming airlines face constant costs in delivering passengers to a market, it is unlikely that differences in the level of demand (as opposed to the price elasticity) contribute to directional fare differences.[6] However, on routes where airlines operate through congested or slot-controlled airports, an airline could face increasing costs of serving passengers. With an upward-sloping cost curve, demand differences could then affect roundtrip fares on a route regardless of which endpoint is costly to serve. Therefore, it is plausible that directional fare differences reflect airlines exploiting differences in passenger price elasticities in addition to differences in the demand

[5] This would be an example of third-degree price discrimination, where airlines charge prices that vary by location. Advance purchase, non-refundability, and minimum stay restrictions are examples of second-degree price discrimination.

[6] An increase in demand will have no effect on price in a competitive industry facing constant costs since the long-run supply curve is horizontal.

for air travel between route endpoints.

A key contribution of this chapter is showing that directional fare differences exist and depend on the incomes of the endpoint cities.[7] Income discrimination arises if airlines deliberately charge passengers beginning roundtrip travel at the higher income endpoint a higher average fare. A price discriminating airline will price in this manner if passengers originating from the higher income endpoint have a lower price elasticity of demand relative to passengers originating from the lower income endpoint. In an airline's yield management strategy, the airline does not need to post higher fares at one endpoint for income discrimination to occur. When selling seats on a flight, an airline will typically allocate a certain number of seats to several different fare levels or "buckets." Once all seats allocated to the cheapest fare bucket are sold, the price rises to the level specified by the next fare bucket. This process continues until all available seats are sold. For simplicity, consider a scenario where an airline operates only a discount and a full fare bucket, charging the same discount and full roundtrip fares no matter what direction is flown. Average roundtrip fares could then directionally differ if the airline allocates fewer seats to the discount fare bucket for passengers beginning roundtrip travel at the endpoint where average income is higher.[8]

After controlling for differences in endpoint populations and differences in the mix of business and leisure passengers between origin and destination markets, this chapter finds evidence that airlines use differences in income to price discriminate on a directional basis. The sensitivity of the income discrimination effect is explored by adding several variables that are correlated with income to the base econometric model. These variables control for the level of competition, whether the roundtrip originates at a hub, differences in the volume of business travelers, differences in the proportion of peak vs. off-peak flights, and differences in

[7] These findings are important because many empirical studies in the existing literature ignore directionality when defining airline markets. For example, LAX-JFK roundtrips and JFK-LAX roundtrips are typically lumped together into a single market.

[8] For more on airline yield management practices, see Belobaba et al. (2015), Belobaba et al. (2009), Donovan (2005), and Talluri and Van Ryzin (2006).

passenger demographics between origin and destination markets. Additional specifications explore how the magnitude of the income discrimination effect differs by distance, time of year, and carrier type (legacy vs. low-cost). After accounting for these factors, the directional difference in roundtrip fares is found to be \$0.18 larger when the difference in average per capita income between the origin cities rises by \$1,000. For example, the average per capita income for residents of San Francisco is more than \$21,000 higher than for Chicago residents, and this finding implies that roundtrip fares are \$3.90 more expensive on average for passengers originating in San Francisco.

The rest of this chapter is organized as follows. Section 1.2 describes the fare data used in the empirical analysis and provides summary measures of directional fare differences in the U.S. airline market. Section 1.3 discusses pricing criteria airlines may utilize in their yield management strategies that potentially contribute to directional roundtrip fare differences. Section 1.4 outlines the econometric model while Section 1.5 discusses results of the empirical analysis. Finally, Section 1.6 concludes.

1.2 Fare Data and Summary Measures of Directional Fare Differences in the U.S. Airline Industry

Ticket and price data are taken from the US Department of Transportation's Airline Origin and Destination Survey (database DB1B). Data from this survey are released quarterly and generated from a 10 percent random sample of all airline tickets that originate in the United States on U.S. based carriers. The analysis sample spans all four quarters of 2015. For a detailed description of DB1B data processing, see Appendix A.

Figure 1.1 displays United's average roundtrip fares for direct travel between its San Francisco (SFO) and Chicago (ORD) hubs, a route with considerable passenger traffic. The solid bars in Figure 1.1 indicate United's average roundtrip fares for flights originating in San Francisco while the patterned bars indicate United's average roundtrip fares for flights originating in Chicago. As illustrated by Figure 1.1, United's fares on the SFO-ORD route are consistently higher when the roundtrip starts in San Francisco. In the first quarter of 2015, average roundtrip fares originating in San Francisco were 6.81% higher than roundtrips originating in Chicago. This directional fare effect is consistent across the second, third, and fourth quarters with average fares originating in San Francisco being 7.6%, 5.68%, and 4.31% higher, respectively.

Figure 1.1: Average Quarterly Roundtrip Fares on United's SFO-ORD Route

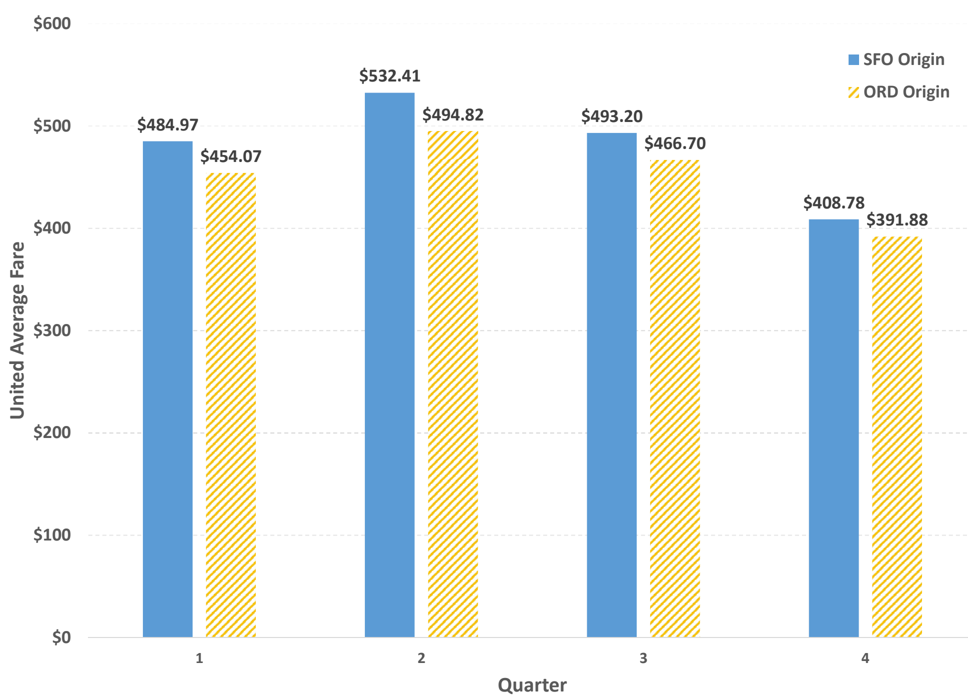

Figure 1.2: Average Quarterly Roundtrip Fares on Virgin America's SFO-ORD Route

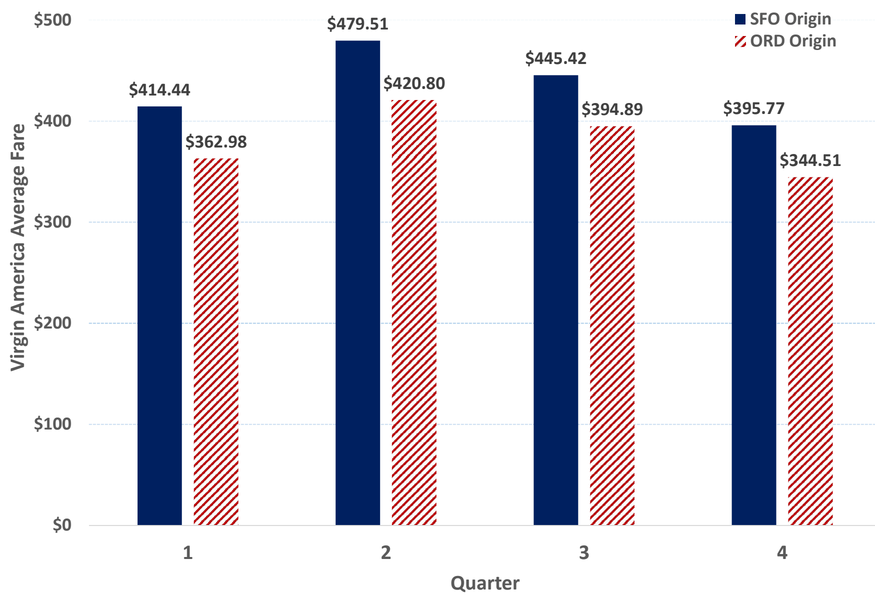

A similar pattern emerges with Virgin America's average fares on the SFO-ORD route. Figure 1.2 displays Virgin's average roundtrip fares for direct travel between San Francisco (SFO) and Chicago (ORD). The solid bars in Figure 1.2 indicate Virgin's average fares for roundtrips originating in San Francisco whereas the patterned bars indicate Virgin's average fares for roundtrips originating in Chicago. Consistent with its status as a low-cost competitor, Virgin's average fares on the SFO-ORD route are lower than United's fares in each quarter regardless of roundtrip origin. However, similar to United, Virgin's average fares on the SFO-ORD route are consistently higher when the roundtrip originates in San Francisco. In the first quarter of 2015, average roundtrip fares originating in San Francisco were 14.18% higher than roundtrips originating in Chicago. This directional fare effect is consistent across all four quarters, with average fares originating in San Francisco being

13.95%, 12.79%, and 14.88% higher for the second, third, and fourth quarters respectively.

Table 1.1 provides a summary of directional fare differences for the top 50 roundtrip routes by nonstop passenger traffic in the fourth quarter of 2015, the most recent quarter of data used for analysis. This table is sorted by the total number of sampled passengers observed in the DB1B data purchasing tickets for nonstop roundtrip travel on the respective route. Column four of Table 1.1 provides the average fare for roundtrip travel originating at the airport indicated in column one. Analogously, column five provides the average fare for roundtrip travel originating at the airport indicated in column two. Column six provides the fare difference while column seven converts this fare difference to a percentage in terms of the average fare in column four.

As demonstrated by Table 1.1, there is substantial variation in directional roundtrip fare differences across the top 50 routes, ranging from a mere 70 cents (.25%) on Southwest's Burbank-Oakland route to $98.74 (over 25%) on JetBlue's New York-Fort Lauderdale route. The next section discusses pricing criteria airlines may utilize in their yield management strategies that potentially contribute to the observed directional roundtrip fare differences.

Table 1.1: Average Fares for Top 50 Direct Roundtrip Routes in Quarter 4 of 2015

Airport 1 [1]	Airport 2 [2]	Carrier [3]	Average Fare (Origin at [1]) [4]	Average Fare (Origin at [2]) [5]	Fare Difference [6]	Percent Difference [7]	Sampled Passengers [8]
Honolulu, HI (HNL)	Kahului, HI (OGG)	Hawaiian	$188.35	$181.12	$7.23	3.84%	9,972
Atlanta, GA (ATL)	New York, NY (LGA)	Delta	$433.83	$443.03	-$9.20	-2.12%	9,677
Burbank, CA (BUR)	Oakland, CA (OAK)	Southwest	$282.69	$283.40	-$0.70	-0.25%	7,277
Seattle, WA (SEA)	Los Angeles, CA (LAX)	Alaska	$272.04	$275.02	-$2.98	-1.09%	7,270
Dallas, TX (DAL)	Houston, TX (HOU)	Southwest	$361.68	$347.83	$13.85	3.83%	6,448
Denver, CO (DEN)	Phoenix, AZ (PHX)	Southwest	$272.06	$280.97	-$8.92	-3.28%	6,168
San Diego, CA (SAN)	San Jose, CA (SJC)	Southwest	$320.49	$300.63	$19.87	6.20%	6,075
San Jose, CA (SJC)	Orange County, CA (SNA)	Southwest	$279.89	$322.79	-$42.90	-15.33%	6,060
Oakland, CA (OAK)	San Diego, CA (SAN)	Southwest	$284.83	$282.39	$2.43	0.85%	6,028
Newark, NJ (EWR)	San Francisco, CA (SFO)	United	$564.92	$590.51	-$25.59	-4.53%	5,993
Sacramento, CA (SMF)	San Diego, CA (SAN)	Southwest	$299.78	$313.68	-$13.90	-4.64%	5,960
Baltimore, MD (BWI)	Orlando, FL (MCO)	Southwest	$330.93	$308.28	$22.65	6.84%	5,878
Oakland, CA (OAK)	Los Angeles, CA (LAX)	Southwest	$250.73	$256.93	-$6.20	-2.47%	5,864
Dallas, TX (DFW)	Chicago, IL (ORD)	American	$338.55	$344.79	-$6.24	-1.84%	5,836
Newark, NJ (EWR)	Orlando, FL (MCO)	United	$360.00	$299.26	$60.74	16.87%	5,796
Atlanta, GA (ATL)	Fort Lauderdale, FL (FLL)	Delta	$277.47	$245.14	$32.33	11.65%	5,790
Chicago, IL (ORD)	New York, NY (LGA)	United	$322.42	$306.66	$15.77	4.89%	5,604
Atlanta, GA (ATL)	Washington, DC (DCA)	Delta	$392.26	$380.74	$11.52	2.94%	5,576
Newark, NJ (EWR)	Los Angeles, CA (LAX)	United	$532.75	$523.74	$9.01	1.69%	5,514
New York, NY (JFK)	Los Angeles, CA (LAX)	Delta	$545.70	$514.34	$31.36	5.75%	5,462
Chicago, IL (ORD)	New York, NY (LGA)	American	$361.05	$354.82	$6.23	1.73%	5,450
Oakland, CA (OAK)	Orange County, CA (SNA)	Southwest	$292.58	$317.87	-$25.29	-8.64%	5,305
Boston, MA (BOS)	Orlando, FL (MCO)	JetBlue	$433.77	$368.95	$64.82	14.94%	5,277
San Francisco, CA (SFO)	Los Angeles, CA (LAX)	Virgin	$250.74	$240.59	$10.15	4.05%	5,273
New York, NY (JFK)	Orlando, FL (MCO)	JetBlue	$360.25	$292.04	$68.22	18.94%	5,272
Detroit, MI (DTW)	Orlando, FL (MCO)	Delta	$384.24	$360.03	$24.21	6.30%	5,256
Dallas, TX (DFW)	Los Angeles, CA (LAX)	American	$333.36	$334.36	-$1.00	-0.30%	5,220
New York, NY (JFK)	Fort Lauderdale, FL (FLL)	JetBlue	$382.87	$284.14	$98.74	25.79%	5,156
Baltimore, MD (BWI)	Fort Lauderdale, FL (FLL)	Southwest	$304.46	$255.56	$48.90	16.06%	5,095
Boston, MA (BOS)	Washington, DC (DCA)	American	$371.43	$370.44	$0.99	0.27%	5,022
Chicago, IL (MDW)	Orlando, FL (MCO)	Southwest	$317.81	$303.26	$14.54	4.58%	4,929
Chicago, IL (ORD)	San Francisco, CA (SFO)	United	$391.88	$408.78	-$16.90	-4.31%	4,921
New York, NY (JFK)	Los Angeles, CA (LAX)	JetBlue	$522.50	$523.30	-$0.80	-0.15%	4,886
Dallas, TX (DFW)	New York, NY (LGA)	American	$406.40	$399.78	$6.62	1.63%	4,850
Chicago, IL (MDW)	Denver, CO (DEN)	Southwest	$302.71	$321.38	-$18.67	-6.17%	4,789
Seattle, WA (SEA)	Las Vegas, NV (LAS)	Alaska	$317.25	$287.95	$29.30	9.23%	4,787
Sacramento, CA (SMF)	Orange County, CA (SNA)	Southwest	$283.14	$321.70	-$38.56	-13.62%	4,776
Denver, CO (DEN)	Las Vegas, NV (LAS)	Southwest	$265.83	$256.42	$9.41	3.54%	4,722
Atlanta, GA (ATL)	Boston, MA (BOS)	Delta	$475.66	$480.27	-$4.61	-0.97%	4,701
Oakland, CA (OAK)	Las Vegas, NV (LAS)	Southwest	$262.21	$236.26	$25.95	9.90%	4,633
New York, NY (JFK)	San Francisco, CA (SFO)	Delta	$521.39	$504.16	$17.24	3.31%	4,618
San Jose, CA (SJC)	Los Angeles, CA (LAX)	Southwest	$261.23	$272.04	-$10.81	-4.14%	4,605
San Jose, CA (SJC)	Las Vegas, NV (LAS)	Southwest	$270.07	$264.55	$5.53	2.05%	4,543
Chicago, IL (MDW)	Las Vegas, NV (LAS)	Southwest	$370.76	$362.88	$7.87	2.12%	4,479
Seattle, WA (SEA)	Orange County, CA (SNA)	Alaska	$303.85	$338.72	-$34.87	-11.48%	4,379
Atlanta, GA (ATL)	Orlando, FL (MCO)	Delta	$304.68	$312.83	-$8.15	-2.68%	4,370
Newark, NJ (EWR)	Chicago, IL (ORD)	United	$442.92	$427.12	$15.81	3.57%	4,356
Sacramento, CA (SMF)	Ontario, CA (ONT)	Southwest	$276.75	$283.34	-$6.59	-2.38%	4,345
New York, NY (LGA)	Miami, FL (MIA)	American	$423.27	$413.51	$9.76	2.31%	4,294
Boston, MA (BOS)	San Francisco, CA (SFO)	United	$551.72	$584.62	-$32.90	-5.96%	4,262

Notes:
[4]: Average roundtrip fare for trip originating at Airport [1] (2015 Q4)
[5]: Average roundtrip fare for trip originating at Airport [2] (2015 Q4)
[6]: [4] - [5]
[7]: [6]/[4]
[8]: Sampled passengers on route (2015 Q4)

1.3 Possible Sources of Directional Fare Differences

There are several pricing criteria airlines may employ in their respective yield management strategies that could explain why average roundtrip fares in a city-pair market differ on a directional basis. One potential criterion is the difference in per capita incomes between origin and destination cities. An airline wishing to maximize revenue may income discriminate by charging passengers beginning roundtrip travel in the higher income locality a higher average fare.[9]

A second criterion is the difference in endpoint populations. Population differences could affect the price elasticity if passengers originating roundtrip travel at the more populous endpoint benefit from greater flight frequency.[10] The endpoint population differential may also result in a demand difference, assuming that metro areas with larger populations have a greater demand for air travel than smaller metro areas. Differences in income can also result in a difference in demand, with higher income endpoints demanding more air travel. However, given the substantial level of competition in the U.S. airline industry and assuming airlines face constant costs in delivering passengers to a market, differences in demand are unlikely to impact the roundtrip fare difference. Yet, it is plausible that airlines face increasing costs in delivering passengers on routes that use congested or slot-controlled airports. If airlines face increasing costs, the demand differential may impact the fare difference, with roundtrip fares higher at the endpoint with greater demand.

[9] For an airline to price discriminate in this manner, passengers residing in the higher income area must have a lower price elasticity of demand for air travel relative to passengers residing in the lower income area, with rich passengers being less sensitive to the cost of travel.

[10] More frequent service increases the probability that passengers are able to purchase tickets with preferred flight times, a benefit passengers are willing to pay more for. If this is the case, the price elasticity of demand is expected to be lower for passengers originating roundtrip travel at the endpoint with the larger population.

Additionally, if the proportion of business and leisure passengers on a route differs by point of origin, airlines may set fares directionally, with roundtrip fares being cheaper in the direction where the mix of passengers is more likely to contain a higher proportion of price-elastic leisure travelers. The following sections describe in more detail how airlines may exploit differences in income, population, and passenger mix between route endpoints when setting roundtrip fares.

1.3.1 Income Discrimination

It is helpful to walk through a simple yield management example outlining how differences in income between endpoints could be reflected in an airline's yield management strategy. For simplicity, consider a route where an airline has 100 available seats in economy class. Two different fare classes are offered. In each direction, a discounted fare is offered for $250 roundtrip and a full fare for $400. Since demand is not high enough to sell all tickets at full fare, the airline's problem is to choose the maximum number of seats that are sold at the discounted fare ("booking limit") to maximize revenue. The remaining seats are sold to higher paying customers. Assuming leisure travelers purchase their seats before most business travelers, the airline must choose a booking limit that leaves a sufficient number of seats for the higher yielding business travelers who buy their tickets closer to the departure date.

In the example outlined above, suppose all seats are sold and the booking limit is chosen to be 60. In this case, the average roundtrip fare in each direction equals $310.[11] However, suppose an airline chooses a booking limit that depends on direction. An airline could then income discriminate by lowering the booking limit for passengers beginning roundtrip travel in the metro area with the higher average income while maintaining a higher booking limit for passengers beginning travel from the opposite endpoint. Continuing with the example

[11] (60*$250 + 40*$400) / 100 = $310.

of the San Francisco (SFO) to Chicago (ORD) route referenced in Figures 1.1 and 1.2, an airline could income discriminate by choosing a lower booking limit for passengers traveling on SFO-ORD roundtrips while maintaining a higher booking limit for passengers on ORD-SFO roundtrips. For instance, suppose the booking limit is chosen to be 50 for SFO-ORD roundtrips and 60 for ORD-SFO roundtrips. Average roundtrip fares for SFO-ORD would be $325 compared to $310 for ORD-SFO.[12]

While the actual values of the $250 discount and $400 full fares do not differ depending on direction flown, the mechanism through which airlines income discriminate is by choosing different booking limits that depend on the difference in average incomes of the route endpoints.[13] Alternatively, airlines could also discriminate by simply charging higher discount or full fares for passengers beginning roundtrip travel from the higher income endpoint. Suppose the booking limit is kept at 60 in each direction but the discount fare is $300 for passengers on SFO-ORD roundtrips compared to $250 for ORD-SFO. The average roundtrip fare for SFO-ORD would then be $340 compared to $310 for ORD-SFO passengers.[14] However, a review of the yield management literature suggests that airlines do not discriminate in this manner. If airlines are income discriminating, it is likely to occur via the airline choosing directional booking limits that depend on the difference in average incomes of the route endpoints.

1.3.2 Differences in Endpoint Populations

As alluded to earlier, differences in endpoint populations may impact the price elasticity of demand if passengers originating at the more populous endpoint benefit from greater flight

[12] (50*$250 + 50*$400) / 100 = $325.

[13] For instance, directional fare differences would arise if an airline offered more promotional (discount) fares at endpoints where incomes are lower. Alaska and Southwest are examples of airlines that send biweekly emails informing passengers about available discount fares. For the SFO-ORD route, roundtrip fares would be higher for SFO originating passengers if more promotional fares are made available for roundtrip passengers originating at ORD.

[14] (60*$300 + 40*$400) / 100 = $340.

frequency. Because the price elasticity of demand is expected to be lower for passengers originating travel at the larger endpoint on a route, airlines may charge passengers originating at these endpoints higher average fares. Referencing the yield management example from Section 1.3.1, this effect could also be captured by an airline choosing a booking limit that takes into account the difference in endpoint populations. By lowering the booking limit for passengers beginning travel at the endpoint with the larger population, resulting average roundtrip fares will be more expensive for passengers beginning travel at that endpoint.

1.3.3 Differences in the Mix of Traveling Passengers

If the proportion of business and leisure passengers on a route differs depending on point of origin, airlines will likely take this difference into account when setting roundtrip fares. In the two-fare-class example from Section 1.3.1, the airline chooses a booking limit that constrains the number of seats sold at the discounted fare. For routes with a high proportion of business travelers, airlines are inclined to lower the booking limit to maximize revenue due to the increased demand from price-inelastic consumers. For instance, consider a roundtrip between Chicago (ORD) and Las Vegas (LAS). Travelers beginning the trip in Chicago are more likely to be leisure travelers compared to passengers beginning the trip in Las Vegas. Therefore, it may make sense for an airline to maintain a lower booking limit for passengers beginning the roundtrip in Las Vegas compared to Chicago. In this case, average roundtrip fares would be higher for passengers on LAS-ORD compared to ORD-LAS.

1.4 Specification of Empirical Model

An observation in the empirical analysis is the average fare for nonstop roundtrip travel from origin airport i to destination airport j on carrier k in quarter l. To analyze the directional effect on roundtrip fares, this average fare is subtracted from the average fare for direct roundtrip travel from origin airport j to destination airport i on the same carrier k in the same quarter l. This fare difference is then regressed on a series of differenced variables to determine what effects differences in income, population, and the mix of traveling passengers between endpoint cities have on the roundtrip fare difference.

Before estimating the empirical model, each airport i and j in the DB1B data is matched to a U.S. Census Bureau metropolitan statistical area (MSA). Per capita income at the MSA level for 2014 is taken from the Bureau of Economic Analysis. The variable DINC is defined to be the per $1,000 difference in per capita incomes between airports i and j. A positive coefficient on DINC would provide evidence of directional income discrimination, since a positive coefficient implies that passengers originating in the higher income metro area are charged higher roundtrip fares.

Population estimates at the MSA level for 2014 are taken from the U.S. Census Bureau's American Community Survey. The variable DPOP is defined as the difference in total population (in millions) between airports i and j. Because demand for air travel rises with income, a positive coefficient on DPOP is expected.

To control for differences in the mix of traveling passengers on a route, the effectiveness of two variables commonly used to identify leisure markets is explored. Foremost, a temperature variable similar to the one used in Brueckner et al. (2013) and Stavins (2001) is constructed using airport weather data from the National Oceanographic and Atmospheric Administration's 1981-2010 climate normals. This variable, DTEMP, is constructed as the difference in the average quarterly high temperatures between airports i and j.[15] A negative value for this variable is likely to indicate a leisure market, where passengers travel from cold to warm climates in the first and fourth quarters. This formulation also assumes

[15] In Brueckner et al. (2013) and Stavins (2001), the absolute difference in the average January high temperatures is used. Here, the actual difference in quarterly average temperatures is used, allowing for negative values of DTEMP.

that leisure markets are warmer than non-leisure markets in the second and third quarters. This assumption seems to hold for popular leisure destinations like Los Angeles, San Diego, Honolulu, Las Vegas, Orlando, and Miami.

The second candidate is similar to the tourist variable employed by Gerardi and Shapiro (2009). This variable, DTOURIST, is defined as the difference in the ratio of accommodation earnings to total nonfarm earnings between the MSAs of airports i and j.[16] A negative value for this variable is likely to indicate a leisure market, as passengers travel to metro areas where earnings from hotel accommodations make up a larger proportion of total earnings.

Theory provides little guidance on the correct functional form for this type of analysis. Therefore, two different forms of the empirical model are estimated. Equations (1.1) and (1.2) below are estimated in levels while equations (1.3) and (1.4) are estimated in constant elasticity or log-log form. All four equations are estimated using ordinary least squares with robust standard errors. Estimates from equations (1.1) and (1.2) are provided in Section 1.5 while estimates from equations (1.3) and (1.4) are provided in Appendix C.[17]

$$FARE_{ijkl} - FARE_{jikl} = \beta_0 + \beta_1 DINC_{i,j} + \beta_2 DPOP_{i,j} + \beta_3 DTEMP_{i,j} + \epsilon_{ijkl} \tag{1.1}$$

$$FARE_{ijkl} - FARE_{jikl} = \beta_0 + \beta_1 DINC_{i,j} + \beta_2 DPOP_{i,j} + \beta_3 DTOURIST_{i,j} + \epsilon_{ijkl} \tag{1.2}$$

[16] MSA level data used to construct DTOURIST comes from the Bureau of Economic Analysis.

[17] Given the empirical framework, any variable that is identical between origin and destination metro areas would be "differenced" away. Therefore, there is little justification for the inclusion of carrier or quarter-of-year fixed effects. Nevertheless, including both carrier and quarter-of-year fixed effects does not significantly change the magnitudes of the coefficient estimates presented in Section 1.5.

$$ln(FARE_{ijkl}) - ln(FARE_{jikl}) = \beta_0 + \beta_1 DLINC_{i,j} + \beta_2 DLPOP_{i,j} + \beta_3 DLTEMP_{i,j} + \epsilon_{ijkl} \tag{1.3}$$

$$ln(FARE_{ijkl}) - ln(FARE_{jikl}) = \beta_0 + \beta_1 DLINC_{i,j} + \beta_2 DLPOP_{i,j} + \beta_3 DLTOURIST_{i,j} + \epsilon_{ijkl} \tag{1.4}$$

The variable DLINC in equations (1.3) and (1.4) is defined as the difference in the logs of average per capita income between cities i and j. Similarly, the variable DLPOP is the difference in the logs of population between cities i and j, DLTEMP is the difference in the logs of the average quarterly high temperatures, and DLTOURIST is the difference in the logs of the ratios of accommodation earnings to total nonfarm earnings.

To prevent small routes from skewing results, an observation is included only if at least 100 sampled passengers are observed traveling in each direction during the quarter. Given that the DB1B data is a 10% sample of tickets sold in a quarter, this restriction implies that only nonstop routes where an airline carries at least 1,000 roundtrip passengers in each direction during the quarter are included. Furthermore, to prevent double counting by including both ij and ji observations in the analysis, an observation is included only if the origin airport code comes after the destination airport code in the alphabet. Referencing the San Francisco (SFO) to Chicago (ORD) route discussed earlier, this restriction implies that the SFO-ORD observation is included in the analysis while the ORD-SFO observation is excluded.[18]

[18] Regression results are approximately the same if both ij and ji observations are included in the analysis with standard errors that are clustered at the origin-destination level.

1.5 Results of Empirical Analysis

Summary statistics for all variables in the empirical analysis are provided in Appendix Table B.1. Using the full sample of observations, the first column of Table 1.2 reports regression results from the model specified by equation (1.1) while the third column reports results from the model specified by equation (1.2). A careful examination of the empirical distribution for the dependent variable DFARE indicates the presence of large outliers (see Table B.1). To alleviate concerns that outliers are driving results, columns two and four of Table 1.2 report estimates from equations (1.1) and (1.2) respectively when the sample is restricted to observations that fall between the 5th and 95th percentiles of DFARE. The large difference in the magnitude of coefficients between columns one and two on DPOP (1.43 to 0.76) and DTEMP (.20 to .08) and between columns three and four on DTOURIST (0.17 to 0.01) indicates that outliers in DFARE have a substantial impact on coefficient estimates. Surprisingly, the small insignificant coefficient on DTOURIST in column four suggests that the variable is not sufficiently capturing the effect of differences in the mix of traveling passengers on the directional fare difference. Therefore, to minimize the impact of outliers and more appropriately control for differences in the mix of traveling passengers, the remaining regressions presented in this section are estimated using equation (1.1) on the restricted sample.

Table 1.2: Directional Price Discrimination Regression Results

	Equation (1.1)		Equation (1.2)	
Analysis Sample:	Full (1)	Restricted (2)	Full (3)	Restricted (4)
DINC ($1,000s)	0.43***	0.35***	0.36***	0.31***
	(0.038)	(0.027)	(0.034)	(0.025)
DPOP (1,000,000s)	1.43***	0.76***	1.43***	0.76***
	(0.065)	(0.043)	(0.065)	(0.043)
DTEMP	0.20***	0.08***		
	(0.037)	(0.024)		
DTOURIST			0.17*	0.01
			(0.098)	(0.082)
Constant	-0.92**	-1.25***	-0.76*	-1.19***
	(0.397)	(0.283)	(0.397)	(0.284)
Observations	7,426	6,684	7,426	6,684
R-squared	0.137	0.109	0.133	0.108

Notes: Robust standard errors are reported in parentheses. The restricted sample removes outliers that fall outside the range encompassing the 5th and 95th percentiles of the dependent variable DFARE.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

The most important result of this chapter is the finding that the coefficient on DINC is positive and significant at any conventional level, providing evidence of income discrimination in the directional pricing of roundtrip flights. Not taking into account the effects of DPOP and DTEMP, the estimated coefficient of 0.35 on DINC in column two of Table 1.2 indicates that the roundtrip fare difference is $0.35 higher on average for each $1,000 difference in income between the endpoint cities. Continuing with the example of the SFO-ORD route, the 2014 average per capita income for residents in the San Francisco metropolitan area was $72,364 whereas the average per capita income for residents in the Chicago metropolitan area was $50,690. Thus, residents in San Francisco have an average per capita income that is \$21,674 higher than residents in Chicago. This difference in per capita income results in roundtrip fares on the SFO-ORD route that are predicted to be \$7.59 higher on average for passengers originating in San Francisco.[19]

To provide more context on the estimated income discrimination effect on the SFO-ORD route, consider United. In the 2015 DB1B data, 8,649 roundtrip passengers are observed traveling the SFO-ORD route on United with an origin at SFO. Since the DB1B data is a 10% sample of all tickets, this implies United carried a total of 86,490 roundtrip passengers on SFO-ORD in 2015. The estimated effect of \$7.59 for each passenger implies that United generated an additional \$656,459 in revenue from these nonstop passengers.[20]

As expected, the coefficients on DPOP in Table 1.2 are positive and significant, indicating that passengers beginning roundtrip travel in the metro area with the larger population pay higher fares on average. As discussed in Section 1.3.2, these results are sensible assuming that passengers beginning roundtrip travel at the more populous endpoint benefit from greater flight frequency. The estimated coefficient of 0.76 on DPOP in column two indicates that roundtrip fares are \$0.76 higher for each 1,000,000 difference in population between origin and destination metro areas. On the Los Angeles (LAX) to Seattle (SEA) route, where the population of the Los Angeles metro area is 9.58 million larger than that of the Seattle metro area, this coefficient implies that roundtrip fares on the LAX-SEA route are \$7.28 higher on average for passengers originating at LAX.[21]

The coefficients on DTEMP have the expected positive signs and are significant at conventional levels. As discussed in Section 1.3.3, airlines potentially take into account the mix of traveling passengers when setting roundtrip fares. Fares on routes with a high proportion of leisure travelers are expected to be cheaper than routes that have a high proportion of price-inelastic business travelers. Since DTEMP is defined as the difference in the average

[19] $21.674*0.35 \approx \$7.59$. This calculation does not take into account the effects of DPOP and DTEMP.
[20] $\$7.59*(8{,}649/.10) = \$656{,}459.10$.
[21] $\$0.76*9.58 = \$7.28$. This calculation does not take into account the effects of DINC or DTEMP.

quarterly high temperatures between origin and destination airports, a large negative value likely indicates travel to a leisure market as passengers travel from cold to warm climates in the first and fourth quarters of the year. For instance, the difference in average quarterly temperatures between New York (JFK) and Miami (MIA) is -35.8, -24.7, -9.4, and -21 degrees Fahrenheit for the first, second, third, and fourth quarters respectively. Not taking into account the effects of DINC and DPOP, the 0.08 coefficient on DTEMP in column two of Table 1.2 implies roundtrip fares on the JFK-MIA route are \$2.86, \$1.98, \$0.75, and \$1.68 cheaper on average for passengers originating at JFK in the first, second, third, and fourth quarters respectively.[22]

1.5.1 Threats to the Identification of the Income Discrimination Effect

The primary threat to the identification of the income discrimination effect is omitted variables that are correlated with the income difference and potentially cause a fare difference. This section discusses five potential omitted variables and evaluates their impact on the magnitude of the DINC coefficient when these variables are added to the specifications presented in Table 1.2.

Hub Premium

Legacy carriers operating hub-and-spoke networks possess substantial market power at their hub airports. Previous studies have shown that this market power results in hub carriers receiving higher fares on routes that depart from their hubs.[23] Given that many hubs are

[22] -35.8*\$0.08 = -\$2.86. -24.7*\$0.08 $\approx$ -\$1.98. -9.4*\$0.08 = \$0.75. -21*\$0.08 = -\$1.68.

[23] For evidence supporting the hub premium, see Borenstein (1989), Borenstein (1991), Evans and Kessides (1993), Lee and Luengo-Prado (2005), Berry et al. (2006), Lederman (2008), and Ciliberto and Williams (2010).

located in cities with high average incomes, it is possible that directional fare differences are driven by route segments originating from hub airports rather than the proposed income discrimination effect. To control for the fare premium a carrier receives on routes departing from their hubs, the variable DHUB is constructed and set equal to one if origin airport i is a hub for carrier k. DHUB is set to minus one if destination airport j is a hub for carrier k. Lastly, DHUB is set equal to zero if airports i and j are both hubs for carrier k.[24]

Differences in the Volume of Business Travelers

While the DTEMP variable is able to identify and distinguish business from leisure markets, it may not sufficiently capture differences in the overall volume of business passengers between endpoints on a route. Clearly, differences in the volume of business passengers will impact the fare difference since these passengers purchase tickets much closer to the departure date when fares are higher. Relative to leisure travelers, business passengers are also more likely to purchase expensive refundable tickets. Therefore, endpoints where the volume of originating business passengers is larger are expected to have higher fares on average. This difference is likely correlated with the income difference since residents of metro areas with large concentrations of business generally have higher average incomes.

To control for differences in the volume of business passengers, two variables are constructed using MSA level data from the Bureau of Labor Statistics' May 2014 Occupational Employment Statistics survey. The first variable, DEMPLOYSHARE, measures the difference in the employed to total population ratio between endpoints. If this ratio is small, it indicates a metro area with a small concentration of business.[25] It is assumed that these types of metro areas have a low volume of business travelers. Conversely, if the employed to total

[24] For a list of hub airports by carrier, see Appendix Table B.2.

[25] Metro areas where the employment to population ratio is small may have a large retired population or a large unemployed population. In either case, these types of metro areas are assumed to have small concentrations of business.

population ratio is large, it indicates a metro area with a large concentration of business. These types of metro areas are expected to have a high volume of business travelers.

The second variable constructed, DMANAGERIAL, measures the difference in the percentage of workers in managerial occupations between endpoints.[26] Considering that workers in managerial occupations are the most likely type of worker to travel for business, DMANAGERIAL, in conjunction with DEMPLOYSHARE, should appropriately control for the effect of differences in the volume of business passengers on the average fare difference between endpoints.

Peak vs Off-Peak Flights

A passenger purchasing a flight leaving at a peak congested time will often pay more than a passenger purchasing a similar flight with a departure leaving at an off-peak time. An example of this type of differential pricing is flights departing in the late afternoon to early evening hours being more expensive than comparable red-eye flights.[27] Thus, an avenue exists for average roundtrip fares to directionally differ if a disproportionate number of flights depart at peak times from one endpoint. If the difference in the proportion of peak vs. off-peak flights on a route is correlated with the income difference, then omitting a variable controlling for this difference will bias the estimate of the income discrimination effect. For instance, it is plausible that passengers residing at higher income endpoints benefit not only from greater flight frequency, but also may benefit from a disproportionate number of flights departing at preferred peak times.

One major limitation of the DB1B data is that exact departure and arrival times of flights

[26] From the Occupational Employment Statistics survey, the proportion of managerial workers for each metro area is constructed as the total number of workers in "Management Occupations" divided by the total number of workers in all occupations. This value is then multiplied by 100 to get the percentage.

[27] During the work week, peak travel times in later half of the day are typically between 4pm and 8pm whereas red-eye flights depart at off-peak hours late in the evening.

are not observed for each itinerary.[28] Therefore, it is difficult to construct a variable that controls for the difference in the proportion of peak and off-peak flights. A next best option is to construct a variable measuring the difference in time zones between origin and destination airports. This variable would be able to control for differences in the proportion of peak flights on routes where the origin and destination airports are in different time zones. However, this variable would be unable to control for differences in the proportion of peak flights on routes within the same time zone. This concern is limited by assuming that flights depart at similar times in both directions on routes within the same time zone. This assumption appears adequate for heavily traveled coastal routes like LAX-SFO and DCA-JFK where flights typically depart at similar times in both directions.

The time zone difference variable constructed for use in the analysis is DTIME. It is defined as the difference in Greenwich Mean Times (GMT) between the origin and destination airports. For instance, JFK is at GMT -5 while LAX is at GMT -8. Therefore, a flight from LAX to JFK would have a DTIME value of -3.[29] Considering the LAX-JFK route, a passenger originating at LAX must depart in the morning if they wish to arrive in New York before close of business, implying that flights originating in western time zones are more constrained to depart at peak travel times. Given this expectation, the coefficient on DTIME is predicted to be negative, resulting in fares that are more expensive for passengers originating roundtrip travel at the more-westerly endpoint on a route.

Differences in Passenger Demographics

Another potential avenue for airlines to directionally price discriminate is to price flights taking into account demographic characteristics like race. Although it is not believed that racial demographics factor into airline pricing, race is highly correlated with other demographics

[28] The dates of travel are also not observed in the DB1B data. Only the quarter of travel is observed.

[29] Analogously, flights from JFK to LAX would have a DTIME value of 3 (-5 - (-8) = 3). Similarly, a flight from LAX to SEA would have a DTIME value of 0 since both airports are in the same time zone.

such as average wealth. Alperovich and Machnes (1994) and Bhadra (2012) find evidence that wealth influences the demand for air travel, with demand decreasing as wealth decreases. Since average wealth is typically lower in metro areas with large black populations, one would expect the price elasticity of demand for air travel to be higher for passengers originating at the endpoint where the proportion of the black population is higher, mirroring the effect of higher income. Differences in wealth may also contribute to a difference in demand between endpoints. Resulting demand differentials may then be reflected in average roundtrip fares, with fares being cheaper at the endpoint where the proportion of the black population is higher.

To control for differences in passenger demographics, the variable DBLACK is constructed as the difference in the fraction of the black population between origin and destination metro areas.[30] It is expected that endpoints where the proportion of the black population is higher will have lower average fares.

[30] Data used to construct DBLACK comes from the 2014 American Community Survey.

Table 1.3: Regression Results with Additional Controls

Analysis Sample:	Restricted (1)
DINC ($1,000s)	0.18***
	(0.032)
DPOP (1,000,000s)	0.60***
	(0.049)
DTEMP	0.14***
	(0.025)
DHUB	7.87***
	(0.475)
DMANAGERIAL	1.45***
	(0.225)
DEMPLOYSHARE	-0.00
	(0.027)
DTIME	-1.33***
	(0.264)
DBLACK	-15.30***
	(3.062)
Constant	-1.65***
	(0.285)
Observations	6,678
R-squared	0.172

Notes: Robust standard errors are reported in parentheses. The restricted sample removes outliers that fall outside the range encompassing the 5th and 95th percentiles of the dependent variable DFARE.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Impact of Additional Controls

Table 1.3 provides results when DHUB, DMANAGERIAL, DEMPLOYSHARE, DTIME, and DBLACK are added to the specification presented in column two of Table 1.2.[31] The

[31] The regression presented in Table 1.3 has 6 fewer observations than the regression presented in column two of Table 1.2 because the data used to construct DMANAGERIAL, DEMPLOYSHARE, and DBLACK

coefficient on DINC remains positive and significant at conventional levels. However, the magnitude has decreased from $0.35 to $0.18 per $1,000 difference in income. For the SFO-ORD route, this coefficient implies that roundtrip fares are $3.90 higher on average for passengers originating in San Francisco.[32] This estimated effect of $3.90 per passenger suggests that United generated an additional $337,311 in revenue from SFO originating nonstop passengers in 2015.[33]

Of the five additional controls added to the specification presented in column two of Table 1.2, only the coefficient on DEMPLOYSHARE is insignificant. As expected, the positive and significant coefficient on DHUB supports the existence of a hub premium, with fares predicted to be $7.87 higher on average when the roundtrip departs from a hub airport.

The coefficient on DMANAGERIAL indicates that fares are $1.45 higher for each 1% difference in the percentage of managerial workers between origin and destination markets. Referencing once more the SFO-ORD route, there are 1.18% more managerial workers in the San Francisco metro area. Not taking into account the effects of other covariates, this coefficient implies that fares are $1.71 higher on average for SFO originating passengers.[34]

The negative coefficient on DTIME indicates that roundtrip flights are cheaper when traversing time zones from east to west. For example, when flying from JFK to LAX, the difference in GMT is plus 3. Thus, the estimated coefficient of -1.33 in column one of Table 1.3 indicates that roundtrip fares are $3.99 cheaper on average for JFK originating passengers on the JFK-LAX route.[35]

The DBLACK coefficient is negative and significant, supporting the expectation that fares are cheaper for passengers originating from metro areas with low levels of wealth. This coef-

was not available for the Walla Walla, WA and Kahului-Wailuka-Lahaina, HI MSAs.

[32] 21.674*$0.18 =$3.90.

[33] $3.90*86,490 = $337,311.

[34] 1.18*$1.45 = $1.71.

[35] -$1.33*3 = -$3.99. This calculation does not take into account the impact of other covariates.

ficient of -15.30 indicates that roundtrip fares are $0.15 cheaper for each 0.01 unit difference in the fraction of the black population between origin and destination metro areas. Continuing with the SFO-ORD example, 17.4% of the Chicago metropolitan area is black compared to 8.4% in San Francisco. This 9% difference corresponds to a roundtrip fare that is $1.38 cheaper for passengers originating in Chicago on the SFO-ORD route.[36]

1.5.2 Asymmetries in the Share of Originating Passengers

Differences in the level of demand between route endpoints may be the driving force behind directional fare differences rather than the proposed income discrimination effect. For example, instead of a yield manager allocating fewer seats to discount fare buckets at high income origins, capacities could be the same regardless of origin point. If the level of demand at the high income endpoint is higher, lower-priced tickets will be sold faster, resulting in more higher-priced tickets being bought. The observed outcome in these markets would be observationally equivalent to a yield manager setting different capacities by point of origin.[37]

To investigate this possibility, the analysis sample is split into two groups. The first group contains routes where 70% or more of the roundtrip passengers originate from one of the endpoints. Routes with large asymmetries in the share of roundtrip passengers are assumed to have levels of demand that are substantially higher at the endpoint with the larger share of originating passengers. Examples of asymmetric routes include city-pairs involving Las Vegas and Honolulu, where most passengers originate from other endpoints. The second group contains all routes that do not meet the cutoff for the first group. If the magnitude of the DINC coefficients are found to be substantially different between groups with the coefficient on DINC larger in the asymmetric sample, then evidence would support the narrative that differences in the level of demand are the primary factor driving directional

[36] 0.09*-$15.30 $\approx$ -$1.38. This calculation does not take into account the effects of other covariates.

[37] I would like to thank the referees for pointing out this possibility.

fare differences.

Table 1.4: Regression Results for Asymmetric and Non-Asymmetric Markets

	Asymmetric Routes		Non-Asymmetric Routes	
Analysis Sample:	Restricted	Restricted	Restricted	Restricted
	(1)	(2)	(3)	(4)
DINC ($1,000s)	0.34***	0.25***	0.35***	0.17***
	(0.067)	(0.080)	(0.030)	(0.034)
DPOP (1,000,000s)	0.65***	0.74***	0.78***	0.53***
	(0.098)	(0.113)	(0.047)	(0.054)
DTEMP	-0.05	0.09*	0.18***	0.18***
	(0.044)	(0.050)	(0.030)	(0.030)
DHUB		3.94***		8.72***
		(1.159)		(0.514)
DMANAGERIAL		1.52***		1.33***
		(0.568)		(0.243)
DEMPLOYSHARE		0.17***		-0.05
		(0.056)		(0.030)
DTIME		1.49***		-2.29***
		(0.577)		(0.292)
DBLACK		-26.56***		-12.10***
		(7.465)		(3.324)
Constant	-0.05	-0.08	-1.56***	-2.21***
	(0.641)	(0.645)	(0.314)	(0.311)
Observations	1,563	1,562	5,121	5,116
R-squared	0.101	0.125	0.118	0.206

Notes: Robust standard errors are reported in parentheses. The restricted sample removes outliers that fall outside the range encompassing the 5th and 95th percentiles of the dependent variable DFARE. Asymmetric routes are city-pair markets where 70% or more of the roundtrip passengers originate from one of the endpoints. Non-Asymmetric routes are all other city-pair markets.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

To determine if this is the case, column one of Table 1.4 provides estimates when equation (1.1) is estimated on asymmetric markets while column three provides estimates for the same equation on non-asymmetric markets. The coefficients on DINC are nearly identical between both samples, with the value in the asymmetric sample equaling \$0.34 compared to \$0.35 in the non-asymmetric sample. After adding the additional control variables in columns two and four, the DINC coefficient decreases to \$0.25 in the asymmetric sample and \$0.17 in the non-asymmetric sample. While the coefficient in the asymmetric sample is larger once additional controls are added, the \$0.08 difference is not large. Furthermore, when equation (1.1) is estimated on the combined sample, the coefficient on DINC is \$0.18 (see Table 1.3), almost identical to the \$0.17 coefficient found in the non-asymmetric sample. Therefore, markets with large asymmetries in the share of originating passengers are not noticeably influencing the magnitude of the coefficient on DINC. This finding suggests that differences in income is the more likely factor explaining directional fare differences than are differences in the levels of demand.

1.5.3 Effects by Route Distance

Table 1.5 investigates whether the income, population, and leisure effects differ by route distance. It is conceivable that the magnitudes of these effects increase with distance. For instance, competition from other transport modes may suppress an airline's ability to price discriminate on routes where distance between origin and destination cities is small. However, when distance between endpoints is large, air travel becomes the only suitable mode of transport.

Supplementing the specification presented in Table 1.3, column one of Table 1.5 adds interaction terms allowing the income, population, and leisure effects to differ with distance. Contrary to expectation, the coefficients on all three interaction terms while having the predicted positive signs, are insignificant.

As an additional sensitivity, the analysis sample is split into three groups by distance. Column two presents estimates when the specification presented in Table 1.3 is estimated on routes where the distance between endpoints is less than 500 miles. Column three estimates the same model on routes between 500 and 1500 miles while column four provides estimates when the distance between endpoints is greater than 1500 miles. As expected, the coefficients on DINC increase between columns two and four, indicating that income discrimination increases with distance. For routes where the distance between origin and destination airports is less than 500 miles, roundtrip fares are \$0.09 higher for each \$1,000 difference in per capita incomes. However, this effect is insignificant. When distance is between 500 and 1500 miles, the estimated effect increases to \$0.13 for each \$1,000 difference in incomes between the endpoint cities. When the distance between origin and destination cities is over 1500 miles, roundtrip fares are \$0.41 higher for each \$1,000 difference in per capita incomes. For the SFO-ORD route, this effect results in a roundtrip fare that is predicted to be \$8.89 higher for passengers originating in San Francisco.[38]

The coefficients on DTEMP also increase with distance. For routes less than 500 miles, roundtrip fares are \$0.07 higher for each one degree difference in the quarterly high temperatures between endpoints. Similar to the coefficient on DINC, this effect is also insignificant, providing suggestive evidence that competition from other transport modes on short haul routes substantially impacts a carrier's ability to directionally price discriminate. For markets where distance is between 500 and 1500 miles, the estimated effect increases to \$0.15 while the effect for routes over 1500 miles is slightly larger at \$0.19 for each one degree difference in the quarterly high temperature between origin and destination airports.

[38] 21.674*\$0.41 $\approx$ \$8.89.

Table 1.5: Regression Results with Differential Effects by Route Distance

Analysis Sample:	Restricted	Restricted < 500 Miles	Restricted Between 500 and 1500 Miles	Restricted > 1500 Miles
	(1)	(2)	(3)	(4)
DINC ($1,000s)	0.14***	0.09	0.13***	0.41***
	(0.047)	(0.063)	(0.042)	(0.091)
DINC*DISTANCE (1,000s miles)	0.05			
	(0.041)			
DPOP (1,000,000s)	0.58***	0.43***	0.70***	0.46***
	(0.079)	(0.093)	(0.072)	(0.110)
DPOP*DISTANCE (1,000s miles)	0.03			
	(0.060)			
DTEMP	0.15***	0.07	0.15***	0.19***
	(0.052)	(0.066)	(0.032)	(0.056)
DTEMP*DISTANCE (1,000s miles)	-0.00			
	(0.043)			
DHUB	7.94***	8.94***	8.69***	5.71***
	(0.476)	(0.881)	(0.663)	(1.080)
DMANAGERIAL	1.45***	1.75***	1.43***	0.32
	(0.226)	(0.436)	(0.298)	(0.637)
DEMPLOYSHARE	0.00	-0.06	0.04	-0.05
	(0.027)	(0.051)	(0.036)	(0.071)
DTIME	-1.37***	-7.51***	-2.10***	-0.06
	(0.268)	(1.022)	(0.438)	(0.573)
DBLACK	-15.12***	-13.17***	-20.47***	-21.02**
	(3.077)	(5.006)	(4.427)	(9.579)
Constant	-1.64***	-0.68	-2.52***	-0.26
	(0.285)	(0.509)	(0.380)	(0.797)
Observations	6,678	1,703	3,792	1,183
R-squared	0.172	0.236	0.172	0.129

Notes: Robust standard errors are reported in parentheses. The restricted sample removes outliers that fall outside the range encompassing the 5th and 95th percentiles of the dependent variable DFARE. The DISTANCE variable is defined as the distance in thousands of miles between origin and destination airports.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

1.5.4 Legacy vs Low-Cost Carriers

Table 1.6 allows for the income, population, and leisure effects to differ depending on the airline's status as a legacy or low-cost carrier.[39] In the DB1B data, the legacy carriers are Alaska, American, Delta, Hawaiian, United, and US Airways while the low-cost carriers are Allegiant, Frontier, JetBlue, Southwest, Spirit, Sun Country, and Virgin America.[40]

The most striking results from the Table 1.6 regressions are the different income effects for legacy and low-cost carriers. Given that legacy carriers typically serve a more price-inelastic customer base than do low-cost carriers, one would expect the income discrimination effect for legacy carriers to be stronger than the effect for low-cost carriers. However, the opposite result is found. For each $1,000 difference in average per capita income between endpoint metro areas, low-cost carriers set average roundtrip fares that are $0.34 higher for passengers originating in the higher income locale compared to $0.04 for legacy carriers.[41] It is conceivable that these results could be driven by routes where low-cost carriers are the only competitor. To investigate this possibility, column two of Table 1.6 restricts the sample to routes where both legacy and low-cost carriers offer nonstop service. Surprisingly, the magnitude of the income discrimination effect is even larger on these routes for low-cost carriers, increasing from $0.34 in column one to $0.56 in column two. One possible explanation for this surprising result is that, even though price elasticities for the typical low-cost passenger are higher than for the typical legacy passenger, elasticities display more variation among the low-cost passenger base than among the legacy passenger base. If this is true, low-cost carriers could exploit this variation to income discriminate more among their price-elastic customer base than legacy carriers.

[39] Since equations (1.1)-(1.4) outline a model using variables with values that are differenced between origin and destination, the legacy indicator variable that does not differ by origin/destination cannot be included by itself as it would be differenced away. To obtain differential effects for legacy carriers, the differenced variables must be interacted with the legacy indicator variable.

[40] Although US Airways and American announced their merger in 2013, flights operated under US Airways flight numbers in the first two quarters of 2015.

[41] $0.34-$0.30 = $0.04.

Table 1.6: Regression Results with Legacy Interactions

Analysis Sample:	Restricted	Competing Routes
	(1)	(2)
DINC ($1,000s)	0.34***	0.56***
	(0.042)	(0.053)
DINC*LEGACY	-0.30***	-0.46***
	(0.053)	(0.070)
DPOP (1,000,000s)	0.71***	0.63***
	(0.061)	(0.079)
DPOP*LEGACY	-0.19**	-0.16
	(0.087)	(0.111)
DTEMP	0.07**	0.10**
	(0.032)	(0.041)
DTEMP*LEGACY	0.16***	0.19***
	(0.047)	(0.063)
DHUB	9.09***	10.48***
	(0.515)	(0.705)
DMANAGERIAL	1.26***	0.33
	(0.226)	(0.291)
DEMPLOYSHARE	-0.00	-0.02
	(0.027)	(0.037)
DTIME	-1.28***	-0.96***
	(0.263)	(0.351)
DBLACK	-15.99***	-14.40***
	(3.040)	(4.421)
Constant	-1.49***	-2.34***
	(0.284)	(0.395)
Observations	6,678	3,390
R-squared	0.184	0.204

Notes: Robust standard errors are reported in parentheses. The restricted sample removes outliers that fall outside the range encompassing the 5th and 95th percentiles of the dependent variable DFARE. The competing routes sample limits the restricted sample to routes where legacy and low-cost carriers both offer nonstop service.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Conversely, since legacy carriers serve a broader mix of price-elastic and price-inelastic passengers, the temperature effect should be stronger for legacy carriers. This expectation is supported by the coefficient on DTEMP*LEGACY in column one of Table 1.6. For each extra one degree difference in the average quarterly high temperature between endpoint cities, legacy carriers set fares that are \$0.23 higher compared to \$0.07 for low-cost carriers.[42] In the sample restricted to competing routes, these effects increase to \$0.29 for legacy carriers and \$0.10 for low-cost carriers.

1.5.5 Effects by Time of Year

Table 1.7 presents estimates when the income, population, and leisure effects are allowed to vary by quarter-of-year. Column one presents estimates when DINC, DPOP, and DTEMP are interacted with indicators specifying whether travel occurred in the first, second, or third quarters. There is substantial heterogeneity in the quarterly effects. The significant and negative coefficients on DINC*Q1, DINC*Q2, and DINC*Q3 indicate that the income discrimination effect is strongest during the fourth quarter. This result is sensible considering that the volume of passengers in the fourth quarter is larger and more likely to contain a higher proportion of leisure travelers due to the Thanksgiving and Christmas holidays. During the fourth quarter, roundtrip fares are \$0.45 higher for each \$1,000 difference in per capita income between the endpoint cities, compared to \$0.08 in the first quarter and \$0.11 in the second quarter.[43] The third quarter effect, while negative, is statistically insignificant.[44]

Similar to the income effects, the population effect is most pronounced in the fourth quarter when the demand for leisure travel is highest. For the fourth quarter, roundtrip fares are \$0.81 higher for each 1,000,000 difference in population between the endpoint cities,

[42] \$0.07+\$0.16 = \$0.23.

[43] \$0.45-\$0.37 = \$0.08 and \$0.45-\$0.34 = \$0.11. These calculations do not take into account the effects of other covariates.

[44] A hypothesis test with the null hypothesis that the coefficients on DINC and DINC*Q3 sum to 0 cannot be rejected at standard significance levels.

compared to \$0.55, \$0.56, and \$0.52 in the first, second, and third quarters, respectively.

In contrast to the income and population effects, the leisure effect is strongest in the first quarter and statistically insignificant in the second, third, and fourth quarters. In the first quarter, roundtrip fares are \$0.42 higher for each one degree difference in the average quarterly high temperature between origin and destination cities.

Column two of Table 1.7 provides estimates when legacy interactions are added to the specification from column one. The coefficient estimates on the legacy interaction terms are robust in the sense that the estimates are similar in magnitude to the estimates from column one of Table 1.6. Therefore, even with the addition of quarter-of-year effects, the income discrimination effect remains stronger for low-cost carriers while the population differential effect remains stronger for legacy carriers.

Table 1.7: Regression Results with Quarter-of-Year Interactions

Analysis Sample:	Restricted (1)	Restricted (2)
DINC ($1,000s)	0.45***	0.61***
	(0.054)	(0.061)
DINC*Q1	-0.37***	-0.38***
	(0.075)	(0.076)
DINC*Q2	-0.34***	-0.34***
	(0.071)	(0.071)
DINC*Q3	-0.46***	-0.47***
	(0.072)	(0.073)
DINC*LEGACY		-0.30***
		(0.052)
DPOP (1,000,000s)	0.81***	0.91***
	(0.080)	(0.090)
DPOP*Q1	-0.26**	-0.24**
	(0.120)	(0.120)
DPOP*Q2	-0.25**	-0.24**
	(0.109)	(0.109)
DPOP*Q3	-0.29***	-0.29***
	(0.110)	(0.110)
DPOP*LEGACY		-0.18**
		(0.085)
DTEMP	-0.07	-0.14***
	(0.047)	(0.051)
DTEMP*Q1	0.42***	0.42***
	(0.060)	(0.059)
DTEMP*Q2	0.08	0.08
	(0.067)	(0.066)
DTEMP*Q3	0.07	0.06
	(0.073)	(0.073)
DTEMP*LEGACY		0.17***
		(0.046)
DHUB	7.87***	9.08***
	(0.468)	(0.508)
DMANAGERIAL	1.41***	1.22***
	(0.223)	(0.223)
DEMPLOYSHARE	0.00	0.00
	(0.027)	(0.026)
DTIME	-1.28***	-1.24***
	(0.263)	(0.263)
DBLACK	-14.29***	-14.99***
	(3.019)	(2.995)
Constant	-1.65***	-1.49***
	(0.282)	(0.281)
Observations	6,678	6,678
R-squared	0.197	0.209

Notes: Robust standard errors are reported in parentheses. The restricted sample removes outliers that fall outside the range encompassing the 5th and 95th percentiles of the dependent variable DFARE.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

1.5.6 Competition and the Income Discrimination Effect

Table 1.8 investigates the impact of the level of competition on the magnitude of the income discrimination effect. Column one presents estimates when the number of carriers offering nonstop service between origin and destination airports is added to the specification in Table 1.3. The coefficient on NCOMPETITORS while negative, is insignificant. This result is expected considering that any variable with values that do not differ by origin or destination would be differenced away given the empirical framework outlined by equations (1.1)-(1.4) in Section 1.4.[45] Column four provides estimates when the Herfindahl-Hirschman Index (HHI) is used to measure the level of competition. Similar to the column one results, the coefficient on HHI is insignificant.

[45] For example, in the first quarter of 2015, NCOMPETITORS equals 3 for SFO-ORD and for ORD-SFO because American, United, and Virgin America offer nonstop service between SFO and ORD in both directions.

Table 1.8: Regression Results with Competition Variables

Analysis Sample:	Restricted (1)	Restricted (2)	Restricted (3)	Restricted (4)	Restricted (5)	Restricted (6)
DINC ($1,000s)	0.18***	0.05	0.07	0.18***	0.38***	0.43***
	(0.032)	(0.054)	(0.055)	(0.032)	(0.067)	(0.069)
DPOP (1,000,000s)	0.60***	0.61***	0.67***	0.60***	0.61***	0.67***
	(0.049)	(0.049)	(0.051)	(0.049)	(0.049)	(0.051)
DTEMP	0.14***	0.15***	0.16***	0.14***	0.15***	0.16***
	(0.025)	(0.026)	(0.026)	(0.025)	(0.025)	(0.026)
DHUB	7.88***	7.97***	8.01***	7.88***	8.01***	8.05***
	(0.475)	(0.477)	(0.476)	(0.475)	(0.478)	(0.476)
DMANAGERIAL	1.45***	1.49***	1.45***	1.45***	1.47***	1.43***
	(0.226)	(0.226)	(0.227)	(0.225)	(0.225)	(0.226)
DEMPLOYSHARE	-0.00	0.01	0.01	-0.00	0.01	0.01
	(0.027)	(0.027)	(0.027)	(0.027)	(0.027)	(0.027)
DTIME	-1.33***	-1.36***	-1.29***	-1.34***	-1.36***	-1.29***
	(0.264)	(0.265)	(0.268)	(0.264)	(0.264)	(0.268)
DBLACK	-15.35***	-14.83***	-13.05***	-15.36***	-14.88***	-13.09***
	(3.064)	(3.073)	(3.124)	(3.064)	(3.072)	(3.123)
NCOMPETITORS	-0.10	-0.15	-0.08			
	(0.235)	(0.236)	(0.237)			
DINC*NCOMPETITORS		0.06***	0.07***			
		(0.020)	(0.020)			
ADJACENT_ORIGIN			-2.38***			-2.33***
			(0.729)			(0.730)
ADJACENT_DEST			1.51**			1.58**
			(0.694)			(0.692)
HHI (1,000s)				0.09	0.10	0.08
				(0.102)	(0.102)	(0.102)
DINC*HHI (1,000s)					-0.03***	-0.03***
					(0.008)	(0.008)
Constant	-1.44**	-1.34**	-1.30**	-2.29***	-2.36***	-2.01**
	(0.572)	(0.573)	(0.625)	(0.769)	(0.768)	(0.804)
Observations	6,678	6,678	6,678	6,678	6,678	6,678
R-squared	0.172	0.173	0.175	0.172	0.173	0.175

Notes: Robust standard errors are reported in parentheses. The restricted sample removes outliers that fall outside the range encompassing the 5th and 95th percentiles of the dependent variable DFARE.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

To obtain income discrimination effects that vary with the level of competition, columns two and five of Table 1.8 present estimates when NCOMPETITORS and HHI are interacted with DINC. The significant positive coefficient on DINC*NCOMPETITORS and the significant negative coefficient on DINC*HHI indicates that roundtrip fares are higher at endpoints where relative incomes are higher, and that these fare differences increase with the level of competition.[46] This result coincides with Borenstein and Rose (1994), who find that the extent of a carrier's price dispersion at the route level increases with the number of competitors.

To illustrate the magnitude of these effects, consider once more the SFO-ORD route serviced by American, United, and Virgin America. The coefficients of 0.05 on DINC and 0.06 on DINC*NCOMPETITORS imply that roundtrip fares are $4.99 higher for passengers originating at SFO.[47] When HHI is used, the coefficients on DINC and DINC*HHI in column five result in a roundtrip fare that is $4.71 higher for SFO originating passengers.[48]

Columns three and six of Table 1.8 provide estimates when variables indicating the presence of competition from nearby airports are added to the specifications in columns two and five. Adopting adjacent market definitions from Brueckner et al. (2014), the list of adjacent airports by metro area is provided in Table 1.9. ADJACENT_ORIGIN is an indicator equaling one if there exists nonstop service from an adjacent airport in the origin market to the final destination. For the SFO-ORD route, ADJACENT_ORIGIN is set to 1 if there exists nonstop service from adjacent origin OAK to destination ORD. Similarly, ADJACENT_DEST is an indicator equaling one if there exists nonstop service from the origin to an adjacent airport in the destination market. For SFO-ORD, ADJACENT_DEST is set to 1 if there exists nonstop service from SFO to adjacent destination MDW.

[46] While this result differs from the standard textbook finding that, ceteris paribus, price discrimination increases with market concentration, Borenstein (1985) and Holmes (1989) find the effect of competition on third-degree price discrimination to be ambiguous in settings of imperfect competition with product differentiation.

[47] 21.674*$0.05 + 21.674*3*$0.06 $\approx$ $4.99. This calculation does not take into account the effects of other covariates.

[48] During the first quarter of 2015, HHI (1000s) from SFO-ORD was equal to 5.42. Therefore, $0.38*21.674 -$0.03*21.674*5.42 $\approx$ $4.71. This calculation does not take into account the effects of other covariates.

Table 1.9: List of Primary and Adjacent Airports

Metro Area	**Primary Airport**	**Adjacent Airport(s)**
Chicago	ORD	MDW
Cincinatti	CVG	DAY
Cleveland	CLE	CAK
Dallas	DFW	DAL
Houston	IAH	HOU
Miami	MIA	FLL
New York	LGA	EWR, JFK
San Francisco	SFO	OAK
Tampa	TPA	PIE
Washington, DC	DCA	IAD, BWI

Notes: These adjacent market definitions are adopted from Table 4 of Brueckner et al. (2014).

The addition of the adjacent competition variables has little impact on the overall magnitude of the income discrimination effect since the coefficients on DINC and DINC*NCOMPETITORS are roughly the same between columns two and three of Table 1.8. The coefficients on DINC and DINC*HHI are also similar between columns five and six. However, the presence of adjacent competition in the origin market has a significant negative effect on the roundtrip fare difference as demonstrated by the coefficients on ADJACENT_ORIGIN in columns three and six. In the model with the number of competitors, roundtrip fares are $2.38 cheaper when adjacent competition exists from an airport in the origin market. In the model with HHI, adjacent origin competition decreases roundtrip fares by $2.33.

The coefficients on ADJACENT_DEST in Table 1.8 indicate that roundtrip fares increase by $1.51-$1.58 when service to an adjacent airport in the destination market is present. While these results are counterintuitive, they suggest that passengers may view service to an adjacent destination airport as less desirable. If this is the case, passengers would be willing to pay a small premium to arrive at their preferred airport in the destination market.

The results presented in Table 1.8 establish that competition has a negligible impact on the level of income discrimination. While this conclusion differs from the standard textbook model predicting that, ceteris paribus, price discrimination decreases with competition, the results are consistent with several theoretical models of price discrimination under oligopoly. For instance, Borenstein (1985) and Holmes (1989) find the effect of competition on third-degree price discrimination to be ambiguous. In some cases, more competition leads to an increase in price discrimination.[49]

1.5.7 Effects at the Market Level

Table 1.8 demonstrates that the level of competition has a trivial impact on the magnitude of the income discrimination effect. Since the same pattern of income discrimination emerges regardless of the number of competitors in a market, there is a concern that observations are being "double counted" in markets served by multiple carriers. To address this concern, the following equation is estimated:

$$FARE_{ijl} - FARE_{jil} = \beta_0 + \beta_1 DINC_{i,j} + \beta_2 DPOP_{i,j} + \beta_3 DTEMP_{i,j} + \epsilon_{ijl} \tag{1.5}$$

where $FARE_{ijl}$ is the average roundrip fare for nonstop travel between origin airport i to destination airport j in quarter l and $FARE_{jil}$ is the average roundtrip fare for nonstop travel between origin j to destination i in quarter l.

The first column of Table 1.10 presents estimates from equation (1.5) using the full sample while column two restricts the sample to observations that fall between the 5th and 95th percentiles of the dependent variable DFARE. The coefficients on DINC are positive, significant, and similar in magnitude between the full and restricted samples. Therefore, evidence of income discrimination in the directional pricing of roundtrip flights remains when analysis

[49] For a comprehensive review of price discrimination under oligopoly, see Stole (2007).

is performed at the market level. The coefficient in column two of Table 1.10 implies that roundtrip fares are $0.29 higher for each $1,000 difference in per capita incomes between the origin and destination cities. When additional controls are added, the precision and magnitude of the coefficient decreases to $0.06.[50]

Table 1.10: Market Level Regressions

Analysis Sample:	Full (1)	Restricted (2)	Restricted (3)
DINC ($1,000s)	0.31***	0.29***	0.06*
	(0.046)	(0.033)	(0.039)
DPOP (1,000,000s)	1.74***	1.02***	1.04***
	(0.075)	(0.051)	(0.058)
DTEMP	0.25***	0.11***	0.18***
	(0.040)	(0.029)	(0.031)
DMANAGERIAL			2.55***
			(0.278)
DEMPLOYSHARE			-0.03
			(0.034)
DTIME			-1.18***
			(0.336)
DBLACK			0.81
			(3.677)
Constant	-2.69***	-2.45***	-2.36***
	(0.469)	(0.336)	(0.342)
Observations	5,200	4,680	4,672
R-squared	0.160	0.144	0.164

Notes: Robust standard errors are reported in parentheses. The restricted sample removes outliers that fall outside the range encompassing the 5th and 95th percentiles of the dependent variable DFARE.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

[50] Since the DHUB variable defined earlier is airline specific, the hub indicator was not added as an additional control in the market level analysis. However, adding a hub indicator that equals one if origin airport i is a hub for at least one carrier does not change the magnitude of the DINC coefficient.

1.6 Conclusion

This chapter has provided the first exploration of an airline's decision to charge passengers different roundtrip fares on a nonstop route depending on trip origin. These directional fare differences cannot be due to differences in cost, since there are no significant cost differences on a roundtrip basis related to the direction of travel. After controlling for differences in endpoint populations and differences in the mix of business and leisure passengers between origin and destination cities, significant evidence is found of directional price discrimination based on income, with passengers at high-income origins charged higher fares.

The analysis of directional price discrimination provided here is certainly not exhaustive. The R-squared in most of the regressions presented is below 0.2, indicating that much of the variation in average roundtrip fare differences remains unexplained. Furthermore, the nature of the DB1B data do not allow one to test whether the income discrimination effect differs by time of day, day of week, or for peak and off-peak flights. The analysis in this chapter focused on nonstop roundtrip flights. It is not clear how the income discrimination effect would carry over to connecting flights. Further empirical work using more detailed fare and itinerary data would help shed light on these questions.

Chapter 2

Are Passengers Compensated for Incurring an Airport Layover? Estimating the Value of Layover Time in the U.S. Airline Industry

2.1 Introduction

In 2015, American, Delta, Southwest, and United accounted for 80% of all domestic U.S. passenger traffic.[1] Three of these carriers, American, Delta, and United, have organized their

[1] See *"2015 U.S.-Based Airline Traffic Data"* released by the U.S. Department of Transportation's Bureau of Transportation Statistics on Thursday March 24th, 2016. This press release may be accessed at `https://www.rita.dot.gov/bts/press_releases/bts018_16`. In addition, see a Dallas News article from October 2015 titled *"American Airlines finally absorbs US Airways' brand"* or an Economist article from April 5th, 2016 titled *"Alaska Airlines' purchase of Virgin America could start a new wave of consolidation."* These articles may be accessed at `https://www.dallasnews.com/business/airlines/2015/10/16/american-airlines-finally-absorbs-us-airways-brand` and `https://www.economist.com/news/business-and-finance/21696326-new-carrier-will-become-americas-fifth-largest-alaska-airlines-purchase-virgin`, respectively.

flight operations into extensive hub-and-spoke networks.[2] These networks typically require passengers originating and concluding travel in non-hub cities to board a connecting flight at a hub on the way to the final destination. An example of the hub-and-spoke network involves a passenger traveling from Orange County, CA (SNA) to Reagan National (DCA) in Washington, DC.[3] Currently, there are no direct options on any carrier from SNA to DCA. Therefore, a passenger must incur a layover when making this trip. This required layover will often occur in one of the carrier's respective hub cities. For instance, a passenger flying American is likely to incur a layover in Dallas (DFW) or Chicago (ORD), two of American's major hubs. Similarly, a passenger flying Delta is likely to incur a layover in Atlanta (ATL), Salt Lake City (SLC), or Minneapolis (MSP), three of Delta's hubs.

From the passenger perspective, layovers are negative drains on utility as the addition to total travel time relative to a nonstop itinerary is a cost incurred by the passenger. Passengers generally do not wish to sit in an airport for an extended period of time waiting for connecting flights. In some city-pair markets where multiple routing options are available to transport passengers from origin to destination, a prospective passenger may be faced with the choice of purchasing an expensive nonstop itinerary or a cheaper one requiring a layover. The difference in fare between nonstop and connecting itineraries indicates that a flight's price is partially a function of total layover time, with flights expected to decrease in price as the amount of layover time increases.[4]

From the carrier perspective, there exists a trade-off between providing convenient flight connections for passengers and reducing airport congestion. By narrowing the gap between flights, a carrier reduces a connecting passenger's layover time. However, narrowing the flight

[2] Southwest primarily operates a point-to-point network.

[3] For detailed descriptions and analysis of hub-and-spoke networks, see Brueckner and Spiller (1991), Brueckner (2004), and Hendricks et al. (1995).

[4] Investigating whether privatizing the San Francisco Bay Area airports would improve passenger welfare and increase airline profits, Yan and Winston (2014) find a large difference in willingness to pay for nonstop and connecting flights. Using a discrete choice model that allows for unobserved preference heterogeneity, Yan and Winston (2014) find that passengers are willing to pay \$212 for nonstop flights compared to \$17 for itineraries with connections that are 90 minutes or less in duration.

gap has the adverse effect of increasing airport congestion (see Brueckner and Lin (2016)).

Taking both the passenger and carrier perspectives into account, it is clear that layover time is a major component that factors into a prospective passenger's purchase decision and an airline's flight scheduling decision. This chapter provides insight into these decisions by providing empirical estimates on the value of layover time in the U.S. airline industry.

One approach to estimate the value of layover time is to obtain data on actual consumer purchases and estimate a multinomial logit (MNL) or multinomial probit model where the ratio of coefficients on the fare and total layover time covariates would provide the appropriate "willingness to pay" measure per minute of layover time. Using a MNL and 1983 data for five randomly selected city-pair markets, Morrison and Winston (1989) finds a value of air travel time of $34 an hour. Reflecting the disutility of waiting for connecting flights, the authors find an even higher value of layover time at $74 an hour. In addition to providing an updated estimate of the average value of layover time, this chapter improves upon the Morrison and Winston (1989) study by exploring heterogeneity in the valuations of layover time.[5] Specifically, this chapter presents estimates that are airline-specific and hub-specific, as well as estimates that differ by purchase date to reflect differences in the valuations of layover time by passenger type.[6]

Instead of relying on consumer purchase data in conjunction with a MNL model, this chapter employs an alternative strategy by constructing a dataset using posted fare and itinerary information from Google flights. With these data, a hedonic regression is estimated where fare is regressed on a vector of flight characteristics, including the total amount of layover time. The estimated hedonic model has two key advantages over the standard MNL model. Foremost, the hedonic model does not require the independence from irrelevant alternatives

[5] An updated estimate of the value of layover time is warranted for several reasons. Foremost, the U.S. airline industry has recently emerged from a period of substantial consolidation. Second, the quality of airports and airline terminals have improved since the early 1980s. Therefore, there are reasons to believe that the value of layover time has changed since 1983.

[6] This approach assumes that business passengers purchase later than leisure passengers.

(IIA) assumption to hold. Most importantly however, the model in this chapter does not have an error term that takes the standard logit form. As discussed in Ackerberg and Rysman (2005), the inflexibility of standard logit errors may lead to biased estimates of price elasticities and other parameters of interest such as substitution patterns and welfare effects.[7]

After collecting data for a period of three weeks for the top 50 domestic U.S. routes where passengers are required to incur a layover (no nonstop option is available), this chapter finds that passengers are compensated with a fare that is approximately $42.74-$47.60 cheaper per hour of layover time. When estimating values for each of the three dominant legacy carriers, United passengers are found to be compensated at an even higher rate of $61.89 per hour.

The rest of this chapter is organized as follows. Section 2.2 provides a brief review of the previous literature concerning layover time. Section 2.3 describes the fare and itinerary data collected for the empirical analysis. Section 2.4 outlines the hedonic regression model while Section 2.5 discusses regression results. Section 2.6 uses the estimates presented in Section 2.5 to evaluate the trade-off between convenient flight connections and airport congestion in a hypothetical scenario applying to Chicago's O'Hare airport. Finally, Section 2.7 concludes.

2.2 Background

Prior to the Airline Deregulation Act of 1978, airlines faced constraints on both fares and route structures. In addition to not being allowed to set fares, airlines had to receive permission from the Civil Aeronautics Board (CAB) to fly new routes or discontinue existing

[7] Other papers discussing the undesirability of standard logit errors include Petrin (2002) and Bajari and Benkard (2005), among others. The shortcomings of using a standard logit error term would not have been known to Morrison and Winston (1989) at the time of their study.

service.[8] After the Airline Deregulation Act became effective in the 1980s, carriers began to organize operations into hub-and-spoke networks.[9] By transporting passengers through a hub, airlines achieve higher traffic densities on all routes. Economies of traffic density (increasing returns at the route level) may then be exploited, resulting in lower cost per passenger.[10] However, incurring a layover is costly from the passenger perspective. Airlines are able to reduce layover time by increasing flight frequency.[11] While increasing frequency benefits passengers because they are more likely to find a flight that leaves at their preferred departure time, increased frequency has the adverse effect of increasing congestion on the tarmac as more flights queue for take-off.

Brueckner and Lin (2016) provided the first theoretical analysis of this trade-off between convenient flight connections and airport congestion. Using a continuous spatial model, Brueckner and Lin (hereafter BL) observed that when the passenger's cost per unit of layover time rises, a monopoly airline will choose to narrow the gap between its flights, resulting in decreased layover times but increased congestion. In a discrete spatial model where flights congest one another only if they operate in the same period, BL find that the monopoly airline will concentrate its flights in one period when the layover costs are high relative to the costs of congestion.

When analyzing congestion at hub airports, Flores-Fillol (2010) finds that both underprovision and overprovision of flight frequency are possible when airlines compete in layover time. Analyzing network structure and airline competition, Brueckner (2004) finds that hub-and-spoke flight frequency, local hub-and-spoke passenger volume, and connecting hub-and-spoke passenger volume all decline when the cost of layover time for connecting pas-

[8] For a detailed history of the U.S. airline industry, see Morrison and Winston (1995).

[9] For a review of the changes in the airline industry brought forth by deregulation, see Levine (1986). Bailey et al. (1985) illustrate this "hubbing" phenomenon during the early 1980s.

[10] Caves et al. (1984) find the the cost per passenger mile falls as traffic density in the airline's network increases.

[11] Morrison and Winston (1995) argue that the adoption of hub-and-spoke networks are partly responsible for the increase in flight frequencies observed in the post-deregulation airline industry.

sengers increases. Using a stated preference experiment, Theis et al. (2006) finds evidence of an n-shaped utility curve with regard to layover time, supporting the hypothesis that passengers are risk averse towards very short connections while also disliking extremely long connections.[12] This chapter provides empirical estimates on the value of layover time, which are relevant to the results in the Brueckner (2004), Brueckner and Lin (2016), Flores-Fillol (2010) and Theis et al. (2006) papers.

This chapter is not the first to provide empirical estimates on the value of layover time. Using data from 1983 for five randomly selected city-pair markets, Morrison and Winston (1989) finds an average value of layover time equal to $74 per hour. This chapter differs from Morrison and Winston (1989) in several respects. Foremost, this chapter utilizes more recent data encompassing fifty markets instead of five. Second, unlike the Morrison and Winston (1989) study, this chapter explores heterogeneity in the valuations of layover time by estimating values that are specific to each of the three dominant legacy carriers (American, Delta, and United) in addition to values that differ with how far in advance airfare is purchased. Third, the model employed in this chapter is different from the multinomial logit (MNL) used in Morrison and Winston (1989).

Instead of a MNL, this chapter estimates a hedonic regression where fare is regressed on a vector of flight characteristics. This approach has two advantages over the MNL model. First, the hedonic model does not require the independence from irrelevant alternatives (IIA) assumption to hold. Most importantly however, the hedonic model does not have an error term that takes the standard logit form. As discussed in Ackerberg and Rysman (2005), the inflexibility of standard logit errors may lead to biased estimates of price elasticities and other parameters of interest such as substitution patterns and welfare effects.[13]

[12] An n-shaped utility curve is similar to an inverted U-shape.

[13] The shortcomings of the standard logit error term would not have been known to Morrison and Winston (1989) at the time of their study. Other papers discussing the undesirability of standard logit errors include Petrin (2002) and Bajari and Benkard (2005).

2.3 Fare and Itinerary Data

To estimate the value of layover time, one approach is to obtain data on actual consumer purchases. With these data, a MNL or multinomial probit (MNP) could be estimated where the ratio of coefficients on the fare and total layover time covariates would provide the appropriate "willingness to pay" measure per minute of layover time.[14] However, obtaining this data is prohibitively expensive. The Airline Origin and Destination Survey (DB1B) released quarterly by the Department of Transportation provides information on actual fares paid. However, the DB1B data do not provide information on specific dates of travel or how long layovers were for each itinerary. With these shortcomings in mind, this chapter employs the next best alternative by constructing a dataset using posted fare and itinerary information. With these data, a hedonic regression is estimated where the flight's price is regressed on a vector of flight characteristics, including the total amount of layover time.

In lieu of collecting posted fare and itinerary data for all possible routes in the domestic U.S. market, the DB1B data is first relied upon to identify the top 50 routes without nonstop service.[15,16] These routes require passengers to use connecting flights en route to their final destination. Connecting routes are ranked by the total number of roundtrip passengers observed traveling the route during the first quarter of 2016.[17] Table 2.1 provides a list of the top 50 routes.

---

[14] If one were to obtain actual consumer purchase data, a Bayesian MNP model (such as the one outlined in McCulloch et al. (2000)) would be ideal for two reasons. Foremost, estimating a MNP model ensures that one is not constrained by the restrictive Independence of Irrelevant Alternatives (IIA) assumption of the MNL model. Secondly, considering that passengers are often choosing between more than six different routing options for travel between origin and destination markets, a Bayesian MNP model will be able to compute the high-dimensional integrals required to estimate choice probabilities, whereas frequentist likelihood-based estimation would be unable to handle such high-dimensional integrals.

[15] Collecting posted fare and itinerary data for all possible origin-destination pairs in the domestic U.S. market would require an immense amount of computing resources. Focusing on the top 50 routes that require passengers to take a connecting flight should provide a reasonable picture on how layover time is valued in the U.S. airline industry.

[16] Throughout this chapter, a route is defined as a unique origin-destination pair. Directionality is not suppressed, meaning that PDX-MCO roundtrips are different from MCO-PDX.

[17] Data from the DB1B survey are released quarterly and generated from a 10 percent random sample of all airline tickets that originate in the United States on U.S. based carriers.

Routes with nonstop service were excluded from the analysis due to concerns regarding self-selection bias. In markets with both nonstop and connecting service, there will be self-selection across routes, with low-value-of-time passengers most likely choosing the connecting option. This self-selection would consequently result in a downward biased estimate of the average value of layover time. Focusing on routes requiring connections rules out this self-selection, leading to better estimates.

Although limiting the analysis to connecting routes removes selection bias on the passenger side, this bias may still play a role if airlines select which city-pairs to serve nonstop and these are routes with a relatively large share of high-value-of-time passengers. While this type of selection is possible, the fifty routes in Table 2.1 were most likely not provided with nonstop service due to insufficient passenger demand. For instance, the PHL-SAT route listed in the fourth row had a total of 8,460 roundtrip passengers in the first quarter of 2016 across all carriers.[18] This traffic equates to 93 passengers per day, insufficient demand to justify daily nonstop service.[19]

[18] Because the DB1B data are a 10% sample, $\frac{846}{.10} = 8,460$.

[19] There were 91 days in the first quarter of 2016. $\frac{8,460}{91} = 92.97$.

Table 2.1: Top 50 Roundtrip Routes Without Nonstop Service in 2016 Q1

Rank	Origin	Destination	DB1B Roundtrip Passengers
1	Portland, OR (PDX)	Orlando, FL (MCO)	1,413
2	Los Angeles, CA (LAX)	New York, NY (LGA)	1,028
3	Sacramento, CA (SMF)	Orlando, FL (MCO)	912
4	Philadelphia, PA (PHL)	San Antonio, TX (SAT)	846
5	Salt Lake City, UT (SLC)	Fort Lauderdale, FL (FLL)	775
6	Oklahoma City, OK (OKC)	Orlando, FL (MCO)	707
7	Jacksonville, FL (JAX)	Las Vegas, NV (LAS)	697
8	Washington, DC (DCA)	San Diego, CA (SAN)	661
9	Hartford County, CT (BDL)	Phoenix, AZ (PHX)	642
10	New York, NY (LGA)	Los Angeles, CA (LAX)	626
11	Indianapolis, IN (IND)	San Diego, CA (SAN)	617
12	Milwaukee, WI (MKE)	San Diego, CA (SAN)	597
13	New York, NY (LGA)	Aspen, CO (ASE)	593
14	San Francisco, CA (SFO)	New York, NY (LGA)	585
15	San Antonio, TX (SAT)	Washington, DC (DCA)	585
16	Raleigh, NC (RDU)	San Diego, CA (SAN)	581
17	Portland, OR (PDX)	Fort Lauderdale, FL (FLL)	553
18	Washington, DC (DCA)	San Antonio, TX (SAT)	553
19	Richmond, VA (RIC)	Las Vegas, NV (LAS)	551
20	Tulsa, OK (TUL)	Orlando, FL (MCO)	548
21	Grand Rapids, MI (GRR)	Phoenix, AZ (PHX)	545
22	Orange County, CA (SNA)	Orlando, FL (MCO)	543
23	Omaha, NE (OMA)	San Diego, CA (SAN)	530
24	Norfolk, VA (ORF)	San Diego, CA (SAN)	529
25	Tampa, FL (TPA)	San Diego, CA (SAN)	527
26	Fort Lauderdale, FL (FLL)	Salt Lake City, UT (SLC)	527
27	Louisville, KY (SDF)	Fort Lauderdale, FL (FLL)	525
28	Minneapolis, MN (MSP)	Jacksonville, FL (JAX)	519
29	Raleigh, NC (RDU)	Austin, TX (AUS)	503
30	San Jose, CA (SJC)	Orlando, FL (MCO)	503
31	Spokane, WA (GEG)	San Diego, CA (SAN)	500
32	Pittsburgh, PA (PIT)	San Diego, CA (SAN)	500
33	Rochester, NY (ROC)	Fort Myers, FL (RSW)	500
34	San Antonio, TX (SAT)	New York, NY (LGA)	499
35	Portland, OR (PDX)	New Orleans, LA (MSY)	498
36	Norfolk, VA (ORF)	Las Vegas, NV (LAS)	496
37	Portland, OR (PDX)	Tampa, FL (TPA)	494
38	St. Louis, MO (STL)	West Palm Beach, FL (PBI)	488
39	Kansas City, KS (MCI)	San Antonio, TX (SAT)	476
40	Hartford County, CT (BDL)	San Francisco, CA (SFO)	475
41	Albuquerque, NM (ABQ)	Washington, DC (DCA)	469
42	Detroit, MI (DTW)	Portland, OR (PDX)	468
43	Tampa, FL (TPA)	Salt Lake City, UT (SLC)	460
44	New York, NY (LGA)	Salt Lake City, UT (SLC)	456
45	Albuquerque, NM (ABQ)	Orlando, FL (MCO)	451
46	Philadelphia, PA (PHL)	Key West, FL (EYW)	449
47	Sacramento, CA (SMF)	Washington, DC (DCA)	443
48	Jacksonville, FL (JAX)	Los Angeles, CA (LAX)	442
49	Manchester, NH (MHT)	Fort Lauderdale, FL (FLL)	442
50	Detroit, MI (DTW)	Orange County, CA (SNA)	440

Notes: Data from the DB1B survey are released quarterly and generated from a 10 percent random sample of all airline tickets that originate in the United States on U.S. based carriers. For instance, if 1,413 roundtrip passengers are observed in the DB1B data traveling between Portland, OR (PDX) and Orlando, FL (MCO) in the first quarter of 2016, we would expect actual passenger traffic to be $\approx \frac{1{,}413}{.10} = 14,130$.

For the purpose of data collection, posted fares and itineraries for the fifty routes listed in Table 2.1 were obtained from Google Flights.[20] Data were collected over a twenty-one day period during the first quarter of 2017, beginning on January 31st and ending on February 20th. Only fares in economy class involving a connection with total layover time less than four hours were collected. This restriction is imposed because it is assumed that most passengers traveling within the continental U.S. do not purchase itineraries involving layovers that are more than four hours in duration.

Fare quotes were obtained daily, for one-way travel between the origin and destination airports listed in Table 2.1.[21] For each route, fares for each of the next thirty-five travel days were collected.[22] For example, on January 31st, fare quotes were obtained for flights departing between February 1st and March 7th. On February 20th, fare quotes were obtained for flights departing between February 21st and March 27th. Focusing on a thirty-five day window should capture business travelers who purchase flights close to the departure date in addition to leisure travelers who purchase flights well in advance of the departure date.

This sampling scheme resulted in a sample of 575,267 observations. Thus, an average of 15.65 fare quotes per query date for each route-departure date combination were retrieved.[23] Figure

[20] Fare quotes for Southwest Airlines are not available on Google Flights (or other travel aggregator websites, such as Kayak and Expedia). However, since Southwest operates primarily a point-to-point network, the impact of omitting Southwest from the analysis should be minimal.

[21] This chapter relies on one-way fares due to the complications involved in gathering roundtrips. The main complication involves specifying trip duration. For any given departure date, there are a substantial number of roundtrip fares that could potentially be gathered, each depending on trip duration. For instance, fares for four-day roundtrips are likely much different from seven-day or ten-day trips. Similar papers using published fare and itinerary data also focus on one-way trips due to this issue. Examples include Bilotkach (2005) and Bilotkach et al. (2010).

[22] Analyzing 90-day fare histories for a variety of routes, Gillen and Mantin (2009) and Mantin and Koo (2009) find average prices are relatively stable during most of the 90-day history and only begin to gradually increase about three-weeks prior to departure. Therefore, extending the five-week horizon for data collection would not appear to add substantial value to the analysis.

[23] 575,267 / (50 routes*35 departure dates per query date*21 query dates) $\approx$ 15.65. This average is similar to the 14.31 average Bilotkach et al. (2010) obtained from querying expedia.com for the New York-London route.

2.1 below summarizes average fares by layover type across the fifty routes. An itinerary with a tight layover involves a connection with 50 or less minutes of layover time. Itineraries with 51-80 minute layovers are considered short while itineraries with 81-120 minutes of layover time are considered long. Finally, itineraries with more than two hours of layover time are considered very long.

Figure 2.1: Average Fares by Layover Type for All Routes

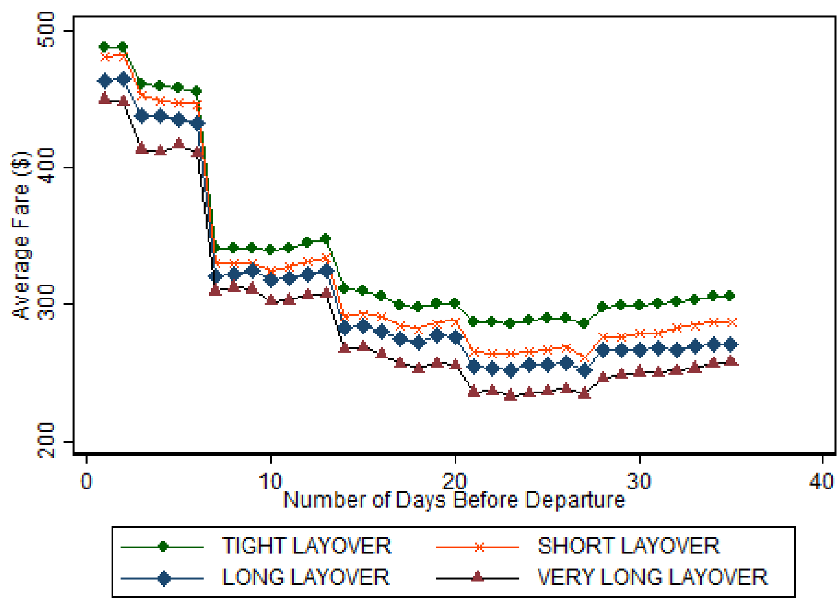

Notes: The data used to produce this graph were collected from Google Flights between January 31st, 2017 and February 20th, 2017. On each day of the sampling period, one-way fares for each of the next thirty-five travel days were collected for the routes listed in Table 2.1. A tight layover involves a connection with 50 or fewer minutes of layover time. Itineraries with 51-80 minute layovers are short while itineraries with 81-120 minute layovers are long. Finally, itineraries with more than two hours of layover time are very long.

Two distinct patterns emerge from Figure 2.1. Foremost, there is clear evidence of advance-purchase discounts as illustrated by the downward sloping trend in average fares for itineraries of all layover types. Fares are highest when "purchased" one-day prior to departure and decline as the number of days before departure increases.[24] Moreover, there are steep declines

[24] Stavins (2001) finds that increasing the advance-purchase requirement by one day results in a \$6.04 decrease in the average ticket price. A ticket with a fourteen day advance-purchase requirement would then cost \$84.56 less than a similar ticket on the same route without the same requirement.

in average fares from six to seven days, thirteen to fourteen days, and twenty to twenty-one days before departure. In other words, airlines sharply increase fares at specific three-week, two-week, and one-week milestones prior to departure.[25] However, this chapter is not the first to observe a pattern of large fare increases within three weeks of departure. Analyzing 90-day fare histories for a variety of routes, Gillen and Mantin (2009) and Mantin and Koo (2009) also find that fares begin to substantially increase two-three weeks prior to departure.

With regard to the present chapter, the most striking takeaway from Figure 2.1 is the decline in average fares by layover type. Regardless of the number of days before departure, average fares are higher for itineraries with tight layovers relative to short layovers. Similarly, average fares are higher for itineraries with short layovers relative to long or very long layovers. Therefore, as layover time increases, average fares decrease. This finding supports the conjecture that passengers are compensated with a discounted fare for purchasing a ticket that requires a layover. The magnitude of the discount depends on total layover time, with the discount increasing as layover time increases.

Figure 2.2 is analogous to Figure 2.1, except that average fares by layover type are computed for the PDX-MCO route instead of across all fifty routes. Similar to Figure 2.1, the downward sloping trend in average fares for itineraries of all layover types is evidence of advance-purchase discounts. Most importantly however, average fares for itineraries with tight layovers are consistently higher than average fares for itineraries with short, long, or very long layovers. This finding reinforces the notion of an inverse relationship between fare and layover time. As illustrated in Figures 2.1 and 2.2, average fares decrease when layover time increases.

[25] Increasing fares at specific milestones before departure are one of many mechanisms airlines employ in their respective yield management strategies. For more on yield management strategies in the airline industry, see Belobaba et al. (2015) and Talluri and Van Ryzin (2006).

Figure 2.2: Average Fares by Layover Type for PDX-MCO Route

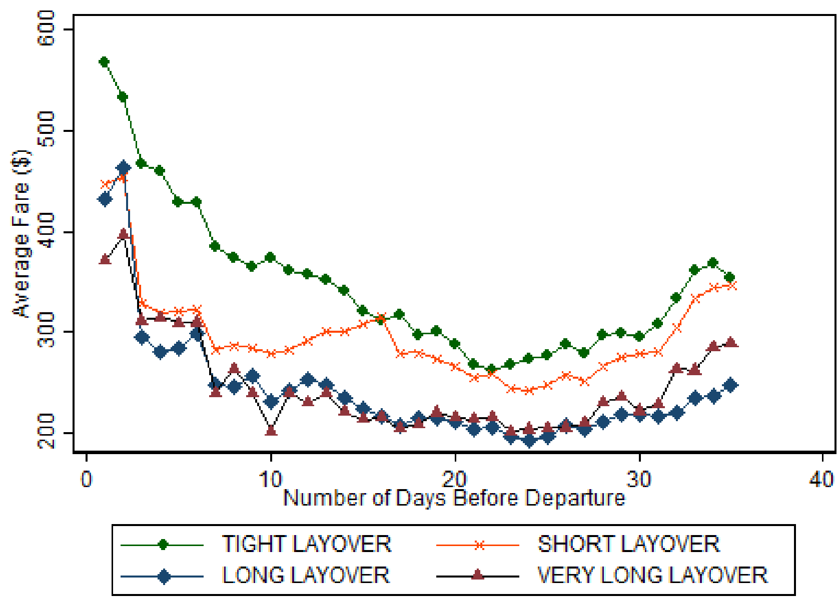

Notes: The data used to produce this graph were collected from Google Flights between January 31st, 2017 and February 20th, 2017. On each day of the sampling period, one-way fares for each of the next thirty-five travel days were collected. A tight layover involves a connection with 50 or fewer minutes of layover time. Itineraries with 51-80 minute layovers are short while itineraries with 81-120 minute layovers are long. Finally, itineraries with more than two hours of layover time are very long.

2.4 Empirical Strategy

A hedonic regression is used to estimate the value of layover time in the U.S. airline industry. This approach, regresses fare on a vector of itinerary characteristics, including the total amount of layover time. Because theory provides little guidance on the correct functional form for this type of analysis, three different forms of the empirical model are estimated. Equation (2.1) below is estimated in levels,

$$Fare_{ijkts} = \beta_0 + \beta_1 \cdot (\text{AIR TRAVEL TIME})_{ijkts} + \beta_2 \cdot (\text{AIR TRAVEL TIME})^2_{ijkts} + \beta_3 \cdot (\text{LAYOVER MINUTES})_{ijkts} + \beta_4 \cdot (\text{LAYOVER MINUTES})^2_{ijkts} + \alpha \cdot X_{ijkts} + \gamma_{ij} + \theta_k + \epsilon_{ijkts} \qquad (2.1),$$

where $Fare_{ijkts}$ is the published fare on date s for one-way travel departing on date t on carrier k between origin airport i and destination airport j. X is a matrix of controls including total itinerary distance in miles, the number of days before departure the fare is observed, and an indicator for travel on a low-cost carrier (LCC). The X matrix also contains fixed-effects controlling for the time-of-day-of-departure[26], the day-of-week-of-departure[27], the day-of-week the fare is observed ("purchased")[28], and the airport where the layover occurs.[29] γ is a city-pair fixed-effect, θ is a carrier-specific fixed-effect, and ϵ is an error term.[30] AIR TRAVEL TIME reflects the total amount of time spent flying, defined as total travel time minus total layover time.[31] A quadratic in air travel time and layover time is included to reflect that the relationship between fare and travel time is not well approximated linearly.[32] The main coefficients of interest are β_3 and β_4 as these coefficients measure the impact on fare for each additional minute of layover time. Of particular interest is the effect on fare of an itinerary requiring a one-hour layover.[33]

[26] To control for the time-of-day-of-departure, departure time of each itinerary's first leg is split into the following five periods: 12:00am-4:59am, 5:00am-8:59am, 9:00am-1:59pm, 2:00pm-7:59pm, and 8:00pm-11:59pm. These fixed-effects control for the possibility that the amount of layover time on an itinerary depends on the time-of-day-of-departure. For instance, itineraries with first leg departures in the morning may have shorter layovers on average relative to itineraries with first leg departures in the evening because flight frequency is higher during daytime hours.

[27] Day-of-week-of-departure fixed effects control for the possibility that passengers traveling on weekends incur longer layovers due to decreased flight frequency.

[28] Puller and Taylor (2012) find that fares purchased on weekends are 5% lower than those purchased on weekdays, supporting the conjecture that airlines price discriminate when the mix of purchasing passengers makes demand more price-elastic.

[29] Omitting a variable controlling for the quality of the layover airport could bias an estimate of the layover time effect if the amount of layover time on an itinerary is correlated with the quality of the layover airport. For example, passengers are more inclined to accept longer layovers at airports with lots of amenities (shopping, airport lounges, restaurants), relative to layovers at airports where the menu of amenities is scarce. Airlines may set fares taking these preferences into account, with fares being cheaper when the itinerary incorporates a layover in a low quality airport (few amenities). Airlines may also adjust their departure schedules so that passengers incurring layovers in low quality airports have short layovers.

[30] Referencing Table 2.1, there are forty-seven city-pairs in the analysis.

[31] AIR TRAVEL TIME is measured in minutes. The magnitude of the layover time effects presented in Section 2.5 do not substantially change if total travel time (inclusive of layover time) enters the model instead of AIR TRAVEL TIME.

[32] As mentioned in Section 2.2, Theis et al. (2006) finds evidence of an n-shaped utility function with regard to layover time, supporting the hypothesis that passengers are risk averse to very short connections (where the possibility of missing the connecting flight is high) while also being averse to connections requiring long layovers.

[33] As discussed in Section 3.4, the data do not include routes with nonstop service due to concerns regarding self-selection bias. Therefore, connecting itineraries are used to estimate what the first hour of layover time

Equation (2.2) is identical to equation (2.1), except that fare is measured in natural logarithm form to further capture non-linearity in the relationship between fare and travel time.[34] The benefit of measuring fare in natural log form is that the impact of layover time can be translated into a percentage change in fare.[35] The natural log transformation also minimizes the impact of potential outliers in fare.

To provide an elasticity interpretation of the results, equation (2.3) below is estimated in logs,

$$ln(Fare)_{ijkts} = \beta_0 + \beta_1 \cdot ln(\text{AIR TRAVEL TIME})_{ijkts} + \beta_2 \cdot ln(\text{AIR TRAVEL TIME})^2_{ijkts} + \beta_3 \cdot ln(\text{LAYOVER MINUTES})_{ijkts} + \beta_4 \cdot ln(\text{LAYOVER MINUTES})^2_{ijkts} + \alpha \cdot X_{ijkts} + \gamma_{ij} + \theta_k + \epsilon_{ijkts} \qquad (2.3),$$

where fare and the travel time covariates enter in natural logarithm form. Results from equation (2.3) are discussed in Appendix D since the coefficients from this equation cannot be used to directly compute the dollar value of a one-hour layover.[36]

All three equations are estimated using ordinary least squares (OLS) with standard errors that are clustered at the city-pair level.[37] To explore heterogeneity in the valuations of layover time, additional specifications present values that are airline-specific and hub-specific in addition to values that differ with the date of purchase (number of days before departure).

is worth.

[34] To provide further evidence of the nonlinear relationship between fare and layover time, Appendix H provides spline regression results that allow the effect of layover time on fare to change throughout the range of layover time values. The spline regressions produce results very similar to those from equations (2.1)-(2.2). Because the models specified by equations (2.1)-(2.2) are more parsimonious and easier to interpret, results from the spline regressions are not presented in the main text.

[35] Since the dependent variable fare is in natural log form and layover minutes are in levels, the marginal effect is interpreted as the $100(e^{\beta_1} - 1)\%$ change in fare.

[36] The coefficients cannot be used to directly compute the dollar value of a one-hour layover because the change from 0 to 60 minutes of layover time cannot be expressed as a percentage.

[37] Fares in different directions within a city-pair market are highly correlated. Therefore, standard errors are clustered at this level to account for this correlation.

2.4.1 Potential Endogeneity of Layover Time

There is a concern that layover times are endogenous if airlines aim to reduce layovers in city-pair markets with a large number of business travelers who are time-sensitive and less price elastic relative to leisure passengers. If layover times are endogenous, then OLS estimates may overstate the true effect of layover time on fares.[38] However, in connecting city-pair markets, it is plausible that fare and layover time are not jointly determined. The argument is similar to one used by Brueckner and Spiller (1994) in measuring the strength of economies of traffic density. After limiting their analysis to connecting markets, Brueckner and Spiller (1994) treated spoke traffic levels as exogenous with respect to fares, arguing that since each connecting city-pair market is relatively small, its passengers constitute a negligible share of total traffic on the spokes along which they travel.

A similar argument is applied here, recognizing that layover time in connecting city-pairs results from a complex hub-and-spoke optimization problem.[39] Within a connecting city-pair, passengers continuing to the final destination constitute a small fraction of total passengers on the flight into the hub. At the hub, these connecting passengers are joined by other passengers who are connecting from other spoke cities or originating at the hub. Since these connecting passengers constitute a small fraction of total traffic on both legs of the trip, layover times are unlikely to be determined by the time sensitivity of passengers in any one (connecting) city-pair market. To further support this argument, Appendix E presents suggestive evidence indicating that layover times result from the hub-and-spoke optimization problem in a manner not consistent with airlines reducing layover times in city-pairs with a large share of time-sensitive business travelers.

[38] To deal with endogeneity, one could estimate equations (2.1)-(2.3) using two-stage least squares (2SLS). However, the use of city-pair, carrier, layover airport, time-of-day-of-departure, day-of-week-of-departure, and day-of-week-of-purchase fixed effects severely limits the pool of candidate instrumental variables.

[39] Alaska, American, Delta, and United account for 96.56% of observations in the analysis sample.

2.5 Results

This section presents results from estimating the models described in the Empirical Strategy Section. First, results from equations (2.1) and (2.2) are presented to establish an average value of layover time. Next, values that are specific to legacy and low-cost carriers are presented.[40] To further explore heterogeneity in the valuations of layover time, separate values are estimated for the three dominant legacy carriers, American, Delta, and United.[41] Further decomposing the dominant legacy carrier results, values specific to each of American, Delta, and United's four largest hub airports are presented. These results are followed by specifications which allow the layover time effect to differ with the number of days before departure that the fare is observed ("purchased").

2.5.1 Average Value of Layover Time

Column one of Table 2.2 provides regression results from the model specified by equation (2.1). All coefficients have their expected signs. Fares are substantially cheaper if travel is on a low-cost carrier while the positive coefficient on TOTAL DISTANCE indicates that fares increase with distance traveled. The coefficient on DAYS BEFORE DEPARTURE indicates that increasing the advance-purchase requirement by one day results in a $5.21 decrease in the average fare. This estimated advance-purchase effect is similar to the $6.04 estimate found in Stavins (2001).

Although only the coefficient on AIR TRAVEL TIME is significant at the 10% level, the negative coefficient on AIR TRAVEL TIME and the small positive coefficient on (AIR TRAVEL TIME)^2 indicate that one-hour of air travel time is valued at $33.92. This estimate is nearly

[40] Low-cost carriers in the data are Frontier, JetBlue, Spirit, Sun Country, and Virgin America. However, only 3.44% of the observations involve travel on a low-cost carrier.

[41] American, Delta, and United operate the largest hub-and-spoke networks and account for 94.44% of observations.

identical to the \$34 value found in Morrison and Winston (1989).[42] In contrast to the coefficients on air travel time, the coefficients on LAYOVER MINUTES and (LAYOVER MINUTES)^2 are both significant at conventional levels. Consistent with the Morrison and Winston (1989) results, this chapter finds an even higher value of layover time, with passengers compensated in the amount of \$47.60 for incurring a one-hour layover.[43] This estimate compares favorably to the hourly figure recommended by the Federal Aviation Administration's (FAA) Office of Aviation Policy and Plans. As of September 2016, the FAA recommends that passengers' time for all travel purposes be valued at \$44.30 per hour (FAA, 2016).

Column two of Table 2.2 presents regression results from the model specified by equation (2.2). The air travel and layover time effects are similar in magnitude to the column one results. However, this model produces a higher R-squared value than the linear approach, suggesting that the absolute effect of travel time on fares depends on fare levels. The coefficients on the travel time covariates indicate that passengers are compensated with a fare that is \$31 (9.75%) and \$42.74 (13.44%) cheaper per hour of air travel and layover time, respectively. To alleviate the concern that high levels of demand on routes originating from larger metro areas are driving results, Appendix F presents estimates when routes originating from the twenty largest Metropolitan Statistical Areas are removed from the analysis. The air travel and layover time estimates presented in Appendix Table F.1 are very similar to the Table 2.2 values, implying that routes originating from large metro areas are not substantially influencing the results.

[42] Using data from the third quarter of 2007, Yan and Winston (2014) find an average value of air travel time equal to \$24 per hour.

[43] As discussed in Section 2.4.1, OLS may overestimate the true effect of layover time on fares if layover times are endogenous. However, the average value found in this chapter is lower than the \$74 average value found in the Morrison and Winston (1989) study.

Table 2.2: Average Value of Layover Time Regression Results

Dependent Variable:	Fare	ln(Fare)
	(1)	(2)
AIR TRAVEL TIME	-0.59*	-0.002**
	(0.30)	(0.00)
(AIR TRAVEL TIME)^2 (100s)	0.05	0.000
	(0.04)	(0.00)
LAYOVER MINUTES	-0.95***	-0.003***
	(0.08)	(0.00)
(LAYOVER MINUTES)^2 (100s)	0.26***	0.001***
	(0.03)	(0.00)
TOTAL DISTANCE (1,000s)	33.14*	0.098*
	(19.04)	(0.05)
DAYS BEFORE DEPARTURE	-5.21***	-0.016***
	(0.34)	(0.00)
LCC	-149.03***	-0.371***
	(42.43)	(0.10)
Constant	732.36***	6.871***
	(62.66)	(0.16)
Observations	575,267	575,267
R-squared	0.39	0.45
City-Pair FE	YES	YES
Carrier FE	YES	YES
Layover Airport FE	YES	YES
Time-of-Departure FE	YES	YES
Day-of-Week-of-Departure FE	YES	YES
Day-of-Week-of-Purchase FE	YES	YES
1-Hour Air Travel Effect (%)	NA	-9.75%
1-Hour Air Travel Effect ($)	-$33.92	-$31.00
1-Hour Layover Effect (%)	NA	-13.44%
1-Hour Layover Effect ($)	-$47.60	-$42.74
2-Hour Layover Effect (%)	NA	-21.80%
2-Hour Layover Effect ($)	-$76.71	-$69.32

Notes: Standard errors clustered at the city-pair level are reported in parentheses. The average fare in the regression sample is $317.99.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

2.5.2 Values of Layover Time by Carrier Type

Table 2.3 presents regression results when separate layover time effects are estimated for legacy and low-cost carriers. In the preferred specification provided in column two, passengers on legacy carriers are found to be compensated in the amount of \$42.58 for incurring a one-hour layover compared to \$46.23 for low-cost carriers. These values are very close to the average value of \$42.74 presented in column two of Table 2.2. While the estimated effect for low-cost carriers is slightly larger than the effect for legacy carriers, the low-cost effect is imprecisely estimated. This imprecise estimate is not surprising considering that 96.56% of all observations collected for the empirical analysis involve travel on a legacy carrier. Since low-cost carriers primarily operate point-to-point networks, almost all travel will occur on legacy carriers in analyses limited to connecting city-pair markets.

Table 2.3: Legacy vs. Low-Cost Regression Results

Dependent Variable:	Fare	ln(Fare)
	(1)	(2)
AIR TRAVEL TIME	-0.60*	-0.002**
	(0.30)	(0.00)
(AIR TRAVEL TIME)^2 (100s)	0.05	0.000
	(0.04)	(0.00)
(LAYOVER MINUTES)*LEGACY	-0.96***	-0.003***
	(0.08)	(0.00)
(LAYOVER MINUTES)^2*LEGACY (100s)	0.26***	0.001***
	(0.03)	(0.00)
(LAYOVER MINUTES)*LCC	-0.81	-0.003*
	(0.49)	(0.00)
(LAYOVER MINUTES)^2*LCC (100s)	0.19	0.001
	(0.17)	(0.00)
TOTAL DISTANCE (1,000s)	33.24*	0.098*
	(19.09)	(0.05)
DAYS BEFORE DEPARTURE	-5.21***	-0.016***
	(0.34)	(0.00)
LCC	-153.64***	-0.339**
	(53.18)	(0.14)
Constant	733.05***	6.871***
	(62.68)	(0.16)
Observations	575,267	575,267
R-squared	0.39	0.45
City-Pair FE	YES	YES
Carrier FE	YES	YES
Layover Airport FE	YES	YES
Time-of-Departure FE	YES	YES
Day-of-Week-of-Departure FE	YES	YES
Day-of-Week-of-Purchase FE	YES	YES
1-Hour Legacy Layover Effect (%)	NA	-13.34%
1-Hour Legacy Layover Effect ($)	-$48.01	-$42.58
1-Hour LCC Layover Effect (%)	NA	-16.27%
1-Hour LCC Layover Effect ($)	-$41.80	-$46.23

Notes: Standard errors clustered at the city-pair level are reported in parentheses. The average legacy fare in the regression sample is $319.20 while the average low-cost fare is $284.16. The legacy carriers are Alaska, American, Delta, and United. The low-cost carriers are Frontier, JetBlue, Spirit, Sun Country, and Virgin America. Legacy carriers account for 96.56% of observations in the analysis sample compared to 3.44% for low-cost carriers.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Given the large disparity in observations by carrier type, Table 2.4 presents regression results when the analysis sample is limited to observations involving travel on American, Delta, or United and separate effects are estimated for each carrier.[44] The coefficients on the layover time variables are all significant at conventional levels. The most striking result from Table 2.4 are the large differences in the valuations of layover time among the three dominant legacy carriers. In the preferred specification presented in column two, the value of a one-hour layover ranges from $38 on American to almost $62 on United. These results indicate that United passengers are compensated at a substantially higher rate for incurring a layover relative to American and Delta passengers. This result may reflect general consumer distaste for United connecting service (or some other underlying difference) that is not being captured by the carrier-specific or layover airport fixed effects. For instance, a March 2017 survey conducted by the American Customer Satisfaction Index (ACSI) found that United had the lowest customer satisfaction level of any legacy carrier.[45] If passengers strictly prefer travel with American or Delta over United, layover discounts must be larger on United to entice passengers to purchase a United connecting ticket when comparable itineraries are also offered by American or Delta.

[44] Only 2.12% of all observations were for travel on Alaska. American, Delta, and United account for 94.44% of all observations.

[45] See a Chicago Tribune article from April 25th, 2017 titled *"United scores lowest among legacy airlines in customer satisfaction; JetBlue tops survey."* This article may be accessed at `http://www.chicagotribune.com/business/ct-airline-passenger-satisfaction-0426-biz-20170425-story.html`.

Table 2.4: Dominant Legacy Carrier Regression Results

Dependent Variable:	Fare	ln(Fare)
	(1)	(2)
AIR TRAVEL TIME	-0.33	-0.001*
	(0.23)	(0.00)
(AIR TRAVEL TIME)^2 (100s)	0.00	0.000
	(0.00)	(0.00)
(LAYOVER MINUTES)*AMERICAN	-0.86***	-0.002***
	(0.15)	(0.00)
(LAYOVER MINUTES)^2*AMERICAN (100s)	0.23***	0.001***
	(0.06)	(0.00)
(LAYOVER MINUTES)*DELTA	-0.87***	-0.003***
	(0.13)	(0.00)
(LAYOVER MINUTES)^2*DELTA (100s)	0.24***	0.001***
	(0.05)	(0.00)
(LAYOVER MINUTES)*UNITED	-1.47***	-0.004***
	(0.14)	(0.00)
(LAYOVER MINUTES)^2*UNITED (100s)	0.45***	0.001***
	(0.05)	(0.00)
TOTAL DISTANCE (1,000s)	17.28	0.056
	(16.44)	(0.05)
DAYS BEFORE DEPARTURE	-5.18***	-0.016***
	(0.36)	(0.00)
Constant	565.35***	6.472***
	(40.70)	(0.11)
Observations	543,308	543,308
R-squared	0.39	0.46
City-Pair FE	YES	YES
Carrier FE	YES	YES
Layover Airport FE	YES	YES
Time-of-Departure FE	YES	YES
Day-of-Week-of-Departure FE	YES	YES
Day-of-Week-of-Purchase FE	YES	YES
1-Hour American (AA) Layover Effect (%)	NA	-12.19%
1-Hour American (AA) Layover Effect ($)	-$43.56	-$38.41
1-Hour Delta (DL) Layover Effect (%)	NA	-13.09%
1-Hour Delta (DL) Layover Effect ($)	-$43.52	-$40.72
1-Hour United (UA) Layover Effect (%)	NA	-18.72%
1-Hour United (UA) Layover Effect ($)	-$72.02	-$61.89

Notes: Standard errors clustered at the city-pair level are reported in parentheses. Average fares for travel on American, Delta, and United in the regression sample are $315.07, $311.09, and $330.60, respectively.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

2.5.3 Values of Layover Time by Hub Airport

An alternative explanation for the difference in layover time values among American, Delta, and United involves the quality of the connecting airports themselves. If passengers perceive that American and Delta's hub airports are of better quality than United's, then layover discounts on United will tend be larger. To investigate this possibility, Table 2.5 presents layover time values specific to each of American, Delta, and United's four largest hub airports.[46] Coefficient estimates used to compute the Table 2.5 values are provided in Appendix Table G.1.

Although there exists considerable heterogeneity in the value of a one-hour layover at each of the hub airports listed in Table 2.5, the pattern of values are consistent with the Table 2.4 results. On average, the value of a one-hour layover at a United hub is substantially larger than the value of a comparable layover at an American or Delta hub. In particular, passengers are compensated at high rates for incurring layovers in Chicago (\$52.02), Denver (\$68.88), and Houston (\$67.55), three of United's major hubs.

[46] For American, the four largest domestic hubs are Dallas/Fort Worth (DFW), Charlotte (CLT), Chicago O'Hare (ORD), and Philadelphia (PHL). The four largest hubs for Delta are Atlanta (ATL), Detroit (DTW), Minneapolis (MSP), and Salt Lake City (SLC). Finally, the four largest domestic hubs for United are Chicago O'Hare (ORD), Houston (IAH), Newark (EWR), and Denver (DEN).

Table 2.5: Value of a 1-Hour Layover by Connecting Hub Airport

Airline (1)	Connecting Hub Airport (2)	1-Hour Layover (3)
American (AA)/United (UA)	Chicago O'Hare (ORD)	-\$52.02 (-15.5%)
American (AA)	Dallas/Fort Worth (DFW)	-\$38.77 (-12.14%)
American (AA)	Charlotte (CLT)	-\$37.98 (-12.68%)
American (AA)	Philadelphia (PHL)	-\$46.78 (-16.21%)
Delta (DL)	Atlanta (ATL)	-\$42.76 (-13.41%)
Delta (DL)	Detroit (DTW)	-\$24.56 (-8.12%)
Delta (DL)	Minneapolis (MSP)	-\$60.77 (-19.48%)
Delta (DL)	Salt Lake City (SLC)	-\$28.98 (-9.33%)
United (UA)	Houston (IAH)	-\$67.55 (-21.42%)
United (UA)	Denver (DEN)	-\$68.88 (-19.16%)
United (UA)	Newark (EWR)	-\$39.46 (-13.82%)
	Other Airports	-\$21.43 (-6.86%)

Notes: The dollar amounts and percentages presented in this table are computed using the coefficient estimates provided in column two of Table G.1. Average fares for each hub airport in the regression sample are: ORD=\$335.62, DFW=\$319.32, CLT=\$299.53, PHL=\$288.59, ATL=\$318.90, DTW=\$302.51, MSP=\$311.94, SLC=\$310.56, IAH=\$315.37, DEN=\$359.50, EWR=\$285.50, Other=\$312.38.

2.5.4 Layover Time and the Number of Days Before Departure

To further explore heterogeneity in the valuations of layover time, Table 2.6 investigates whether these values differ with the number of days before departure. Assuming business travelers purchase closer to the departure date than leisure passengers, this table sheds light on the compensation these two different groups receive for itineraries with comparable amounts of layover time. Since business travelers typically value their time at much higher rates, it is possible that layover discounts are larger when airfare is purchased close to the departure date. However, business travelers are also more price-inelastic relative to leisure passengers. Therefore, conditional on having a layover, business travelers are expected to pay more for airfare per layover minute, suggesting that the discount received for incurring a layover may be lower for business travel.[47]

To determine if the value of layover time differs with the date of purchase (and the direction of the effect), two interaction terms are added to the specifications presented in Table 2.2. The coefficients on (LAYOVER MINUTES)*DAYS BEFORE DEPARTURE and (LAYOVER MINUTES)^2*DAYS BEFORE DEPARTURE in Table 2.6 are significant at conventional levels, indicating that the value of layover time does depend on the number of days before departure. Using the preferred specification in column two of Table 2.6, Table 2.7 presents layover time estimates for passengers purchasing airfare 1, 7, 14, 21, 28, or 35 days before departure. The Table 2.7 values confirm that passengers purchasing airfare closer to the departure date are compensated at lower rates for the layovers they incur. The value of a one-hour layover ranges from $25.50 (8.02%) one-day before departure to over $59 (18.74%)

[47] Consider a business passenger and leisure passenger traveling on the same connecting itinerary. Because they are on the same pair of flights (first flight into the hub, second flight to final destination), both passengers incur the same amount of layover time. However, if the business passenger purchases airfare closer to the departure date, the business passenger will pay more for the same connecting itinerary than the leisure passenger. Since both are incurring the same amount of layover time, the business passenger pays more for airfare per layover minute. Therefore, finding that layover discounts are lower closer to the date of departure may reflect "dominance" of the advance-purchase fare effect in the (LAYOVER MINUTES*DAYS BEFORE DEPARTURE) interaction term.

five weeks before departure.

Table 2.6: Days Before Departure Regression Results

Dependent Variable:	Fare	ln(Fare)
	(1)	(2)
AIR TRAVEL TIME	-0.59*	-0.002**
	(0.30)	(0.00)
(AIR TRAVEL TIME)^2 (100s)	0.05	0.000
	(0.04)	(0.00)
LAYOVER MINUTES	-0.56***	-0.001***
	(0.18)	(0.00)
(LAYOVER MINUTES)*DAYS BEFORE DEPARTURE	-0.02**	-0.000***
	(0.01)	(0.00)
(LAYOVER MINUTES)^2 (100s)	0.07	0.000
	(0.07)	(0.00)
(LAYOVER MINUTES)^2*DAYS BEFORE DEPARTURE	0.01***	0.000***
	(0.00)	(0.00)
DAYS BEFORE DEPARTURE	-4.37***	-0.013***
	(0.44)	(0.00)
TOTAL DISTANCE (1,000s)	33.42*	0.099*
	(19.07)	(0.05)
LCC	-150.01***	-0.372***
	(42.39)	(0.10)
Constant	715.63***	6.810***
	(62.94)	(0.16)
Observations	575,267	575,267
R-squared	0.39	0.45
City-Pair FE	YES	YES
Carrier FE	YES	YES
Layover Airport FE	YES	YES
Time-of-Departure FE	YES	YES
Day-of-Week-of-Departure FE	YES	YES
Day-of-Week-of-Purchase FE	YES	YES

Notes: Standard errors clustered at the city-pair level are reported in parentheses.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Table 2.7: Value of Layover Time by Purchase Date

Days Before Departure (1)	1-Hour Layover (2)	2-Hour Layover (3)
1	-\$25.50 (-8.02%)	-\$46.49 (-14.62%)
7	-\$31.55 (-9.92%)	-\$54.54 (-17.15%)
14	-\$38.54 (-12.12%)	-\$63.92 (-20.10%)
21	-\$45.57 (-14.33%)	-\$73.33 (-23.06%)
28	-\$52.60 (-16.54%)	-\$82.71 (-26.01%)
35	-\$59.59 (-18.74%)	-\$92.09 (-28.96%)

Notes: The dollar amounts and percentages presented in this table are computed using the coefficient estimates provided in column two of Table 2.6. The average fare in the regression sample is \$317.99.

2.6 Hypothetical Impact of Implementing Tighter Connections at Chicago's O'Hare Airport

The estimates presented in the previous section provide a foundation for evaluating the trade-off between convenient flight connections and airport congestion. For example, the total fare gain from an airline reducing passenger layover time may be directly computed from the estimates in Tables 2.2, 2.3, 2.4, or 2.6. This gain could then be compared to the loss resulting from increased airport congestion. However, making this comparison requires estimates of congestion costs. These costs are expected to vary by time of day. Periods with high flight volume will have higher congestion costs relative to periods with low flight volume.

As a major hub for American and United, Chicago's O'Hare Airport (ORD) is one of the most congested airports in the United States. Currently, 60% of arriving ORD passengers board connecting flights.[48] Given the immense volume of connecting traffic, ORD provides an excellent test case for evaluating the trade-off between reducing passenger layover time and airport congestion. Estimates of congestion costs at ORD are also directly available from Johnson and Savage (2006).[49]

To evaluate this trade-off, the total fare gain from reducing passenger layover time will be compared to the extra congestion cost the airline incurs when the connecting flight is moved to an earlier (more congested) part of the day. Since congestion costs vary with the volume of airport traffic, the extra cost the airline incurs will depend on the difference in the volume of flights departing between the new and previously scheduled departure times. This extra congestion cost will be positive if the length of the departure queue at the new time is longer than the length at the previously scheduled departure time.

In addition to the difference in queue lengths, the extra congestion cost also depends on whether the airline takes into account how the schedule change impacts other flights in the departure queue. If an airline ignores the effects of their scheduling decisions on other departing flights, then the airline is "atomistic." However, carriers operating a large fraction of the airport's flights will not ignore how their scheduling impacts other departing flights. For instance, a hub carrier moving a flight to a congested period will not only delay the flights of other carriers, it will also delay its own planes in the departure queue. Therefore, when determining flight schedules at ORD, American and United should "internalize" the

[48] In 1993, 56% of passengers at O'Hare were making connecting flights (Morrison and Winston, 1995). This number has increased to about 60% today (see `http://www.airport-technology.com/projects/chicago/`).

[49] Average congestion costs at 27 major U.S. airports, including O'Hare, are provided in Daniel and Harback (2008) and Daniel and Harback (2009). The analysis in this chapter relies on the estimates in Johnson and Savage (2006) because rather than reporting an average congestion cost, the authors present congestion costs that differ depending on the length of the departure queue or by time of day.

congestion they impose on their other departing flights.[50] This internalizing behavior results in congestion costs that are higher than analogous costs for atomistic carriers.

For atomistic carriers, Johnson and Savage (2006) estimate congestion costs using the following equation,

$$ATOMISTIC\ COST\ (AC) = \$92.97 * e^{(2.498+0.00051D^2)} \qquad (2.4),$$

where D is the length of the departure queue. Applying this equation, the atomistic congestion cost of entering a queue with 5 or 10 flights is estimated to be $1,145 and $1,189, respectively.[51]

For United, congestion costs are estimated using the following equation,

$$UNITED\ COST\ (UC) = AC + \frac{\partial AC}{\partial D} * UA's\ flight\ share * D \qquad (2.5),$$

where AC is the atomistic cost from equation (2.4), *"UA's flight share"* is the fraction of United flights in the departure queue, and D is the length of the departure queue.[52][53] Currently, United accounts for about 40% of the traffic at ORD.[54] Thus, after fixing *UA's flight share* in equation (2.5) at 40%, United's congestion cost of entering a departure queue with 10 flights would be $1,238, a $49 increase over the atomistic cost.[55]

Equation (2.4) may be used to calculate the increase in cost that an atomistic carrier would incur if the airline altered its flight schedule such that the new departure occurs when the queue is at a longer length.[56] For example, consider an atomistic carrier entering a departure queue that is 10 flights long. Now suppose that to reduce passenger layover time, the carrier

[50] Brueckner (2002a), Brueckner (2002b), and Brueckner (2005) show that in oligopolistic markets, airlines will only internalize the congestion they impose on themselves.

[51] $\$92.97 * e^{(2.498+0.00051*5^2)} = \$1,144.85$. $\$92.97 * e^{(2.498+0.00051*10^2)} = \$1,189.49$.

[52] Similarly, American's congestion cost would be $AC + \frac{\partial AC}{\partial D}$*$AA$*'s flight share*$D$.

[53] After differentiating, equation (2.5) becomes:
$UC = \$92.97 * e^{(2.498+0.00051D^2)} + 1.15295 * e^{(0.00051D^2)} * D * UA's\ flight\ share * D$.

[54] In Johnson and Savage (2006), United's market share was 40.5%.

[55] $\$92.97 * e^{(2.498+0.00051*10^2)} + 1.15295 * e^{(0.00051*10^2)} * 10^2 * .40 = \$1,238.02$.

[56] Equivalently, equation (2.5) may be used to calculate the increase in cost for United in a similar scenario.

moves up the departure by twenty minutes. However, instead of entering a queue with 10 flights, the flight now enters a departure queue with 20 flights. Referencing equation (2.4), the carrier would incur a congestion cost of $1,386, a $197 increase from the $1,189 cost it would previously have incurred. Therefore, for the atomistic carrier to permanently move the departure, the gain from reducing passenger layover time by twenty minutes must be at least as large as the $197 increase in cost.

To compare the gain from tighter connections to the loss resulting from increased airport congestion, assumptions must first be made regarding the share of connecting passengers and their layover times. Focusing on ORD, 12.71% of the 575,267 itineraries in the Table 2.2 regression sample had a layover at ORD. Of these connecting itineraries, the most common city-pair market in the analysis is the Los Angeles (LAX)-New York (LGA) route with 9.6% of the ORD connections.

Table 2.8 provides the mean and standard deviation of LAYOVER MINUTES for the three carriers observed flying the LAX-LGA route in the regression sample. For a passenger traveling with American or United, the average layover is between 72 and 75 minutes. However, if a passenger flies with Spirit, the average layover is over two hours. Considering that LAX-LGA passengers incur layovers that are over one hour on average regardless of carrier, there are large gains that could be realized if airlines are able to implement tighter connections for these passengers.

Table 2.8: Average LAYOVER MINUTES for ORD connecting itineraries in the LAX-LGA city-pair market

Airline	Observations	Mean
American	2,121	72.79
		(39.10)
Spirit	499	129.59
		(2.45)
United	4,402	75.21
		(41.05)
Total	7,022	78.34
		(41.48)

Notes: Standard deviation reported in parentheses.

Table 2.9 summarizes the convenient flight connection and airport congestion trade-off in a hypothetical scenario where Spirit (treated as an atomistic carrier), decreases layover time at ORD by twenty minutes for LAX-LGA passengers. In February 2017, American, Spirit, and United transported an average of 162 passengers per departure on the LAX-ORD segment.[57] Considering that 60% of ORD arriving passengers make connections, 97 of these passengers are expected to board connecting flights. For simplicity, assume that 40 of these passengers continue to LGA.[58] Using estimates from column two of Table 2.3, layover time is valued at \$46.23 an hour for low-cost carriers. Using these estimates, tightening connections by twenty minutes results in a gain of \$17.45 per passenger. Therefore, tightening connections for 40 passengers results in a total gain of \$698.[59] The increase in congestion cost incurred when Spirit enters a longer departure queue is determined from equation (2.4). These costs

[57] This average may be computed from the U.S. Department of Transportation's T-100 Domestic Segment database.

[58] This number is certainly too high considering that connecting passengers on any given flight are likely connecting to a number of other spoke cities. However, a flight departing for LGA at ORD will also include passengers arriving from cities other than Los Angeles. Thus, moving up a LGA departure by twenty minutes also reduces the amount of layover time for passengers connecting from these other spoke cities.

[59] $\$17.45 * 40 = \$698$.

are then subtracted from the \$698 gain to produce the numbers presented in column four of Table 2.9.

Table 2.9: Hypothetical impact of tightening ORD connections by 20 minutes on LAX-LGA for Spirit passengers (Layover time valued at \$46.23 per hour)

Change in queue length at takeoff (1)	Increase in congestion cost (2)	Gain from tighter connection (3)	Difference (4)
From 0 to 5	\$15	\$698	\$683
From 5 to 10	\$45	\$698	\$653
From 10 to 20	\$197	\$698	\$501
From 20 to 30	\$403	\$698	\$295
From 30 to 40	\$767	\$698	-\$69

Notes: Layover time for low-cost carriers is valued at \$46.23 per hour referencing the results in column two of Table 2.3. Therefore, decreasing layover time by twenty minutes for 40 connecting LAX-LGA passengers results in a total revenue gain of approximately \$698. The increase in congestion cost incurred when entering a longer departure queue is determined from equation (2.4). The costs in column (2) are then subtracted from the \$698 gain to produce the numbers presented in column (4).

Interpreting the Table 2.9 results, Spirit will benefit from tightening connections by twenty minutes for 40 connecting passengers in four of the five outlined scenarios. Foremost, if the schedule change moves the flight's departure from a queue with 0 flights to one with 5 flights, the overall gain is \$683. Moving the departure from a queue with 5 to 10 or 10 to 20 flights results in gains of \$653 and \$501, respectively. Finally, if the schedule change moves the flight from a queue with 20 flights to one with 30 flights, the gain is \$295. However, the overall gain is negative if the flight is moved to a departure queue with 40 other flights.

Table 2.10 repeats the analysis presented in Table 2.9 for United. However, using the estimates from column two of Table 2.4, passenger layover time on United is valued at \$61.89 per hour. Using these estimates, tightening connections by twenty minutes results in a gain of \$23.67 per passenger. Even though United internalizes the congestion imposed on its own flights, the higher value of layover time leads to results that are similar to the atomistic results in Table 2.9. It is profitable for United to shorten connections by twenty minutes if the connecting flight is moved from a departure queue with 0 to 5, 5 to 10, 10 to 20, or 20 to 30 fights. However, the overall gain is negative if the flight is moved to a departure queue with 40 other flights. In this scenario, United suffers a loss of $832.

Table 2.10: Hypothetical impact of tightening ORD connections by 20 minutes on LAX-LGA for United passengers (Layover time valued at $61.89 per hour)

Change in queue length at takeoff (1)	Increase in congestion cost (2)	Gain from tighter connection (3)	Difference (4)
From 0 to 5	$26	$947	$921
From 5 to 10	$81	$947	$866
From 10 to 20	$374	$947	$573
From 20 to 30	$833	$947	$114
From 30 to 40	$1,779	$947	-$832

Notes: Layover time is valued at $61.89 per hour referencing the results in column two of Table 2.4. Therefore, decreasing layover time by twenty minutes for 40 connecting LAX-LGA passengers results in a total revenue gain of approximately $947. The increase in congestion cost incurred when entering a longer departure queue is determined from equation (2.5). The costs in column (2) are then subtracted from the $947 gain to produce the numbers presented in column (4).

The hypothetical exercise described in this section outlines how one might use the layover time estimates provided in this chapter to evaluate the trade-off between convenient flight connections and airport congestion. However, a more accurate evaluation of this trade-off will involve congestion costs that differ by time of day and under various weather conditions.[60] One must also consider that flights departing from a hub are populated with passengers connecting from multiple spoke cities. Moving up the departure of one flight will likely induce the carrier to alter the schedule of other flights within the carrier's network. Therefore, an accurate evaluation of this trade-off will involve a much more complex analysis than the simple hypothetical scenario analyzed here.

[60] Johnson and Savage (2006) estimate congestion costs in bad weather that are significantly higher than analogous congestion costs in good weather. Congestion costs at ORD are also highest in the late afternoon and early evening hours.

2.7 Conclusion

This chapter provides empirical estimates on the value of layover time in the U.S. airline industry. Utilizing published fare and itinerary data from Google Flights for the top 50 routes without nonstop service, significant evidence is found of a negative relationship between fare and layover time. As the amount of layover time on an itinerary increases, fares are found to decrease, supporting the conjecture that passengers are compensated with a discounted fare for purchasing a ticket that requires a layover. The magnitude of the discount depends on total layover time, with the discount increasing as layover time increases. Of the three dominant legacy carriers, United passengers are found to be compensated at the highest rate for the layovers they incur.

As discussed in Section 2.4.1, there is a concern that layover times are endogenous if airlines aim to reduce layovers in city-pair markets with a large number of time-sensitive passengers. If layover times are endogenous, then estimates presented in this chapter may overstate the true effect of layover time on fares. However, because this chapter limits the analysis to city-pair markets where nonstop service is not available, an argument similar to the one employed in Brueckner and Spiller (1994) is used to argue that layover times in these connecting markets may be treated as exogenous.

There are a few limitations with regard to the present study. Foremost, the data collected represent published fares, not actual consumer purchases. Second, data for Southwest is not included since Southwest fares are not published on Google Flights (or other travel aggregator websites such as Kayak and Expedia). Finally, the analysis focuses on a small subset of connecting domestic routes. Therefore, the data cannot be used to determine whether the value of layover time differs on routes where both nonstop and connecting service are available. It is also not clear if the value of layover time differs on international routes. Further empirical work using detailed fare and itinerary data encompassing a wider variety of routes would help shed light on these questions.

Chapter 3

Capacity constraints and service quality: Do airport slot controls reduce flight delays?

3.1 Introduction

The economic cost of flight delays have been shown to be substantial for both passengers and airlines. Forbes (2008) finds that airfares fall by \$1.42 for each additional minute of flight delay while Peterson et al. (2013) estimate that U.S. net welfare would increase by \$17.6 billion if flight delays are reduced by 10%.[1] To mitigate persistent congestion and delays, regulatory authorities have typically placed restrictions on the number of flights that are allowed to takeoff and land at capacity constrained airports (slot controls). While the effect of slot constraints on airfares has received considerable attention, very few papers have

[1] Regarding the value of travel time, Luttmann (2019) finds that an hour spent on a plane (air travel time) is valued at \$31-\$34 while an hour spent at an airport on a connection (layover time) is valued at \$42-\$47.

studied whether slot controls are effective at accomplishing their primary goal of reducing flight delays.[2] In a future where passenger demand is expected to surpass current network and airport capacity (the International Air Transport Association (IATA) estimates that as many as 100 additional airports worldwide will need to be placed under slot control in the next 10 years), determining how effective slot controls are at reducing airport congestion is an important empirical question.[3] This chapter fills that gap by taking advantage of a regulatory change applying to John F. Kennedy (JFK) and Newark (EWR) airports. On March 30, 2008, the Federal Aviation Administration (FAA) implemented slot controls at JFK, capping the number of arrivals and departures to 81 per hour. A similar policy went into effect at EWR on June 20, 2008.

While one would expect the introduction of slot controls to translate to a decrease in flight delays (resulting from a drop in flight traffic), the actual relationship is more complicated. Examining the link between airport concentration and the length of high-volume flight periods (known as flight banks), Ater (2012) finds that hub airlines choose longer arrival and departure banks as their share of flights in the bank increases. These longer flight banks are shown to be associated with shorter flight delays, indicating that hub airlines internalize the congestion they impose on their own flights. With respect to delay and congestion management, the findings in Ater (2012) suggest that policies aimed at reducing congestion at highly concentrated airports will only have a limited impact because dominant airlines already internalize congestion. In this respect, EWR and JFK provide interesting case studies. When slot controls were introduced in 2008, Continental accounted for 72% of flight traffic at EWR while Delta accounted for 31% of traffic at JFK.[4]

[2] Abramowitz and Brown (1993), Borenstein (1989), Ciliberto and Williams (2010), Dresner et al. (2002), Evans and Kessides (1993), Fukui (2019), Gayle and Wu (2013), Morrison (2001), Peteraf and Reed (1994), Snider and Williams (2015), and Van Dender et al. (2007) all find that fares are higher at slot-controlled airports.

[3] Current forecasts call for 1.9% per year growth of enplanements over the next two decades, resulting in 46% more passengers in the U.S. market by 2038 (FAA, 2018). Worldwide, IATA predicts that as many as 100 additional airports will need to be placed under slot control in the next 10 years (see `https://www.travelweekly.com/Travel-News/Airline-News/More-and-more-airports-running-out-of-space`).

[4] United and Continental announced and completed their merger in 2010.

Using a differences-in-differences (DD) design, this chapter finds that these controls were not effective at reducing the incidence or severity of arrival and departure delays at EWR and JFK. In the months following the introduction of slot controls, the average arrival delay actually increased by 7 minutes at EWR (after controlling for weather, seasonality, competition, and unobserved differences in on-time performance across airlines and routes).

The DD design employed in this chapter differs from the typical approach where the treatment group indicator does not vary over time. Because the FAA implemented slot controls at JFK and EWR in an effort to either reduce or ensure that delays did not increase relative to summer 2007 levels, flights at both airports in the year prior to slot control enactment (2007) are utilized as a control group, resulting in both the treatment and pre-post indicators varying over time. The advantage of this approach is any seasonal difference in flight-level on-time performance will be accounted for. This approach is necessary because a clean counterfactual airport for EWR and JFK is not readily available. Although most U.S. airports were not slot-controlled in 2008, daily traffic volumes at EWR and JFK are larger than most other U.S. airports and the airspace in New York is generally more congested than other multiple airport metro areas. In addition, since weather is a major contributing factor to flight delays, a good counterfactual airport should have similar weather.

To determine why slot controls were not effective at both airports, three potential mechanisms are explored. These mechanisms are a decrease in the length of Continental and Delta's arrival and departure banks, a decrease in scheduled flight times, and an increase in market concentration. During the slot control period, the length of Delta's departure banks at JFK decreased by about 2 minutes with no substantive change in the length of Continental's arrival or departure banks at EWR. Further, the scheduled time for flights arriving into or departing from EWR decreased by 1.5-2.2 minutes while the scheduled time for flights departing from JFK decreased by over 3.5 minutes after slot controls were enforced. Finally, no evidence of an increase in market concentration resulting from the slot allocation process is found. Collectively, these findings highlight the need for policymakers to carefully consider how the allocation of takeoff and landing slots will impact scheduling decisions (both of individual flights and the length of the dominant carrier's arrival and departure banks) and market concentration when implementing similar policies.

The remainder of this chapter is organized as follows. Section 3.2 discusses the costs associated with airport congestion and delays while Section 3.3 provides a history on slot controls in the U.S. and discusses why the FAA implemented slot controls at JFK and EWR in 2008. Section 3.4 describes the flight-level on-time performance, market concentration, and weather data used in the empirical analysis. Section 3.5 outlines the differences-in-differences estimation strategy while Section 3.6 presents results of the empirical analysis. Section 3.7 provides a further decomposition of the results by examining three potential mechanisms that may explain why slot controls were not effective at accomplishing their primary goal of reducing delays. Finally, Section 3.8 concludes.

3.2 Airport Congestion and Delays

The costs associated with congestion and delays are substantial for passengers and airlines. Exploiting an exogenous shock to delays resulting from a legislative change in takeoff and landing restrictions at LaGuardia, Forbes (2008) finds that fares decrease by \$1.42 on average for each additional minute of flight delay.[5] Examining the impact of delays on consumer and producer welfare for a sample of U.S. routes, Britto et al. (2012) find delays raise fares and reduce demand. A 10% reduction in delays implies a benefit of \$1.50-\$2.50 per passenger while the welfare gain for airlines is nearly three times larger at \$4.44 per passenger.[6] In

[5] Forbes (2008) finds the estimated fare response in competitive markets to be substantially larger at \$2.44 per additional minute of delay.

[6] In a similar study incorporating a larger sample of routes, Zou and Hansen (2014) find decreasing delays by 25% results in a fare reduction of \$3 per passenger. This estimate is similar to the \$3.06 number found in Britto et al. (2012) when delays fall by 20%.

another study, Peterson et al. (2013) estimate that U.S. net welfare would increase by $17.6 billion if flight delays are reduced by 10%.[7]

Collectively, these studies establish that both passengers and airlines reap substantial benefits from a reduction in flight delays. When faced with an inefficient market outcome, economists typically advocate for a Pigouvian tax to correct for the negative externality.[8] However, in the context of airport congestion, there exists substantial debate concerning the optimal congestion-based pricing scheme. While some researchers advocate for a uniform congestion toll regardless of a carrier's market share, others argue that tolls should only reflect the congestion each airline imposes on flights it does not operate.

Contrary to the analysis of highway congestion where individual users are atomistic (they do not consider the congestion their vehicle imposes on other drivers), Daniel (1995) was the first to recognize that airlines at heavily concentrated airports may internalize congestion they impose on their own flights. If airlines exhibit internalizing behavior, then optimal congestion tolls should be reduced (relative to the atomistic case) to reflect only the congestion imposed on other carriers.[9] Investigating this possibility, Brueckner (2002a) presents a theoretical model showing that congestion is fully internalized when an airport is dominated by a monopolist. However, under Cournot oligopoly, carriers internalize only the congestion they impose on themselves.

Testing the internalization hypothesis empirically, Brueckner (2002a), Mayer and Sinai (2003), Ater (2012), and Miranda and Oliveira (2018) find evidence that delays decrease as airport dominance increases, suggesting that dominant airlines internalize congestion at heavily con-

[7] In a more ambitious scenario where flight delays are reduced by 30%, Peterson et al. (2013) estimate a U.S. net welfare increase of $38.5 billion.

[8] Currently, airlines are charged weight-based landing and departure fees that do not take into account the amount of congestion on the tarmac.

[9] Pels and Verhoef (2004) and Brueckner (2005) find optimal congestion tolls are much lower than atomistic tolls when carriers internalize self-imposed congestion. Interestingly, Pels and Verhoef (2004) find that pure congestion tolls, even when corrected for self-imposed congestion, may cause a decrease in welfare when there is significant market power.

centrated airports. Using a different delay measure, Rupp (2009) finds the opposite result, concluding that carriers do not internalize. Although Daniel (1995) was the first to suggest internalizing behavior, he argued that it is optimal for dominant carriers not to internalize congestion when faced with competitive pressure from fringe carriers. For instance, if the dominant carrier internalized congestion costs by reducing its number of flights, fringe carriers would likely fill the void with new service. Supporting this theory, Daniel and Harback (2008) present empirical evidence rejecting the internalization hypothesis at most major US airports. Seeking to unite conflicting results from the Brueckner (2002a) and Daniel (1995) studies, Brueckner and Van Dender (2008) provide theoretical results that show carriers in a Stackelberg competitive environment behave in an atomistic manner while carriers in a Cournot competitive environment exhibit internalizing behavior.

Although an optimal congestion-based tolling scheme could be administered, the Federal Aviation Administration (FAA) has typically addressed congestion at capacity constrained airports by implementing slot controls restricting the number of flights. These controls are currently in place at Washington National (DCA), LaGuardia (LGA), and John F. Kennedy (JFK) airports. Slot controls have previously been implemented at Chicago O'Hare (ORD) and Newark (EWR). This chapter investigates whether slot controls are effective at reducing congestion and delays at large airports by looking at the months prior to and after the implementation of slot controls at EWR and JFK in 2008.

EWR and JFK provide interesting case studies. Foremost, EWR was a major hub for Continental airlines while Delta operates a medium-sized hub at JFK.[10] At the time slot controls were implemented, Continental controlled 72% of EWR traffic while Delta controlled 31% of JFK traffic. If Continental and Delta already internalized self-imposed congestion prior to the implementation of slot controls, few congestion-related benefits may have been realized. Conversely, implementing slot controls at hub airports has the potential to exacerbate

[10] When the slot controls were initially implemented in 2008, EWR was a hub for Continental Airlines. United and Continental announced and completed their merger in 2010.

congestion if the policy results in the hub carrier moving from self-internalizing to atomistic behavior with respect to flight scheduling. Analyzing the impact of slot controls at EWR and JFK will provide insight into how the effectiveness of these controls potentially differs with airport concentration.

3.3 History of Slot Controls in the United States

Airports around the world are designated at three different levels according to their degree of congestion. According to the FAA, Level 1 airports have sufficient capacity to meet demand while Level 2 airports have some periods when demand approaches capacity limits, but a voluntary schedule-facilitation process prevents systemic delays.[11] Finally, Level 3 airports have demand that far exceeds the airport's capacity and without controls, would have unacceptable levels of delays. Airlines require approval to operate during slot-controlled hours at all Level 3 airports.

Level 3 slot controls were first implemented in the U.S. in the late 1960s. To reduce congestion and delays at John F. Kennedy (JFK), Newark (EWR), LaGuardia (LGA), Washington National (DCA), and Chicago O'Hare (ORD), the FAA adopted the High Density Rule (HDR). Effective in 1969, the HDR capped the number of hourly arrivals and departures permitted at these five airports while also requiring airlines to acquire takeoff and landing slots to operate during slot-controlled hours. While DCA has remained slot-controlled under the HDR since 1969, there have been substantial changes to the application of the HDR at the four other airports. After determining that airport capacity could meet demand, the FAA suspended slot controls at EWR in 1970. Then, on April 5, 2000, Congress passed the Wendall H. Ford Aviation Investment and Reform Act of the 21st Century (AIR-21) requiring the HDR to be phased out at ORD by July 1, 2002 and at JFK and LGA by

[11] See `https://www.faa.gov/about/office_org/headquarters_offices/ato/service_units/systemops/perf_analysis/slot_administration/slot_definition/`.

January 1, 2007. In phasing out the HDR, AIR-21 permitted the Secretary of the Department of Transportation to grant two types of exemptions from the HDR's flight restrictions at LGA. To increase competition at slot-controlled airports, the first exemption permitted the Secretary to grant slot exemptions to a new entrant or an incumbent carrier holding fewer than 20 slots (DOT, 2015). The second exemption aimed to improve service to small communities by requiring exemptions to be granted to a carrier operating an aircraft with less than 71 seats to small-hub or non-hub airports for an unrestricted number of flights (DOT, 2015).

Although the HDR was to be phased out at LGA on January 1, 2007, the exemptions permitted by AIR-21 resulted in a substantial increase in delays. In 2000, the average arrival delay at LGA increased 144% from 15.52 minutes in March 2000 to 37.86 minutes in September 2000 (DOT, 2015). Based on this experience, the FAA determined that lifting the HDR at LGA would result in a substantial increase in delays. Therefore, the FAA issued an order in 2006 limiting the number of scheduled operations at LGA to 75 per hour, the same hourly limit in place under the HDR.

3.3.1 John F. Kennedy Airport After AIR-21

After the HDR was phased out at JFK on January 1, 2007, average daily operations increased by 21% over 2006 levels. As a result of the substantial increase in operations, on-time gate arrivals (defined as arriving within 15 minutes of scheduled arrival time) declined from 68.5% in 2006 to 62.19% in 2007 (DOT, 2015). Additionally, average daily arrival delays exceeding one hour increased by 87% relative to 2006 levels while taxi-out delays increased by 15%.[12] To address projected increases in demand for summer 2008 and the previous over-scheduling in summer 2007 when JFK lacked scheduling limits, the FAA issued an order on January 18, 2008 capping the number of hourly scheduled flight operations to 81 per hour. This order

[12] Taxi-out delays measure the time that aircraft wait prior to departing the runway.

went into effect at 6:00am on March 30, 2008. The FAA's goal was to reduce peak evening departure delays from the summer 2007 average of nearly 80 minutes.

3.3.2 Newark Liberty International Airport After AIR-21

Since 2000, Newark has been one of the most delay-prone airports in the U.S. The percentage of on-time gate arrivals decreased from 70.66% in 2000, to 63.97% in 2006 and 61.71% in 2007. Over the same period, the average daily counts of flights with arrival delays greater than one hour were 54 in 2000, 79 in 2006, and 93 in 2007 (DOT, 2008a). To alleviate persistent congestion and delays at EWR, the FAA proposed placing temporary limits on scheduled flight operations on March 18, 2008. The FAA indicated these limits would be necessary due to the projected increase in flight delays that were likely to occur during summer 2008 if proposed flight schedules were implemented as requested by the carriers. In particular, U.S. and foreign carriers requested almost 100 new operations, adding to the schedules that produced pronounced delays during the summer of 2007. According to the FAA, EWR is capable of handling an average of 83 operations per hour. However, if proposed summer 2008 schedules were implemented, peak periods would have experienced operations in excess of 90 per hour.

Following through with their proposal on March 18th, the FAA issued an order on May 21, 2008 officially establishing an hourly limit of 81 scheduled operations between the hours of 6:00am and 10:59pm Eastern Time (DOT, 2008b). The order went into effect on June 20, 2008 with an initial expiration of October 24, 2009.[13] However, the order was amended four times between 2008 and 2014, pushing the eventual expiration date of the slot controls to October 30, 2016.[14] The FAA established a goal of no increase in delays relative to summer

[13] Similar orders were also issued for the other New York area airports in 2007 and 2008. Hourly limits went into effect at LaGuardia Airport (75 scheduled operations per hour) on January 1, 2007, and at John F. Kennedy Airport (81 scheduled operations per hour) on March 30, 2008.

[14] The May 21st, 2008 order was amended on October 7th, 2009, April 4th, 2011, May 14th, 2013, and

2007 levels while permitting additional operations to the extent possible.

To enforce the hourly limit of 81 scheduled operations, the FAA assigned arrival and departure slots in half-hour intervals by day of week.[15] Once allocated a slot, airlines are required to utilize the slot 80% of the time or risk having their slot revoked by the FAA.[16] While the 80% usage requirement was established to provide airlines with a buffer to suspend flights for operational reasons (such as cancellations for weather or ground delays), this usage requirement may have adversely impacted competition at EWR and JFK. By not requiring airlines to schedule flights for 100% of their slots, some existing airport capacity will go unused. This unused capacity could be allocated to other carriers (or new entrants) to stimulate competition. Furthermore, rather than evaluating each slot individually, the FAA allows airlines to apply the 80% requirement to their pool of slots within each half-hour period.[17] This enforcement policy would seem to provide larger carriers with an incentive to obtain as many slots as possible to increase their operational flexibility without necessarily increasing their number of operations.[18] For the sake of holding onto valuable slot holdings, this "use-it-or-lost-it" policy potentially creates an incentive for airlines to operate higher frequencies with smaller aircraft (Fageda and Flores-Fillol, 2017).

Overall, the slot control policies implemented by the FAA at EWR and JFK were very

March 26th, 2014.

[15] To provide an example of how slots are allocated by the FAA, a list of the arrival and departure slots initially assigned to Continental and Delta at EWR and JFK are provided in Tables I.1, I.2, J.1, J.2, and J.3.

[16] Slots are typically assigned by the FAA for use over a three-month period.

[17] The FAA also does not require airlines to report the same flight in a specific slot over the course of a reporting period (typically three months). For example, if Continental was awarded a Monday 8:30am departure slot for the summer of 2008, Continental would not be required to fly between the same two endpoints for each 8:30am flight it operates on Mondays during the summer 2008 reporting period. Therefore, Continental could use the slot on Monday June 20th, 2008 to fly between EWR and LAX. The same departure slot could then be used by Continental to fly between EWR and SFO on Monday June 27th, 2008.

[18] This "hoarding" behavior would further result in an inefficient use of airport capacity while potentially hindering competition. For instance, in a 2012 Government Accountability Office (GAO) report reviewing slot-controlled airports, the GAO suggests that the 80% usage requirement restricted entry and resulted in large carriers "hoarding" slots by flying excessive frequencies using small planes. Specifically, the GAO found that scheduled passenger flights at slot-controlled airports were 75% more likely to be scheduled by airlines using an aircraft with fewer than 100 seats than flights at other like-sized airports that are not slot-controlled.

similar. Scheduled flight operations were limited to 81 per hour at both airports between the hours of 6:00am and 10:59pm Eastern Time. These controls also went into effect within three months of each other (Monday March 30th, 2008 at JFK and Monday June 20th, 2008 at EWR). In addition, although arrival and departure slots were initially assigned by the FAA in fifteen minute intervals at JFK and thirty minute intervals at EWR, the FAA enforced the 80% slot utilization requirement in half-hour intervals at both airports.

3.4 Data

3.4.1 Flight-level On-Time Performance (OTP)

Since 1987, all U.S. carriers with at least one percent of total domestic traffic have been required to report on-time performance (OTP) data to the Bureau of Transportation Statistics (BTS). For each domestic flight, the OTP data includes the scheduled arrival and departure time, the actual arrival and departure time, time spent taxiing from the gate to the runway, time spent taxiing to the gate after landing, the flight and tail numbers of the aircraft, whether the flight was canceled or diverted, and indicators specifying whether the flight departed or arrived more than fifteen minutes past the scheduled departure or arrival time.[19]

This dataset has been used in several empirical studies of the airline industry. For instance, Mazzeo (2003), Rupp et al. (2006), and Greenfield (2014) find that OTP improves in less concentrated airport-pair markets. In a related study, Prince and Simon (2014) find that incumbent OTP actually worsens in response to entry and entry threats from low-cost carriers. Other papers relying on these data have answered questions concerning schedule padding

[19] The U.S. Department of Transportation defines a late arrival as any flight that arrives to the gate fifteen or more minutes past the scheduled arrival time. In other words, a flight that arrives at the gate fourteen minutes past the scheduled arrival time is considered "on-time."

(Forbes et al., 2019b), information disclosure requirements (Forbes et al., 2019a), mergers and quality provision (Prince and Simon, 2017), the effect of the internet on scheduled flight times (Ater and Orlov, 2015), multimarket contact and service quality (Prince and Simon, 2009), and the determinants of flight cancellations (Rupp and Holmes, 2006).

Figure 3.1 summarizes the OTP data for all Newark (EWR) domestic flights between January 1st, 2007 and December 31st, 2008. Panel A plots the daily number of arriving flights along with the pre-slot (January 1st, 2007 through June 19th, 2008) and slot control (June 20th, 2008 through December 31st, 2008) averages. Foremost, there exists clear within-week variation in the pattern of arrivals, with traffic on weekends noticeably lower than weekdays. As illustrated in Panel A, the average number of domestic arrivals falls from 415 in the pre-slot control period to 359.77 after slot controls are enacted. To see if this reduction in arriving traffic reduced the daily percentage of late arrivals (defined by the Department of Transportation as a flight arriving fifteen or more minutes late), Panel B plots the daily fraction of flights arriving late along with the pre-slot and slot control averages. The reduction in arriving traffic depicted in Panel A generated a modest decrease in the daily percentage of late arrivals from an average of 37.03% in the pre-slot control period to 34.46% in the slot control period.

Looking at the flight-level, Panel C of Figure 3.1 plots the daily average arrival delay along with the pre-slot and slot control averages. Although there was a large reduction in the daily average number of arrivals at EWR, there was not a large reduction in the average arrival delay. In the pre-slot control period, the daily average arrival delay was 22.55 minutes compared to 21.70 minutes in the slot control period.

Panels D-F of Figure 3.1 are similar to Panels A-C, except that the variables plotted are the daily number of domestic departures, the daily fraction of flights departing late, and the daily average departure delay. Similar to Panel A, Panel D of Figure 3.1 illustrates a substantial reduction in the daily number of departures from an average of 415.05 in the pre-slot control period to 359.82 after slot controls are enacted. This reduction in departing traffic led to a small decrease in the daily percentage of late departures. Depicted in Panel E, the daily percentage of late departures fell from an average of 29.54% in the pre-slot control period to 27.95% in the slot control period. However, the reduction in departing traffic did not translate to a large decrease in the average departure delay. Demonstrated in Panel F, the daily average departure delay slightly decreased from a pre-slot control average of 19.44 minutes to 18.54 minutes in the slot control period.

Figure 3.1: Newark Airport (EWR) Flight Traffic and Delay Summary

Notes: The data used to construct Panels A-F come from the Bureau of Transportation Statistics Airline On-Time Performance (OTP) database. The pre-slot control period is from January 1st, 2007 thorough June 19th, 2008 while the slot control period is from June 20th, 2008 thorough December 31st, 2008. Adopting the Department of Transportation definition, a flight in Panel B is considered to arrive late if the flight arrives 15 or more minutes past the scheduled arrival time. Similarly, a flight is considered to depart late in Panel E if the flight leaves 15 or more minutes past the scheduled departure time.

Figure 3.2 summarizes the OTP data for all JFK domestic flights between January 1st, 2007 and December 31st, 2008. Panel A plots the daily number of arriving flights along with the pre-slot (January 1st, 2007 through March 29th, 2008) and slot control (March 30th, 2008 through December 31st, 2008) averages. Illustrated in Panel A, the average number of domestic arrivals falls from 343.17 in the pre-slot control period to 323.23 after slot controls are enacted. To see if this reduction in arriving traffic reduced the daily percentage of late arrivals, Panel B plots the daily fraction of flights arriving late along with the pre-slot and slot control averages. Depicted in Panel B, there was a sizable decrease in the daily percentage of late arrivals from an average of 33.53% in the pre-slot control period to 29.56% in the slot control period. Turning to the individual flight-level, Panel C plots the daily average arrival delay along with the pre-slot and slot control averages. The reduction in arriving JFK traffic also resulted in a two minute decrease in the average arrival delay. In the pre-slot control period, the daily average arrival delay was 17.90 minutes compared to 15.26 minutes in the slot control period.

Panels D-F of Figure 3.2 are similar to Panels A-C, except that the variables plotted are the daily number of domestic departures, the daily percentage of flights departing late, and the daily average departure delay. Similar to Panel A, Panel D of Figure 3.2 illustrates a large reduction in the daily average number of departures from 343.24 in the pre-slot control period to 323.23 after slot controls are enacted. This reduction in daily departures also resulted in decreases in the daily percentage of late departures and the daily average departure delay. Depicted in Panel E, the daily fraction of late departures fell from an average of 27.72% in the pre-slot control period to 22.50% in the slot control period. Furthermore, the daily average departure delay decreased over three minutes from a pre-slot control average of 17.46 minutes to 14.04 minutes in the slot control period (see Panel F).

Figure 3.2: John F. Kennedy Airport (JFK) Flight Traffic and Delay Summary

Notes: The data used to construct Panels A-F come from the Bureau of Transportation Statistics Airline On-Time Performance (OTP) database. The pre-slot control period is from January 1st, 2007 thorough March 29th, 2008 while the slot control period is from March 30th, 2008 thorough December 31st, 2008. Adopting the Department of Transportation definition, a flight in Panel B is considered to arrive late if the flight arrives 15 or more minutes past the scheduled arrival time. Similarly, a flight is considered to depart late in Panel E if the flight leaves 15 or more minutes past the scheduled departure time.

At face value, Figures 3.1 and 3.2 suggest that the implementation of slot controls at EWR and JFK in 2008 slightly reduced the frequency and magnitude of arrival and departure delays at both airports. However, these figures should not be interpreted as strong evidence to the effectiveness of slot controls since the figures do not control for other factors that influence flight delays such as weather, seasonality, and competition. In addition, delays may have decreased in 2008 for reasons not related to the implementation of slot controls.

For instance, passenger traffic significantly fell in the latter half of 2008 as a result of the financial crisis. Therefore, OTP performance could have improved due to the financial crisis (and not due to slot controls) as airplanes take less time to load and unload with the decrease in passenger traffic. Section 3.5 outlines an empirical model that controls for weather, seasonality, competition, the 2008 financial crisis, and unobserved differences in OTP across airlines and routes while Section 3.6 presents results from that model.

3.4.2 Market Concentration

The data used to construct monthly route-level market concentration (Herfindahl-Hirschman Index) are derived from the U.S. Department of Transportation's (DOT) T-100 service-segment database. Information contained in the T-100 database is derived from the Form 41, which commercial air carriers have been required to submit since 1990. In addition to financial operating characteristics, Form 41 contains information on the number of departures, passengers served, and available seats for each nonstop U.S. route segment on a specific plane type in a given month.

To construct the monthly Herfindahl-Hirschman Index for each route with EWR or JFK as an endpoint, the number of passengers is first aggregated to the origin-destination-airline-year-month level. These passenger sums are then divided by the total number of passengers traveling on the origin-destination segment across all carriers in that year-month. These passenger shares are then squared and the resulting sum of these squared shares is the Herfindahl-Hirschman Index (HHI).[20]

[20] HHI is measured at the origin-destination-year-month level and ranges from 1 (highly competitive) to 10,000 (monopoly).

3.4.3 Weather Data

Weather data for EWR and JFK come from the National Oceanographic and Atmospheric Administration's (NOAA) National Climatic Data Center. Recorded at the hourly level, the NOAA data include information on precipitation (inches), humidity, temperature (degrees Fahrenheit), and windspeed (miles per hour). Figure 3.3 plots the daily high and low temperatures in addition to total precipitation levels for EWR and JFK. As expected, there is clear seasonality in both temperature (Panels A and C) and precipitation levels (Panels B and D), with the lowest temperatures occurring during the winter and highest precipitation levels occurring during the summer.

Figure 3.3: Daily Temperature and Precipitation Summary for EWR and JFK

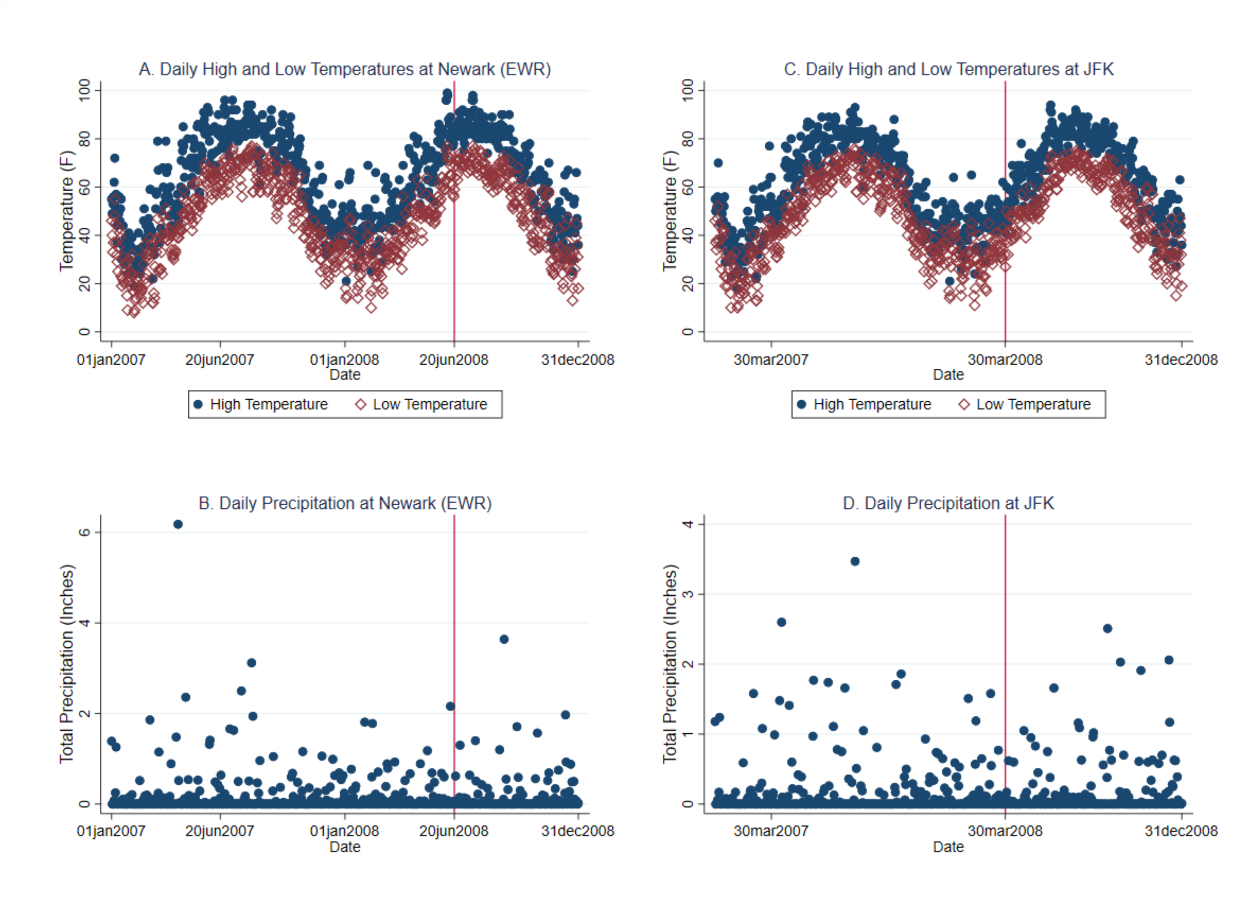

Notes: The data used to construct Panels A-D come from the National Oceanographic and Atmospheric Administration's (NOAA) National Climatic Data Center. The vertical line in Panels A and B indicate the date slot controls (June 20th, 2008) were implemented at Newark (EWR) airport. The vertical line in Panels C and D indicate the date (March 30th, 2008) slot controls were implemented at John F. Kennedy (JFK) airport.

3.5 Empirical Strategy

The goal of the empirical analysis is to identify the effect of slot controls on the incidence and severity of arrival and departure delays at JFK and EWR in 2008. Outcomes of interest include departure and arrival delay minutes, the probability that a flight departs or arrives fifteen or more minutes late, and the probability that a flight departs or arrives sixty or more minutes late. Identifying the effect of slot controls on delay measures at these airports is difficult because a clean counterfactual group for each airport is not readily available. Although the majority of airports in the U.S. were not slot-controlled at the time these controls were implemented, daily traffic volumes at JFK and EWR are larger than most other U.S. airports and the airspace in New York is generally more congested than other multiple airport metro areas. Furthermore, because weather is a major contributing factor to delays, a good counterfactual airport should have similar weather. To overcome this issue, the empirical strategy takes advantage of the reason slot controls were implemented in the first place. At JFK, the goal of these controls were to reduce delays relative to summer 2007 levels. At EWR, the goal of these controls were to ensure that delays did not increase relative to summer 2007 levels. Given that these policies were implemented to reduce or ensure that delays did not increase relative to summer 2007 levels, summer 2007 provides an ideal counterfactual group for use in a differences-in-differences (DD) design. This model uses flights at the same airport in the year prior to the implementation of slot controls as a control group. This strategy assumes that, in the absence of slot controls, delays at EWR and JFK would have followed similar trends in 2007 and 2008. Visual evidence of this assumption is illustrated in Figure 3.1 for EWR and Figure 3.2 for JFK.

The estimating equation for EWR is,

$$y_{ijkt} = \beta_0 + \beta_1 \cdot \mathrm{TREAT}_t + \beta_2 \cdot \mathrm{POST}_t + \beta_3 \cdot \mathrm{TREAT}_t \cdot \mathrm{POST}_t + \beta_4 \cdot \mathrm{HHI}$$

$$+\beta_5 \cdot \mathrm{FINANCIAL\ CRISIS} + \gamma \cdot W_{jt} + \delta_{kj} + \theta_t + \epsilon_{ijkt} \qquad (3.1),$$

where y_{ijkt} is one of the measures of OTP for individual flight i on directional route j operated by carrier k on date t.[21] TREAT is a dummy variable equal to one in 2008 and zero in 2007, POST is a dummy variable equal to one if the day of the year is on or after June

[21] As mentioned earlier, outcomes of interest include departure and arrival delay minutes, the probability that a flight departs or arrives fifteen or more minutes late, and the probability that a flight departs or arrives sixty or more minutes late.

20th and zero otherwise, and β_3 is the DD coefficient of interest.[22] HHI is the Herfindahl-Hirschman Index while FINANCIAL CRISIS is an indicator equal to one if the day is on or after the collapse of Lehman Brothers on September 15th, 2008. The estimated sign of the coefficient on FINANCIAL CRISIS is ambiguous. If airlines responded to the financial crisis by cutting airport staff (for example, gate agents and baggage handlers), then OTP may have declined with less staff available to process checked luggage. Conversely, the drop in passenger traffic resulting from the financial crisis may have led to an improvement in OTP as the airport becomes less congested and planes take less time to load and unload. To flexibly control for weather-related shocks to OTP, W is a matrix of airport-level weather controls containing quartics in hourly temperature, windspeed, precipitation, and humidity in addition to one-hour lags of these variables. δ_{kj} are airline-route fixed effects that control for unobserved differences in OTP across airlines and routes while θ is a matrix containing a linear time trend in addition to month-of-year, day-of-week, and hour-of-day fixed effects. The month and day-of-week fixed effects control for seasonality while the hour-of-day fixed effects control for variation in airport traffic and congestion levels throughout a typical day. Since delay profiles differ across days, standard errors are clustered at the daily level to account for this correlation.

The estimating equation for JFK, equation (3.2), is similar to equation (3.1), except that POST is a dummy variable equal to one if the day of the year is on or after March 30th and zero otherwise.

[22] This DD design differs from the typical approach where the treatment group indicator does not vary over time. Because the FAA implemented slot controls at JFK and EWR in an effort to either reduce or ensure that delays did not increase relative to summer 2007 levels, flights at both airports in 2007 (year prior to slot control enactment) are utilized as a control group, resulting in both the treatment and POST indicators varying over time. The advantage of this approach is any seasonal difference in flight-level on-time performance will be accounted for.

3.5.1 EWR Estimation Sample

Since the empirical strategy uses the year prior to the implementation of slot controls as a control group, the estimation sample for EWR includes the universe of arriving and departing EWR domestic flights between January 1st, 2007 and December 31st, 2008. To limit the impact of the 2008 financial crisis, an additional robustness check limits the analysis sample to the symmetric 60-day window before and after slot controls were implemented on June 20th, 2008 and compares the OTP performance of these 2008 flights to comparable flights in 2007. This restriction limits the sample to the universe of arriving and departing domestic flights from April 21st, 2007 to August 18th, 2007 and from April 21st, 2008 to August 18th, 2008.

3.5.2 JFK Estimation Sample

The estimation sample for JFK includes the universe of JFK arriving and departing domestic flights from January 1st, 2007 to December 31st, 2008. To limit the impact of the 2008 financial crisis, an additional robustness check limits the analysis sample to the symmetric 60-day window before and after slot controls were implemented on March 30th, 2008 and compares the OTP performance of these 2008 flights to comparable flights in 2007. This restriction limits the analysis sample to the universe of arriving and departing domestic flight from January 30th, 2007 to May 28th, 2007 and from January 30th, 2008 to May 28th, 2008.

3.6 Results

This section presents results from estimating the models described in the Empirical Strategy Section. First, results from equations (3.1) and (3.2) are presented in Section 3.6.1. Then, the analysis sample is restricted to the symmetric 60-day window before and after slot controls were implemented in 2008 (March 30th for JFK and June 20th for EWR) and compares the OTP performance of these flights to comparable flights in 2007. Results from this restricted sample are presented in Section 3.6.2.

3.6.1 Full Sample Results

Column (1) of Table 3.1 provides results from equation (3.1) when arrival delay minutes is the outcome of interest and the analysis sample contains all EWR domestic arrivals between January 1st, 2007 and December 31st, 2008. Consistent with Mazzeo (2003) and Greenfield (2014), the positive and significant coefficient on HHI (1000s) indicates that OTP declines in more concentrated airport-pair markets. Although insignificant, the coefficient on FINANCIAL CRISIS suggests that the average flight-level arrival delay decreased by over three and a half minutes after the collapse of Lehman Brothers on September 15th, 2008. As discussed in the Empirical Strategy section, this negative effect supports the conjecture that OTP improved during the FINANCIAL CRISIS due to the large decline in passenger traffic. Most importantly however, the positive and significant coefficient on TREAT*POST (Slot Control) indicates that the average flight-level arrival delay at EWR **increased** by 7.13 minutes after slot controls were implemented on June 20th, 2008. Thus, after controlling for weather, seasonality, the 2008 financial crisis, and unobserved differences in OTP across airlines and routes, the implementation of slot controls adversely impacted the OTP of EWR arriving flights between June 20th, 2008 and December 31st, 2008 relative to comparable EWR arriving flights over the same time period in 2007. Section 3.7 evaluates a few of the potential reasons why OTP at EWR suffered once slot controls were implemented.

Column (2) of Table 3.1 provides results from equation (3.1) when the probability that a flight arrives sixty or more minutes late is the outcome of interest while column (3) provides results when the probability that a flight arrives fifteen or more minutes late is the outcome of interest. The positive and significant coefficient on TREAT*POST (Slot Control) in column (2) indicates that the probability that a flight arrives into EWR sixty or more minutes late increased by 4% once slot controls were enforced. Similarly, the significant coefficient on TREAT*POST (Slot Control) in column (3) indicates that the probability that a flight arrives into EWR fifteen or more minutes late increased by 6% after slot controls were implemented on June 20th, 2008.

Columns (4)-(6) of Table 3.1 are similar to columns (1)-(3), except that the outcomes of interest are departure delay minutes, the probability that a flight departs sixty or more minutes late, and the probability that a flight departs fifteen or more minutes late. The coefficient on TREAT*POST (Slot Control) in column (4) while insignificant, suggests that the average EWR departure delay increased by over one and half minutes once slot controls were enforced. Similarly, the insignificant positive coefficient on TREAT*POST (Slot Control) in column (5) suggests that the probability that a flight departs from EWR sixty or more minutes late increased by 1% once slot controls were enforced while the insignificant 0.02 coefficient in column (6) suggests that the probability that a flight departs from EWR fifteen or more minutes late increased by 2% after the implementation of slot controls.

Table 3.1: Differences-in-Differences Regression Results for Newark (June 20th, 2007 through December 31st, 2007 Counterfactual)

	Arriving Flights			Departing Flights		
Dependent Variable:	Delay Minutes	Delay > 60	Delay > 15	Delay Minutes	Delay > 60	Delay > 15
	(1)	(2)	(3)	(4)	(5)	(6)
TREAT*POST (Slot Control)	7.13**	0.04**	0.06***	1.57	0.01	0.02
	(3.09)	(0.02)	(0.02)	(1.67)	(0.01)	(0.02)
POST (June 20th-December 31st)	0.57	0.00	-0.01	2.97	0.00	0.04*
	(4.79)	(0.03)	(0.03)	(2.17)	(0.01)	(0.02)
TREAT (Year = 2008)	-20.52	-0.08	-0.01	-12.15	-0.06	-0.04
	(25.68)	(0.15)	(0.20)	(12.87)	(0.09)	(0.13)
HHI (1000s)	1.31***	0.01***	0.01***	0.89***	0.00**	0.01***
	(0.34)	(0.00)	(0.00)	(0.20)	(0.00)	(0.00)
FINANCIAL CRISIS	-3.63	-0.03	-0.03	-2.89	-0.02*	-0.04**
	(3.39)	(0.02)	(0.03)	(1.87)	(0.01)	(0.02)
Observations	279,895	279,895	279,895	281,398	281,398	281,398
R-squared	0.28	0.24	0.25	0.40	0.29	0.26
Time Trend	YES	YES	YES	YES	YES	YES
Airline-Route FE	YES	YES	YES	YES	YES	YES
Month-of-Year FE	YES	YES	YES	YES	YES	YES
Day-of-week FE	YES	YES	YES	YES	YES	YES
Hour-of-day FE	YES	YES	YES	YES	YES	YES
Weather Controls	YES	YES	YES	YES	YES	YES

Notes: Standard errors clustered at the daily level are reported in parentheses. The sampling period is January 1st, 2007 through December 31st, 2008.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Column (1) of Table 3.2 provides results from equation (3.2) when arrival delay minutes is the outcome of interest and the analysis sample contains all JFK domestic arrivals between January 1st, 2007 and December 31st, 2008. Consistent with Table 3.1, the positive (although insignificant) coefficient on HHI (1000s) indicates that OTP worsens in more concentrated airport-pair markets. However, the coefficient on HHI changes signs in columns (2)-(6) with different outcomes of interest. Although insignificant, the coefficient on FINANCIAL CRISIS suggests that the average flight-level arrival delay at JFK increased by over three minutes after the collapse of Lehman Brothers on September 15th, 2008. While in contrast to the EWR results in Table 3.1, the positive coefficient supports the conjecture that OTP may have declined during the FINANCIAL CRISIS if airlines responded by laying off airport staff such as gate attendants and baggage handlers. The positive coefficient on TREAT*POST

(Slot Control) in column (1) of Table 3.2 suggests that the average flight-level arrival delay at JFK increased by 3.18 minutes after slot controls were implemented on March 30th, 2008. However, this estimated effect is not significant at conventional levels.

Column (2) of Table 3.2 provides results from equation (3.2) when the probability that a flight arrives sixty or more minutes late is the outcome of interest while column (3) provides results when the probability that a flight arrives fifteen or more minutes late is the outcome of interest. The coefficient on TREAT*POST (Slot Control) in column (2) suggests that the probability that a JFK arriving flight arrives sixty or more minutes late increased by 1% once slot controls were enforced. However, in line with the column (1) results, the estimated effect is insignificant at conventional levels. Similarly, the marginally significant coefficient on TREAT*POST (Slot Control) in column (3) indicates that the probability that a flight arrives into JFK fifteen or more minutes late increased by 3% after the implementation of slot controls on March 30th, 2008.

Columns (4)-(6) of Table 3.2 are similar to columns (1)-(3), except that the outcomes of interest are departure delay minutes, the probability that a flight departs sixty or more minutes late, and the probability that a flight departs fifteen or more minutes late. The coefficient on TREAT*POST (Slot Control) in column (4) while insignificant, suggests that the average JFK departure delay increased by over two minutes once slot controls were enforced. Similarly, the insignificant coefficient on TREAT*POST (Slot Control) in column (5) suggests that the probability that a flight departs from JFK sixty or more minutes late increased by 1% once slot controls were enforced while the insignificant coefficient in column (6) suggests that the probability that a flight departs from JFK fifteen or more minutes late increased by 2% after slot controls were put in place.

Table 3.2: Differences-in-Differences Regression Results for JFK
(March 30th, 2007 through December 31st, 2007 Counterfactual)

	Arriving Flights			Departing Flights		
Dependent Variable:	Delay Minutes	Delay > 60	Delay > 15	Delay Minutes	Delay > 60	Delay > 15
	(1)	(2)	(3)	(4)	(5)	(6)
TREAT*POST (Slot Control)	3.18	0.01	0.03*	2.04	0.01	0.02
	(2.57)	(0.01)	(0.02)	(2.02)	(0.01)	(0.02)
POST (March 30th-December 31st)	3.05	-0.00	0.07	2.66	0.00	0.05*
	(3.69)	(0.02)	(0.04)	(2.37)	(0.02)	(0.03)
TREAT (Year = 2008)	-43.50**	-0.20*	-0.28*	-29.23**	-0.13*	-0.24*
	(21.09)	(0.12)	(0.16)	(13.26)	(0.08)	(0.13)
HHI (1000s)	0.10	-0.00	-0.00	-0.67***	-0.00	-0.01***
	(0.20)	(0.00)	(0.00)	(0.15)	(0.00)	(0.00)
FINANCIAL CRISIS	3.17	0.01	0.03*	3.04*	0.01	0.04***
	(2.39)	(0.01)	(0.02)	(1.59)	(0.01)	(0.01)
Observations	234,597	234,597	234,597	235,975	235,975	235,975
R-squared	0.19	0.15	0.17	0.29	0.20	0.20
Time Trend	YES	YES	YES	YES	YES	YES
Airline-Route FE	YES	YES	YES	YES	YES	YES
Month-of-Year FE	YES	YES	YES	YES	YES	YES
Day-of-week FE	YES	YES	YES	YES	YES	YES
Hour-of-day FE	YES	YES	YES	YES	YES	YES
Weather Controls	YES	YES	YES	YES	YES	YES

Notes: Standard errors clustered at the daily level are reported in parentheses. The sampling period is January 1st, 2007 through December 31st, 2008.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

3.6.2 Short-Time Horizon (60-day window)

After controlling for several factors that influence flight delays such as weather, seasonality, and competition, the results in Tables 3.1 and 3.2 reveal that the implementation of slot controls at EWR and JFK in 2008 did not result in a measurable decrease in the average flight-level arrival or departure delay at both airports. Remarkably, the OTP of flights arriving into EWR declined during the slot control period. However, the slot control period at both airports coincided with the 2008 financial crisis. To alleviate concerns that results are being driven by the financial crisis, Tables 3.3 and 3.4 restrict the analysis sample to the symmetric 60-day window before and after slot controls were implemented at both airports and compares the OTP of those 2008 flights to comparable flights in 2007. For EWR, the analysis sample is restricted to all domestic flights from April 21st, 2007 to August 18th, 2007 and from April 21st, 2008 to August 18th, 2008. For JFK, the analysis sample is restricted to all domestic flights from January 30th, 2007 to May 28th, 2007 and from January 30th, 2008 to May 28th, 2008.

Columns (1)-(3) of Table 3.3 present results from equation (3.1) when analysis is performed on the restricted sample and outcomes of interest are arrival delay minutes, the probability that a flight arrives sixty or more minutes late, and the probability that a flight arrives fifteen or more minutes late. The coefficients on TREAT*POST (Slot Control) in columns (1)-(3) are positive, significant, and similar in magnitude to the Table 3.1 results. In column (1), the coefficient indicates that the average arrival delay at EWR increased by over nine minutes once slot controls were enforced. In columns (2) and (3), the probability that a flight arrives sixty or more minutes late increased by 4% while the probability that a flight arrives fifteen or more minutes late increased by 9% during the slot control period.

The results presented in columns (4)-(6) of Table 3.3 are also consistent with columns (4)-(6) of Table 3.1. Although insignificant, the coefficient on TREAT*POST (Slot Control) in column (4) indicates that the average departure delay increased by over three minutes after slot controls were implemented while in column (5), the coefficient suggests that the probability that a flight departs from EWR sixty or more minutes late increased by 2% during the slot control period. In column (6), the coefficient on TREAT*POST (Slot Control) is significant at the 5% level and indicates that the probability that a flight departs EWR fifteen or more minutes late increased by 4% during the slot control period.

Table 3.3: Short-Time Horizon Regression Results for Newark
(June 20th, 2007 through August 18th, 2007 Counterfactual)

	Arriving Flights			Departing Flights		
Dependent Variable:	Delay Minutes	Delay > 60	Delay > 15	Delay Minutes	Delay > 60	Delay > 15
	(1)	(2)	(3)	(4)	(5)	(6)
TREAT*POST (Slot Control)	9.28**	0.04*	0.09***	3.09	0.02	0.04**
	(4.33)	(0.02)	(0.03)	(2.12)	(0.01)	(0.02)
POST (June 20th-August 19th)	-1.75	-0.00	-0.04	0.39	-0.01	0.01
	(4.96)	(0.03)	(0.03)	(2.26)	(0.01)	(0.02)
TREAT (Year = 2008)	-68.56	-0.22	-0.30	-71.71***	-0.30*	-0.56**
	(54.12)	(0.31)	(0.37)	(26.33)	(0.17)	(0.24)
HHI (1000s)	-2.26***	-0.01	-0.00	0.78*	0.00	0.01
	(0.79)	(0.00)	(0.01)	(0.47)	(0.00)	(0.01)
Observations	94,060	94,060	94,060	94,839	94,839	94,839
R-squared	0.28	0.24	0.26	0.45	0.32	0.29
Time Trend	YES	YES	YES	YES	YES	YES
Airline-Route FE	YES	YES	YES	YES	YES	YES
Month-of-Year FE	YES	YES	YES	YES	YES	YES
Day-of-week FE	YES	YES	YES	YES	YES	YES
Hour-of-day FE	YES	YES	YES	YES	YES	YES
Weather Controls	YES	YES	YES	YES	YES	YES

Notes: Standard errors clustered at the daily level are reported in parentheses. The sampling period is April 21st, 2007 to August 18th, 2007 and from April 21st, 2008 to August 18th, 2008.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Columns (1)-(3) of Table 3.4 present results from equation (3.2) when analysis is performed on the restricted sample and outcomes of interest are arrival delay minutes, the probability that a flight arrives sixty or more minutes late, and the probability that a flight arrives fifteen or more minutes late. Consistent with Table 3.2, the coefficients on TREAT*POST (Slot Control) are insignificant in columns (1)-(3). However, in contrast to Table 3.2, the estimated effects are negative, suggesting that the implementation of slot controls decreased the incidence and severity of arrival delays at JFK. For instance, the coefficient in column (1) suggests that the average arrival delay decreased by over two and a half minutes during the slot control period while the coefficient in column (3) suggests that the probability that a flight arrives fifteen or more minutes late decreased by 3% once slot controls were enforced.

Columns (4)-(6) of Table 3.4 present results from equation (3.2) when the outcomes of interest are departure delay minutes, the probability that a flight departs sixty or more minutes late, and the probability that a flight departs fifteen or more minutes late. Consistent with columns (4)-(6) of Table 3.2, the coefficients on TREAT*POST (Slot Control) are all insignificant, suggesting that the implementation of slot controls at JFK in 2008 did not have any measurable impact on departure delays.

Table 3.4: Short-Time Horizon Regression Results for JFK
(March 30th, 2007 through May 29th, 2007 Counterfactual)

	Arriving Flights			**Departing Flights**		
Dependent Variable:	Delay Minutes	Delay > 60	Delay > 15	Delay Minutes	Delay > 60	Delay > 15
	(1)	(2)	(3)	(4)	(5)	(6)
TREAT*POST (Slot Control)	-2.78	-0.02	-0.03	1.09	0.00	-0.00
	(2.98)	(0.02)	(0.03)	(2.36)	(0.01)	(0.02)
POST (March 30th-May 29th)	4.98	0.02	0.10*	3.57	0.01	0.06*
	(4.18)	(0.03)	(0.05)	(3.09)	(0.02)	(0.03)
TREAT (Year = 2008)	-42.09	-0.11	-0.12	-23.73	-0.05	-0.10
	(34.62)	(0.22)	(0.28)	(26.81)	(0.16)	(0.23)
HHI (1000s)	-1.04***	-0.01***	-0.01***	-1.18***	-0.00	-0.02***
	(0.33)	(0.00)	(0.00)	(0.31)	(0.00)	(0.00)
Observations	78,378	78,378	78,378	78,692	78,692	78,692
R-squared	0.19	0.15	0.17	0.26	0.19	0.21
Time Trend	YES	YES	YES	YES	YES	YES
Airline-Route FE	YES	YES	YES	YES	YES	YES
Month-of-Year FE	YES	YES	YES	YES	YES	YES
Day-of-week FE	YES	YES	YES	YES	YES	YES
Hour-of-day FE	YES	YES	YES	YES	YES	YES
Weather Controls	YES	YES	YES	YES	YES	YES

Notes: Standard errors clustered at the daily level are reported in parentheses. The sampling period is January 30th, 2007 to May 28th, 2007 and from January 30th, 2008 to May 28th, 2008.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

3.7 Discussion

Collectively, the results presented in Tables 3.1, 3.2, 3.3, and 3.4 indicate that slot controls were not effective at reducing arrival and departure delays at EWR and JFK. Notably, the establishment of slot controls at EWR led to a decline in the on-time performance (OTP) of arriving flights. This section investigates several mechanisms that may explain why slot controls were not effective at reducing congestion and delays at EWR and JFK. These mechanisms include a decrease in scheduled flight time, an increase in market concentration, and changes in the length of arrival and departure banks of flights.

3.7.1 Scheduled Flight Time and Schedule Padding

As mentioned in Section 3.4, airlines with at least 1% of domestic scheduled passenger revenues have been required to submit monthly information to the U.S. Department of Transportation about the OTP of their flights. The goal of this program has been to provide consumers with better information about expected OTP while also creating incentives for airlines to invest in their OTP (Forbes et al., 2019b). However, the DOT considers a flight "on-time" if it arrives at the gate less than fifteen minutes after its scheduled arrival time. As discussed in Forbes et al. (2019b), this definition may result in "schedule padding" behavior. Because the DOT measures delays by comparing a flight's actual arrival time to its scheduled arrival time, an airline can increase the likelihood that a flight arrives "on-time" by increasing the amount of time scheduled for the flight.

With regard to slot constraints, if airlines perceive that slot controls will be effective at reducing airport congestion, airlines may respond by decreasing the scheduled time of their flights. If this behavior occurs, then it is plausible that implementing slot controls would result in no change to published OTP as the potential gains from slot controls are captured by the decrease in scheduled flight time. To determine if this "reverse schedule padding" behavior occurred at EWR and/or JFK, the following equation is estimated,

$$ScheduledTime_{ijkt} = \beta_0 + \beta_1 \cdot \mathrm{TREAT}_t + \beta_2 \cdot \mathrm{POST}_t + \beta_3 \cdot \mathrm{TREAT}_t \cdot \mathrm{POST}_t + \beta_4 \cdot \mathrm{HHI}_{jt}$$

$$+\beta_5 \cdot \text{FINANCIAL CRISIS}_t + \gamma \cdot \text{X}_{ijkt} + \epsilon_{ijkt} \qquad (3.3),$$

where $ScheduledTime_{ijkt}$ is the scheduled time in minutes of individual flight i on directional route j operated by carrier k on date t. X is a matrix containing aircraft, airline-route, month-of-year, day-of-week, and hour-of-day fixed effects in addition to a linear time trend. Individual aircraft fixed effects are included since different plane types fly at different speeds while hour-of-day fixed effects are included to control for variation in airport traffic and congestion levels throughout a typical day.[23] Because delays tend to propagate through the day, these hour-of-day fixed effects also capture if flights scheduled for later in the day are allocated more time. Standard errors are clustered at the airline-route level.

Column (1) of Table 3.5 presents results from equation (3.3) when the analysis sample contains all EWR domestic arrivals between January 1st, 2007 and December 31st, 2008. The negative and significant coefficient on TREAT*POST (Slot Control) indicates that on average, the scheduled time of flights arriving into EWR decreased by over two minutes during the slot control period. To determine if there was a similar change in the actual time it takes to complete a flight, column (2) presents results when the dependent variable is actual elapsed time in minutes instead of scheduled time. The significant positive coefficient on TREAT*POST (Slot Control) in column (2) reveals that on average, the actual time it takes to complete a EWR arriving flight increased by over one and a half minutes after slot controls went into effect. Together, the column (1) and (2) results support the decline in OTP of EWR arrivals observed in Tables 3.1 and 3.3. The decline in OTP appears to be partially due to a decrease in the scheduled time alloted for flights arriving into EWR.

The column (1)-(2) results also reveal a potential unintended consequence of how slots are allocated at capacity constrained airports. Because takeoff and landing slots were allocated by the FAA in thirty minute intervals at EWR, the decrease in scheduled time of arriving flights could be a result of the allocation process. For instance, consider an airline that is

[23] Unique aircraft are identified by tail number in the OTP dataset.

awarded a landing slot between 2:00pm-2:29pm but not one between 2:30pm-2:59pm. To improve the OTP performance of an arriving flight, the airline may wish to set the arrival time for 2:35pm. However, because the airline was not allocated an arrival slot during the 2:30pm-2:59pm interval, the airline would set the scheduled arrival time for 2:29pm.

Column (3) of Table 3.5 presents results from equation (3.3) when the analysis sample contains all EWR domestic departures between January 1st, 2007 and December 31st, 2008. The significant coefficient on TREAT*POST (Slot Control) indicates that on average, the scheduled time for flights departing from EWR decreased by one and half minutes during the slot control period. However, when the dependent variable is changed to actual elapsed time in column (4), there does not appear to be any statistically significant difference in the actual time it takes to complete a flight departing from EWR between the pre-slot and slot control periods. Regardless, the decrease in scheduled time observed in column (3) is consistent with the results in Tables 3.1 and 3.3 that find no significant reduction in departure delays at EWR during the slot control period.

Table 3.5: Scheduled and Actual Time Regressions for EWR and JFK (2007 Counterfactual)

	EWR Flights				JFK Flights			
	Arrivals		Departures		Arrivals		Departures	
Dependent Variable:	SchedTime	ActualTime	SchedTime	ActualTime	SchedTime	ActualTime	SchedTime	ActualTime
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
TREAT*POST (Slot Control)	-2.24***	1.59***	-1.52***	0.30	-0.24	-1.41***	-3.62***	1.38**
	(0.37)	(0.40)	(0.44)	(0.32)	(0.44)	(0.52)	(0.48)	(0.56)
POST	1.84***	-1.03**	0.87***	-3.23***	0.27	2.70***	1.16***	-3.49***
	(0.20)	(0.44)	(0.26)	(0.53)	(0.28)	(0.60)	(0.30)	(0.72)
TREAT (Year = 2008)	0.82	7.06***	1.80***	-2.15	-0.55	-6.86***	4.03***	12.12***
	(0.65)	(1.79)	(0.55)	(1.76)	(0.88)	(1.55)	(0.80)	(2.30)
HHI (1000s)	0.13	0.12	-0.25	0.27	-0.58**	-0.25	0.15	-0.16
	(0.46)	(0.41)	(0.56)	(0.48)	(0.28)	(0.24)	(0.33)	(0.32)
FINANCIAL CRISIS	-1.63***	-0.85*	1.50***	-1.11***	-2.89***	2.16***	0.08	-6.07***
	(0.45)	(0.45)	(0.37)	(0.37)	(0.48)	(0.58)	(0.61)	(0.89)
Observations	283,150	280,734	283,384	281,525	240,623	236,882	240,691	237,381
R-squared	0.99	0.93	0.99	0.95	0.99	0.96	0.99	0.96
Time Trend	YES	YES	YES	YES	YES	YES	YES	YES
Aircraft FE	YES	YES	YES	YES	YES	YES	YES	YES
Airline-Route FE	YES	YES	YES	YES	YES	YES	YES	YES
Month-of-Year FE	YES	YES	YES	YES	YES	YES	YES	YES
Day-of-week FE	YES	YES	YES	YES	YES	YES	YES	YES
Hour-of-day FE	YES	YES	YES	YES	YES	YES	YES	YES

Notes: Standard errors clustered at the airline-route level are reported in parentheses. The sampling period is January 1st, 2007 through December 31st, 2008. In columns (1)-(4), POST is an indicator equal to one if the day of the year is on or after June 20th. In columns (5)-(8), POST is an indicator equal to one if the day of the year is on or after March 30th.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Column (5) of Table 3.5 presents results from equation (3.3) when the analysis sample contains all JFK domestic arrivals between January 1st, 2007 and December 31st, 2008. The insignificant negative coefficient on TREAT*POST (Slot Control) indicates that there was no statistically significant difference in the scheduled time of JFK arriving flights between the pre-slot and slot control periods. To determine if there was a change in the actual time it takes to complete a flight, column (6) presents results when the dependent variable is actual elapsed time instead of scheduled time. The significant negative coefficient on TREAT*POST (Slot Control) in column (6) reveals that on average, the actual time it takes to complete a JFK arriving flight decreased by over one minute once slot controls went into effect. Nonetheless, the results in columns (5)-(6) are generally consistent with the results from Tables 3.2 and 3.4 that find no statistically significant improvement in the OTP of JFK arriving flights in the presence of slot controls.

Column (7) of Table 3.5 presents results from equation (3.3) when the analysis sample contains all JFK domestic departures between January 1st, 2007 and December 31st, 2008. The significant coefficient on TREAT*POST (Slot Control) in column (7) indicates that on average, the scheduled time for flights departing from JFK decreased by over three and a half minutes during the slot control period while the coefficient in column (8) indicates that the actual time it takes to complete a flight that departs from JFK increased by over one minute once slot controls went into effect. This increase in actual time coupled with the decrease in scheduled time are consistent with the results from Tables 3.2 and 3.4 that find no statistically significant improvement in the OTP of JFK departing flights.

Overall, the results in Table 3.5 reveal that a decrease in scheduled flight times helps explain why OTP did not improve at EWR and JFK after slot controls were introduced. In particular, a decrease in scheduled time combined with an increase in the actual time it takes to complete an arriving flight likely explains why the OTP of flights arriving into EWR declined after slot controls were implemented (see columns (1)-(2)). Similarly, a decrease in scheduled time and an increase in actual time also helps explain why a reduction in departure delays was not observed at JFK (see columns (7)-(8)).

3.7.2 Competition

To determine if an increase in market concentration may explain why OTP did not improve once slot controls were implemented at EWR and JFK, the following equation is estimated,

$$ln(HHI)_{jt} = \beta_0 + \beta_1 \cdot \text{TREAT}_t + \beta_2 \cdot \text{POST}_t + \beta_3 \cdot \text{TREAT}_t \cdot \text{POST}_t + \beta_4 \cdot \text{FINANCIAL CRISIS}_t +$$

$\gamma_j + \delta_t + \epsilon_{jt}$ (3.4),

where $ln(HHI)_{jt}$ is the natural logarithm of the Herfindahl-Hirschman Index on directional route j in month t. γ is a route fixed effect while δ_t is a month-of-year fixed effect. Standard errors are clustered at the route level.

Columns (1) and (2) of Table 3.6 present results from equation (3.4) when the analysis sample is restricted to EWR arriving and departing routes, respectively. The marginally significant coefficient in the first two columns indicate that market concentration decreased by 1% on average after slot controls went into effect at EWR. Therefore, the decline in OTP of flights arriving into EWR does not appear to be driven by an increase in market concentration during the slot control period.

Table 3.6: HHI Regressions for EWR and JFK
(2007 Counterfactual)

	EWR		**JFK**	
	Arrivals	Departures	Arrivals	Departures
Dependent Variable:	ln(HHI)	ln(HHI)	ln(HHI)	ln(HHI)
	(1)	(2)	(3)	(4)
TREAT*POST (Slot Control)	-0.01*	-0.01*	0.03	0.02
	(0.01)	(0.01)	(0.03)	(0.03)
POST	0.11**	0.09	-0.29	-2.27
	(0.05)	(0.06)	(1.06)	(1.38)
TREAT (Year = 2008)	0.66	0.21	-0.27	-1.05
	(0.57)	(0.62)	(1.68)	(1.22)
FINANCIAL CRISIS	0.02	0.03	0.02	0.02
	(0.01)	(0.02)	(0.02)	(0.02)
Observations	2,021	2,025	1,448	1,433
R-squared	0.97	0.97	0.92	0.92
Time Trend	YES	YES	YES	YES
Route FE	YES	YES	YES	YES
Month-of-Year FE	YES	YES	YES	YES

Notes: Standard errors clustered at the route level are reported in parentheses. The sampling period is January, 2007 through December, 2008. In columns (1)-(2), POST is an indicator equal to one if the month is July-December. In columns (3)-(4), POST is an indicator equal to one if the month is April-December.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

The last two columns of Table 3.6 present results from equation (3.4) when the analysis sample is restricted to JFK arriving and departing routes, respectively. Although insignificant, the positive coefficients on TREAT*POST (Slot Control) in columns (3) and (4) suggest that market concentration increased by 2-3% at JFK during the slot control period. If market concentration increased at JFK, then finding that the OTP performance of flights did not improve after slot controls were implemented (as the results in Tables 3.2 and 3.4 indicate) is not surprising.[24]

3.7.3 Length of Arrival and Departure Banks

Studying the relationship between airport concentration and the length of high-volume periods (known as flight banks), Ater (2012) finds that hub airlines choose longer arrival and departure banks as their share of flights in the bank increases. These longer flight banks are also shown to be associated with shorter flight delays, indicating that hub airlines internalize the congestion they impose on their own flights. With respect to delay and congestion management, the findings in Ater (2012) suggest that policies aimed at reducing congestion at highly concentrated airports will only have a limited impact because dominant airlines already internalize congestion. The key finding in this chapter (that the implementation of slot controls at EWR and JFK did not result in a decrease in arrival and departure delays at both airports) supports this expectation. At the time slot controls were enacted, Continental accounted for 72% of the flight traffic at EWR while Delta controlled 31% of the traffic at JFK.

The results in Ater (2012) also provide a potential explanation for why flight delays may increase under a slot control scheme at highly concentrated airports. For example, if the allocation of slots results in the hub airline reducing the length of its arrival and departure

[24] The overall impact of market concentration on delays is inconclusive at JFK. Depending on specification, the sign of HHI is either positive or negative in Tables 3.2 and 3.4. Moreover, the HHI coefficients are also insignificant in most specifications.

banks, then delays may increase. To determine if this behavior occurred at EWR and/or JFK, the following equation is estimated,

$$\text{BANK LENGTH}_{it} = \beta_0 + \beta_1 \cdot \text{TREAT}_t + \beta_2 \cdot \text{POST}_t + \beta_3 \cdot \text{TREAT}_t \cdot \text{POST}_t + \beta_4 \cdot \text{BANK FLIGHTS}_{it} + \beta_5 \cdot \text{FINANCIAL CRISIS}_t + \gamma \cdot X_t + \epsilon_{it} \qquad (3.5),$$

where BANK LENGTH$_{it}$ is the length in minutes of arrival or departure bank i on date t.[25] BANK FLIGHTS is the total number of flights scheduled to operate during bank i on date t while X is a matrix containing month-of-year and day-of-week fixed effects in addition to a linear time trend. Standard errors are clustered at the daily level.

[25] To identity the length of arrival and departure banks, the procedure outlined in Ater (2012) is followed. The identification of flight banks relies on flights operated by the relevant hub carrier only (Continental for EWR and Delta for JFK). To identify arriving banks, the number of arriving flights for each minute on every day is derived for EWR and JFK. Then, a moving average of the number of flights arriving at the airport in a 21-minute time window is computed. This moving average is used to compute a daily moving average and its standard deviation. An arriving bank threshold is defined as the moving average at one standard deviation above the daily moving average. An arriving bank period occurs when the 21-min moving average is higher than the threshold. The structure of departing banks is derived in a similar manner.

Table 3.7: Bank Length Regressions for EWR and JFK
(2007 Counterfactual)

	EWR		JFK	
Dependent Variable:	Arrival Bank Length (Minutes)	Departure Bank Length (Minutes)	Arrival Bank Length (Minutes)	Departure Bank Length (Minutes)
	(1)	(2)	(3)	(4)
TREAT*POST (Slot Control)	0.03	0.26	-0.02	-1.65***
	(0.18)	(0.20)	(0.40)	(0.35)
POST	-0.65**	-0.54**	0.12	0.65
	(0.27)	(0.22)	(1.32)	(1.05)
TREAT (Year = 2008)	-2.67	0.99	7.01**	3.34
	(2.04)	(2.21)	(2.90)	(3.18)
BANK FLIGHTS	1.60***	1.21***	1.37***	1.36***
	(0.01)	(0.01)	(0.01)	(0.01)
FINANCIAL CRISIS	0.63**	1.83***	-0.44	0.79***
	(0.26)	(0.27)	(0.29)	(0.30)
Observations	8,005	5,989	4,279	3,773
R-squared	0.90	0.96	0.85	0.83
Time Trend	YES	YES	YES	YES
Month-of-Year FE	YES	YES	YES	YES
Day-of-Week FE	YES	YES	YES	YES

Notes: Standard errors clustered at the daily level are reported in parentheses. The sampling period is January 1st, 2007 through December 31st, 2008. In columns (1)-(2), POST is an indicator equal to one if the day of the year is on or after June 20th. In columns (3)-(4), POST is an indicator equal to one if the day of the year is on or after March 30th.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Columns (1) and (2) of Table 3.7 present regression results for EWR when equation (3.5) is estimated separately for arrival and departure banks. The insignificant positive coefficients on TREAT*POST (Slot Control) in the first two columns indicate that the length of arrival and departure banks at EWR did not significantly decrease after slot controls were enacted. Therefore, it does not appear that a decrease in the length of Continental's arrival or departure banks explains the increase in arrival delays found in Tables 3.1 and 3.3.

The last two columns of Table 3.7 present results when equation (3.5) is estimated separately for JFK's arrival and departure banks. The negative and significant coefficient on

TREAT*POST (Slot Control) in column (4) indicates that the length of Delta's departure banks decreased by over one and a half minutes on average after slot controls were implemented at JFK. Although the goal of slot controls at JFK were to reduce the incidence and severity of departure delays (relative to summer 2007 levels), it appears that a decrease in the length of departure banks at JFK helps explain why slot controls at JFK did not accomplish their primary goal.

Of the three mechanisms that potentially explain why OTP did not improve at both airports, it appears that a reduction in scheduled flight times is most likely responsible as opposed to an increase in market concentration or a decrease in the length of arrival and departure flight banks (with the exception of the length of Delta's departure banks at JFK). Airlines likely decreased scheduled flight times for one of two reasons. If airlines perceive that slot controls are effective at reducing congestion, then airlines may respond to slot controls by decreasing flight times. If this behavior occurred, then any potential gains from slot controls would be captured by the reduction in flight times. Alternatively, the decrease in scheduled times may result from the slot allocation process. For example, an airline awarded a landing slot between 2:00pm-2:29pm but not one between 2:30pm-2:59pm would be forced to set the arrival time of a flight for 2:29pm even if it wished to schedule the arrival for 2:35pm to improve the flight's expected OTP.

3.8 Conclusion

This chapter has examined whether slot controls implemented at John F. Kennedy (JFK) on March 30th, 2008 and Newark (EWR) on June 20th, 2008 were effective at reducing the incidence and severity of flight delays at both airports. At JFK, slot controls were implemented with the primary goal of reducing peak evening departure delays from the summer 2007 average of 80 minutes. At EWR, the goal was to ensure that arrival and departure delays did not increase relative summer 2007 levels. Using flight-level on-time performance (OTP) data for the universe of domestic flights between January 1st, 2007 and December 31st, 2008, these controls are found to have not been effective at reducing the incidence or severity of arrival and departure delays at both airports. Surprisingly, the implementation of these controls resulted in a 7 minute increase in the average arrival delay at EWR (6% increase in the probability that a flight arrives late). These results support Ater (2012), who suggested that policies aimed at reducing congestion and delays at highly concentrated airports will only have a limited impact because dominant airlines already internalize congestion. At the time slot controls were implemented, Delta accounted for 31% of flight traffic at JFK while Continental controlled 72% of flights at EWR.

Examining three potential mechanisms that may explain why slot controls were not effective, the length of Delta's departure banks at JFK were found to have decreased by about 2 minutes with no change in the length of Continental's arrival or departure banks at EWR. Further, the scheduled time for flights arriving into or departing from EWR decreased by 1.5-2.2 minutes while the scheduled time for flights departing from JFK decreased by over 3.5 minutes during the slot control period. Finally, no evidence of an increase in market concentration is found at both airports as a result of the slot allocation process.

Overall, it appears that OTP did not improve at both airports because airlines responded to slot controls by reducing scheduled flight times. As discussed in Section 3.7.1, airlines likely decreased scheduled times for one of two reasons. If airlines perceived that these controls would be effective at reducing congestion, then airlines could have responded by decreasing scheduled fight times. This explanation suggests that any gains from slot controls were captured by the decrease in scheduled times. Alternatively, because slots were allocated by the FAA in thirty (fifteen) minute intervals at EWR (JFK), the decrease in scheduled times could be a result of the slot allocation process. For instance, consider an airline awarded a landing slot between 2:00pm-2:29pm but not one between 2:30pm-2:59pm. To improve OTP, the airline may wish to schedule the arrival for 2:35pm. However, because the airline was not awarded a 2:30pm-2:59pm arrival slot, the airline would schedule the flight to arrive at 2:29pm.

In a future where passenger demand is forecast to surpass current network and airport capacity, the results in this chapter highlight the need for policymakers to carefully consider how the allocation of takeoff and landing slots will impact scheduling decisions and market concentration when implementing similar policies. Foremost, policymakers should be aware that attaching specific time intervals to airport slots may impact the internalizing behavior of dominant carriers by altering the length of the carrier's arrival and departure banks in addition to scheduled flight times. Second, because OTP has been shown to decline with market concentration, policymakers should also consider how the allocation of slots will impact competition. These concerns would be mitigated if policymakers were to switch to a congestion-based price schedule in lieu of slot controls. However, congestion pricing may not be politically feasible.

More empirical evidence is needed on the effectiveness of slot controls. Future research could focus on European airports where slot controls are more prevalent. The key issue is identifying an appropriate counterfactual airport for use in a differences-in-differences (DD) design. While this chapter relied on a DD model with flights at the same airport in the year prior to slot control enactment serving as the counterfactual, other approaches that could be applied in the future include event study and synthetic control methods.

Bibliography

Abramowitz, A. D. and Brown, S. M. (1993). Market share and price determination in the contemporary airline industry. *Review of Industrial Organization*, 8(4):419–433.

Ackerberg, D. A. and Rysman, M. (2005). Unobserved product differentiation in discrete-choice models: Estimating price elasticities and welfare effects. *The Rand Journal of Economics*, 36(4):771.

Alperovich, G. and Machnes, Y. (1994). The role of wealth in the demand for international air travel. *Journal of Transport Economics and Policy*, pages 163–173.

Ater, I. (2012). Internalization of congestion at US hub airports. *Journal of Urban Economics*, 72(2-3):196–209.

Ater, I. and Orlov, E. (2015). The effect of the internet on performance and quality: Evidence from the airline industry. *Review of Economics and Statistics*, 97(1):180–194.

Bailey, E. E., Graham, D. R., and Kaplan, D. P. (1985). *Deregulating the airlines*, volume 10. MIT press.

Bajari, P. and Benkard, C. L. (2005). Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach. *Journal of Political Economy*, 113(6):1239–1276.

Belobaba, P., Odoni, A., and Barnhart, C. (2009). Fundamentals of pricing and revenue management. *The Global Airline Industry*, page 73.

Belobaba, P., Odoni, A., and Barnhart, C. (2015). *The global airline industry.* John Wiley & Sons.

Berry, S., Carnall, M., and Spiller, P. T. (2006). Airline hubs: Costs, markups and the implications of customer heterogeneity. *Competition policy and antitrust.*

Bhadra, D. (2012). Disappearance of american wealth and its impact on air travel: An empirical investigation. In *Journal of the Transportation Research Forum*, volume 51.

Bilotkach, V. (2005). Understanding price dispersion in the airline industry: Capacity constraints and consumer heterogeneity. *Unpublished Manuscript.*

Bilotkach, V., Gorodnichenko, Y., and Talavera, O. (2010). Are airlines' price-setting strategies different? *Journal of Air Transport Management*, 16(1):1–6.

Borenstein, S. (1985). Price discrimination in free-entry markets. *The RAND Journal of Economics*, pages 380–397.

Borenstein, S. (1989). Hubs and high fares: Dominance and market power in the US airline industry. *The RAND Journal of Economics*, pages 344–365.

Borenstein, S. (1991). The dominant-firm advantage in multiproduct industries: Evidence from the US airlines. *The Quarterly Journal of Economics*, 106(4):1237–1266.

Borenstein, S. and Rose, N. L. (1994). Competition and price dispersion in the US airline industry. *Journal of Political Economy*, 102(4):653–683.

Britto, R., Dresner, M., and Voltes, A. (2012). The impact of flight delays on passenger demand and societal welfare. *Transportation Research Part E: Logistics and Transportation Review*, 48(2):460–469.

Brueckner, J. K. (2002a). Airport congestion when carriers have market power. *American Economic Review*, 92(5):1357–1375.

Brueckner, J. K. (2002b). Internalization of airport congestion. *Journal of Air Transport Management*, 8(3):141–147.

Brueckner, J. K. (2004). Network structure and airline scheduling. *The Journal of Industrial Economics*, 52(2):291–312.

Brueckner, J. K. (2005). Internalization of airport congestion: A network analysis. *International Journal of Industrial Organization*, 23(7):599–614.

Brueckner, J. K., Lee, D., and Singer, E. (2014). City-pairs versus airport-pairs: A market-definition methodology for the airline industry. *Review of Industrial Organization*, 44(1):1–25.

Brueckner, J. K., Lee, D., and Singer, E. S. (2013). Airline competition and domestic US airfares: A comprehensive reappraisal. *Economics of Transportation*, 2(1):1–17.

Brueckner, J. K. and Lin, M. H. (2016). Convenient flight connections vs. airport congestion: Modeling the rolling hub. *International Journal of Industrial Organization*, 48:118 – 142.

Brueckner, J. K. and Spiller, P. T. (1991). Competition and mergers in airline networks. *International Journal of Industrial Organization*, 9(3):323–342.

Brueckner, J. K. and Spiller, P. T. (1994). Economies of traffic density in the deregulated airline industry. *The Journal of Law and Economics*, 37(2):379–415.

Brueckner, J. K. and Van Dender, K. (2008). Atomistic congestion tolls at concentrated airports? Seeking a unified view in the internalization debate. *Journal of Urban Economics*, 64(2):288–295.

Caves, D. W., Christensen, L. R., and Tretheway, M. W. (1984). Economies of density versus economies of scale: Why trunk and local service airline costs differ. *The RAND Journal of Economics*, pages 471–489.

Ciliberto, F. and Williams, J. W. (2010). Limited access to airport facilities and market power in the airline industry. *The Journal of Law and Economics*, 53(3):467–495.

Dana, Jr, J. D. (1998). Advance-purchase discounts and price discrimination in competitive markets. *Journal of Political Economy*, 106(2):395–422.

Daniel, J. I. (1995). Congestion pricing and capacity of large hub airports: A bottleneck model with stochastic queues. *Econometrica: Journal of the Econometric Society*, pages 327–370.

Daniel, J. I. and Harback, K. T. (2008). (When) Do hub airlines internalize their self-imposed congestion delays? *Journal of Urban Economics*, 63(2):583–612.

Daniel, J. I. and Harback, K. T. (2009). Pricing the major US hub airports. *Journal of Urban Economics*, 66(1):33–56.

Donovan, A. W. (2005). Yield management in the airline industry. *Journal of Aviation/Aerospace Education & Research*, 14(3).

DOT (January 2015). Federal Register January 8, 2015. *Federal Register*, 80(5).

DOT (March 2008a). Federal Register March 18, 2008. *Federal Register*, 73(53).

DOT (May 2008b). Federal Register May 21, 2008. *Federal Register*, 73(99).

Dresner, M., Windle, R., and Yao, Y. (2002). Airport barriers to entry in the US. *Journal of Transport Economics and Policy (JTEP)*, 36(3):389–405.

Durrleman, S. and Simon, R. (1989). Flexible regression models with cubic splines. *Statistics in medicine*, 8(5):551–561.

Evans, W. N. and Kessides, I. N. (1993). Localized market power in the US airline industry. *The Review of Economics and Statistics*, pages 66–75.

FAA (2016). *Economic Values for the Evaluation of FAA Investment and Regulatory Decisions: A Revised Guide*. Federal Aviation Administration, Washington, DC.

FAA (2018). *FAA Aerospace Forecast: Fiscal Years 2018-2038*. Federal Aviation Administration, Washington, DC.

Fageda, X. and Flores-Fillol, R. (2017). Airport congestion and airline network structure. In *The Economics of Airport Operations*, pages 335–359. Emerald Publishing Limited.

Flores-Fillol, R. (2010). Congested hubs. *Transportation Research Part B: Methodological*, 44(3):358 – 370. Economic Analysis of Airport Congestion.

Forbes, S. J. (2008). The effect of air traffic delays on airline prices. *International journal of industrial organization*, 26(5):1218–1232.

Forbes, S. J., Lederman, M., and Wither, M. J. (2019a). Quality disclosure when firms set their own quality targets. *International Journal of Industrial Organization*, 62:228–250.

Forbes, S. J., Lederman, M., and Yuan, Z. (2019b). Do airlines pad their schedules? *Review of Industrial Organization*, pages 1–22.

Fukui, H. (2019). How do slot restrictions affect airfares? New evidence from the US airline industry. *Economics of Transportation*, 17:51–71.

Gale, I. L. and Holmes, T. J. (1993). Advance-purchase discounts and monopoly allocation of capacity. *The American Economic Review*, pages 135–146.

Gayle, P. G. and Wu, C.-Y. (2013). A re-examination of incumbents response to the threat of entry: Evidence from the airline industry. *Economics of Transportation*, 2(4):119–130.

Gerardi, K. S. and Shapiro, A. H. (2009). Does competition reduce price dispersion? New evidence from the airline industry. *Journal of Political Economy*, 117(1):1–37.

Gillen, D. and Mantin, B. (2009). Price volatility in the airline markets. *Transportation Research Part E: Logistics and Transportation Review*, 45(5):693–709.

Greenfield, D. (2014). Competition and service quality: New evidence from the airline industry. *Economics of Transportation*, 3(1):80–89.

Harrell Jr, F. E. (2015). *Regression modeling strategies: with applications to linear models, logistic and ordinal regression, and survival analysis.* Springer.

Hendricks, K., Piccione, M., and Tan, G. (1995). The economics of hubs: The case of monopoly. *The Review of Economic Studies*, 62(1):83–99.

Holmes, T. J. (1989). The effects of third-degree price discrimination in oligopoly. *The American Economic Review*, 79(1):244–250.

Johnson, T. and Savage, I. (2006). Departure delays, the pricing of congestion, and expansion proposals at Chicago O'Hare Airport. *Journal of Air Transport Management*, 12(4):182–190.

Lederman, M. (2008). Are frequent-flyer programs a cause of the hub premium? *Journal of Economics & Management Strategy*, 17(1):35–66.

Lee, D. and Luengo-Prado, M. J. (2005). The impact of passenger mix on reported "hub premiums" in the US airline industry. *Southern Economic Journal*, pages 372–394.

Levine, M. E. (1986). Airline competition in deregulated markets: Theory, firm strategy, and public policy. *Yale J. on Reg.*, 4:393.

Luttmann, A. (2019). Are passengers compensated for incurring an airport layover? Estimating the value of layover time in the US airline industry. *Economics of Transportation*, 17:1–13.

Mantin, B. and Koo, B. (2009). Dynamic price dispersion in airline markets. *Transportation Research Part E: Logistics and Transportation Review*, 45(6):1020–1029.

Mayer, C. and Sinai, T. (2003). Network effects, congestion externalities, and air traffic delays: Or why not all delays are evil. *American Economic Review*, 93(4):1194–1215.

Mazzeo, M. J. (2003). Competition and service quality in the US airline industry. *Review of industrial Organization*, 22(4):275–296.

McCulloch, R. E., Polson, N. G., and Rossi, P. E. (2000). A bayesian analysis of the multinomial probit model with fully identified parameters. *Journal of Econometrics*, 99(1):173 – 193.

Miranda, V. A. and Oliveira, A. V. (2018). Airport slots and the internalization of congestion by airlines: An empirical model of integrated flight disruption management in brazil. *Transportation Research Part A: Policy and Practice*, 116:201–219.

Morrison, S. A. (2001). Actual, adjacent, and potential competition: Estimating the full effect of Southwest Airlines. *Journal of Transport Economics and Policy (JTEP)*, 35(2):239–256.

Morrison, S. A. and Winston, C. (1989). Enhancing the performance of the deregulated air transportation system. *Brookings Papers on Economic Activity. Microeconomics*, 1989:61–123.

Morrison, S. A. and Winston, C. (1995). *The evolution of the airline industry*. Brookings Institution Press.

Netessine, S. and Shumsky, R. (2002). Introduction to the theory and practice of yield management. *INFORMS Transactions on Education*, 3(1):34–44.

Pai, V. (2010). On the factors that affect airline flight frequency and aircraft size. *Journal of Air Transport Management*, 16(4):169–177.

Pels, E. and Verhoef, E. T. (2004). The economics of airport congestion pricing. *Journal of Urban Economics*, 55(2):257–277.

Peteraf, M. A. and Reed, R. (1994). Pricing and performance in monopoly airline markets. *The Journal of Law and Economics*, 37(1):193–213.

Peterson, E. B., Neels, K., Barczi, N., and Graham, T. (2013). The economic cost of airline flight delay. *Journal of Transport Economics and Policy (JTEP)*, 47(1):107–121.

Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110(4):705–729.

Poirier, D. J. (1976). The econometrics of structural change; with special emphasis on spline functions. Technical report.

Prince, J. T. and Simon, D. H. (2009). Multimarket contact and service quality: Evidence from on-time performance in the US airline industry. *Academy of Management Journal*, 52(2):336–354.

Prince, J. T. and Simon, D. H. (2014). Do incumbents improve service quality in response to entry? Evidence from airlines' on-time performance. *Management Science*, 61(2):372–390.

Prince, J. T. and Simon, D. H. (2017). The impact of mergers on quality provision: Evidence from the airline industry. *The Journal of Industrial Economics*, 65(2):336–362.

Puller, S. L. and Taylor, L. M. (2012). Price discrimination by day-of-week of purchase: Evidence from the US airline industry. *Journal of Economic Behavior & Organization*, 84(3):801–812.

Rupp, N., Owens, D., Plumly, L., et al. (2006). Does competition influence airline on-time performance. *Advances in airline economics*, 1:251–272.

Rupp, N. G. (2009). Do carriers internalize congestion costs? Empirical evidence on the internalization question. *Journal of Urban Economics*, 65(1):24–37.

Rupp, N. G. and Holmes, G. M. (2006). An investigation into the determinants of flight cancellations. *Economica*, 73(292):749–783.

Snider, C. and Williams, J. W. (2015). Barriers to entry in the airline industry: A multidimensional regression-discontinuity analysis of AIR-21. *Review of Economics and Statistics*, 97(5):1002–1022.

Stavins, J. (2001). Price discrimination in the airline market: The effect of market concentration. *Review of Economics and Statistics*, 83(1):200–202.

Stole, L. A. (2007). Price discrimination and competition. *Handbook of industrial organization*, 3:2221–2299.

Stone, C. J. and Koo, C.-Y. (1986). Additive splines in statistics. In *Proceedings of the statistical computing section*, volume 45, page 48.

Talluri, K. T. and Van Ryzin, G. J. (2006). *The theory and practice of revenue management*, volume 68. Springer Science & Business Media.

Theis, G., Adler, T., Clarke, J.-P., and Ben-Akiva, M. (2006). Risk aversion to short connections in airline itinerary choice. *Transportation Research Record: Journal of the Transportation Research Board*, 1951:28–36.

Van Dender, K. et al. (2007). Determinants of fares and operating revenues at US airports. *Journal of Urban Economics*, 62(2):317–336.

Yan, J. and Winston, C. (2014). Can private airport competition improve runway pricing? The case of San Francisco Bay area airports. *Journal of Public Economics*, 115:146–157.

Zou, B. and Hansen, M. (2014). Flight delay impact on airfare and flight frequency: A comprehensive assessment. *Transportation Research Part E: Logistics and Transportation Review*, 69:54–74.

Appendix A

Description of DB1B Data Processing

As mentioned in Section 1.2, ticket and price data for all four quarters of 2015 are taken from the US Department of Transportation's Airline Origin and Destination Survey (referred to as database DB1B). Data from this survey are released quarterly and generated from a 10 percent random sample of all airline tickets that originate in the United States on U.S. based carriers.

To produce the final analysis sample, several processing steps are applied. Foremost, only tickets involving direct or nonstop roundtrip travel are included.[1] Tickets involving a connection are excluded to directly compare how legacy and low-cost carriers differ in their directional pricing of nonstop roundtrip flights. This restriction is imposed because legacy and low-cost carriers differ in their network structures, with legacy carriers operating hub-and-spoke networks that typically require passengers originating and concluding travel in non-hub cities to board a connecting flight at a hub on the way to the final destination. In contrast, low-cost carriers generally deemphasize hub-and-spoke operations, instead operating point-to-point networks.

[1] One-way tickets are excluded. In the 2015 DB1B data, 74.39% of all passengers are roundtrip passengers. Moreover, 35.74% of all passengers purchase roundtrips that are nonstop in both directions (these types of tickets have two coupons in the DB1B data).

Second, similar to Brueckner et al. (2013), roundtrip fares less than $25 are excluded since these fares are assumed to be frequent flier tickets.[2] Third, only tickets involving travel in coach class (unrestricted or restricted) between airport-pairs in the U.S. are included. Tickets in first-class are excluded because these types of tickets entail a significantly different quality of service than travel in coach. Moreover, including first-class tickets in an average roundtrip fare calculation would upwardly bias an estimate of the average fare paid by most consumers on a route.[3]

Next, the operating carrier variable was recoded for regional carriers that are owned by or have a partnership with a major airline. To assign regional carriers to major airlines, a similar approach to the one outlined in Pai (2010) was undertaken. In most cases, the operating carrier was replaced with the ticketing carrier for tickets involving travel on a regional carrier. Table A.1 provides the regional carrier assignments that were determined after careful analysis of annual 10K reports filed with the Securities and Exchange Commission for the major airlines. These assignments were cross-checked with regional carrier route maps for accuracy.

[2] When redeeming frequent flier miles, a passenger must still pay associated ticket taxes and fees. For example, the September 11th security fee is $5.60 each way or $11.20 roundtrip.

[3] Including first-class tickets does not significantly change the results presented throughout this chapter.

Table A.1: Regional Carrier Assignments

Major Carrier	Regional Carrier
Alaska (AS)	Horizon Air*
	Peninsula Airways
	SkyWest Airlines
American (AA)	Air Wisconsin
	American Eagle
	Cape Air
	Compass Air
	ExpressJet
	Mesa Airlines[a]
	Piedmont Airlines*
	PSA Airlines*
	Republic Airlines
	SkyWest Airlines
	Trans States Airlines
Delta (DL)	Comair*
	Compass Air
	Endeavor Air*
	ExpressJet
	GoJet Airlines
	Shuttle America
	SkyWest Airlines
Hawaiian Airlines (HA)	Empire Airlines
United (UA)	Commutair
	ExpressJet
	GoJet Airlines
	Mesa Airlines[b]
	Republic Airlines
	Shuttle America
	SkyWest Airlines
	Trans States Airlines

Notes: Regional carriers that serve multiple major airlines are assigned in the DB1B data according to the ticketing carrier.
* Asterisks indicate that the regional carrier is owned by the major carrier.
a Only Routes involving Phoenix (PHX) or Dallas (DFW).
b Only routes involving Houston (IAH) or Washington, DC (IAD).

After assigning regional carriers to a major carrier, observations are then aggregated to produce an average fare. However, in the DB1B data, each itinerary contains a passenger count corresponding to the number of sampled passengers observed paying that particular fare for travel on a given carrier during the quarter. Therefore, different observed fares on a given carrier for travel in the same market generate separate observations. Using the passenger counts as weights, observations are aggregated to the origin-destination-carrier-quarter level to produce a weighted average fare. Resulting weighted average fares are then compared with the corresponding weighted average fare for roundtrip travel in the opposite direction on the same carrier. To prevent small routes from skewing results, an observation is included only if at least 100 sampled passengers are observed traveling in each direction during the quarter.[4] Given that the DB1B is a 10% sample of tickets sold in a quarter, this restriction implies that only nonstop routes where an airline carries at least 1,000 roundtrip passengers in each direction during the quarter are included.

Finally, to prevent double counting by including both ij and ji observations in the analysis, an observation is included only if the origin airport code comes after the destination airport code in the alphabet. This final restriction reduces the dataset to 7,426 observations encompassing 1,452 city-pair markets.

[4] This restriction is similar to the one employed by Gerardi and Shapiro (2009).

Appendix B

Summary Statistics and List of Hub Airports

Table B.1: Summary Statistics for Variables in Empirical Analysis

Variable	N	Mean	Standard Deviation	Minimum	5th Percentile	25th Percentile	Median	75th Percentile	95th Percentile	Maximum
Differenced Variables										
DFARE	7,426	-$1.77	$36.65	-$300.06	-$58.19	-$21.15	-$1.88	$17.35	$55.35	$252.36
DINC ($1,000s)	7,426	-$0.22	$12.58	-$39.11	-$19.81	-$9.39	-$1.08	$7.46	$21.61	$45.36
DPOP (1,000,000s)	7,426	-0.64	7.54	-19.72	-16.60	-4.19	-0.68	3.15	13.14	19.88
DTEMP	7,426	0.80	13.64	-50.80	-20.70	-8.30	0.30	9.00	23.80	49.00
DTOURIST	7,426	-0.10	3.41	-18.14	-2.24	-0.20	0.02	0.31	2.41	16.00
Legacy Interactions										
DINC*LEGACY	7,426	$0.21	$9.25	-$32.52	-$15.65	-$1.53	$0.00	$1.80	$17.97	$40.88
DPOP*LEGACY	7,426	-0.35	6.01	-19.72	-11.30	-1.34	0.00	0.00	9.15	19.88
DTEMP*LEGACY	7,426	0.02	9.68	-50.80	-16.80	-1.90	0.00	0.90	18.70	49.00
Quarter Interactions										
DINC*Q1	7,426	-$0.06	$5.87	-$39.11	-$9.98	$0.00	$0.00	$0.00	$8.05	$45.36
DINC*Q2	7,426	-$0.06	$6.38	-$39.11	-$11.35	$0.00	$0.00	$0.00	$10.32	$45.36
DINC*Q3	7,426	-$0.06	$6.43	-$39.11	-$11.75	$0.00	$0.00	$0.00	$10.74	$45.36
DPOP*Q1	7,426	-0.13	3.50	-19.72	-4.61	0.00	0.00	0.00	3.56	19.88
DPOP*Q2	7,426	-0.18	3.82	-19.72	-5.07	0.00	0.00	0.00	4.04	19.88
DPOP*Q3	7,426	-0.17	3.86	-19.72	-5.12	0.00	0.00	0.00	4.15	19.88
DTEMP*Q1	7,426	0.46	8.61	-50.80	-12.00	0.00	0.00	0.00	16.60	49.00
DTEMP*Q2	7,426	0.22	6.62	-29.20	-11.30	0.00	0.00	0.00	12.90	38.10
DTEMP*Q3	7,426	-0.20	5.32	-34.60	-9.00	0.00	0.00	0.00	7.60	33.70
Other Control Variables										
DHUB	7,426	-0.04	0.69	-1.00	-1.00	-1.00	0.00	0.00	1.00	1.00
DMANAGERIAL	7,418	-0.17	1.72	-5.21	-2.94	-1.39	-0.16	0.98	2.65	4.60
DEMPLOYSHARE	7,418	-0.12	13.09	-31.50	-21.40	-8.68	-0.18	8.85	22.26	43.61
DTIME	7,426	-0.31	1.24	-5.00	-3.00	-1.00	0.00	0.00	2.00	5.00
DBLACK	7,418	-0.03	0.11	-0.38	-0.21	-0.10	-0.02	0.04	0.15	0.40
Competition Variables										
NCOMPETITORS	7,426	2.14	1.15	1.00	1.00	1.00	2.00	3.00	4.00	6.00
HHI (1,000s)	7,426	6.95	2.66	2.08	3.05	4.91	6.26	10.00	10.00	10.00
ADJACENT_ORIGIN	7,426	0.22	0.41	0.00	0.00	0.00	0.00	0.00	1.00	1.00
ADJACENT_DEST	7,426	0.29	0.45	0.00	0.00	0.00	0.00	1.00	1.00	1.00
DINC*NCOMPETITORS	7,426	-$0.21	$30.08	-$121.70	-$48.17	-$16.12	-$1.28	$14.16	$54.34	$164.16
DINC*HHI (1,000s)	7,426	-$1.51	$95.34	-$391.06	-$152.27	-$59.25	-$6.74	$47.42	$169.19	$453.57
Indicator Variables										
LEGACY	7,426	0.55	0.50	0	0	0	1	1	1	1
Quarter 1	7,426	0.22	0.41	0	0	0	0	0	1	1
Quarter 2	7,426	0.26	0.44	0	0	0	0	1	1	1
Quarter 3	7,426	0.26	0.44	0	0	0	0	1	1	1
Quarter 4	7,426	0.27	0.44	0	0	0	0	1	1	1

Table B.2: List of Hub Airports by Carrier

Carrier	Hub Airports
Alaska (AS)	Anchorage, AK (ANC)
	Los Angeles, CA (LAX)
	Portland, OR (PDX)
	San Francisco, CA (SFO)
	Seattle, WA (SEA)
American (AA)	Charlotte, NC (CLT)
	Chicago, IL (ORD)
	Dallas, TX (DFW)
	Los Angeles, CA (LAX)
	Miami, FL (MIA)
	New York, NY (JFK)
	New York, NY (LGA)
	Philadelphia, PA (PHL)
	Phoenix, AZ (PHX)
	Washington, DC (DCA)
Delta (DL)	Atlanta, GA (ATL)
	Boston, MA (BOS)
	Cincinnati, OH (CIN)
	Detroit, MI (DFW)
	Minneapolis, MN (MSP)
	Los Angeles, CA (LAX)
	New York, NY (JFK)
	New York, NY (LGA)
	Salt Lake City, UT (SLC)
	Seattle, WA (SEA)
Hawaiian (HA)	Honolulu, HI (HNL)
	Kahului, HI (OGG)
United (UA)	Chicago, IL (ORD)
	Denver, CO (DNL)
	Houston, TX (IAH)
	Los Angeles, CA (LAX)
	Newark, NJ (EWR)
	San Francisco, CA (SFO)
	Washington, DC (IAD)
US Airways (US)*	Charlotte, NC (CLT)
	Philadelphia, PA (PHL)
	Phoenix, AZ (PHX)
	Washington, DC (DCA)

Notes: * Although US Airways and American announced their merger in 2013, flights operated under US Airways flight numbers in the first two quarters of 2015.

Appendix C

Regression Results for Equations (1.3) and (1.4)

Table C.1: Difference in Logs Regression Results

	Equation (1.3)		Equation (1.4)	
Analysis Sample:	Full	Restricted	Full	Restricted
	(1)	(2)	(3)	(4)
DLINC	0.071***	0.060***	0.068***	0.054***
	(0.005)	(0.005)	(0.005)	(0.005)
DLPOP	0.018***	0.012***	0.019***	0.013***
	(0.001)	(0.001)	(0.001)	(0.001)
DLTEMP	0.009	-0.001	0.014**	0.007
	(0.006)	(0.005)	(0.006)	(0.005)
DLTOURIST			-0.003**	-0.005***
			(0.001)	(0.001)
Constant	-0.002**	-0.004***	-0.002**	-0.004***
	(0.001)	(0.001)	(0.001)	(0.001)
Observations	7,426	6,684	7,426	6,684
R-squared	0.161	0.133	0.162	0.137

Notes: Robust standard errors are reported in parentheses. The restricted sample removes outliers that fall outside the range encompassing the 5th and 95th percentiles of the dependent variable DLFARE.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Using the full sample of observations, the first column of Table C.1 reports regression results from the model specified by equation (1.3) while the third column reports results from the model specified by equation (1.4). To alleviate concerns that outliers are driving results, columns two and four of Table C.1 report estimates from equations (1.3) and (1.4) respectively when the sample is restricted to observations that fall between the 5th and 95th percentiles of the dependent variable DLFARE. The estimated coefficient of 0.06 on DLINC in column two indicates that roundtrip fares are 0.06% higher for each 1% difference in average per capita income between the endpoint cities. For the SFO-ORD route where San Francisco residents have average incomes that are 42.7% higher Chicago residents, this coefficient implies a roundtrip fare that is 2.56% higher for SFO originating passengers.[1]

Table C.2: Difference in Logs with Additional Controls

Analysis Sample:	Restricted
	(1)
DLINC	0.044***
	(0.005)
DLPOP	0.012***
	(0.001)
DLTEMP	0.012**
	(0.005)
DHUB	0.012***
	(0.001)
DMANAGERIAL	0.004***
	(0.001)
DEMPLOYSHARE	0.000***
	(0.000)
DTIME	-0.002***
	(0.001)
DBLACK	-0.067***
	(0.010)
Constant	-0.005***
	(0.001)
Observations	6,678
R-squared	0.160

Notes: Robust standard errors are reported in parentheses. The restricted sample removes outliers that fall outside the range encompassing the 5th and 95th percentiles of the dependent variable DLFARE.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Table C.2 provides results when DHUB, DMANAGERIAL, DEMPLOYSHARE, DTIME, and DBLACK are added to the specification presented in column two of Table C.1.[2] The

[1] $\frac{\$72{,}364-\$50{,}690}{\$50{,}690} * 100 = 42.7\%$. 0.06%*42.7 = 2.56%.

[2] The regression presented in Table C.2 has 6 fewer observations than the regression presented in column

coefficient on DLINC remains positive and significant at conventional levels. However, the magnitude has decreased from 0.06% to 0.044% per 1% difference in income. For the SFO-ORD route, this coefficient implies that roundtrip fares are 1.88% higher on average for passengers originating in San Francisco.[3]

two of Table C.1 because the data used to construct DMANAGERIAL, DEMPLOYSHARE, and DBLACK was not available for the Walla Walla, WA and Kahului-Wailuka-Lahaina, HI MSAs.

[3] 0.044*42.7 $\approx$ 1.88%.

Appendix D

Equation (2.3) Regression Results

Table D.1 reports regression results from the model specified by equation (2.3). Similar to the results from equations (2.1) and (2.2) presented in Table 2.2, the coefficients on the layover time covariates are significant at conventional levels. Although the coefficients in Table D.1 cannot be used to directly compute the dollar value of a one-hour layover, the elasticity interpretation is still useful.[1]

The last three rows of Table D.1 present the predicted percent change in fare from increasing an itinerary's layover time from 60 to 80 minutes (33.33% increase), 60 to 90 minutes (50% increase), and from 60 to 120 minutes (100% increase), respectively. Increasing layover time from 60 to 80 minutes results in a fare decrease of -4.31% ($13.71) while increasing an itinerary's layover time from one to two hours results in a fare decrease of -12.93% ($41.12). These estimated layover time effects are similar in magnitude to the values presented in Table 2.2.

[1] The coefficients cannot be used to directly compute the dollar value of a one-hour layover because the change from 0 to 60 minutes of layover time cannot be expressed as a percentage.

Table D.1: Log-Log Regression Results

Dependent Variable:	ln(Fare)
	(1)
ln(AIR TRAVEL TIME)	0.21
	(1.08)
$ln(\text{AIR TRAVEL TIME})^2$	-0.04
	(0.10)
ln(LAYOVER MINUTES)	-0.41***
	(0.09)
$ln(\text{LAYOVER MINUTES})^2$	0.03***
	(0.01)
TOTAL DISTANCE (1,000s)	0.10*
	(0.05)
DAYS BEFORE DEPARTURE	-0.02***
	(0.00)
LCC	-0.37***
	(0.10)
Constant	7.63**
	(3.09)
Observations	575,267
R-squared	0.45
City-Pair FE	YES
Carrier FE	YES
Layover Airport FE	YES
Time-of-Departure FE	YES
Day-of-Week-of-Departure FE	YES
Day-of-Week-of-Purchase FE	YES
Effect of Increasing Layover from 60 to 80 Minutes	-4.31%
	(-$13.71)
Effect of Increasing Layover from 60 to 90 Minutes	-6.47%
	(-$20.57)
Effect of Increasing Layover from 60 to 120 Minutes	-12.93%
	(-$41.12)

Notes: Standard errors clustered at the city-pair level are reported in parentheses. The average fare in the regression sample is $317.99.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Appendix E

Evidence in Support of the Exogeneity of Layover Times

As discussed in Section 2.4.1, there is a concern that layover times are endogenous if airlines aim to reduce layovers in city-pair markets with a large share of business travelers who are time-sensitive and less price elastic relative to leisure travelers. If this type of scheduling behavior occurs, then routes where the origin metro area contains a large fraction of managerial workers should experience lower layover times on average.[1] To investigate if this pattern occurs in the data, Figure E.1 plots the average number of layover minutes for each of the fifty routes listed in Table 2.1 along with the fraction of managerial workers in the origin airport Metropolitan Statistical Area (MSA). If airlines are reducing layover times in markets with a large share of time-sensitive business travelers, then a significant downward trend should be observed in Figure E.1.

The points plotted in Figure E.1 do not display a clear downward trend. Therefore, the figure supports the argument outlined in Section 2.4.1. Namely, that layover times are unlikely to

[1] Workers in managerial occupations are the most likely type of worker to travel for business.

be determined by the time sensitivity of passengers in any one connecting city-pair market. Furthermore, the pairwise correlation between these average layover times and the fraction of managerial workers while slightly negative, is statistically insignificant.[2]

Figure E.1: Average Layover Times by the Fraction of Managerial Workers in the Origin MSA

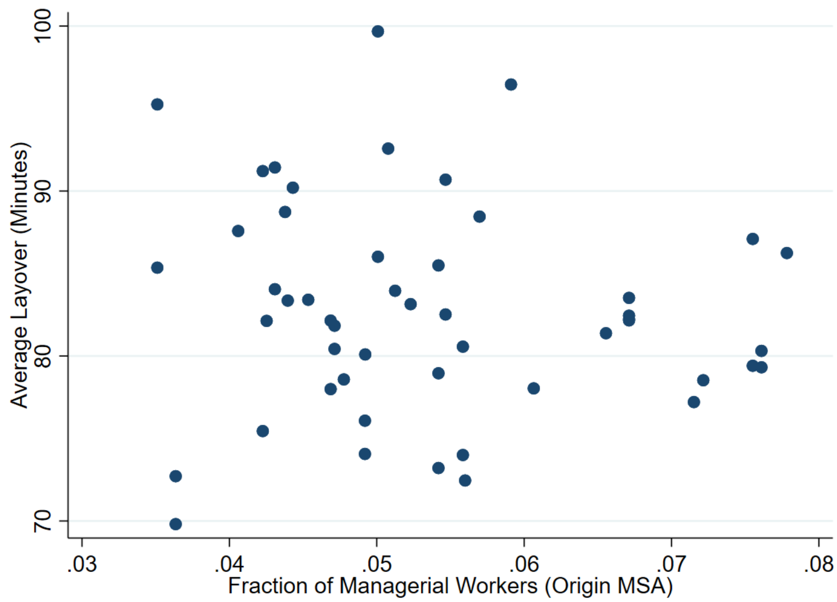

Notes: Data for the fraction of managerial workers comes from the May 2017 Bureau of Labor Statistics (BLS) Occupational and Employment Statistics survey. For each Metropolitan Statistical Area (MSA), the fraction of managerial workers is defined as the number of workers in management occupations divided by the total number of employed persons.

[2] The pairwise correlation is -0.09. However, the p-value of a Pearson's correlation test with a null hypothesis of no correlation ($H_0 : \rho = 0$) cannot be rejected at standard significance levels given that the p-value of the test is 0.5167.

Appendix F

Small Cities Subsample Regression Results

Considering that metro areas with larger populations have a greater demand for air travel relative to smaller metro areas, there is a concern that high levels of demand on routes originating from larger metro areas are driving the estimated air travel time and layover time effects presented in Table 2.2. While the city-pair fixed effects included in equations (2.1)-(2.3) should appropriately control for any differences in demand across routes, Table F.1 repeats the analysis presented in Table 2.2 when routes originating from the twenty most populous Metropolitan Statistical Areas (MSA) are removed. This restriction removes from the analysis sample routes originating from the Atlanta (ATL), Boston (MHT), Detroit (DTW), Los Angeles (LAX and SNA), Miami (FLL), Minneapolis (MSP), New York (LGA), Philadelphia (PHL), San Diego (SAN), San Francisco (SFO), Tampa (TPA), and Washington DC (DCA) metro areas. After these routes are removed, the air travel and layover time estimates presented in Table F.1 are still very similar to the Table 2.2 values. Therefore, routes originating from large metropolitan areas do not appear to be substantially influencing the results.

Table F.1: Average Value of Layover Time Regression Results
(Small Cities Subsample)

Dependent Variable:	Fare	*ln*(Fare)
	(1)	(2)
AIR TRAVEL TIME	-0.66	-0.002
	(0.46)	(0.001)
(AIR TRAVEL TIME)^2 (100s)	0.03	0.000
	(0.05)	(0.00)
LAYOVER MINUTES	-0.92***	-0.002***
	(0.09)	(0.00)
(LAYOVER MINUTES)^2 (100s)	0.25***	0.001***
	(0.03)	(0.00)
TOTAL DISTANCE (1,000s)	65.47**	0.17**
	(30.64)	(0.08)
DAYS BEFORE DEPARTURE	-5.59***	-0.017***
	(0.38)	(0.00)
LCC	-165.64***	-0.456***
	(56.63)	(0.13)
Constant	737.80***	6.879***
	(82.62)	(0.21)
Observations	380,800	380,800
R-squared	0.35	0.41
City-Pair FE	YES	YES
Carrier FE	YES	YES
Layover Airport FE	YES	YES
Time-of-Departure FE	YES	YES
Day-of-Week-of-Departure FE	YES	YES
Day-of-Week-of-Purchase FE	YES	YES
1-Hour Air Travel Effect (%)	NA	-11.02%
1-Hour Air Travel Effect ($)	-$38.63	-$35.22
1-Hour Layover Effect (%)	NA	-12.14%
1-Hour Layover Effect ($)	-$46.15	-$38.80
2-Hour Layover Effect (%)	NA	-19.82%
2-Hour Layover Effect ($)	-$74.26	-$63.35

Notes: Standard errors clustered at the city-pair level are reported in parentheses. The average fare in the regression sample is $319.64. The analysis sample includes the Table 2.1 routes that do not originate from the Atlanta (ATL), Boston (MHT), Detroit (DTW), Los Angeles (LAX and SNA), Miami (FLL), Minneapolis (MSP), New York (LGA), Philadelphia (PHL), San Diego (SAN), San Francisco (SFO), Tampa (TPA), and Washington DC (DCA) metro areas.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Appendix G

Hub Airport Regression Results

Table G.1: Hub Airport Regression Results

Dependent Variable:	Fare	ln(Fare)
	(1)	(2)
AIR TRAVEL TIME	-0.57*	-0.002**
	(0.30)	(0.001)
(AIR TRAVEL TIME)^2 (100s)	0.04	0.000
	(0.04)	(0.000)
(LAYOVER MINUTES)*(ORD Layover)	-1.05***	-0.003***
	(0.22)	(0.001)
(LAYOVER MINUTES)^2*(ORD Layover) (100s)	0.30***	0.001***
	(0.08)	(0.000)
(LAYOVER MINUTES)*(IAH Layover)	-1.63***	-0.004***
	(0.26)	(0.001)
(LAYOVER MINUTES)^2*(IAH Layover) (100s)	0.48***	0.001***
	(0.10)	(0.000)
(LAYOVER MINUTES)*(DEN Layover)	-1.61***	-0.004***
	(0.37)	(0.001)
(LAYOVER MINUTES)^2*(DEN Layover) (100s)	0.48***	0.001**
	(0.14)	(0.000)
(LAYOVER MINUTES)*(EWR Layover)	-1.06**	-0.003**
	(0.43)	(0.001)
(LAYOVER MINUTES)^2*(EWR Layover) (100s)	0.29*	0.001
	(0.16)	(0.000)
(LAYOVER MINUTES)*(DFW Layover)	-0.80***	-0.002***
	(0.24)	(0.001)
(LAYOVER MINUTES)^2*(DFW Layover) (100s)	0.21**	0.001**
	(0.09)	(0.000)
(LAYOVER MINUTES)*(CLT Layover)	-0.95***	-0.003***
	(0.27)	(0.001)
(LAYOVER MINUTES)^2*(CLT Layover) (100s)	0.28**	0.001*
	(0.11)	(0.000)
(LAYOVER MINUTES)*(PHL Layover)	-1.13***	-0.003***
	(0.40)	(0.001)
(LAYOVER MINUTES)^2*(PHL Layover) (100s)	0.29**	0.001**
	(0.13)	(0.000)
(LAYOVER MINUTES)*(ATL Layover)	-0.95***	-0.003***
	(0.18)	(0.001)
(LAYOVER MINUTES)^2*(ATL Layover) (100s)	0.27***	0.001***
	(0.07)	(0.000)
(LAYOVER MINUTES)*(DTW Layover)	-0.39	-0.001*
	(0.27)	(0.001)
(LAYOVER MINUTES)^2*(DTW Layover) (100s)	0.05	0.000
	(0.11)	(0.000)
(LAYOVER MINUTES)*(MSP Layover)	-1.38***	-0.004***
	(0.30)	(0.001)
(LAYOVER MINUTES)^2*(MSP Layover) (100s)	0.42***	0.001***
	(0.11)	(0.000)
(LAYOVER MINUTES)*(SLC Layover)	-0.60	-0.002
	(0.56)	(0.002)
(LAYOVER MINUTES)^2*(SLC Layover) (100s)	0.12	0.000
	(0.20)	(0.001)
(LAYOVER MINUTES)*(Other Layover Airport)	-0.42**	-0.001**
	(0.20)	(0.001)
(LAYOVER MINUTES)^2*(Other Layover Airport) (100s)	0.06	0.000
	(0.07)	(0.000)
TOTAL DISTANCE (1,000s)	32.70*	0.096*
	(18.48)	(0.050)
DAYS BEFORE DEPARTURE	-5.20***	-0.016***
	(0.34)	(0.001)
LCC	-145.74***	-0.361***
	(42.78)	(0.100)
Constant	691.95***	6.77***
	(64.47)	(0.17)
Observations	575,267	575,267
R-squared	0.40	0.45
City-Pair FE	YES	YES
Carrier FE	YES	YES
Layover Airport FE	YES	YES
Time-of-Departure FE	YES	YES
Day-of-Week-of-Departure FE	YES	YES
Day-of-Week-of-Purchase FE	YES	YES

Notes: Standard errors clustered at the city-pair level are reported in parentheses.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Appendix H

Spline Regression Results

To provide further evidence of the nonlinear relationship between fare and layover time, this appendix presents regression results from an alternative identification strategy that allows the effect of layover time on fare to change throughout the range of layover time values. A typical way to model structural change in the relationship between a dependent variable y and an independent variable x, is to regress y on a spline function of x.[1] One option is a cubic spline, where y is modeled as a continuous piecewise third-order polynomial in x. However, the points of structural change must first be specified before estimating a spline regression.[2] When analyzing the distribution of layover time, it is possible to classify layovers as either tight, short, long, or very long.

Figure H.1 below displays the distribution of layover minutes in the Table 2.2 regression sample. The vertical red lines are placed at 50, 80, and 120 minutes of layover time. These values split layover time into four categories. Itineraries with less than 50 minutes of layover time are considered to have tight connections. Passengers may be concerned when purchasing this type of itinerary because of the high likelihood of missing the connecting flight if the first

[1] For a review of regression splines, see Poirier (1976).
[2] Points of structural change are referred to as the "knots" in the spline regression function.

leg arrives late. Itineraries with 51-80 layover minutes are considered short while itineraries with 81-120 minutes are considered long. Finally, itineraries with more than two hours of layover time are considered very long. With these classifications, 22.60% of the regression sample have tight layovers, 40.69% have short layovers, 20.46% have long layovers, and 16.25% have very long layovers.

Figure H.1: Distribution of Layover Minutes

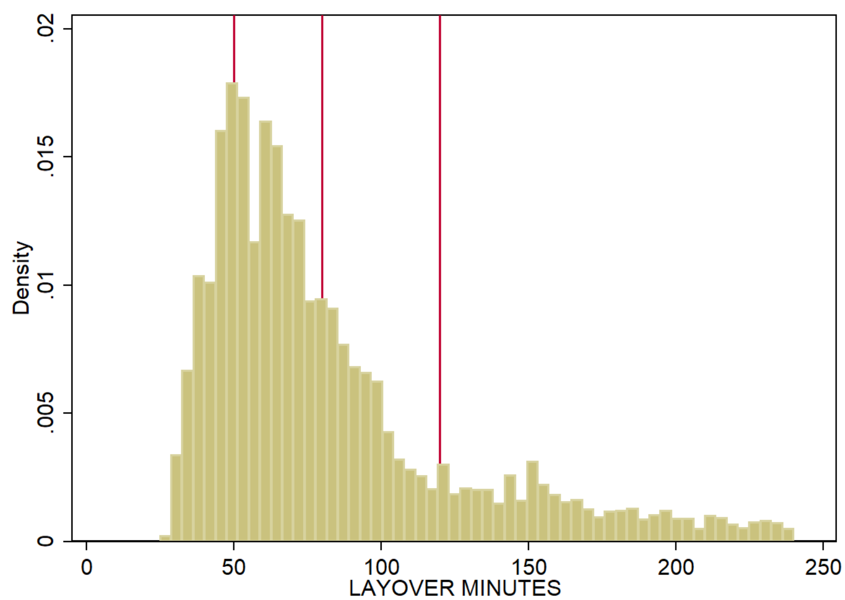

Notes: Vertical red lines are placed at 50, 80, and 120 minutes.

Table H.1 presents regression results when the layover minute covariates in equations (2.1) and (2.2) are replaced by a restricted cubic spline function with knots placed at 50, 80, and 120 minutes of layover time. Cubic splines are frequently used to model nonlinear relationships because they provide great flexibility for fitting the data and are visually smooth.[3] However, cubic spline functions are often poorly behaved in the tails.[4] To address this issue, Stone and Koo (1986) advocate for the use of restricted cubic splines, which are constrained

[3] Cubic splines are visually smooth because they have continuous first and second derivatives.
[4] Here, tails refer to the parts of the distribution before the first knot and after the last knot.

to be linear in the tails.[5]

Column one of Table H.1 presents estimates from the model specified by equation (2.1), when LAYOVER MINUTES and (LAYOVER MINUTES)^2 are replaced by a restricted cubic spline function with knots placed at 50, 80, and 120 minutes of layover time. The coefficients on the spline coefficents are significant at conventional levels and the results indicate that passengers are compensated with a fare that is \$53.26 cheaper for incurring a one-hour layover. This value is similar in magnitude to the \$47.60 estimate found in column one of Table 2.2.

Column two of Table H.1 provides estimates from the model specified by equation (2.2), when the layover minute covariates are replaced by a restricted cubic spline function with knots placed at 50, 80, and 120 minutes of layover time. The column two results indicate that passengers are compensated with a fare that is \$48.59 or \$72.66 (15.28% or 22.85%) cheaper for incurring a one or two-hour layover, respectively. These values are also similar in magnitude to the one and two-hour effects presented in column two of Table 2.2.

[5] For more on restricted cubic splines, see Durrleman and Simon (1989) and Harrell Jr (2015). To force linearity when LAYOVER MINUTES is less than the first knot of 50, quadratic and cubic terms in LAYOVER MINUTES are omitted. Furthermore, with knots placed at 50, 80, and 120 minutes of layover time, the additional variable in the restricted cubic spline is defined as follows:

$$Add'l\ Var = (X-50)^3_+ - (X-80)^3_+ * \frac{(120-50)}{(120-80)} + (X-120)^3_+ * \frac{(80-50)}{(120-80)}$$

where X = LAYOVER MINUTES. For numerical behavior, the $Add'l\ Var$ is divided by the term $\tau = (120-50)^2$. Therefore, at 60 minutes of layover time, the additional variable equals:

$$\frac{(60-50)^3}{(120-50)^2} = \frac{1000}{4900} \approx 0.2041.$$

Table H.1: Restricted Cubic Spline Regressions

Dependent Variable:	Fare	ln(Fare)
	(1)	(2)
AIR TRAVEL TIME	-0.59*	-0.002**
	(0.30)	(0.00)
(AIR TRAVEL TIME)^2 (100s)	0.05	0.000
	(0.04)	(0.00)
LAYOVER MINUTES	-0.89***	-0.003***
	(0.08)	(0.00)
LAYOVER MINUTES (*Add'l Var*)	0.58***	0.002***
	(0.07)	(0.00)
TOTAL DISTANCE (1,000s)	32.68*	0.097*
	(19.04)	(0.05)
DAYS BEFORE DEPARTURE	-5.22***	-0.016***
	(0.34)	(0.00)
LCC	-150.24***	-0.375***
	(42.27)	(0.10)
Constant	738.44***	6.891***
	(61.61)	(0.16)
Observations	575,267	575,267
R-squared	0.39	0.45
City-Pair FE	YES	YES
Carrier FE	YES	YES
Layover Airport FE	YES	YES
Time-of-Departure FE	YES	YES
Day-of-Week-of-Departure FE	YES	YES
Day-of-Week-of-Purchase FE	YES	YES
1-Hour Air Travel Effect (%)	NA	-9.67%
1-Hour Air Travel Effect ($)	-$33.65	-$30.75
1-Hour Layover Effect (%)	NA	-15.28%
1-Hour Layover Effect ($)	-$53.26	-$48.59
2-Hour Layover Effect (%)	NA	-22.85%
2-Hour Layover Effect ($)	-$79.37	-$72.66

Notes: Standard errors clustered at the city-pair level are reported in parentheses. The average fare in the regression sample is $317.99. Knots are placed at 50, 80, and 120 minutes of layover time. With these knots, LAYOVER MINUTES (*Add'l Var*) is defined as follows:

$$Add'l\ Var = (X-50)^3_+ - (X-80)^3_+ * \frac{(120-50)}{(120-80)} + (X-120)^3_+ * \frac{(80-50)}{(120-80)}$$

where X = LAYOVER MINUTES. For numerical behavior, the $Add'l\ Var$ is divided by the term $\tau = (120-50)^2$.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Appendix I

Initial Slot Allocation for Delta and Continental at Newark (EWR)

Table I.1: Initial Slot Allocation at Newark (EWR) for Delta Airlines

Carrier	Time	Arr/Dep	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Delta	6:00 AM	Departure	2	2	2	2	2	2	2
Delta	7:00 AM	Departure	2	2	2	2	2	2	2
Delta	7:30 AM	Departure	1	1	1	1	1	1	1
Delta	8:00 AM	Departure	1	1	1	1	1	1	1
Delta	9:00 AM	Arrival	1	1	1	1	1	1	1
Delta	9:30 AM	Arrival	1	1	1	1	1	1	1
Delta	10:00 AM	Departure	2	2	2	2	2	2	2
Delta	10:30 AM	Arrival	1	1	1	1	1	1	1
Delta	11:00 AM	Departure	1	1	1	1	1	1	1
Delta	12:00 PM	Arrival	1	1	1	1	1	1	1
Delta	12:30 PM	Arrival	1	1	1	1	1	1	1
Delta	1:00 PM	Departure	1	1	1	1	1	1	1
Delta	1:30 PM	Departure	1	1	1	1	1	1	1
Delta	2:30 PM	Arrival	1	1	1	1	1	1	1
Delta	3:00 PM	Departure	1	1	1	1	1	1	1
Delta	4:00 PM	Arrival	1	1	1	1	1	1	1
Delta	4:30 PM	Arrival	1	1	1	1	1	1	1
Delta	4:30 PM	Departure	1	1	1	1	1	1	1
Delta	5:00 PM	Arrival	1	1	1	1	1	1	1
Delta	5:00 PM	Departure	1	1	1	1	1	1	1
Delta	5:30 PM	Arrival	1	1	1	1	1	1	1
Delta	5:30 PM	Departure	1	1	1	1	1	1	1
Delta	6:00 PM	Departure	1	1	1	1	1	1	1
Delta	6:30 PM	Arrival	1	1	1	1	1	1	1
Delta	7:00 PM	Departure	1	1	1	1	1	1	1
Delta	7:30 PM	Arrival	1	1	1	1	1	1	1
Delta	9:00 PM	Arrival	1	1	1	1	1	1	1
Delta	9:30 PM	Arrival	1	1	1	1	1	1	1
Delta	10:30 PM	Arrival	1	1	1	1	1	1	1

Source: Federal Register Volume 73, No. 99, Wednesday, May 21, 2008. Beginning on June 20th, 2008, slot controls were in effect between 6:00AM and 10:59PM.

Table I.2: Initial Slot Allocation at Newark (EWR) for Continental Airlines

Carrier	Time	Arr/Dep	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Continental	6:00 AM	Arrival	4	4	4	4	4	4	4
Continental	6:00 AM	Departure	9	9	9	9	9	9	9
Continental	6:30 AM	Arrival	1	1	1	1	1	1	1
Continental	6:30 AM	Departure	20	20	20	20	20	20	20
Continental	7:00 AM	Arrival	6	6	6	6	6	6	6
Continental	7:00 AM	Departure	15	15	15	15	15	15	15
Continental	7:30 AM	Arrival	19	19	19	19	19	19	19
Continental	7:30 AM	Departure	17	17	17	17	17	17	17
Continental	8:00 AM	Arrival	14	14	14	14	14	14	14
Continental	8:00 AM	Departure	16	16	16	16	16	16	16
Continental	8:30 AM	Arrival	7	7	7	7	7	7	7
Continental	8:30 AM	Departure	24	24	24	24	24	24	24
Continental	9:00 AM	Arrival	6	6	6	6	6	6	6
Continental	9:00 AM	Departure	21	21	21	21	21	21	21
Continental	9:30 AM	Arrival	9	9	9	9	9	9	9
Continental	9:30 AM	Departure	4	4	4	4	4	4	4
Continental	10:00 AM	Arrival	10	10	10	10	10	10	10
Continental	10:00 AM	Departure	7	7	7	7	7	7	7
Continental	10:30 AM	Arrival	7	7	7	7	7	7	7
Continental	10:30 AM	Departure	12	12	12	12	12	12	12
Continental	11:00 AM	Arrival	18	18	18	18	18	18	18
Continental	11:00 AM	Departure	6	6	6	6	6	6	6
Continental	11:30 AM	Arrival	14	14	14	14	14	14	14
Continental	11:30 AM	Departure	13	13	13	13	13	13	13
Continental	12:00 PM	Arrival	14	14	14	14	14	14	14
Continental	12:00 PM	Departure	18	18	18	18	18	18	18
Continental	12:30 PM	Arrival	19	19	19	19	19	19	19
Continental	12:30 PM	Departure	8	8	8	8	8	8	8
Continental	1:00 PM	Arrival	10	10	10	10	10	10	10
Continental	1:00 PM	Departure	18	18	18	18	18	18	18
Continental	1:30 PM	Arrival	20	20	20	20	20	20	20
Continental	1:30 PM	Departure	9	9	9	9	9	9	9
Continental	2:00 PM	Arrival	17	17	17	17	17	17	17
Continental	2:00 PM	Departure	5	5	5	5	5	5	5
Continental	2:30 PM	Arrival	17	17	17	17	17	17	17
Continental	2:30 PM	Departure	20	20	20	20	20	20	20
Continental	3:00 PM	Arrival	17	17	17	17	17	17	17
Continental	3:00 PM	Departure	21	21	21	21	21	21	21
Continental	3:30 PM	Arrival	19	19	19	19	19	19	19
Continental	3:30 PM	Departure	10	10	10	10	10	10	10
Continental	4:00 PM	Arrival	11	11	11	11	11	11	11
Continental	4:00 PM	Departure	11	11	11	11	11	11	11
Continental	4:30 PM	Arrival	19	19	19	19	19	19	19
Continental	4:30 PM	Departure	13	13	13	13	13	13	13
Continental	5:00 PM	Arrival	13	13	13	13	13	13	13
Continental	5:00 PM	Departure	18	18	18	18	18	18	18
Continental	5:30 PM	Arrival	6	6	6	6	6	6	6
Continental	5:30 PM	Departure	13	13	13	13	13	13	13
Continental	6:00 PM	Arrival	16	16	16	16	16	16	16
Continental	6:00 PM	Departure	15	15	15	15	15	15	15
Continental	6:30 PM	Arrival	13	13	13	13	13	13	13
Continental	6:30 PM	Departure	8	8	8	8	8	8	8
Continental	7:00 PM	Arrival	20	20	20	20	20	20	20
Continental	7:00 PM	Departure	17	17	17	17	17	17	17
Continental	7:30 PM	Arrival	8	8	8	8	8	8	8
Continental	7:30 PM	Departure	14	14	14	14	14	14	14
Continental	8:00 PM	Arrival	18	18	18	18	18	18	18
Continental	8:00 PM	Departure	26	26	26	26	26	26	26
Continental	8:30 PM	Arrival	9	9	9	9	9	9	9
Continental	8:30 PM	Departure	13	13	13	13	13	13	13
Continental	9:00 PM	Arrival	8	8	8	8	8	8	8
Continental	9:00 PM	Departure	17	17	17	17	17	17	17
Continental	9:30 PM	Arrival	16	16	16	16	16	16	16
Continental	9:30 PM	Departure	11	11	11	11	11	11	11
Continental	10:00 PM	Arrival	11	11	11	11	11	11	11
Continental	10:00 PM	Departure	6	6	6	6	6	6	6
Continental	10:30 PM	Arrival	9	9	9	9	9	9	9
Continental	10:30 PM	Departure	1	1	1	1	1	1	1

Source: Federal Register Volume 73, No. 99, Wednesday, May 21, 2008. Continental was initially allocated 72% of the available slots at Newark (EWR). Beginning on June 20th, 2008, slot controls were in effect between 6:00AM and 10:59PM.

Appendix J

Initial Slot Allocation for Delta and Continental at John F. Kennedy (JFK)

Table J.1: Initial Slot Allocation at John F. Kennedy (JFK) for Continental Airlines

Carrier	Time	Arr/Dep	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Continental	7:30 AM	Departure	0	1	1	1	1	1	0
Continental	1:15 PM	Departure	1	1	1	1	1	1	0
Continental	2:45 PM	Departure	0	0	0	0	0	0	1
Continental	4:15 PM	Departure	1	1	1	1	1	1	0
Continental	5:45 PM	Departure	1	1	1	1	1	1	0
Continental	8:30 PM	Arrival	1	1	1	1	1	1	1

Source: Federal Register Volume 73, No. 13, Friday, January 18, 2008. Beginning on March 30th, 2008, slot controls were in effect between 6:00AM and 10:59PM.

Table J.2: Initial Arrival Slot Allocation at John F. Kennedy (JFK) for Delta Airlines

Carrier	Time	Arr/Dep	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Delta	6:00 AM	Arrival	2	2	2	2	2	2	2
Delta	6:15 AM	Arrival	2	2	2	2	2	2	2
Delta	6:30 AM	Arrival	2	2	2	2	2	2	2
Delta	7:00 AM	Arrival	2	2	2	2	2	2	2
Delta	7:15 AM	Arrival	5	5	5	5	5	5	5
Delta	7:30 AM	Arrival	7	7	7	7	7	7	7
Delta	7:45 AM	Arrival	5	5	5	5	5	5	5
Delta	8:15 AM	Arrival	5	5	5	5	5	5	5
Delta	8:30 AM	Arrival	1	1	1	1	1	1	1
Delta	8:45 AM	Arrival	1	1	1	1	1	1	1
Delta	9:15 AM	Arrival	1	1	1	1	1	1	1
Delta	9:30 AM	Arrival	1	1	1	1	1	1	1
Delta	10:00 AM	Arrival	1	1	1	1	1	1	1
Delta	10:15 AM	Arrival	2	2	2	2	2	2	2
Delta	10:30 AM	Arrival	3	3	3	3	3	3	3
Delta	11:00 AM	Arrival	1	1	1	1	1	1	1
Delta	11:15 AM	Arrival	3	3	3	3	3	3	3
Delta	11:30 AM	Arrival	3	3	3	3	3	3	3
Delta	11:45 AM	Arrival	5	5	5	5	5	5	5
Delta	12:00 PM	Arrival	2	2	2	2	2	2	2
Delta	12:15 PM	Arrival	2	2	2	2	2	2	2
Delta	12:30 PM	Arrival	2	2	2	2	2	2	2
Delta	12:45 PM	Arrival	11	11	11	11	11	11	11
Delta	1:00 PM	Arrival	1	1	1	1	1	1	1
Delta	1:15 PM	Arrival	1	1	1	1	1	1	1
Delta	1:30 PM	Arrival	4	4	4	4	4	4	4
Delta	1:45 PM	Arrival	10	10	10	10	10	10	10
Delta	2:00 PM	Arrival	1	1	1	1	1	1	1
Delta	2:15 PM	Arrival	4	4	4	4	4	4	4
Delta	2:30 PM	Arrival	2	2	2	2	2	2	2
Delta	2:45 PM	Arrival	10	10	10	10	10	10	10
Delta	3:00 PM	Arrival	3	3	3	3	3	3	3
Delta	3:15 PM	Arrival	7	7	7	7	7	7	7
Delta	3:30 PM	Arrival	6	6	6	6	6	6	6
Delta	3:45 PM	Arrival	4	4	4	4	4	4	4
Delta	4:15 PM	Arrival	2	2	2	2	2	2	2
Delta	4:30 PM	Arrival	2	2	2	2	2	2	2
Delta	4:45 PM	Arrival	7	7	7	7	7	7	7
Delta	5:00 PM	Arrival	2	2	2	2	2	2	2
Delta	5:15 PM	Arrival	3	3	3	3	3	3	3
Delta	5:30 PM	Arrival	3	3	3	3	3	3	3
Delta	5:45 PM	Arrival	4	4	4	4	4	4	4
Delta	6:00 PM	Arrival	4	4	4	4	4	4	4
Delta	6:15 PM	Arrival	6	6	6	6	6	6	6
Delta	6:30 PM	Arrival	2	2	2	2	2	2	2
Delta	6:45 PM	Arrival	6	6	6	6	6	6	6
Delta	7:00 PM	Arrival	2	2	2	2	2	2	2
Delta	7:30 PM	Arrival	3	3	3	3	3	3	3
Delta	7:45 PM	Arrival	2	2	2	2	2	2	2
Delta	8:00 PM	Arrival	6	6	6	6	6	6	6
Delta	8:15 PM	Arrival	2	2	2	2	2	2	2
Delta	8:30 PM	Arrival	3	3	3	3	3	3	3
Delta	8:45 PM	Arrival	1	1	1	1	1	1	1
Delta	9:00 PM	Arrival	1	1	1	1	1	1	1
Delta	10:00 PM	Arrival	9	9	9	9	9	9	9
Delta	10:15 PM	Arrival	2	2	2	2	2	2	2
Delta	10:30 PM	Arrival	1	1	1	1	1	1	1
Delta	11:15 PM	Arrival	1	1	1	1	1	1	1
Delta	11:45 PM	Arrival	1	1	1	1	1	1	1

Source: Federal Register Volume 73, No. 13, Friday, January 18, 2008. Beginning on March 30th, 2008, slot controls were in effect between 6:00AM and 10:59PM.

Table J.3: Initial Departure Slot Allocation at John F. Kennedy (JFK) for Delta Airlines

Carrier	Time	Arr/Dep	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Delta	6:00 AM	Departure	2	2	2	2	2	2	2
Delta	6:45 AM	Departure	1	1	1	1	1	1	1
Delta	7:00 AM	Departure	3	3	3	3	3	3	3
Delta	7:45 AM	Departure	2	2	2	2	2	2	2
Delta	8:00 AM	Departure	4	4	4	4	4	4	4
Delta	8:15 AM	Departure	7	7	7	7	7	7	7
Delta	8:30 AM	Departure	8	8	8	8	8	8	8
Delta	8:45 AM	Departure	5	5	5	5	5	5	5
Delta	9:00 AM	Departure	8	8	8	8	8	8	8
Delta	9:15 AM	Departure	3	3	3	3	3	3	3
Delta	9:30 AM	Departure	3	3	3	3	3	3	3
Delta	9:45 AM	Departure	5	5	5	5	5	5	5
Delta	10:15 AM	Departure	1	1	1	1	1	1	1
Delta	10:45 AM	Departure	3	3	3	3	3	3	3
Delta	11:00 AM	Departure	3	3	3	3	3	3	3
Delta	11:15 AM	Departure	2	2	2	2	2	2	2
Delta	11:45 AM	Departure	2	2	2	2	2	2	2
Delta	12:00 PM	Departure	1	1	1	1	1	1	1
Delta	12:15 PM	Departure	2	2	2	2	2	2	2
Delta	12:30 PM	Departure	3	3	3	3	3	3	3
Delta	12:45 PM	Departure	4	4	4	4	4	4	4
Delta	1:00 PM	Departure	1	1	1	1	1	1	1
Delta	1:15 PM	Departure	3	3	3	3	3	3	3
Delta	1:30 PM	Departure	4	4	4	4	4	4	4
Delta	1:45 PM	Departure	7	7	7	7	7	7	7
Delta	2:15 PM	Departure	1	1	1	1	1	1	1
Delta	2:30 PM	Departure	1	1	1	1	1	1	1
Delta	2:45 PM	Departure	8	8	8	8	8	8	8
Delta	3:45 PM	Departure	6	6	6	6	6	6	6
Delta	4:00 PM	Departure	5	5	5	5	5	5	5
Delta	4:15 PM	Departure	7	7	7	7	7	7	7
Delta	4:30 PM	Departure	6	6	6	6	6	6	6
Delta	4:45 PM	Departure	2	2	2	2	2	2	2
Delta	5:00 PM	Departure	4	4	4	4	4	4	4
Delta	5:15 PM	Departure	1	1	1	1	1	1	1
Delta	5:30 PM	Departure	4	4	4	4	4	4	4
Delta	5:45 PM	Departure	5	5	5	5	5	5	5
Delta	6:00 PM	Departure	1	1	1	1	1	1	1
Delta	6:30 PM	Departure	4	4	4	4	4	4	4
Delta	6:45 PM	Departure	5	5	5	5	5	5	5
Delta	7:00 PM	Departure	7	7	7	7	7	7	7
Delta	7:15 PM	Departure	6	6	6	6	6	6	6
Delta	7:30 PM	Departure	5	5	5	5	5	5	5
Delta	7:45 PM	Departure	5	5	5	5	5	5	5
Delta	8:00 PM	Departure	4	4	4	4	4	4	4
Delta	8:30 PM	Departure	2	2	2	2	2	2	2
Delta	8:45 PM	Departure	5	5	5	5	5	5	5
Delta	9:00 PM	Departure	9	9	9	9	9	9	9
Delta	9:15 PM	Departure	4	4	4	4	4	4	4
Delta	10:00 PM	Departure	5	5	5	5	5	5	5
Delta	10:15 PM	Departure	2	2	2	2	2	2	2
Delta	11:30 PM	Departure	1	1	1	1	1	1	1
Delta	11:45 PM	Departure	1	1	1	1	1	1	1

Source: Federal Register Volume 73, No. 13, Friday, January 18, 2008. Beginning on March 30th, 2008, slot controls were in effect between 6:00AM and 10:59PM.