THE EFFECT OF AIR TRAFFIC DELAYS ON AIRLINE PRICES*

Silke I. Januszewski
University of California, San Diego

July 2004

Abstract

The relationship between air traffic delays and competition has received much attention recently. I estimate the price responses to longer flight delays in competitive and non-competitive markets, using a policy change to instrument for potentially endogenous flight delays in the price regressions. I find that prices fall substantially in competitive markets but not so in non-competitive markets. Quantile regressions on the price distribution imply that prices fall most for business travelers who have sufficient flexibility to switch to more restricted tickets while the price decrease is smallest for the least elastic travelers at the top of the price distribution.

JEL Codes: L11, L15, L93

*I would like to thank Severin Borenstein for providing me with the DB1A data, Amadeo Odoni and the ICAT lab at MIT for giving me access to the ASQP database, Nancy Rose and Glenn Ellison for their guidance and advice, Sara Ellison and Paul Joskow for detailed comments on an earlier draft, Chris Hansen and Victor Chernozhukov for sharing their code for estimating Instrumental Variables Quantile Regression, Mara Berman, Stephane Bratu, Andrew Hertzberg, Chris Mayer, Andrew Sweeting and various seminar participants for helpful comments and discussions, and the Alfred P. Sloan Foundation for financial support through the MIT Global Airline Industry Program.

I Introduction

Airport congestion and flight delays have received increased attention during the past few years. Mayer and Sinai [2003a, 2003b], Mazzeo [2003], and Rupp *et al.* [2003] study the relationship between flight delays and the level of competition on a route, testing theories about the effect of competition on service quality which go back at least to Spence [1975, 1976]. Concern about flight delays also comes from policymakers and airline executives. Air traffic delays are the most common source of customer complaints by airline passengers (see, for example, 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, 2001, has alleviated congestion concerns for the short term but according to the most recent projections by the Federal Aviation Administration (FAA) the demand for air travel is expected to return to pre-September 2001 levels in 2005.

In evaluating policies regarding flight delays and airport congestion, an important consideration is how passengers react to longer flight delays – for example, by reducing their demand for flights with long delays and choosing to fly from less congested airports. Much of the policymakers' concern regards the harm to consumers rather than to the airlines. However, if consumers get compensated for longer flight delays by paying lower prices for their flights in equilibrium then the harm to consumers is potentially quite small.

I study the effect of increased flight delays on prices paid by air travelers using an exogenous policy change created by the Aviation Investment and Reform Act for the 21st Century

(AIR-21) as an instrument for flight delays. The AIR-21 Act affected takeoff and landing restrictions at La Guardia Airport in New York City over the March 2000 to January 2001 period. Restrictions were lifted on routes to non-hub and small hub airports, but not on routes to larger hub airports. The new legislation allowed entry only into some markets, but the additional flights increased congestion – and, as a result, delays – for *all* flights at La Guardia. I study a subset of routes on which no new flights were allowed to be added but which were nevertheless affected by the increase in airport congestion.

We would expect flight delays to be endogenous to the pricing decision of an airline for several reasons. Demand shocks that are observable to the airline and to consumers but unobservable to the researcher might lead us to find 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. The existing literature on hub networks suggests another reason why we might find a positive correlation between air traffic delays and ticket prices: Mayer and Sinai [2003a] show that flight delays are significantly longer at hub airports than at non-hub airports.[1] Borenstein [1989, 1991] and Evans and Kessides [1993] demonstrate that airlines charge higher prices at their hubs and at more highly concentrated airports, controlling for route-specific demand and cost side variables. The legislative change at La Guardia 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 reaction to an increase in flight delays. My point estimates imply a reduction in fares of $1.07 per additional minute of average flight delays. If I focus on the reliability of arriving on time, I find a reduction in price of $0.64 for a one percentage point increase in flights arriving at least fifteen minutes behind schedule. When I estimate the effect separately for flights with competition from nearby airports which were not affected by the policy change, I find a larger price response to delays on flights which face such competition. The implied fare reduction on competitive routes is $2.92 per additional minute of delay. Interestingly, I find no statistically significant reduction in fares on non-competitive routes. The likely explanation for this finding is that at the same time as passengers' willingness-to-pay was reduced because of longer flight delays, airlines also experienced an increase in marginal costs of serving passengers so that there was no net change in prices. The policy conclusion is that there is a substantial loss in consumer surplus from longer flight delays on routes without direct competition, but a much smaller or no loss in consumer surplus if airlines face competition from less delayed airlines. This has especially interesting implications in light of the fact that many airlines now face increasing competition from new entrants who often choose to fly from less congested airports located in the periphery of metropolitan areas. To my knowledge, such empirical evidence on the question of whether the price response to changes in product quality is larger in competitive markets than in monopoly markets has not been previously presented in the literature.

In order to investigate whether flight delays have a differential impact on the prices paid by different types of passengers, such as "leisure" and "business" travelers, I also estimate quantile regressions on the price distribution, using a new instrumental variables technique developed by Chernozhukov and Hansen [forthcoming]. I find that prices decreased for all travelers. The evidence suggests that the largest decrease was experienced by business travelers who were sufficiently flexible to switch to more restricted tickets. The next largest decrease in prices was for leisure travelers, while prices for very inelastic business travelers at the top of the price distribution decreased by the smallest amount even though this group likely has the highest value of their time.

The remainder of the paper is organized as follows: The next section presents a simple theoretical model of the price response to variations in product quality. Section III provides more detail on the changes in the legislation for takeoff and landing restrictions at La Guardia Airport. Section IV describes the data sources and measures of flight delay used in the empirical estimation of section V, and section VI concludes.

II A Simple Model of the Relationship Between Flight Delays and Airline Prices

In an early model of the airline industry, Douglas and Miller [1974] show how under price regulation competition in service quality can dissipate rents. This section provides a simple model of the relationship between price and quality when both are choice variables. I will show under which conditions decreased product quality causes prices to fall, both in the single-firm case and for the case of a duopoly. This form of the exposition is motivated by the observation that in most of the markets in my sample only one or two airlines offer direct service.[2] The extension of the duopoly case to an oligopoly with more than two firms is straightforward.

Spence [1975] points out that a monopolist sets the level of product quality equal to the valuation of the marginal customer while a social planner would choose product quality based on the average valuation for quality of all consumers. Therefore, a monopolist may over-provide or under-provide quality relative to the social optimum depending on whether the valuation for quality of the marginal customer is greater or smaller than the valuation of the average customer. Mayer and Sinai [2003a, 2003b], Mazzeo [2003], and Rupp *et al.* [2003] study the relationship between service quality and competition in airline markets on the example of flight delays. Their results are mixed: Mayer and Sinai find in their work that delays are longer, i.e. service quality is lower, in more competitive markets. In different specifications, Mazzeo and Rupp *et al.* find the opposite effect. To my knowledge, empirical evidence integrating the effect on prices and investigating the question of whether the price reaction to changes in product quality is larger in competitive markets than in monopoly markets has not been presented previously in the literature.[3]

Airline demand is determined by the number of people who want to travel on a particular route. This depends on the population at both endpoints and on factors such as the distance of the route and whether or not one endpoint is a tourist destination.[4] Furthermore, the demand on the route and the allocation of demand between carriers depend on the quality of flight service. Among the dimensions of service quality which affect demand are flight frequency, on-time arrivals, and in-flight service. In the empirical application, I focus on on-time arrivals for several reasons. Variations in flight frequency are very small on the routes that I study. Other dimensions of service quality, such as in-flight service, are not observable to me. Therefore, I will assume that these are constant within airline and control for these with airline fixed effects.

Airline costs can be separated into airport-, route-, flight-, and passenger-level costs. I will focus on the passenger as the unit of analysis and therefore consider costs on the airport-, route-, and flight-level as fixed costs, and passenger-level costs as marginal costs. The presence of airport- and route-level fixed costs creates barriers to entry so that many routes are only served by a small number of airlines. Berry's [1992] finding that entry into a route is more likely if the entrant is already present at one or both of the endpoint airports provides evidence for the existence of such barriers to entry. Flight-level fixed costs influence the firm's decision of how many flights to offer and which aircraft to fly. This decision determines the airline's capacity on a route. I will assume for simplicity that marginal costs are constant as long as the number of tickets sold on a route is below the total capacity.

Flight delays affect the demand for air travel as well as its costs. Higher expected delay and reduced reliability of on-time arrivals will reduce travelers' willingness-to-pay. On the cost side, air traffic delays affect both the fixed costs of operating a flight and the marginal costs of transporting a passenger. In my data, I observe no reduction in the number of flights on a route when flight delays increase. It is likely that airlines did not reduce their flights at La Guardia because of the option value associated with maintaining the rights to using their allocated time-slots. Increases in marginal or passenger-level costs of delays are due to increased time that passengers spend at the airport or in the plane and require service by the airlines.

If there were no changes in marginal costs, we would be able to interpret the price effect entirely as demand effect because seat capacity stayed constant throughout the time period we observe. It is likely that marginal cost increases were non-zero but fairly small. Unfortunately, we are not able to observe marginal costs directly. Given this, the price reactions we observe are a lower bound to the reduction in willingness-to-pay.

Airlines can influence their flight delays through scheduling decisions, maintenance and crew management, choice of aircraft, and gate management. However, these variables cannot be adjusted in the short run. As a result, expected flight delays are predetermined in the short run but decision variables to the airline in the long run. During the episode studied here, flight delays increased as a result of an exogenous shock to airport congestion. Since the entire airport was affected during all times of the day except the early morning hours, airlines were unable to avoid longer delays by changing their schedules. We observe, for example, that increases in flight delays during this period are not significantly different across markets or airlines.

I model the expected change in prices in response to an exogenous shock to service quality assuming that prices are flexible in the short run but that service quality is not. Assume that the demand for air travel, $D(p, s)$ is a function of price p and service quality s. Further assume that firms maximize their profits by choosing their price and their quality level in a two-staged game, choosing quality in the first stage and price in the second stage. I will take the quality decision of the first stage as given and examine the effect of an exogenous shock to flight delays on the pricing decision in the second stage.

Consider the case of a monopolist with a profit function of the form

$$\Pi(p, s) = D(p, s) * (p - c) - F, \tag{1}$$

where p, s and $D(p, s)$ are price, service quality and demand, as above, c are the marginal costs, and F are the fixed costs. The firm's optimal price is determined by the first-order condition of profit with respect to price:

$$\Pi_p = D(p, s) + D_p(p, s) * (p - c) = 0 \tag{2}$$

where X_i indicates the derivative of variable X with respect to i.

Taking the total derivative of Π_p yields a condition that describes how the monopolist's equilibrium price responds to a change in the quality of the product:

$$\frac{dp}{ds} = -\frac{D_s + D_{ps} * (p - c)}{2D_p + D_{pp} * (p - c)} \tag{3}$$

where again $X_i \equiv \frac{dX}{di}$ and $X_{ij} \equiv \frac{d^2X}{didj}$. The arguments of the demand functions are suppressed.

Assume that demand increases with product quality , $D_s > 0$, and decreases with price, $D_p < 0$, and that price is greater than marginal costs, $p - c > 0$. Then the equilibrium price will increase with the level of product quality if $D_{ps} \geq 0$ and $D_{pp} \leq 0$, i.e. if the derivative of demand with respect to price is increasing in quality and the demand function is convex. If either $D_{ps} < 0$ or $D_{pp} > 0$, then the effect depends on the relative magnitudes of the variables. The price will be increasing in quality if the either the second derivatives or $(p - c)$ are small relative to the first derivatives.

In the more general case of an oligopoly where a firm's demand function depends both on its own price and product quality and on the competitors' prices and quality, the relationship of a firm's equilibrium price to its quality is given as the solution to an implicit function that depends on first and second derivatives of the demand function with respect to prices and quality levels:

$$dp_i(2D_{p_i} + D_{p_ip_i} * (p_i - c)) + dp_{-i}(D_{p_{-i}} + D_{p_ip_{-i}} * (p_i - c)) \\ + ds_i(D_{s_i} + D_{p_is_i} * (p_i - c)) + ds_{-i}(D_{s_{-i}} + D_{p_is_{-i}} * (p_i - c)) = 0 \quad (4)$$

Here, p_i and s_i refer to the firm's own price and quality, respectively, p_{-i} and s_{-i} are the same variables for the competitors. If we assume that the market is served by a duopoly and that the demand function is linear in all of its arguments, this expression simplifies to the following:

$$-(2D_{p_i} + D_{p_{-i}}) * dp_i + D_{p_{-i}} * (dp_i - dp_{-i})) = (D_{s_i} + D_{s_{-i}}) * ds_i - D_{s_{-i}} * (ds_i - ds_{-i}) \quad (5)$$

For well-behaved demand functions, $D_{p_i} < 0$, $D_{p_{-i}} > 0$, $D_{s_i} > 0$, and $D_{s_{-i}} < 0$. If the shock to quality is symmetric such that relative quality levels, $(ds_i - ds_{-i})$, and relative prices, $(dp_i - dp_{-i})$, do not change, then the overall effect on prices is

$$\frac{dp_i}{ds_i} = \frac{-(D_{s_i} + D_{s_{-i}})}{2D_{p_i} + D_{p_{-i}}}, \tag{6}$$

$i = 1, 2$. If the products of the two firms are imperfect substitutes, then we would expect $|D_{p_i}| > |D_{p_{-i}}|$ and $|D_{s_i}| > |D_{s_{-i}}|$, i.e. a firm's demand reacts more strongly to changes in its own price and quality than to changes in its competitor's price and quality. Under these conditions, a firm's equilibrium price moves in the same direction as its quality so that prices decrease in response to an aggregate negative shock to quality.

The partial deregulation of slot controls at La Guardia was a symmetric shock to flights out of that airport. However, there was no shock to the quality at the competing airports in New York City, John F. Kennedy and Newark Airports. For this case, equation 5 above implies that on the routes on which airlines flying to La Guardia and John F. Kennedy and Newark Airports compete, we should see a negative effect on prices at La Guardia and a positive effect on prices at the other airports. On the routes out of La Guardia which are predominantly served by one airline, the model predicts that prices should fall. For general demand functions, the price effect can be smaller or larger on monopoly routes than on duopoly routes depending on the cross-derivatives of the demand function, $D_{p_{-i}}$ and $D_{s_{-i}}$.

III The La Guardia "Experiment"

La Guardia Airport in New York City is one of four airports in the United States which are governed by the so-called High Density Rule. The High Density Rule was imposed in 1969 to limit the number of takeoffs and landings at highly congested airports. For La Guardia, 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 or sell the rights to using these time slots.

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

After the passage of the AIR-21 bill, airlines filed for almost 600 applications for exemption slots at La Guardia. 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, La Guardia accounted for 25 percent of all flight delays within the United States. Only 44.5 percent of the flights to or from La Guardia arrived within 15 minutes of their scheduled arrival, down from 75.2 percent in the same month of the previous year. In comparison, the numbers for the entire U.S. domestic system were 78.9 percent and 79.3 percent, respectively. In response to this dramatic increase in flight delays, the Federal Aviation Administration 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 La Guardia to 75 an hour starting January 31, 2001.[5] 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.

This attempt at deregulation of slot controls at La Guardia Airport provides an interesting setting to study the effect of flight delays on airline prices because delays increased due to a exogenous change in legislation. On routes which experienced entry, prices may decrease both for competitive reasons and because service quality got worse. However, no new flights were allowed to be scheduled on routes to large and medium-sized hub airports, so that entry was restricted to small hubs and non-hubs. We can therefore estimate the price response to increased flight delays on the routes to large and medium-sized hubs which experienced a delay shock but no change in their entry regulations.

Another advantage of this setting is that the flight delays at La Guardia were highly publicized at the time. 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. The large scale of flight delays attracted much attention so that travelers were likely to be quite aware of the fact that expected delays at La Guardia airport had increased.

IV Data and Measurement Issues

IV.1 Data Sources

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 percent sample of all 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, the price paid for the ticket, and a code indicating first-class or coach-class travel. However, there is no information on the day or time of travel other than the quarter in which the flight was taken. Also, 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.

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 at the flight level. All air carriers that accounted for at least one percent 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 DB1A.

Information on flight delays is available for routes to eighteen airports which did not fall under the AIR-21 exemption rules because they were either large or medium-sized hubs or were not "under-served". The total number of passenger observations for routes between these airports and La Guardia is over 1.1 million for the sample period, with observations per airport ranging from 13,000 to 194,000. The smallest markets are the routes between La Guardia and Memphis, TN, Dulles Airport in Washington, DC, and Raleigh-Durham, NC. The largest markets are the routes to and from Chicago O'Hare, Boston, Atlanta, and Washington National Airport, DC.

Following the existing literature on the airline industry, a market is defined as a pair of origin and destination airports, or a *route*. The estimation is restricted to round-trip and coach-class tickets, and non-stop flights. Ticket observations with one or more stops are excluded because our data do not provide us with sufficient information to link a passenger's itinerary with flight schedules to determine the delay experienced by connecting passengers.[6] First-class tickets are not included in the base specification because one would expect the relationship between price and quality for those travelers to be different from the one for coach-class travelers. Robustness checks including first-class tickets are reported in section V.1.

IV.2 Measuring Flight Delays

I construct delay measures based on the arrival delay of a flight, i.e. the difference between scheduled and actual arrival time. Since passengers value their time, their willingness-to-pay should depend negatively on the expected arrival delay of a flight.

As mentioned in the previous section, we need to aggregate delay statistics over all flights in a quarter to match them with the available price data. All aggregated measures are constructed at the level of route, airline, and time period. The first aggregated measure I compute are the average minutes of delay during the quarter. For the standard version of this variable, I count early arrivals as negative delays. The first column of table 2 shows how the mean of this variable develops over time. For the year 1999, the average flight delays range from 7.6 to 14.6 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 16.9 minutes in the second and 23.0 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 22.1 minutes. In appendix A.2 I present some results using an alternative definition of this variable, averaging only over late arrivals and counting early arrivals as zero delays. All results presented in the paper are robust to using this alternative variable definition.

In addition to the average minutes of delay, I use as a second measure the fraction of flights in a quarter which are more than fifteen 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 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 fifteen minutes late but not from arriving *less* than fifteen minutes late. As a robustness check, I have also estimated all results presented here with the fraction of flights arriving more than thirty or forty-five minutes late. Some of these results can be found in appendix A.2. All results are robust to using these alternative definitions. Table 2 presents summary statistics on the flights delayed over fifteen minutes. It shows that in 1999, when La Guardia was already among the most congested airports in the United States, 22.1 to 28.5 percent of all flights were at least 15 minutes late. Over the year 2000, these numbers increased up to 43.1 percent 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 further robustness checks, recomputing arrival delays as the difference between the actual 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. Again, appendix A.2 shows some results based on this delay measure. The results are generally robust to using this definition for the delay variable.

V Empirical estimation

I estimate a price equation for the routes between La Guardia and large and medium-sized hubs in my sample for the years 1999 and 2000. The slot controls for some routes to La Guardia 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 2nd and 3rd 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-specific 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.

I checked the results 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. I found no substantial difference in the results. 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, our 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 we are 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.[7]. I use two alternative definitions of delay, the log of mean delay and the fraction of flights delayed over fifteen 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:

$$ln(p_{it}) = \beta_0 + \beta_1 delay_{jlt} + \beta_2 ln(p\_index_{jlt}) + \sum_{k=1}^{K} \gamma_k r_{jlk} + \epsilon_{it} \tag{7}$$

where p_{it} is an individual ticket price on route j and airline l in time period t, $delay_{jlt}$ is a measure of delay , p_index_{jlt} is the price index, and r_{jlk} are airline-route fixed effects with $r_{jlk} = 1$ if passenger i travels with airline l on route j and zero otherwise. ϵ_{it} is an unobservable error term. Note that all observable regressors are aggregated at the level of carrier, route, and time period so that only the unobservable can explain variations in prices within a route-carrier-time group.

The price index is a predicted value from a hedonic price regression for the forty largest U.S. airports, excluding the New York City airports which are studied here. 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 equation 7 to control for changes in demand and cost conditions that are common to all of the largest U.S. airports. For example, fuel prices are rising substantially over the time period studied here. Details on the construction of the price index can be found in appendix A.1.

Each period, we 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 3 reports mean fares and their standard deviation for all markets. The table also shows the 20th, 50th, and 80th percentile of the price distributions. 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 booking, and time of travel, we cannot observe any of these. Berry *et al.* [1997] address this problem by assuming a bimodal distribution for the error term in a random coefficients model of airline demand. That approach reflects the hypothesis that travelers fall into two main groups, so-called business and leisure travelers. While the empirical price distribution in some of the markets studied here is indeed bimodal, not all of the markets share this feature.

I will proceed by first estimating the effect of flight delays on mean fares using least squares estimation techniques. These estimations are reported in section V.1. Section V.2 will explore the effects on different types of travelers with quantile regression techniques. Under a simple sorting assumption, the quantile regression approach allows us to estimate different effects for passengers who buy high-priced tickets compared to low price travelers.

V.1 Least squares estimation results

This section describes the least squares regressions based on equation 7. We start with table 4 which shows the first-stage regression results of the instrumental variables estimation. The table reports the coefficients for the instruments, a dummy for the expansionary period, *post*1, and a dummy for the containment period, *post*2. The results are reported for the two preferred measures of delay, the log of mean delay and the fraction of flights over fifteen minutes delayed. Table 4 shows for both dependent variables that the instruments have a large positive effect on delays. The R-squared is quite high with 0.71 and 0.81, respectively.

Table 5 contains the main results of this section. Panel A reports the results for log mean delay; panel B contains the results for flights delayed over fifteen minutes. Column 1 starts with an ordinary least squares (OLS) regression, reporting the coefficients for flight delay and for the hedonic price index. Carrier-route level fixed effects are included in the regression but not reported in table 5. The base specification only includes coach class tickets. Panel A shows that the OLS point estimate for the effect of delay on prices is negative but is estimated with a large standard error. Note that throughout this section the reported standard errors are clustered at the airline-route-time period level. Column 2 proceeds to the instrumental variables (IV) estimation. The first specification includes only the delay measure and the hedonic price index as regressors; it excludes the fixed effects. Here, we find a large negative and statistically significant effect of flight delays on airline prices. The estimated elasticity is –0.1027. The elasticity of prices with respect to the hedonic price index is 0.867 and we cannot statistically reject that it is equal to 1, implying that the estimated hedonic price index is a good predictor of actual prices.

Column 3 has the results for our preferred specification including again carrier-route fixed effects to control for unobserved differences across carrier-route pairs and instrumenting for flight delays with our policy dummies. We find a negative and statistically significant effect of flight delays on airline prices. The point estimate for the elasticity is –0.0704. At the sample averages, this implies that an additional minute of delay reduces the price by $1.07.

During the time in which the policy of relaxed slot controls was in effect at La Guardia, average flight delays increased by 7.3 minutes compared to the same period of the prior year, implying a total price reaction attributable to longer flight delays of approximately \$7.70 or 4 percent of the average ticket price in 1999 on the routes in the sample. This is a considerable effect in the airline industry which tends to have very small margins on its revenue.

The comparison with the OLS estimate in column 1 shows that the coefficient from the IV estimation is much larger in magnitude. This implies that the OLS results are biased towards zero.

Finally, the last column of table 4 shows the results from an estimation that includes not only coach but also first-class tickets. Including these tickets has very small effect on the results. The estimated coefficient on log mean delay is now –0.0714 and again it is precisely estimated. Only a little over 2 percent of the passengers in this sample travel first class, and on some routes there are no first-class travelers at all. Column 4 shows that the estimation results are robust to including these passengers.

Panel B of table 5 presents a very similar picture. The delay variable has a negative effect on fares in all specifications. In the OLS estimation, the coefficient is estimated more precisely than before but compared to the IV results it still seems to be biased towards zero. As before the effect is estimated to be larger when carrier-route fixed effects are excluded, and including first-class passengers has little effect on the results. The preferred specification of column 3 gives a coefficient estimate of –0.3209 on the fraction of flights delayed over fifteen minutes. At the sample means, this implies that an increase in delayed flights by one percentage point reduces prices by \$0.64. For the period in which the policy was in effect, these delays increased by over ten percentage points, implying an average price decrease due to delays of \$6.46 or 3.3 percent of the average ticket price in 1999.

The overall picture we get from these results is that longer flight delays reduce the willingness-to-pay of airline passengers and in this case they reduce the price that airlines charge even in a period in which the costs that airlines face may be rising. The estimated effects are substantial and robust to using a variety of different specifications. In the remainder of this section, we will explore how the price effect is related to the degree of competition on the route before we get to the quantile regression results in the next section.

To investigate the effect of competition on the size of the price reaction, I classify routes as "competitive" or "non-competitive" based on whether a dominant airline has a passenger share of greater than 50 percent in the New York City market at the beginning of 1999, where the New York City market is defined as all passengers traveling from another city to La Guardia, John F. Kennedy or Newark Airports. Panel A of table 6 lists these market shares for the airlines flying to La Guardia for all time periods in our sample. The markets classified as competitive are routes to Boston, Chicago, Raleigh/Durham, and Washington, DC. The table shows that flight shares for the airlines flying to La Guardia are very stable over time. Panel B lists the Herfindahl index for all routes to La Guardia based on city-level flight shares. Market concentration varies only slightly more than the La Guardia flight shares. Therefore, I will treat the competitiveness of a route as predetermined for the period studied here. I use the classification into "competitive" and "non-competitive" routes to construct a dummy variable that I will interact with the effect of flight delays to investigate whether the price reactions are systematically different across these routes. Panel B of table 6 shows that when classified by the Herfindahl index rather than by flight shares, Miami might be classified as a competitive route rather than a non-competitive route. I test the sensitivity of my classification both by moving Miami into the competitive group. The results reported below are robust to this modification.

Table 7 presents the results of regressions in which the flight delay variables are interacted with variables for the competitiveness of the route. Panel A shows the results for mean delay, panel B shows them for the fraction of flights delayed more than fifteen minutes as the delay variable. All regressions use the base specification as in column 3 of table 5 and add an interaction term for flight delays.[8] In column 1, the interaction is with a dummy variable that is equal to one for the routes which are classified as competitive. The results show that the direct effect of flight delays is not distinguishable from zero. The interaction term with the dummy for competitive routes has a large negative coefficient and is statistically highly significant. The point estimate is –0.1791. At the sample averages, this implies a price reaction on these routes of \$ –2.92 for each additional minute of delay, almost three times as large as the average price reduction implied for all routes. This means that the net price effect of the increase in flight delays is zero on non-competitive routes while we find a large negative effect on competitive routes.

All other results in table 7 support this conclusion. Column 2 includes an interaction term with one minus the flight share of the airline. The results are very similar as before. The direct effect of flight delays is now estimated slightly larger than before but we cannot reject that the effect is equal to zero. The effect of the interaction term is now –0.1571. Lastly, column 3 shows results including an interaction term with one minus the Herfindahl index of the route. The direct effect is still indistinguishable from zero while the interaction effect is somewhat smaller than before but still negative and statistically significant at the 5 percent level in a one-sided test. Panel B shows the results for these regressions using the alternative definition of the flight delay variable and supports the same conclusions that we got from panel A. The implied price response here is \$ –1.31 in competitive markets for a one percentage point increase in the flights delayed over fifteen minutes.

To further investigate the source of the price response in competitive markets, I estimate specifications in which I include a firm's own flight delays at La Guardia and its competitors' delays at the other airports. 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. Table 8 contains these results. Again, results are reported for the two different delay measures in panel A and panel B, respectively. We first look at the effect on fares at La Guardia airport and then turn to the price effects at the other New York City airports. All estimations reported here are ordinary least squares regressions because we have no variables to instrument for delay at the airports which were not affected by the policy change. The previously reported results from table 5 imply that we should expect the OLS coefficients to be biased towards zero.

Column 1 of table 8 shows the results for fares at La Guardia, restricting the estimation to competitive routes. When controlling for competitors' delay, the elasticity with respect to the airlines' own mean delay is estimated to be –0.0564. The competitors' delay has a significantly positive effect on prices. The elasticity is estimated to be 0.0456. It is only slightly smaller in magnitude than the effect of own delays, indicating that these products are close substitutes.

Next, we turn to the effect on fares at the other New York City airports. First, column 2 shows the effect of airlines' own delays at John F. Kennedy and Newark Airports only for comparison with the results reported for La Guardia in table 5. The effect is negative and statistically significant but it is indeed smaller than the coefficient estimated in the instrumental variables regressions for La Guardia. Once we control for the competitors' delays in column 3, the effect of airlines' own delays on fares doubles to –0.0663. For competitors' delays, we estimate a coefficient of 0.0572. Again, this point estimate is only slightly smaller in magnitude than the one for firms' own delays.

Comparing columns 1 and 3 in panel A, we see that the estimated effects of own delays and of the competitors' delays on fares at La Guardia and fares at the other two airports are very similar. Statistically, we cannot reject that the estimated coefficients are equal. Panel B, which again uses the alternative delay measure, confirms this result. Here, we also find a negative effect of a firm's own delays and a positive effect of the competitors' delays. This supports the hypothesis that a firm's prices are increasing in its own quality and decreasing in the competitors' quality. In these regressions, we find smaller point estimates for the effects at John F. Kennedy and Newark Airports than at La Guardia but again we cannot statistically reject that the effects are equal.

Finally, in table 9 I report results for the effect of increased delays on the revenues from the La Guardia routes in my sample. Column 1 shows results of an OLS regression of log revenue on the dummy variables for the time periods after the policy change, $post1$ and $post2$, controlling for carrier-route fixed effects. The effect is negative and statistically highly significant. The implied effect of the policy change is a loss of $20742 in revenue per quarter for the 21 carrier-route pairs in the sample or $69141 for the time during which the policy was in effect. This calculation does not include other routes and airlines at La Guardia which are excluded from the estimation here because the change in regulation allowed new entry on those routes. Columns 2 and 3 of table 9 show that both delay measures have a large negative effect on revenues.

V.2 Quantile regression results

We now turn to the quantile regression results. The purpose of these regressions is to explore how the price reductions in response to longer delays vary along the price distribution.

Airline prices vary systematically with restrictions imposed on the ticket, such as advance booking and Saturday night stay-over requirements. Unfortunately, we cannot observe these restrictions in the publicly available data. Instead, I rely on the fact that tickets with more restrictions are sold at lower prices that tickets with fewer restrictions. This ranking itself is not affected by the increases in flight delays at La Guardia. Instead, the relative prices of restricted tickets compared to unrestricted tickets may have changed. If this is indeed the case, then we should find different effects of delays at the bottom compared to the top of the price distribution.

Assume that there are different types of consumers who buy airline tickets, such as different types of business travelers who buy tickets with few or no restrictions and leisure travelers who buy tickets with more restrictions. Assume further that these consumer types can be ranked by the restrictiveness of their tickets and that this ranking is not affected by the length of flight delays. Then, we can infer differences in the price responses of these types from differences in the delay coefficients along the price distribution.

Following Koenker and Bassett [1978], the quantile regression analog to the classical least squares regression function, the conditional quantile function, can be written as

$$y_i = x_i'\beta_\theta + u_{i\theta}, \quad i = 1, ..., n \tag{8}$$

where y is the dependent variable, x are the regressors and β_θ is the regression coefficient for the θ'th quantile. The identifying assumption is that the θ'th quantile of y_i conditional on x_i is equal to $x_i'\beta_\theta$,

$$Q_\theta(y_i|x_i) = x_i'\beta_\theta \tag{9}$$

Koenker and Bassett describe algorithms for the computation of the conditional quantile functions.

In my estimations, I use an instrumental variables (IV) estimator for quantile regression functions proposed by Chernozhukov and Hansen [2001, 2002]. Their method is based on an assumption of *similarity* which means that the error term is equal in distribution for all values of the endogenous variable. This assumption is a weaker version of a rank invariance assumption. In our context, this assumption requires that the expected ranking of passengers by the prices that they pay does not change with variations in flight delays. Passengers at different points of the price distribution can react differently to flight delays, but the assumption rules out that a business traveler who pays a higher price in expectation than a tourist traveler when flight delays are low, pays a lower price in expectation than the tourist traveler when flight delays are high. The model also requires that conditional on the endogenous variable X, the error term is independent of the instruments.

Chernozhukov and Hansen derive an estimation equation of the form

$$Pr(Y \leq q(X, \theta)|Z) = \theta \tag{10}$$

for the θ'th quantile, where X is the endogenous variable, Z is the instrument, and $q(.)$ is an unknown function. They demonstrate that the function $q(.)$ can be solved for by inverse quantile regression methods and develop an estimator based on that. The estimates presented here are from IV quantile regressions using that estimator. The standard errors are computed using methods developed by Powell [1984] and Buchinsky [1995].

I specify a conditional quantile function which is analogous to equation 7. The log of ticket price is assumed to be a function of the predicted price index, carrier-route fixed effects, and log mean flight delays, instrumented for with a dummy for the post-policy change period. Results for the alternative delay variable are not reported but are quite similar to the results for log mean delay. I estimate the instrumental quantile regressions for the 20th to 80th percentile of the price distribution. The estimation is restricted to the routes which were identified as competitive routes in the previous section. These are the routes for which the least squares regressions found a significantly negative price effect. The number of passengers stays fairly constant on these routes over time so that the results are unlikely to be contaminated by large changes in the types of travelers on the route. Only coach-class tickets are included in the estimation.

Figure 1 presents a graph of the estimated coefficients on log mean delay. Table 10 reports the coefficients and standard errors for the 2nd through 8th deciles. The results show that the size of the effect ranges from –0.0633 to –0.5826. Interestingly, the size of the effect is not monotonic along the price distribution. Instead, the effect increases in size from –0.1254 to –0.5826 between the 20th and 60th percentile and then starts to decrease to -0.0633 at the 80th percentile. This means that in the lower range of the price distribution, up to the 60th percentile, ticket prices of passengers who paid higher prices declined relatively more than prices for travelers with less expensive tickets. However, at the top of the price distribution we see a smaller decline in prices relative to the middle of the distribution, i.e. for the passengers who bought the most expensive tickets prices declined relatively less.

These results indicate a change in the mix of tickets bought by passengers on the route. The fact that the price effect of flight delays is relatively small at the highest quantiles suggests that prices for unrestricted tickets did not decline very much. These are the passengers which tend to be most inelastic in their travel demand and for whom we would expect the highest mark-ups. The largest price effect is found around the 60th percentile of the price distribution. This effect can be explained by the bimodal nature of the price distributions in most markets. The distributions have one mode at a lower price, presumably the standard restricted coach fare, and another mode at a higher price, presumably the standard unrestricted coach fare. We observe that after the policy change, the mode at the higher fares decreases relative to the modes at the lower fares. This means that the mix of tickets which airlines sell changes such that relatively more low-price tickets are sold. As a result, we find very large price effects for the quantiles which shift from the higher mode to the lower mode of the distribution.

This shift may well be driven by an endogenous reaction of the airlines flying to La Guardia, e.g. by an increase in special promotions in reaction to the decrease in service quality through longer delays. This means that while we find a fairly small price reduction for the most inelastic travelers at the very top of the distribution, passengers who used to buy high-priced tickets but are more elastic in their demand are able to shift to substantially lower-priced tickets. Finally, for the travelers at the bottom of the price distribution who bought restricted tickets before and after the policy change, we find a price reduction but since they remain with the same type of ticket the price reduction is smaller than in the middle of the price distribution.

An interesting implication of this pattern is that while business travelers with unrestricted tickets may have the highest value of their time, they are also the travelers with the most inelastic demand. Therefore, increased flight delays which should hurt these travelers most do not translate into large reductions in fares. In contrast, high-price travelers who are more flexible in their demand are compensated most for longer flight delays through decreases in their fares. Low-price travelers who tend to be leisure travelers and have a lower value of their time also experience a significant price reduction. Interestingly, this price decrease is larger in magnitude than for the highest-price travelers.

V.3 Discussion of the results

I find that prices for air travel fall in response to a deterioration in service quality in the form of longer flight delays. The estimated price elasticity of –0.0704 with respect to flight delays is likely to be a lower bound on the effect on demand for air travel for two reasons. First, we would expect at least a small increase in marginal costs of operation during this period. This increase works against the price decrease from reduced demand. This is also the likely reason for why the net price effect that we find on non-competitive routes is small and cannot statistically be distinguished from zero. Second, traffic from passengers connecting between flights at La Guardia grew slightly from 1999 to 2000. On average, 1.2 percent more connecting passengers traveled in the second through fourth quarter of 2000 on the routes studied here than in the same quarters of the year before. These additional connecting passengers probably increased the opportunity costs for seats on these routes which is another reason why the estimates presented here are likely to underestimate the full demand response. However, neither the marginal cost increase nor the effect of additional connecting passengers are likely to be large. We can therefore regard the estimates presented here as a useful lower bound to the true demand response.

The results of my estimation are also quite consistent with previous findings in the literature by Morrison and Winston [1989]. They 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 \$0.64 in 2000 dollars. In another 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 a survey conducted by the Air Transport Association of America [Federal Aviation Administration, 1997]. Considering that my study focuses on a market which is likely to have a higher share of business travelers than the national average, these findings are quite in line with the price reaction of \$ –1.07 per minute which I find here. Suzuki [2000] finds that market shares are positively influenced by on-time performance which is consistent with the interpretation that demand reacts negatively to flight delays.

Berry *et al.* [1997] estimate airline demand in a discrete choice framework with a random coefficients model. In that model, consumers can differ in their valuations for product characteristics which avoids the unreasonable implications for substitution patterns of the traditional logit model. Berry *et al.* also use the Department of Transportation's DB1A ticket data for their estimation, as do Morrison and Winston. The problem that these data pose for discrete choice estimation is that the consumer's choice set is unobservable. The database reports tickets on a quarterly basis. Within the quarter, there is no information on the date of purchase or date of flight. In addition to that, ticket restrictions are unobservable. Especially for travelers who pay high ticket prices, we would expect that many of the observed low-price tickets are not in their choice set at the time of purchase. As the estimates of the preference parameters are likely to be influenced by the assumptions on the choice set, I restrict my analysis here to a direct estimation of the price effects.

The unobservability of ticket restrictions creates a further difficulty for discrete choice demand estimation because consumers' preferences for price and for product attributes may be correlated. Berry *et al.* address this problem by allowing for correlation in consumer tastes. Specifically, they assume a bimodal distribution of consumer tastes and show that a tri-modal distribution can also be identified by the data. This approach has the advantage that it can identify taste parameters but it requires an assumption on the functional form of the correlation between consumer tastes. In that regard, it is more restrictive than the quantile regression approach followed here. Berry *et al.* find evidence that "business travelers" are willing to pay more for characteristics associated with hub size and for flight frequency than "tourist travelers". In addition, they find that business travelers at hubs are much more likely to fly with the hub airline than with another carrier even at a considerable price premium. This suggests that these travelers have high costs of switching to another airline. These results are consistent with my finding of relatively small price decreases for high-priced tickets in response to longer flight delays. My findings also suggest that there are more than two types of consumers in airline markets as seen by the different price reactions along the price distribution.

We can use the estimates generated here to approximate the welfare effects of the policy change at La Guardia airport. The average delay on all routes at La Guardia increase from 12.49 minutes in the 2nd to 4th quarters of 1999 to 18.93 minutes in the same period of 2000. Using the price elasticity estimated for all routes, this translates into a price response of $ –6.52 per traveler. The DB1A database, which samples 10 percent of all passengers, reports over 1.1 million travelers at La Guardia in the 2nd to 4th quarters of 2000. This translates into a welfare loss from these travelers of over $74 million. If we use the much larger price elasticities estimated for the competitive routes, the implied reduction in willingness-to-pay from longer delays is $ –16.58 per passenger or a welfare loss of $163 million for all passengers in the 2nd to 4th quarters of 2000.

However, the additional flights allocated to the routes to non-hub and small hub airports allowed new entry on these routes. The expectation that prices would fall under increased competition on these routes was indeed the motivation for passing the legislation to lift slot controls. To assess the effect of the legislation on the routes on which entry was now allowed, I perform regressions of the log fare on these routes on the predicted value of the price index and an indicator for the time periods after the legislation was passed, controlling for route fixed effects. The estimated coefficient for the post-legislation time periods is −0.1052. At the mean fare on these routes of $156.27, this translates into a total price effect of $ −14.51 per traveler. This price effect can be decomposed into a reduced willingness-to-pay because of higher delays and into the competitive effect. Using the estimates obtained for the routes which did not experience entry, $5.67 of the price reduction can be attributed to longer flight delays. The remaining $8.84 can be interpreted as the price reduction due to increased competition. With a little over 480,000 travelers recorded in the DB1A database on these routes in the 2nd to 4th quarters of 2000, the estimated welfare gain from increased competition on these routes is about $42.5 million.

These welfare calculations do not include the additional costs to airlines generated by the increased delays. These costs create an additional welfare loss. On the passenger side, there is additional willingness-to-pay for higher flight frequency [Richard, 1998]. However, most of the flights were added by airlines who entered a route which they did not serve before. The frequency of service by carriers with existing service on a route did not increase on average during the post-legislation period. In fact, average flight frequency declined because new entrants offered on average fewer flights than incumbents. So, there is no added welfare gain from increased flight frequency. In sum, the welfare loss incurred by this legislation is between $74–163 million, not counting additional costs to airlines, while the welfare benefit from increased competition is approximately $42.5 million. The net welfare loss is therefore at least $31.5 million for the time period during which the legislation was in effect, and over $120 million if we use the larger estimates.

VI Conclusion

A legislative change at La Guardia Airport in New York City provides an interesting opportunity to study the effect of an exogenous shock to product quality on prices in the airline industry. I find that prices fall as quality deteriorates, but that the price reaction depends on the degree of competition. Routes which only have direct service by one carrier see almost no change in price while prices fall substantially on competitive routes. In the airline industry, in which a many small to mid-sized markets only have one carrier which provides direct service, this means that changes in service quality affect the price paid by consumers in those markets very differently than in larger and more competitive markets.

I find that in the entire sample prices fall by $1.07 for each additional minute of delay. On competitive routes, the implied price response is $2.92 per minute. The theoretical predictions for whether the price reaction to an exogenous change in service quality should be greater for competitive markets than for monopoly markets are ambiguous. To my knowledge, this paper is the first to present empirical evidence on this question.

Simple welfare calculations imply a net welfare loss of at least $31.5 million from the policy change. According to my estimates, Congress imposed a large welfare loss from longer flight delays on all passengers traveling through La Guardia Airport while the new policy was in effect, in an effort to benefit consumers traveling from New York City to a set of small airports. The policy was eventually revoked by the Federal Aviation Administration which cited safety concerns. My findings suggest that the FAA's action also improved social welfare greatly by reducing delay times.

In the aftermath of September 11, 2001, the airline industry has suffered from numerous disruptions to operations and a sharp decrease in demand. It is expected, however, that demand for air travel will return to pre-September 2001 levels by 2005. As a consequence, the issue of flight delays and airport congestion is likely to re-emerge as an important policy question. The effect of air traffic delays on prices in competitive and non-competitive markets should play an important role in the policy considerations.

A Appendix

A.1 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 sections V.1 and V.2. The index is based on a hedonic price regression for the 40 largest U.S. airports as measured by domestic passenger enplanements, excluding La Guardia, John F. Kennedy and Newark Airports. I estimate the log of the average coach-class fare 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. The estimation result is used to predict the mean price for the routes to La Guardia, John F. Kennedy and Newark Airports. This prediction is included as a control variable in equation 7. 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 sections V.1 and V.2.

A.2 Results using alternative delay measures

This appendix explores the robustness of the main result in table 5 to changing the definition of the delay variable. Table A.2 shows the results for the base regression of table 5 column 3 using four alternative measures of flight delays. We 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 that using this definition of the delay variable we find a negative and statistically significant effect on prices. The estimated coefficient is slightly larger than the one found in table 5 for mean delays but, statistically, we cannot reject that the two coefficients are equal. The implied price effects are very close to the ones reported in section V.1.

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. We see that again the effect of flight delays is estimated to be negative and close to the previous estimates. Here, the point estimate is smaller than before but we cannot reject that the coefficients of this estimation and of the preferred specification in panel A of table 5, column 3, are equal. The results in section V.1 appear to be quite robust to changing the definition of the delay variable.

Finally, in columns 3 and 4 we 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' willingness-to-pay.

Here, we vary the length of that delay from fifteen minutes to thirty or forty-five minutes. We expect that longer delays should have a larger effect on prices and the results that we find 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.

References

[1] Berry, S. (1992), "Estimation of a Model of Entry in the Airline Industry", *Econometrica* 60, 889-917.

[2] Berry, S., M. Carnall and P.T. Spiller (1997), "Airline Hubs: Costs, Markups and the implications of customer heterogeneity", *mimeo*, Yale University.

[3] 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.

[4] Borenstein, S. (1991), "The Dominant Firm Advantage in Multiproduct Industries: Evidence from U.S. Airlines", *Quarterly Journal of Economics* 106, 1237-66.

[5] Bratu, S. and C. Barnhart (2002), "A Study of Passenger Delay for a Major Hub-and-Spoke Airline", *mimeo*, Center for Transportation and Logistics, MIT.

[6] Buchinsky, M. (1995), "Estimating the asymptotic covariance matrix for quantile regression models. A Monte Carlo Study", *Journal of Econometrics* 68, 303-338.

[7] Chernozhukov, V. and C. Hansen (forthcoming), "An IV Model of Quantile Treatment Effects", *Econometrica*.

[8] Chernozhukov, V. and C. Hansen (forthcoming), "Instrumental Quantile Regression Inference for Structural and Treatment Effect Models", *Journal of Econometrics*.

[9] Department of Transportation (2000), *Air Travel Consumer Report*, various issues.

[10] Douglas, G.W. and J.C. Miller (1974), "Quality Competition, Industry Equilibrium, and Efficiency in the Price-Constrained Airline Market", *American Economic Review* 64(4), 657-69.

[11] Evans, W.N. and I. Kessides (1993), "Localized Market Power in the U.S. Airline Industry", *Review of Economics and Statistics* 75(1), 66-75.

[12] Federal Aviation Administration (1997), *Treatment of Values of Passenger Time in Air Travel*, http://api.hq.faa.gov/economic/742SECT1.PDF.

[13] Koenker, R. and Basset (1978), "Regression Quantiles", *Econometrica* 46(1), 33-50.

[14] Mayer, C. and T. Sinai (2003a), “Network Effects, Congestion Externalities, and Air Traffic Delays: Or Why All Delays Are Not Evil”, *American Economic Review*, 93(4), 1194-1215.

[15] Mayer, C. and T. Sinai (2003b), ”Why Do Airlines Systematically Schedule Their Flights to Arrive Late?”, *mimeo*, Wharton School, University of Pennsylvania.

[16] Mazzeo, M.J. (2003), “Competition and Service Quality in the U.S. Airline Industry”, *Review of Industrial Organization*, 22(4), 275-296.

[17] Morrison, S. and C. Winston (1989), ”Enhancing the Performance of the Deregulated Air Transportation System”, *Brookings Papers on Economic Activity*, 61-123.

[18] Powell, J. (1984), ”Least absolute deviation estimation for the censored regression model”, *Journal of Econometrics* 32, 143-155.

[19] Richard, O. M (1998) ”An Empirical Analysis of Pricing and Flight Frequency on Airline Routes”, *University of Rochester Financial Research and Policy Working Paper* 97-21.

[20] Rupp, N.G., D.H. Ownes, and L.W. Plumly (2003), ”Does Competition Influence Airline On-Time Performance?”, *mimeo*, East Carolina University.

[21] Spence, A.M. (1975), Monopoly, Quality, and Regulation, *Bell Journal of Economics* 6(2), 417-429.

[22] Spence, A.M. (1976), Product Differentiation and Welfare, *American Economic Review* 66, 407-414.

[23] Suzuki, Y. (2000), “The Relationship Between On-Time Performance and Airline Market Share: A New Approach”, *Transportation Research Part E* 36, 139-54.

Notes

[1]Mazzeo [2003] finds, in contrast, that flights out of and especially into hubs experience *shorter* delays but also finds that larger market shares are correlated with *longer* flight delays.

[2]A market is defined as an origin and destination pair, as is standard in the literature on the airline industry.

[3]Rupp *et al.* include price as well as the level of competition as explanatory variables in their service quality regressions but do not account for the endogeneity of price in those regressions.

[4]See, for example, Borenstein [1989] and Evans and Kessides [1993].

[5]This increased the limit by 10 flights per hour as compared to the rule before March 2000.

[6]Bratu and Barnhart [2001] show with data for connecting passengers of one major airline that the delay experienced by those passengers is much higher than the available flight-level statistics for direct flights would suggest.

[7]The routes are defined here as airport pairs rather than directional routes, e.g. the effect for the Boston-La Guardia market is restricted to be the same as the effect for the La Guardia-Boston market. We 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 airport-pair definition allows us to reduce the number of fixed effects in the estimation and to increase efficiency.

[8]The direct effect of the competition variables is absorbed by the route fixed effects.

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 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.

Table 2: Mean values of delay measures

Year	Quarter	Mean delay (in minutes)	Fraction of flights delayed more than 15
1999	1	9.2	22.4%
	2	14.6	28.5%
	3	14.1	26.8%
	4	7.6	22.1%
2000	1	8.9	23.1%
	2	16.9	29.4%
	3	23.0	39.2%
	4	22.1	43.1%

Source: ASQP database.

Table 3: Summary statistics of the fare distribution

Origin/Destination	Fare summary statistics				
	Mean	Standard deviation	20th percentile	Median	80th percentile
Atlanta (ATL)	188	151	98	118	257
Boston (BOS)	127	46	86	117	182
Cleveland (CLE)	239	171	99	142	468
Charlotte (CLT)	256	147	108	195	402
Cincinnati (CVG)	242	160	107	143	450
Washington (DCA)	119	50	75	113	157
Denver (DEN)	284	231	151	196	348
Dallas/Ft. Worth (DFW)	334	297	120	155	697
Detroit (DFW)	167	109	102	125	203
Houston (HOU)	221	213	118	131	237
Washington (IAD)	120	56	60	120	169
Memphis (MEM)	207	174	102	121	321
Miami (MIA)	197	166	98	144	225
Minneapolis/St.Paul (MSP)	325	228	140	166	606
Chicago (ORD)	250	183	105	151	444
Pittsburgh (PIT)	201	116	83	153	314
Raleigh-Durham (RDU)	156	128	73	90	228
St. Louis (STL)	360	252	129	215	647

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

Table 4: First-stage regression of delay at LGA on instruments.

Dependent variable	Ln (Mean delay) (1)	Fraction over 15 min delayed (2)
Post1	0.4699	0.0880
	(0.0010)	(0.0001)
Post2	0.7431	0.1718
	(0.0013)	(0.0001)
R squared	0.7145	0.8084
Observations	1097000	1097000

Additional regressors that are not reported here are
a hedonic price index and carrier-route fixed effects.
Standard errors in parentheses.

Table 5: Results for fares at La Guardia

Panel A:

Dependent variable	Ln (fare)	Ln (fare)	Ln (fare)	Ln (fare)
	(OLS)	(IV)	(IV)	(IV)
	(1)	(2)	(3)	(4)
Ln (Mean delay at LGA)	-0.0102	-0.1027	-0.0704	-0.0714
	(0.0113)	(0.0562)	(0.0227)	(0.0225)
Ln (Price index)	0.4115	0.8670	0.4565	0.4595
	(0.1219)	(0.0906)	(0.1122)	(0.1116)
Carrier-route effects	yes	no	yes	yes
First-class included	no	no	no	yes
R squared	0.1714	0.1188	0.1693	0.1737
Observations	1097000	1097000	1097000	1123863

Panel B:

Dependent variable	Ln (fare)	Ln (fare)	Ln (fare)	Ln (fare)
	(OLS)	(IV)	(IV)	(IV)
	(1)	(2)	(3)	(4)
Fraction over 15 minutes	-0.1991	-0.4384	-0.3209	-0.3268
delayed	(0.0858)	(0.2210)	(0.1023)	(0.1010)
Ln (Price index)	0.4481	0.8096	0.4751	0.4769
	(0.0788)	(0.0519)	(0.0814)	(0.0804)
Carrier-route effects	yes	no	yes	yes
First-class included	no	no	no	yes
R squared	0.1718	0.126	0.1717	0.1759
Observations	1097000	1097000	1097000	1097000

Standard errors, clustered at the carrier-route-time period level, in parentheses.

Table 6

Panel A: Share of flights on the route for airlines serving LGA (defined within MSA).

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

Panel B: Herfindahl index for the routes to LGA (defined for share of flights within MSA).

Origin/ destination	1999				2000				Competitive
	1st qr.	2nd qr.	3rd qr.	4th qr.	1st qr.	2nd qr.	3rd qr.	4th qr.	
Atlanta	0.4941	0.5379	0.5	0.5008	0.4875	0.4889	0.4926	0.4795	no
Boston	0.209	0.2174	0.2139	0.1956	0.2024	0.2485	0.2212	0.2212	yes
Chicago	0.2806	0.3047	0.3071	0.3086	0.3124	0.325	0.3211	0.3206	yes
Cleveland	0.6976	0.8766	0.871	0.8131	1	1	1	1	no
Charlotte	0.7123	0.7131	0.7096	0.6992	0.6735	0.6838	0.6905	0.7131	no
Cincinnati	0.6442	0.6597	0.6619	0.6657	0.6576	1	1	1	no
Denver	0.4595	0.4582	0.4324	0.4324	0.4336	0.452	0.497	0.5144	no
Dallas	0.589	0.5926	0.563	0.5197	0.5556	0.5568	0.5586	0.5655	no
Detroit	0.5738	0.5746	0.5931	0.5357	0.68	0.688	0.6866	0.6523	no
Houston	0.693	0.7128	0.7174	0.7136	0.6895	0.6913	0.7654	0.7813	no
Memphis	1	0.7732	0.7773	0.7732	1	1	1	1	no
Miami	0.366	0.3636	0.3982	0.4226	0.4517	0.4305	0.4638	0.4471	no
Minneapolis/ St.Paul	0.6499	0.666	0.6269	0.6857	0.6781	0.6819	0.6938	0.6921	no
Pittsburgh	0.638	1	0.6114	0.6085	0.6138	0.6068	1	1	no
Raleigh/ Durham	0.3019	0.4132	0.3944	0.3426	0.3794	0.3966	0.4038	0.41	yes
St. Louis	0.7439	0.7496	0.7099	0.7199	0.7535	0.7689	0.7099	0.7192	no
Washington, DC	0.1929	0.2058	0.2054	0.2385	0.287	0.3187	0.2895	0.2891	yes

Table 7: Delay effect by level of competition.

Panel A:

Dependent variable	Ln (Fare)	Ln (Fare)	Ln (Fare)
	(IV)	(IV)	(IV)
	(1)	(2)	(3)
Ln (Mean delay at LGA)	-0.0033	-0.0200	-0.0252
	(0.0191)	(0.0305)	(0.0294)
Ln (Mean delay at LGA)*	-0.1791		
Competitive	(0.0607)		
Ln (Mean delay at LGA)*		-0.1571	
(1 - flightshare)		(0.0721)	
Ln (Mean delay at LGA)*			-0.1152
(1 - route HHI)			(0.0657)
Ln (Price index)	0.4332	0.3765	0.4016
	(0.1715)	(0.1431)	(0.1373)
R squared	0.1659	0.1665	0.1671
Observations	1097000	1097000	1097000

Panel B:

Dependent variable	Ln (Fare)	Ln (Fare)	Ln (Fare)
	(IV)	(IV)	(IV)
	(1)	(2)	(3)
Fraction over 15 minutes delayed	-0.0637	-0.0608	-0.0731
	(0.0920)	(0.1641)	(0.1756)
Fraction over 15 minutes delayed*	-0.7840		
Competitive	(0.2003)		
Fraction over 15 minutes delayed*		-0.7566	
(1 - flightshare)		(0.3513)	
Fraction over 15 minutes delayed*			-0.6303
(1 - route HHI)			(0.3362)
Ln (Price index)	0.5491	0.4574	0.4660
	(0.0906)	(0.0868)	(0.0855)
R squared	0.1723	0.1713	0.1712
Observations	1097000	1097000	1097000

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.

Table 8: Results including competitors' delays.

Panel A:

Dependent variable	Ln (LGA fare)	Ln (EWR / JFK fare)	Ln (EWR / JFK fare)
	(OLS) (1)	(OLS) (2)	(OLS) (3)
Ln (Own mean delay)	-0.0564	-0.0329	-0.0663
	(0.0289)	(0.0187)	(0.0268)
Ln (Competitors' mean delay at other airports)	0.0456	---	0.0572
	(0.0245)		(0.0273)
Ln (Price index)	0.1781	0.7934	0.7632
	(0.1491)	(0.1692)	(0.1691)
R squared	0.1822	0.1956	0.203
Observations	468433	538747	511188

Panel B:

Dependent variable	Ln (LGA fare)	Ln (EWR / JFK fare)	Ln (EWR / JFK fare)
	(OLS) (1)	(OLS) (2)	(OLS) (3)
Own fligths over 15 minutes delayed	-0.6735	-0.2297	-0.3153
	(0.1617)	(0.1334)	(0.1377)
Competitors' fligths over 15 minutes delayed	0.6770	---	0.2546
	(0.2585)		(0.1389)
Ln (Price index)	0.3653	0.7437	0.6936
	(0.1525)	(0.1657)	(0.1616)
R squared	0.1866	0.1978	0.2071
Observations	482004	545115	514152

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

LGA: La Guardia
EWR: Newark
JFK: John F. Kennedy

Table 9: Revenue effects at La Guardia

Dependent variable	Ln (Revenue)	Ln (Revenue)	Ln (Revenue)
	(OLS)	(IV)	(IV)
	(1)	(2)	(3)
Post	-0.0476		
	(0.0157)		
Ln (Mean delay at LGA)		-0.0866	
		(0.0252)	
Fraction over 15 minutes delayed			-0.4554
			(0.1127)
R squared	0.9369	0.9288	0.9327
Observations	336	336	336

Additional regressors that are not reported here are carrier-route fixed effects.
Robust standard errors in parentheses.

Table 10: Quantile regression results, competitive routes only.

Percentile	Coefficients on Ln(Mean Delay)
20	-0.1254 (0.0039)
30	-0.2527 (0.0042)
40	-0.3512 (0.0048)
50	-0.4374 (0.0043)
60	-0.5826 (0.0081)
70	-0.3332 (0.0118)
80	-0.0633 (0.0065)

Additional regressors that are not reported here are carrier-route fixed effects and a price index. Standard errors in parentheses.

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

Dependent variable	Ln (Mean fare)
Ln (Mean Population)	-0.1331
	(0.0090)
Hub	0.3336
	(0.0082)
Tourist destination	-0.2209
	(0.0079)
Slot controls	0.0114
	(0.0105)
Distance (1000 miles)	0.6201
	(0.0671)
Distance (1000 miles) ^{2}	-0.0998
	(0.0236)
Route HHI	0.1622
	(0.0163)
R squared	0.3258
Observations	14116

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.

Table A.2: La Guardia results, varying delay definition

Dependent variable	Ln (Fare) (IV) (1)	Ln (Fare) (IV) (2)	Ln (Fare) (IV) (3)	Ln (Fare) (IV) (4)
Ln (Mean delay at LGA), positive	-0.0968			
	(0.0299)			
Ln (Mean delay at LGA), schedule-adjusted		-0.0554		
		(0.0291)		
Fraction over 30 minutes delayed			-0.3950	
			(0.1208)	
Fraction over 45 minutes delayed				-0.5725
				(0.1739)
Ln (Price index)	0.4873	0.5575	0.4925	0.4896
	(0.0886)	(0.1135)	(0.0824)	(0.0845)
R squared	0.1711	0.1721	0.1717	0.1716
Observations	1097000	1097000	1097000	1097000

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.

Figure 1: Quantile regression coefficients, 20th - 80th percentile