Available online at www.sciencedirect.com
ScienceDirect
International Journal of Industrial Organization
26 (2008) 1218 – 1232
International Journal of Industrial Organization
www.elsevier.com/locate/econbase

The effect of air traffic delays on airline prices

Silke J. Forbes

Department of Economics, University of California at San Diego, 9500 Gilman Drive #0508, La Jolla, CA 92093-0508, United States

Received 29 January 2007; received in revised form 12 December 2007; accepted 14 December 2007
Available online 30 January 2008

Abstract

A legislative change in takeoff and landing restrictions at LaGuardia Airport provides an opportunity to study the effect of an exogenous shock to product quality on prices in the airline industry. I test how the price response varies with the degree of competition in the market. I find that prices fall by $1.42 on average for each additional minute of flight delay, and that the price response is substantially larger in more competitive markets.

JEL Classifications: L11; L15; L93

Keywords: Airline pricing; Flight delays; Product quality; Competition

1. Introduction

Economists have long proposed to solve congestion problems at capacity-constrained airports through peak-load pricing. However, legislators have not embraced this solution. Instead, since 1969 four of the most heavily congested airports in the United States have been subject to takeoff and landing constraints, which effectively impose entry restrictions at these airports. In 2000, Congress made an attempt to lift these restrictions and to allow new entry into the affected airports, but without providing a mechanism to alleviate congestion. While some routes experienced entry and significantly lower prices due to the new competition, all routes suffered from a dramatic increase in congestion and flight delays.

I estimate the effect of the increase in air traffic delays on airline prices. In order to identify this relationship, I exploit the exogenous shock to the length of flight delays created by the Aviation Investment and Reform Act for the 21st Century (AIR-21). The AIR-21 Act allowed new entry on a subset of routes at LaGuardia Airport. I estimate the effect of longer flight delays on the *other* routes, for which entry restrictions were left unchanged, on ticket prices. Within that set of routes, I also test how the price response to longer flight delays varies with the degree of competition in the market.

Airport congestion and air traffic delays have received much attention from researchers and policy-makers during recent years. Air traffic delays are the most common source of customer complaints by airline passengers (U.S. Department of Transportation, 2000). Statistics from the Department of Transportation on the on-time performance of major U.S. carriers show that flight delays have increased significantly since the data were first collected in 1987. The United States Congress discussed proposals to impose minimum standards on the quality of airline service in 1999. Increases in air traffic delays and their impact on consumers played an important role in that debate. The sharp decrease in demand for air travel after the events of September 11,

E-mail address: sjanuszewski@ucsd.edu.

0167-7187/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.ijindorg.2007.12.004

2001, alleviated congestion concerns for the short term, but air traffic returned to pre-September 2001 levels in 2005 and is expected to continue to grow, according to the most recent forecasts by the Federal Aviation Administration.

The relationship between competition and service quality is theoretically ambiguous (Spence 1975, 1976), but has been studied empirically for airline markets by Mayer and Sinai (2003a,b), Mazzeo (2003), and Rupp et al. (2006). Mazzeo and Rupp et al. find that there are shorter and fewer flight delays on routes and at airports with more competition. Mayer and Sinai (2003a) focus on the role of airport concentration in internalizing the externalities that flight delays impose on other flights at the airport. They find that hubs have longer delays than non-hubs, but airport concentration has a negative effect on the length of delays. Studies on the relationship between competition and product quality in other industries include Hoxby (2000) for education and Dranove and White (1994), who review the evidence for hospital markets.

The AIR-21 Act affected takeoff and landing restrictions at LaGuardia Airport in New York City from March 2000 to January 2001. Restrictions were lifted on routes to non-hub and small hub airports, but not on routes to larger hub airports. While the new legislation allowed entry only into some markets, the additional flights increased congestion and flight delays for all flights at LaGuardia. I study the price response on the routes for which entry restrictions were left unchanged. In the estimation, I use the policy change to instrument for flight delays, which are potentially endogenous to the price on the route.

The endogeneity of flight delays to airline prices may come from two sources. First, demand shocks that are observable to the airline and to consumers but unobservable to the researcher could lead to a positive correlation between flight delays and prices because not only prices, but also the number of flights and, as a result, flight delays may respond positively to an increase in demand. Second, the existing literature on hub networks also suggests a positive correlation between air traffic delays and ticket prices. While Mayer and Sinai (2003a) find that flight delays are significantly longer at hub airports than at non-hub airports, Borenstein (1989, 1991) and Evans and Kessides (1993) demonstrate that airlines charge higher prices at their hubs and at more concentrated airports. The legislative change at LaGuardia Airport provides a predictable but exogenous variation in flight delays, and can therefore be used as an instrument in the estimation of the price response.

I find that prices fall in response to longer flight delays. My point estimates imply a price reduction of \$1.42 on average for direct passengers per additional minute of flight delay. The findings are robust to various ways of measuring delays and of defining the sample. The price response is substantially smaller for connecting passengers with about \$0.77 per minute of additional delay. Given that these passengers are potentially even more inconvenienced by flight delays than direct passengers, because the delays may cause them to miss connections, this suggests a much lower value of time for this group of passengers than for the direct passengers.

When I interact the effect of flight delays with measures of the competitiveness of the route, I find a significantly larger price response on competitive routes than on non-competitive routes. On competitive routes, the implied price response to an additional minute of delay is \$−2.44 for direct passengers. My findings suggest that a decrease in quality has a strong negative effect on the market price in competitive markets, whereas there is a much weaker effect in markets with low levels of competition. In my setting the extent of the decrease in service quality is exogenous to the markets I study, and thus does not allow firms to endogenously vary the extent to which flight delays increase with the level of competition. My findings on the price responses, however, suggest that firms in competitive markets are hurt much more by a decrease in quality than firms in less competitive markets.

The remainder of the paper is organized as follows: Section 2 provides more detail on the policy change in takeoff and landing restrictions. Section 3 describes the data sources and measures of flight delay and competition used in the empirical estimation of Section 4. Section 5 concludes.

2. The LaGuardia "experiment"

LaGuardia Airport in New York City is one of the four airports in the United States which are governed by the so-called High Density Rule (HDR).[1] The HDR was imposed in 1969 to limit the number of takeoffs and landings at highly congested airports. For LaGuardia, the number of flights was limited to 65 per hour. Takeoff and landing slots were originally allocated based on existing service. Subsequently, it was possible for airlines to trade the rights to using these time slots.

The HDR not only manages the capacity at congested airports, but also has the effect of restricting entry and thereby potentially keeping prices above the competitive level. In the case of the New York City airports, policy-

[1] The other airports are Chicago O'Hare, John F. Kennedy in New York, and Reagan National in Washington, DC.

makers from communities in upstate New York repeatedly voiced concerns that the slot controls limited access to New York City and hampered the economic activity of their communities. Responding in part to these concerns, Congress decided in March 2000 to phase out the HDR by 2007 as part of the Aviation Investment and Reform Act for the 21st Century (AIR-21). A special provision was made for LaGuardia Airport: Here, new service to non-hub and small hub airports in "underserved" communities was to be exempt from the HDR under certain conditions starting April 1, 2000.

After the passage of the AIR-21 bill, airlines filed almost 600 applications for exemption slots at LaGuardia. Previously, the airport had handled 1064 flights daily. By September 2000, about 300 flights a day had been added under the exemptions. As a result, the number of flights surpassed the airport's capacity even under good weather conditions. In September 2000, LaGuardia accounted for 25% of all flight delays within the United States. Only 44.5% of the flights to or from LaGuardia arrived within 15 minutes of their scheduled arrival, down from 75.2% in the same month of the previous year. In comparison, the numbers for the entire U.S. domestic system were 78.9% and 79.3%, respectively. The flight delays at LaGuardia were highly publicized and travelers were likely to be quite aware of them. Among others, the *New York Times*, a national newspaper, reported several news stories in prominent places about this event throughout the spring, summer, and fall of 2000.

In response to the dramatic increase in flight delays at LaGuardia, the Federal Aviation Administration (FAA) imposed a moratorium on any new exemptions for flights during the peak traffic hours at the end of September 2000. In November of the same year, the FAA announced that it would limit the number of takeoffs and landings at LaGuardia to 75 an hour starting January 31, 2001, increasing the limit by 10 flights per hour as compared to the rule before March 2000. The exemption slots were allocated by a lottery which was held on December 4, 2000. Table 1 provides an overview of the timing of these events.

Table 1
Overview of events

February 9, 1999	The Clinton administration announces plans to phase out the High Density Rule.
March 17, 2000	Congress passes a bill to keep the High Density Rule in effect until 2007, but lift the restrictions on planes with less than 71 seats traveling between small hub or non-hub airports and La Guardia starting April 1, 2000.
September 22, 2000	The FAA issues a moratorium on additional flights during the hours of 8–10 a.m. and 5:30–8:30 p.m. Any new flights already scheduled for these hours need to be rescheduled to another time of the day.
November 11, 2000	The FAA announces that it will limit the number of flights at La Guardia to 75 per hour—an increase of 10 flights per hour compared to the regulations before March, 2000. A lottery will be held to allocate exemption slots among the carriers which have already applied for exemptions from the High Density Rule.
December 4, 2000	FAA holds its lottery. The new slot rules are to be in effect starting January 31, 2001.

This attempt at a deregulation of slot controls at LaGuardia Airport provides an interesting setting to study the effect of flight delays on airline prices. Flight delays, which should be considered endogenous to the prices on a route, increased due to an exogenous change in legislation. In many contexts, the effect of flight delays on airline prices would be difficult to estimate because – if longer delays were caused by more flights at the airport – prices might decrease both due to an increase in competition and because service quality deteriorated. These effects would typically be hard to separate. However in this setting, the new regulation only allowed new flights to be scheduled on routes that were considered 'underserved' and were going to small hubs or non-hubs. These routes experienced significant entry after the deregulation. Examples of destinations that were entered from LaGuardia are Burlington, VT, Buffalo, NY, and Rochester, NY. On all other routes that were *not* considered to be 'underserved' – mostly medium-sized and large hubs – no new flights were permitted. As a result, these routes experienced an exogenous increase in flight delays but the routes were not subject to any increase in the level of competition. In the empirical estimation, I investigate the price response to increased flight delays on this subset of routes. Since the new regulatory rules were not designed to affect any of these routes directly and had no effect on these routes except through the longer flight delays, I use the timing of the regulatory change to instrument for the length of flight delays.

One potential concern in using this instrumenting strategy, which relies on-time period fixed effects as instruments, is that it no longer allows me to control for time fixed effects directly in the estimation. This will be important if there is an underlying time trend in prices, for example overall price increases that are due to higher fuel costs, or if there is an underlying time trend in delays, for example due to the expansion in the overall number of domestic flights. I address this concern in the following ways. First, in order to control for price trends that are common to all airports, I estimate a predicted price index for each carrier-route in my sample based on prices on other routes which were not affected by the change in regulation at LaGuardia. Second, I re-estimate my specifications adding routes from Reagan National Airport in Washington, DC as a control group. Reagan National Airport is probably the airport that is most comparable to LaGuardia. Like LaGuardia, it is a slot-controlled airport, it is located in a large metropolitan area with multiple airports and, of those airports, it is the one that is located closest to downtown. It is also geographically close to LaGuardia. The control group specifications allow me to control explicitly for time fixed effects in the relationship between flight delays and airline prices.

An overall time trend in delays during this period, which would have led passengers to expect part of the increase in delays that happened at LaGuardia after the policy change, does not appear to be much of a concern in the data. An analysis of flight delay data for all flights by major U.S. carriers from 1996–2000 shows an annual trend in flight delays that is in fact slightly downward, with a reduction in the average delay of a flight of about 6 seconds per year.[2] Compared to the national trend, LaGuardia experienced a significant increase in the average length of flight delays after the deregulation, ranging from an additional 6.8 minutes in the first quarter after the regulatory change to an additional 18.4 minutes in the third quarter after the change. If the analysis is limited to flights between the largest 40 airports in the U.S., there is no evidence of any significant annual trend in delays. LaGuardia still experiences a significant positive shock to its flight delays, here ranging from an additional 7.6 minutes in the first quarter after the deregulation to an additional 18.1 minutes in the third quarter after deregulation.

Another potential concern with the analysis is that airlines might have responded to the longer delays at LaGuardia by rescheduling their flights away from the most congested time periods. However, the slot controls at LaGuardia – which remained in effect for all flights in my sample – meant that there were severe limitations on rescheduling any flights. An airline which would have wanted to change the departure or arrival time for one of its flights would have had to switch the flight to another time slot for which the airline owned unused takeoff or landing rights. In addition, the airlines' flexibility in moving to a less congested time slot was severely limited by the fact that all time slots outside the early morning hours and the late evening hours were equally congested with similar numbers of flights during all half-hour periods from 7:30 am to 8:30 pm. This is another side-effect of the slot controls. Airlines trying to avoid these congested times would therefore have had to schedule their flights in the early morning or late evening, inconveniencing their travelers and likely finding that their passengers would be willing to pay less as a results of that. Furthermore, changes in airline schedules require coordination with the rest of the airline's network and are quite costly. Such changes are typically done with at least 2 months prior notice. In this case, it was unclear to the airlines how long the deregulation of slot controls would remain in effect at LaGuardia, for example because the FAA opposed the deregulation from the very beginning. It is therefore likely that airlines were reluctant to implement schedule changes that might only be temporary.

3. Data and descriptive statistics

3.1. Data sources and sample

This paper draws on two main sources of data: The first is the Department of Transportation's Origin and Destination Data Bank 1A (DB1A), a 10% sample of airline tickets sold in the United States. These data are collected quarterly and include the full itinerary of each passenger, the carrier operating the flight and the price paid for the ticket. The itinerary information allows to distinguish between passengers on direct flights and passengers with connecting flights who change planes during their trip. However, there is no information on the day or time of travel other than the quarter in which the flight was taken. Furthermore, this database does not include the day of purchase or any restrictions imposed on the ticket, such as advance purchase or Saturday night stay-over requirements.

Each period, I observe a large number of different ticket prices in each of the markets in the estimation. The number of observed prices ranges from 140 to 1276 for a market in a given time period. Table 2 reports mean fares and their standard deviation for all markets. The table also shows the 20th, 50th, and 80th percentile of the price distributions.[3] The price distributions are left-skewed with the median fare less than the mean fare. Although we know that prices vary systematically with product attributes such as ticket restrictions, time of

[2] The analysis controls for carrier-route fixed effects and for quarter fixed effects to address seasonality.

[3] An earlier version of this paper included quantile regressions estimating separately the effect of flight delays on all deciles of the price distribution. The results are available upon request.

Table 2
Summary statistics of the fare distribution

Origin/destination	Fare summary statistics				
	Mean	Standard deviation	20th percentile	Median	80th percentile
Atlanta (ATL)	174	130	97	117	229
Boston (BOS)	127	46	86	117	182
Cleveland (CLE)	233	166	99	138	468
Charlotte (CLT)	249	144	105	175	402
Cincinnati (CVG)	232	154	105	140	423
Washington (DCA)	119	50	75	113	157
Denver (DEN)	260	174	151	193	329
Dallas/Ft. Worth (DFW)	311	275	120	149	650
Detroit (DFW)	162	101	101	124	203
Houston (HOU)	207	187	117	131	207
Washington (IAD)	118	52	60	118	164
Memphis (MEM)	203	169	101	120	269
Miami (MIA)	177	119	96	142	215
Minneapolis/ St. Paul (MSP)	318	222	140	166	606
Chicago (ORD)	241	170	105	148	429
Pittsburgh (PIT)	197	113	83	145	311
Raleigh-Durham (RDU)	153	125	73	90	211
St. Louis (STL)	320	236	125	189	609
All routes	200	162	98	135	274

Source: DB1A. Fares are reported as one half of a round-trip fare and in U.S. dollars.

booking, and time of travel, we cannot observe any of these.[4]

The second data source is the Department of Transportation's Airline Service Quality Performance (ASQP) database which contains daily information about scheduled and actual departure and arrival times for each individual flight. These data identify the departure and arrival airport as well as the carrier. All air carriers that accounted for at least 1% of domestic passenger revenues in the year prior are required to report their flights at all airports that account for at least one percent of domestic passenger enplanements. I construct various aggregate measures of delay from these data and match them to the price data from the DB1A.

The sample for the main empirical estimation consists of 18 routes between LaGuardia and airports which did not fall under the AIR-21 exemption rules – i.e. routes on which no new flights were allowed to be added – and for which information on flight delays is available. In addition to these routes, some of the empirical specifications include routes from Reagan National Airport in Washington, DC, as a control group. Following the existing literature on the airline industry, I define a market as a pair of origin and destination cities or a (non-directional) route.[5] I restrict the sample to round-trip, coach-class tickets. The empirical estimation separately considers the effect of flight delays on passengers with direct flights and on passengers who connect between flights.

3.2. Flight delay measures

Since the price data from the DB1A identify the route and carrier that a passenger traveled on, but not the date or time of the flight the passenger took, other than the quarter in which the flight was taken, I need to aggregate the delay statistics over all flights by a carrier on a route in a quarter to match them with the available price data. I construct flight delay measures based on the arrival delay of a flight, i.e. the difference between scheduled and actual arrival time. I aggregate the flight delays in several ways to test whether the results are robust to the aggregation method.

The first aggregated measure I compute are the average minutes of delay per flight during the quarter.[6] Table 3, column 1 shows the mean values of this variable for the observations in my sample for each quarter in 1999 and 2000. In 1999, the average flight delays range from 7.2 to 12.9 minutes. Delays are highest in the second and third quarters of the year. The value for the first quarter of 2000, the last time period before the new legislation went into effect, is very close to the value of the same quarter of the prior year. After the slot controls on other routes were lifted, delays increased to an average of 14.1 minutes in the second and 20.5 minutes in the third quarter. After the passage of the FAA's moratorium on new flights, the average delay declined slightly in the fourth quarter of 2000 to 19.6 minutes.

[4] Berry et al. (1996) address this problem by assuming a bimodal distribution for the error term in a random coefficients model of airline demand. While the empirical price distribution in some of the markets studied here is indeed bimodal, not all of the markets share this feature.

[5] The routes are defined here as city pairs rather than directional routes, e.g. the effect for the Boston–New York market is restricted to be the same as the effect for the New York–Boston market. I cannot reject the equality of the price distributions for round-trip tickets originating at either endpoint in this sample, using Kolmogorov–Smirnov tests. Using the city pair definition allows me to reduce the number of fixed effects in the empirical estimation and to increase efficiency.

[6] Early arrivals are counted as negative delays. Appendix B presents estimation results counting early arrivals as zero delays instead. All results presented in the paper are robust to using that alternative definition.

Table 3
Mean values of delay measures

Year	Quarter	Mean delay (in minutes)	Percentage of flights delayed more than 15 minutes
1999	1	8.3	20.9
	2	12.9	25.9
	3	12.6	24.9
	4	7.2	20.8
2000	1	8.2	21.3
	2	14.1	26.1
	3	20.5	36.2
	4	19.6	39.5

Source: ASQP database.

In addition to the average minutes of delay, I use as a second measure the fraction of flights by a carrier on a route in a quarter which are more than 15 minutes delayed. This is the measure that the FAA uses in its widely published on-time statistics. This variable can be interpreted as a measure of the reliability of the flight time because it measures the probability of arriving very late rather than the expected length of the delay. Its definition implicitly assumes that passengers receive disutility from arriving more than 15 minutes late but not from arriving less than 15 minutes late.[7] Table 3 presents summary statistics on the percentage of flights delayed over 15 minutes for the observations in my sample. In 1999, when LaGuardia was already among the most congested airports in the United States, 20.8 to 25.9% of all flights were at least 15 minutes late. Over the year 2000, these numbers steadily increased and reached 39.5% in the fourth quarter.

Airlines can control arrival delays to some extent by choosing the scheduled duration of the flight. In order to account for this effect I perform a robustness check in which I recompute arrival delays as the difference between the time spent by the traveler from scheduled departure to actual arrival and a 'normal' flight duration computed as the 15th percentile of the actual flight duration for any airline on that route in the corresponding month of 1998. Appendix B shows some results based on this delay measure. The estimation results are robust to using this alternative definition for the delay variable.

3.3. Competition measures

In the empirical estimation, I will test whether the effect of flight delays on airline prices differs with the degree of competition in the market. Table 4 lists the market shares for all airlines and routes in my sample. These market shares are computed based on the airlines' numbers of direct flights on the route. When an endpoint is located in a metropolitan area with multiple airports – this includes LaGuardia – the market share is defined based on all airports in the metropolitan area.[8]

Airline markets are highly concentrated. Typically, no more than three airlines offer direct service on a route, and it is quite common for routes to have direct service by only a single airline. Table 4 shows that most routes in my sample are dominated by one carrier. 13 of the 17 origin and destination pairs[9] have one airline that consistently has over 50% market share, and the market shares are quite stable over time. There are two reasons why the airlines would not have reacted to the change in slot controls by changing their number of flights. First, as discussed above, the typical time lag for changing airline schedules is several months, and there was a lot of uncertainty about how long the legislation would be in effect. Second, not only LaGuardia but also JFK airport is governed by slot constraints. With these constraints in place, there is an option value to keeping slots which reduces the airlines' incentives to decrease their number of flights. I will therefore treat the market shares as predetermined in the empirical analysis. I also define a dummy variable for a market being 'competitive'. This dummy is set equal to one if there is no dominant airline with a market share of more than fifty percent. The markets in this 'competitive' category are Boston, Chicago, Raleigh-Durham, and Washington, DC.[10]

4. Empirical estimation and results

4.1. Empirical framework

I estimate the effect of flight delays on fares for 18 routes from LaGuardia Airport. The sample period is the years 1999 and 2000. Slot controls for some routes to LaGuardia

[7] As a robustness check, I have also estimated all results presented here with the fraction of flights arriving more than 30 or 45 minutes late. Some of these results can be found in Appendix B. All results are robust to using these alternative definitions.

[8] This definition implies that airports within a metropolitan area are close substitutes. In the case of New York City, JFK airport is located 12 miles from LaGuardia; Newark airport is 25 miles away. I tested alternative definitions of market share based only on flights at the endpoint airports and found qualitatively similar results to the ones presented here.

[9] Note that there are two Washington, DC, airports in my sample, Reagan National and Dulles. Table 3 shows summary statistics on fares for both airports separately, whereas Table 4 combines those observations when presenting market shares. In the empirical estimation, I allow for separate fixed effects for routes going to Reagan National and to Dulles.

[10] This measure of route-level competition is highly correlated with the airline's dominance at the non-LaGuardia endpoint of the route.

Table 4
Share of flights on the route for airlines serving LGA (defined within metropolitan area)

Origin/destination	Carrier	1999				2000				Competitive
		1st qr.	2nd qr.	3rd qr.	4th qr.	1st qr.	2nd qr.	3rd qr.	4th qr.	
Atlanta	Delta	0.66	0.70	0.67	0.67	0.66	0.66	0.66	0.64	no
Boston	Delta	0.26	0.28	0.28	0.30	0.32	0.38	0.33	0.33	yes
Chicago	American	0.34	0.37	0.37	0.36	0.37	0.38	0.38	0.39	yes
	United	0.36	0.37	0.37	0.38	0.37	0.39	0.38	0.38	yes
Cleveland	Continental	0.83	0.93	0.93	0.90	1.00	1.00	1.00	1.00	no
Charlotte	US Airways	0.83	0.83	0.82	0.82	0.79	0.80	0.81	0.83	no
Cincinnati	Delta	0.77	0.78	0.78	0.79	0.78	1.00	1.00	1.00	no
Denver	United	0.63	0.63	0.61	0.61	0.61	0.63	0.67	0.69	no
Dallas	American	0.75	0.75	0.72	0.68	0.72	0.72	0.72	0.73	no
	Delta	0.11	0.12	0.12	0.11	0.11	0.11	0.11	0.11	no
Detroit	Northwest	0.73	0.73	0.74	0.71	0.80	0.81	0.81	0.79	no
Houston	American	0.13	0.12	0.12	0.12	0.13	0.13	0.14	0.13	no
	Continental	0.82	0.83	0.84	0.83	0.82	0.82	0.86	0.88	no
Memphis	Northwest	1.00	0.87	0.87	0.87	1.00	1.00	1.00	1.00	no
Miami	American	0.56	0.55	0.59	0.61	0.63	0.61	0.65	0.63	no
	United	0.06	0.06	0.03	0.04	0.04	0.03	0.04	0.04	no
Minneapolis/St. Paul	Northwest	0.77	0.79	0.75	0.80	0.80	0.80	0.81	0.81	no
Pittsburgh	US Airways	0.76	1.00	0.74	0.73	0.74	0.73	1.00	1.00	no
Raleigh/ Durham	US Airways	0.17	0.19	0.22	0.26	0.26	0.22	0.22	0.25	yes
St. Louis	TWA	0.85	0.85	0.82	0.83	0.86	0.87	0.82	0.83	no
Washington, DC	Delta	0.28	0.26	0.26	0.30	0.33	0.45	0.40	0.40	yes
	United	0.16	0.21	0.21	0.15	0.22	0.29	0.25	0.25	yes
	US Airways	N/A	0.20	0.19	0.31	0.34	0.12	0.22	0.22	yes

Source: OAG database.

were lifted in April 2000, so that the year 1999 and the first quarter of 2000 were still unaffected by the deregulation. The second through fourth quarter of 2000 are the time periods which were affected by the policy change. We can distinguish between an expansionary period in the second and third quarters of 2000 and a containment period in the fourth quarter of 2000, after the moratorium imposed by the FAA. The delay effect is identified by the time-series variation over these periods as well as the cross-sectional variation across routes and airlines, controlling for route-airline fixed effects and a hedonic price index for each route in each time period. Since flight delays are potentially endogenous to the pricing decision, I instrument for those using indicator variables for the two post-deregulation periods, the expansionary period in the second and third quarter of 2000 and the containment period in the last quarter of 2000.[11] The estimation does not include any observations from the year 2001 for several reasons. First, the reduction of takeoff and landing slots according to the results of the "slot lottery" went into effect at the end of January, 2001. Observations from the first quarter fall therefore under two different regulatory regimes which we cannot separate in the data. Second, my price data have a large number of missing observations in the second quarter of 2001. Finally, the events of September 11, 2001, caused a large decline in airline demand and an increase in costs related to security measures which affected prices and flight delays in the third quarter of 2001 and thereafter for reasons other than the ones that I am interested in here.

To determine the effect of flight delays on ticket prices, I estimate log price as a function of flight delays controlling for route and airline-specific demand and cost components by including fixed effects for each airline-route pair. I use two alternative definitions of delay, the log of mean delay and the fraction of flights delayed over 15 minutes, in all of the following regressions. In order to control for changes in costs and demand patterns over time, I also include a price index for the airline and the route as a control variable in the estimation. The estimated price equation is as follows:

$$\ln(p_{ijt}) = \beta_0 + \beta_1 \text{delay}_{jt} + \beta_2 \ln(p\_\text{index}_{jt}) + \sum_k \gamma_k r_k + \varepsilon_{ijt} \quad (1)$$

where p_{ijt} is an individual ticket price on airline-route pair j in time period t, delay_{jt} is a measure of delay for airline-

[11] The results were checked for robustness using the years 1998 and 2000 in the estimation instead of 1999 and 2000 to account for the possibility that airlines may have anticipated some of the post-deregulation effects of increased flight delays during the discussions of various legislative proposals in 1999. There was no substantial difference in the results when using this alternate sample period.

route pair j in time period t, p_index_{jt} is the price index, and r_k are airline-route fixed effects with $r_k=1$ if $j=k$ and zero otherwise. ε_{ijt} is an unobservable error term. In the estimation, the standard errors are clustered at the airline-route-time period level. I instrument for the potentially endogenous variable delay_{jt} with *post*1, a dummy which equals one for the second and third quarters of 2000, and *post*2, a dummy which equals one for the fourth quarter of 2000.

The price index is a predicted value from a hedonic price regression for routes between the forty largest U. S. airports, excluding the New York City airports, using data for 1999 and 2000. This regression includes as explanatory variables the distance of the route and its square, the geometric mean of the populations at both endpoint airports, indicator variables for tourist destinations, slot-controlled airports, and hub airports, the Herfindahl index for the route, and time fixed effects. This price index is included in Eq. (1) to control for changes in demand and cost conditions over the sample period that are common to all of the largest U.S. airports. For example, fuel prices rose substantially over the time period studied here. Details on the construction of the price index can be found in Appendix A.

The main specification considers all passengers who flew directly between LaGuardia and one of the other airports in my sample during the years 1999 and 2000. This specification estimates how flight delays affect the prices paid by passengers on direct flights. The advantage of this specification is that the flight delays encountered by these passengers only come from routes involving LaGuardia Airport, and are not confounded by possible changes in delays on other segments of the passengers' flights.

A second set of empirical specifications considers the effect of flight delays on the prices paid by connecting passengers.[12] These passengers tend to have a lower willingness-to-pay for quality than passengers on direct flights, but they may also be inconvenienced to a larger extent by flight delays than direct passengers if the delays cause them to miss their flight connections. It is therefore, *a priori*, unclear if the price effect on connecting passengers would be larger or smaller than that on direct passengers. One could consider passengers who originate at LaGuardia and connect at another airport, as well as passengers who connect at LaGuardia. I focus on the first set of passengers here. In particular, I include in my sample all passengers who originate at LaGuardia, connect to another flight at one of the 18 airports in my sample, and travel on an itinerary on which at least 70 passengers travel during each of the quarters in my sample period.[13] I do not include passengers who connect at LaGuardia because only a single route, Boston to Reagan National Airport in Washington, DC, has a sufficient number of passengers connecting at LaGuardia.

This set of specifications estimates the effect of flight delays on the *first* segment of the flight, from LaGuardia to the hub at which the passenger connects, on the price paid by connecting passengers. This implicitly assumes that there is no systematic change in flight delays on the second segment of the route over time that might affect the ticket price.[14] As for direct passengers, the estimations on connecting passengers also include controls for a carrier-route-time period specific price index and carrier-route fixed effects. The price index is predicted for each origin–destination pair. It does not differ by the airport through which the passenger connected.[15] The carrier-route fixed effects, however, are specific to the connecting airport and the destination airport.

A final set of specifications uses routes between Reagan National Airport and the other airports in my sample, except LaGuardia, as a set of control groups.[16] These specifications include passengers on direct flights only. The use of a control group allows me to include time fixed effects to control for effects that are common to all routes in the treatment and control groups in each time period. Reagan National is a slot-controlled airport like LaGuardia. The slot controls ensure that – as in the case of LaGuardia – entry into routes to and from Reagan National is very limited and, as a result, the level of competition in these markets does not fluctuate much from quarter to quarter. In addition, Reagan National is also located in a large

[12] I would like to thank an anonymous referee for this suggestion.

[13] This corresponds to at least 7 passengers in the DB1A data, and translates to 5 passengers per week on average.

[14] The available data do not allow me to calculate passenger-level delays for connecting passengers because I cannot identify which flights a passenger was booked on and if she missed her connection or not. I use the delay on the first segment as a proxy for the delay on the passenger's entire trip. Increased delays on flights out of LaGuardia likely led to a higher rate of missed connections and longer delays on the passengers' entire trips.

[15] For example, the same value for the price index is predicted for a passenger traveling LaGuardia–Dallas–New Orleans and a passenger traveling LaGuardia–Atlanta–New Orleans.

[16] I would like to thank an anonymous referee for this suggestion.

metropolitan area with multiple large airports and one might expect that the passenger mix at these two airports is fairly similar.

4.2. Estimation results

Table 5 presents the results of the first stage of the instrumental variables (IV) regressions. The table reports the coefficients for *post*1, a dummy for the second and third quarters of 2000, and *post*2, a dummy for the fourth quarter of 2000. The results are reported for my two preferred measures of delay, the log of mean delay and the fraction of flights over 15 minutes delayed. The table is divided into three panels. The first panel shows the results for the sample of direct passengers on LaGuardia routes, the second panel shows the results for connecting passengers, and the final panel shows the sample that includes the routes from Reagan National as control groups. In that final sample, the instruments in the first stage are *post*1 and *post*2 interacted with a dummy that equals one for the LaGuardia routes and zero for the Reagan National routes. Table 5 shows that the instruments have a large positive effect on flight delays, using either delay measure.[17] The *R*-squared is quite high in all specifications, ranging between 0.63 and 0.82.

Table 5
First-stage regressions of delays on instruments

Dependent variable	Ln (Mean delay)	Fraction over 15 min delayed
	(1)	(2)
Panel A: Direct passengers, LaGuardia routes only		
Post1	0.6240 (0.0012)**	0.1049 (0.0001)**
Post2	0.9978 (0.0016)**	0.2012 (0.0002)**
Carrier-route fixed effects	Yes	Yes
R-squared	0.7266	0.8150
Observations	1,096,556	1,096,556
Panel B: Connecting passengers, LaGuardia routes only		
Post1	0.7587 (0.0075)**	0.1094 (0.0009)**
Post2	0.9055 (0.0093)**	0.1958 (0.0011)**
Carrier-route fixed effects	Yes	Yes
R-squared	0.6349	0.7644
Observations	16,190	16,190
Panel C: Direct passengers, including routes to Reagan National Airport as control group		
Post1*LaGuardia	0.5178 (0.0020)**	0.0783 (0.0002)**
Post2*LaGuardia	0.7468 (0.0027)**	0.1695 (0.0003)**
Carrier-route fixed effects	Yes	Yes
R-squared	0.6357	0.6915
Observations	676,791	692,137

Clustered standard errors in parentheses. ** significant at 1%.
Additional regressor, not reported here, is the hedonic price index.

Table 6 reports the second-stage results. Again, the table is divided into three panels for the three different samples. Panel A shows the results for direct passengers on the LaGuardia routes. Column 1 starts with an ordinary least squares (OLS) regression for comparison with the instrumental variables results. This first specification does not include any fixed effects. The effect of flight delays on prices is not statistically distinguishable from zero. The price index has a positive and highly significant coefficient here and in all other specifications, indicating that it is highly predictive of the prices on the route. All standard errors that are reported here and in later tables are clustered at the airline-route-time period level.

The two-stage least squares results are reported in columns 2 through 5. I begin by showing the instrumental variables results without airline-route fixed effects. The coefficient on log mean delays is negative and statistically significant. The estimated elasticity is −0.1465. The comparison with column 1, in which this effect was much smaller and statistically not significant, suggests that the OLS estimate is indeed biased towards zero and instrumenting for flight delays is important. Once I add carrier-route fixed effects in column 3, the estimated coefficient on log mean delays falls to −0.0936 but remains statistically significant. This suggests that heterogeneity across routes and carriers accounts for some of the observed price differences, but there is still a sizeable and significant effect that can be attributed to flight delays. At the sample means, this estimate implies that each additional minute of delay reduces the price for direct airline travel by $1.42. During the time in which the policy of relaxed slot controls was in effect at LaGuardia, average flight delays increased by 7.2 minutes compared to the same period of the prior year, implying a total price reaction attributable to longer flight delays of approximately $10.28 or 5.2% of the average ticket price in 1999 on the routes in the sample. This is a considerable effect, especially given that the airline industry tends to have very small margins on its revenue.

In column 4, I show results for the alternate delay variable, the fraction of flights delayed more than 15 minutes, again including carrier-route fixed effects. I find a large and highly significant negative coefficient of

[17] Note that in the last panel, there are slightly fewer observations for the log mean delay measure because some of the routes in the control group have negative mean delays and therefore have missing values for log mean delay.

Table 6
Results for fares at La Guardia (Dependent variable is log fare.)

	(OLS)	(IV)	(IV)	(IV)	(IV)
	(1)	(2)	(3)	(4)	(5)
Panel A: Direct passengers, LaGuardia routes only					
Ln (Mean delay at LGA)	−0.0018 (0.0166)	−0.1465 (0.0718)*	−0.0936 (0.0207)**		−0.0555 (0.0252)*
Fraction over 15 min delayed				−0.4863 (0.1052)**	
Ln (Price index)	0.5480 (0.0269)**	0.6996 (0.0915)**	0.3559 (0.1038)**	0.4394 (0.0822)**	0.1527 (0.1121)
Carrier-route fixed effects	no	no	yes	yes	yes
R-squared	0.12	0.1	0.17	0.17	0.24
Observations	1,096,556	1,096,556	1,096,556	1,096,556	491,743
Panel B: Connecting passengers, LaGuardia routes only					
Ln (Mean delay at LGA)	−0.1328 (0.0351)**	−0.1338 (0.0548)*	−0.0533 (0.0202)**		
Fraction over 15 min delayed				−0.3037 (0.1073)**	
Ln (Price index)	1.0787 (0.0972)**	1.0786 (0.0969)**	0.2111 (0.1100)+	0.1561 (0.1052)	
Carrier-route fixed effects	no	no	yes	yes	
R-squared	0.11	0.11	0.3110	0.31	
Observations	16,190	16,190	16,190	16,190	
Panel C: Direct passengers, including routes to Reagan National Airport as control group					
Ln (Mean delay at LGA)	−0.0086 (0.0122)	−0.0442 (0.0356)	−0.1377 (0.0327)**		−0.1315 (0.0776)+
Fraction over 15 min delayed				−0.6440 (0.1532)**	
Ln (Price index)	0.4866 (0.0272)**	0.5030 (0.0295)**	0.5956 (0.1377)**	0.5969 (0.1316)**	
Carrier-route fixed effects	no	no	yes	yes	yes
Time period fixed effects	no	no	no	no	yes
R-squared	0.06	0.06	0.1088	0.12	0.11
Observations	676,791	676,791	676,791	692,137	676,791

Clustered standard errors in parentheses. ** significant at 1%, * significant at 5%, + significant at 10%.

−0.4863. At the sample means, this implies that an increase in delayed flights by one percentage point reduces prices by \$0.97. For the period in which the policy was in effect, this delay measure increased by 10 percentage points, implying an average price decrease due to delays of about \$9.72 or 5.0% of the average ticket price in 1999.

Some of the routes in my sample, while not serving an airport that was entered from LaGuardia after the deregulation, have some connecting passengers that travel this route as part of a connecting trip to one of the airports that were eventually entered directly from LaGuardia. In order to ensure that my results are not picking up an effect that comes through competition for these connecting travelers, I estimate a specification in which I drop all routes on which more than 3% of the passengers were connecting passengers traveling to on of the airports which were eventually entered from LaGuardia. The results of this specification are shown in column 5 of Panel A. I find a smaller point estimate in this specification of −0.0555, but I cannot reject equality of this coefficient with the one found for the full sample in column 3. As in the full sample, the coefficient is significantly less than zero. The implied price response of this effect, evaluated at the means of this reduced sample, is \$−1.06 per minute of delay.

Panel B of Table 6 presents the results for passengers who originate at LaGuardia and connect to another flight at one of the 18 airports in my sample. Columns 1 and 2 again show the OLS and instrumental variables results, respectively, in specifications without carrier-route fixed effects. Both sets of results are very similar. The estimated elasticity in the IV specification is −0.1338. Once I add carrier-route fixed effects in column 3, this estimate falls to −0.0533, implying a price response at the sample means of \$−0.77 per additional minute of delay. This is about half the size of the price response that I find for direct passengers. The estimated coefficient on the fraction of flights that are over 15 minutes delayed is −0.3037. Both of these estimates are smaller in magnitude than the ones I find for direct passengers, but as in the sample of direct passengers these estimates are statistically highly significant.

Panel C shows the results for direct passengers in the larger sample that includes the control groups. In this sample, I find no significant effect in the specifications without carrier-route fixed effects, but once I add those fixed effects in column 3 I again find a large and significant negative coefficient on log mean delays of −0.1377. The coefficient on the fraction of flights delayed more than 15 minutes is −0.6440 in this sample. Both of these estimates are larger in magnitude than the corresponding results from Panel A which did not include the control groups. However, I cannot reject that the results in Panel A and in Panel C are equal.

The advantage of using the sample with the control groups is that it allows me to add fixed effects for each time period while instrumenting for delay with the dummies for the time periods of the policy change interacted with a dummy for the LaGuardia routes. In column 5 of Panel C, I show a specification that includes time period fixed effects in addition to the carrier-route fixed effects. This specification does not include the price index. I find a point estimate on log mean delays that is very close to the one in the equivalent specification in column 3. The standard error on the estimate, however, is larger in this specification, as one would expect given that there are fewer degrees of freedom once the time fixed effects are included. This indicates that replacing the price index, which is supposed to control for changes in average demand and cost conditions over time, with simple time fixed effects does not substantially affect the estimate on flight delays. This finding is important because it alleviates concerns with the specifications in Panel A, which include only LaGuardia routes and in which the instrumenting strategy, which relies on-time effects, does not allow to include time fixed effects separately.

Table 7 presents the results of regressions in which the flight delay variables are interacted with measures of the competitiveness of the route. Panel A shows the results for the sample of direct passengers on the LaGuardia routes. All regressions are based on the preferred specification from column 3 of Table 6, which includes carrier-route fixed effects, and add an interaction term for flight delays. In column 1, the interaction is with the dummy variable for 'competitive' routes.[18] The direct effect of flight delays is negative with a point estimate of −0.0351. The standard error on the estimate is rather large, with a *p*-value of the estimate of 0.13. The interaction term with the dummy for competitive routes has a large negative coefficient of −0.1600 and is statistically highly significant. At the sample averages, this implies a price reaction on these routes of \$−2.44 for each additional minute of delay, a substantially larger effect than the average price reduction implied for all routes.

In column 2, I include an interaction term for flight delays with one minus the market share of the airline. I also control for the direct effect of one minus the market share. As in column 1, the point estimate on the direct effect of flight delays is negative but estimated with a large standard error. Here, the *p*-value is 0.16. The interaction term of one minus the market share with log mean delays is again negative with a point estimate of −0.1418. This effect is significant at the 10% level. The direct effect of the market share is not distinguishable from zero. This is not surprising since the regression includes route-carrier fixed effects so that the effect on market share is only identified by the very small changes in market shares within routes and carriers over time.

Columns 3 and 4 show the results for the alternative definition of the flight delay variable, the fraction of flights delayed more than 15 minutes. Column 3 shows negative and highly significant coefficients for non-competitive markets and for competitive markets, with point estimates of −0.2949 and −0.6802, respectively. The implied price response at the sample means for a one percentage point increase in flights delayed over 15 minutes is \$−1.36 in competitive markets and \$ −0.59 in non-competitive markets. In column 4, I interact the fraction of flights delayed more than 15 minutes with one minus the market share. The results are very similar to the specification with the competition dummy, but the standard errors on the estimated effects are substantially larger.

To further investigate the source of the price response in competitive markets, I estimate specifications in which I include an airline's own flight delays and its competitors' delays at the other New York City airports. I restrict the estimation to competitive routes. The results are presented in columns 5 and 6. These regressions test the hypothesis that a firm's price is increasing in its own quality but, controlling for its own quality, price is also decreasing in the competitors' quality. I estimate these specifications with OLS because I have no variables to instrument for delays at the airports which were not affected by the policy change. The previously reported OLS results imply that we should expect the OLS coefficients to be biased towards zero, i.e. the true effects are likely to be larger than the ones I find here. As expected, I find that there is a negative effect of the airline's own delays on the route and a positive effect of the competitors' delays. The elasticity with respect to the airline's own mean delay is estimated to be −0.0538, and the elasticity with respect to the competitors' delays is 0.0414, only slightly smaller in magnitude than the effect of own delays. I find qualitatively similar results using the alternate delay measure in column 6.

[18] In this specification, the direct effect of the competition dummy is absorbed by the route fixed effects.

Table 7
Delay effect by level of competition (Dependent variable is log fare)

	(IV)	(IV)	(IV)	(IV)	(OLS)	(OLS)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Direct passengers, LaGuardia routes only						
Ln (Mean delay at LGA)	−0.0351 (0.0232)	−0.0426 (0.0303)			−0.0538 (0.0258)*	
Fraction over 15 min delayed			−0.2949 (0.1011)**	−0.2537 (0.1652)		−0.7107 (0.1813)**
Ln (Mean delay at LGA)* competitive	−0.1600 (0.0683)*					
Fraction over 15 min delayed* competitive			−0.6802 (0.2125)**			
Ln (Mean delay at LGA)* (1 — flightshare)		−0.1418 (0.0728)+				
Fraction over 15 min delayed* (1 — flightshare)				−0.6769 (0.3457)+		
(1 — flightshare)		0.1167 (0.1505)		0.0750 (0.1114)		
Ln (Competitors' mean delay at other airports)					0.0414 (0.0239)+	
Competitors' flights over 15 min delayed						0.7255 (0.2765)*
Ln (Price index)	0.3460 (0.1516)*	0.3185 (0.1176)**	0.5115 (0.0899)**	0.4539 (0.0911)**	−0.1099 (0.1487)	0.2578 (0.1743)
Carrier-route fixed effects	yes	yes	yes	yes	yes	yes
R-squared	0.16	0.16	0.17	0.17	0.1820	0.1860
Observations	1,096,556	1,096,556	1,096,556	1,096,556	468,433	482,004
Panel B: Direct passengers, including routes to Reagan National Airport as control group						
Ln (Mean delay at LGA)	−0.0782 (0.0321)*	−0.0511 (0.0720)				
Fraction over 15 min delayed			−0.3676 (0.1403)*	0.0845 (0.2756)		
Ln (Mean delay at LGA)* competitive	−0.2102 (0.0803)**	−0.1677 (0.0645)**				
Fraction over 15 min delayed* competitive			−0.8580 (0.2352)**	−0.6300 (0.2036)**		
Ln (Price index)	0.6947 (0.1822)**		0.7892 (0.1992)**			
Carrier-route fixed effects	yes	yes	yes	yes		
Time fixed effects	no	yes	no	yes		
R-squared	0.1	0.11	0.1110	0.1183		
Observations	676,395	676,395	691,741	691,741		

Clustered standard errors in parentheses. ** significant at 1%, * significant at 5%, + significant at 10%.

Panel B of Table 7 shows specifications with interactions of the flight delay variables with the dummy for competitive routes for the sample that includes the control groups. Column 1 starts by re-estimating the specification from column 1 in Panel A with log mean delays and carrier-route fixed effects on this larger sample. The coefficients are slightly larger than before, and both the direct effect of delays and the interaction effect are statistically highly significant. In the next column, I replace the price index with time fixed effects. The coefficient on the direct effect on delays now loses significance, but the magnitude of both the direct and the interacted effects is similar to the previous estimates. I then repeat both of these specifications with the alternate delay measure and again find similar results as before. Overall, the results from the larger sample in Panel B suggest that the previous findings are highly robust to including the control groups. I do not present results for connecting passengers in this table because all markets for connecting flights are competitive according to my definition.

4.3. Discussion of the results

I find that airline prices fall in response to longer flight delays. The price elasticity with respect to flight delays that I estimate in this paper is likely to be a lower bound on the effect, since one would expect at least a small increase in marginal costs of operation to airlines as a result of the increased level of congestion at LaGuardia. This increase works against the price decrease from reduced demand. I find a smaller effect for connecting passengers than for direct passengers. This indicates that those passengers are less sensitive to flight delays, even though they may be inconvenienced to a greater extent by delays if they miss their connecting flights. The results are robust to including routes to and from Reagan National Airport as a control group. This finding is important because including the control groups allow me to control explicitly for time fixed effects. I find that including those fixed effects rather than the price index, which is predicted for each carrier-route-time period based on observations from other large U.S. airports, does significantly alter my coefficient estimates.

The results of my estimation are quite consistent with previous findings in the literature. Morrison and Winston (1989) estimate the effect of flight delays on airline demand in a simple logit model and find that an increase of one percentage point in the share of flights delayed more than 15 minutes reduces passengers' willingness-to-pay by $0.61, measured in 1983 dollars. This is equivalent to $1.05 in 2000 when inflated by the Consumer Price Index. My findings imply a price reaction of $0.97 in 2000 dollars. In a different study, the United States Department of Transportation recommends a "[value] for aviation passenger travel time" of $0.50 per minute in 1995 or $0.55 per minute in 2000 dollars based on survey results (Federal Aviation Administration, 1997). I find a substantially larger effect of $1.42 per minute, but my sample likely contains travelers with a higher value of time, including a much larger fraction of business travelers, compared to a national sample of travelers.

I find a big difference in the price response between competitive and non-competitive markets. The result that prices fall much more in competitive markets sheds some light on the link between service quality and competition in airline markets. Previous studies by Mazzeo and Rupp et al. have found that on-time performance is better on routes and at airports with more competition, and Suzuki (2000) finds that market shares are positively related to on-time performance. My findings suggest that a decrease in quality has a strong negative effect on the market price in competitive markets, whereas there is a much weaker effect in markets with low levels of competition. In my setting the extent of the decrease in service quality is exogenous to the markets I study, and thus does not allow firms to endogenously vary the extent to which flight delays increase with the level of competition. My findings on the price responses, however, suggest that firms in competitive markets are hurt much more by a decrease in quality than firms in less competitive markets. This is in line with their findings of Mazzeo and Rupp et al. that, at least in this industry, firms who can choose their level of quality will choose a higher level of quality in markets where they face more competition.

5. Conclusion

A legislative change in takeoff and landing restrictions at LaGuardia Airport provides an opportunity to study the effect of an exogenous shock to product quality on prices in the airline industry. I find that prices fall as flight delays increase, and that the price decrease is substantially larger in more competitive markets. Prices fall by $1.42 on average for direct passengers and by $0.77 on average for connecting passengers for each additional minute of delay. On competitive routes, the implied price response is substantially larger with $2.44 per minute for direct passengers.

In the aftermath of September 11, 2001, the airline industry has suffered from numerous disruptions to operations and a sharp decrease in demand. However, demand for air travel returned to pre-September 2001 levels in 2005. As a consequence, the issue of flight delays and airport congestion has re-emerged as an important policy question. The lesson from the policy "experiment" at LaGuardia is that increased airport congestion leads to substantial reduction in the prices that passengers are willing to pay. This is especially true on competitive routes, while on less competitive the price elasticity is much smaller, suggesting that on those routes airlines can shift a larger proportion of the welfare cost of longer flight delays to consumers.

Appendix A. Construction of the hedonic price index

This section details the construction of the price index which is used as a control variable in the regressions of Section 4. The index is based on a hedonic price regression for routes between the 40 largest U.S. airports as measured by domestic passenger enplanements, excluding LaGuardia, John F. Kennedy and Newark Airports. I estimate the log of the coach-class fares on route j with carrier l in time period t as a function of the geometric mean of the populations of the endpoint cities of the route, indicator variables for at least one of the endpoints of the route being a hub airport, a tourist destination, or a slot-controlled airport, respectively, the Herfindahl index of the route based on the share of passengers, fixed effects for each carrier and each time period. I also include the distance of the route and its square and interactions of the distance and the squared distance with the time fixed effects. Table A.1 shows the results of this estimation. This specification is chosen to control for route-level observables which influence the price. I checked the estimation results for robustness using different functional forms and including the log of fuel prices instead of time fixed effects. None of these variations had a significant effect on the results in Section 4. The results of this regression are then used to predict a price index for the LaGuardia routes in my sample, and for the control routes involving Reagan National Airport.

Table A.1
Construction of the route-level price index

Dependent variable	Ln (Fare)
Ln (Mean population)	−0.1357 (0.0003)**
Hub	0.3598 (0.0003)**
Tourist destination	−0.2084 (0.0002)**
Slot controls	0.1431 (0.0002)**
Route HHI	0.1821 (0.0004)**
R-squared	0.7
Observations	6,648,007

Results from an ordinary least squares regression. Additional regressors that are not reported time fixed effects interacted with distance and distance squared and carrier fixed effects.
Robust standard errors in parentheses.

Appendix B. Results using alternative delay measures

This appendix explores the robustness of the main result in Table 6 to changing the definition of the delay variable. Table A.2 shows the results for the base regression of Table 6 column 3 using four alternative measures of flight delays. I start in column 1 with the log of mean delays, where the mean is computed averaging over positive delays only counting early arrivals as a delay of zero minutes rather than as a negative delay. This definition would be appropriate if passengers derived disutility from late arrivals but no utility from early arrivals. The estimation results show a negative and statistically significant effect on prices. The estimated coefficient is larger than the one found in Table 6 for mean delays but, statistically, I cannot reject that the two coefficients are equal.

Column 2 shows results for mean delays adjusted for schedule changes. This variable does not simply take the arrival delays reported by the airlines. Instead, it computes delays as the difference between the actual flight time and a reasonably achievable flight time. The latter is defined as the 15th percentile of the flight time distribution on the route in 1998. Here, I find again a negative point estimate for the delay effect, but it is not statistically significant in this specification.

Finally, in columns 3 and 4 I explore the effects of alternative definitions of our other delay variable, the fraction of flights delayed. As explained above, using this variable implicitly assumes that only delays over a certain length affect passengers' utility. Here, I vary the length of that delay from 15 minutes to 30 or 45 minutes. On would expect that longer delays should have a larger effect on prices and the results suggest that this is indeed the case. The point estimates are negative and significantly different from zero for both the flights delayed over 30 minutes and over 45 minutes. The coefficient estimates are larger in magnitude for longer delays.

Table A.2
La Guardia results, varying delay definition (dependent variable is log fare.)

	(IV)	(IV)	(IV)	(IV)
	(1)	(2)	(3)	(4)
Ln (Mean delay at LGA), positive	−0.1353 (0.0296)**			
Ln (Mean delay at LGA), schedule-adjusted		−0.0406 (0.0318)		
Fraction over 30 min delayed			−0.6141 (0.1274)**	
Fraction over 45 min delayed				−0.8592 (0.1823)**
Ln (Price index)	0.4057 (0.0863)**	0.2599 (0.1199)*	0.4713 (0.0855)**	0.4497 (0.0873)**
R-squared	0.17	0.17	0.17	0.17
Observations	1,096,556	1,090,121	1,096,556	1,096,556

Additional regressors that are not reported here are carrier-route fixed effects.
Standard errors in parentheses. Standard errors are clustered at the carrier-route-time period level.

References

Berry, S., Carnall, M., Spiller, P.T., 1996. Airline Hubs: Costs, Markups and the Implications of Customer Heterogeneity, NBER Working Paper No. 5561.

Borenstein, S., 1989. Hubs and high fares: airport dominance and market power in the U.S. airline industry. RAND Journal of Economics 20 (1), 44–65.

Borenstein, S., 1991. The dominant firm advantage in multiproduct industries: evidence from U.S. airlines. Quarterly Journal of Economics 106, 1237–1266.

Dranove, D., White, W.D., 1994. Recent theory and evidence on competition in hospital markets. Journal of Economics and Management Strategy 3 (1), 169–209.

Evan s, W.N., Kessides, I., 1993. Localized market power in the U.S. airline industry. Review of Economics and Statistics 75 (1), 66–75.

Federal Aviation Administration, 1997. Treatment of values of passenger time in air travel. http://api.hq.faa.gov/economic/742SECT1.PDF.

Hoxby, C.M., 2000. Does competition among public schools benefit students and taxpayers? American Economic Review 90 (5), 1209–1238.

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

Mayer, C., Sinai, T., 2003b, Why do airlines systematically schedule their flights to arrive late? mimeo, Wharton School, University of Pennsylvania.

Mazzeo, M.J., 2003. Competition and service quality in the U.S. airline industry. Review of Industrial Organization 22 (4), 275–296.

Morrison, S., Winston, C., 1989. Enhancing the performance of the deregulated air transportation system. Brookings Papers on Economic Activity 61–123.

Rupp, N.G., Owens, D.H., Plumly, L.W., 2006. Does competition influence airline on-time performance? In: Lee, D. (Ed.), Advances in Airline Economics, vol. I. Elsevier.

Spence, A.M., 1975. Monopoly, quality, and regulation. Bell Journal of Economics 6 (2), 417–429.

Spence, A.M., 1976. Product differentiation and welfare. American Economic Review 66, 407–414.

Suzuki, Y., 2000. The relationship between on-time performance and airline market share: a new approach. Transportation Research Part E, Logistics and Transportation Review 36, 139–154.

U.S. Department of Transportation, 2000. Air Travel Consumer Report, various issues.