

Testing Theories of Scarcity Pricing and Price Dispersion in the Airline Industry

Steven L. Puller, Anirban Sengupta, and Steven N. Wiggins¹

October 2008

Abstract

This paper uses a unique new dataset – ticket transaction data – to test between two broad classes of theories regarding airline pricing. The first group of theories, as advanced by Dana (1999b) and Gale and Holmes (1993), postulates that airlines practice scarcity based pricing and predict that variation in ticket prices is driven by differences between high demand and low demand states. Dana's theory predicts that airlines sell tickets with higher and more dispersed prices in unexpectedly high demand states; Gale and Holmes predict that more discounted "advance purchase" seats are sold in off-peak demand periods. Both of these groups of theories predict substantially higher shares of low price tickets in off-peak versus peak flights. The second group of theories, as advanced in the yield management literature, indicate that fare variation is driven by differences in ticket characteristics as associated with price discrimination. We use a census of ticket transactions from one of the major computer reservation systems to study relationships between fares, ticket characteristics, and flight load factors. The central advantage of our dataset is that it contains additional variables not previously available. These variables measure both the ticket characteristics central to the price discrimination theory and information on load factor and peak/off-peak travel times needed to test the scarcity pricing theory. We find only modest support for the scarcity pricing theories – the fraction of discounted advance purchase seats is only slightly higher on off-peak flights and fare dispersion is nearly the same. However, ticket characteristics that are associated with second-degree price discrimination drive much of the variation in ticket pricing.

¹ Texas A&M University. Puller: puller@econmail.tamu.edu. Sengupta: asengupta@econmail.tamu.edu Wiggins: swiggins@tamu.edu We thank Severin Borenstein, Diego Escobari, Li Gan, Jim Griffin, Nancy Rose, Adam Shapiro, and Manuelita Ureta for helpful discussions, and seminar participants at MIT, Texas A&M, Yale and the International Industrial Organization Conference. Manuel Hernandez provided excellent research assistance.

1. Introduction

It is well-known that airline prices exhibit substantial price dispersion. Borenstein and Rose (1994) use 1986 data to show that two randomly selected passengers on the same airline and route will pay an expected difference of 36% of the average ticket price. Fares also vary based upon a wide variety of ticket characteristics, such as refundability, advance purchase discounts, Saturday night stays, and various travel and stay restrictions. Sengupta and Wiggins (2006) establish that these characteristics account for roughly 80 percent of the variation in fares.²

There are two major groups of theories that are used to explain this price dispersion. This paper use new, unique data to test between these theories. The first group posits that airline prices are set to allocate capacity in the context of a market where demand is uncertain and capacity is costly and perishable; that is, a seat on an airline is costly to provide but loses its value if not filled at departure. The leading models of this view are Dana (1999a, 1999b) and Gale and Holmes (1993).³ Under these models, airline seats are priced so that higher prices reflect a lower probability of sale (Dana), or advance purchase discounts are used to encourage travelers with low opportunity cost of time to fly in off-peak periods (Gale and Holmes). Both theories predict higher price dispersion in high demand states because there will be a higher proportion of high fares observed for transacted tickets. In Dana this comparative static is driven by airlines running out of low price seats in high demand states. In Gale and Holmes, airlines offer more discounted fares on low demand flights, resulting in a greater proportion of high price fares on high demand flights.

An alternative theory of airline pricing comes from the revenue management literature, which focuses primary attention on the use of ticket restrictions to engage in a form of second-degree price discrimination. In this literature, ticket restrictions such as non-refundability or stay restrictions are used to create fencing devices between customers with different valuations. In these theories, customers sort based on their willingness to accept restrictions, but restrictions are not intended to move certain customers to off-peak times as in Gale and Holmes.

² For a detailed survey of the history of airline regulation and pricing, see Borenstein and Rose (2007).

³ Dana's model builds on the pioneering analyses of Prescott (1975) and Eden (1990).

These theories are not mutually exclusive. Nevertheless, the revenue management literature focuses primary attention on variation in prices associated with different customer groups, and does not provide the sharp predictions regarding the allocation of capacity in high demand states found in either Dana or Gale and Holmes. This paper tests between these groups of theories.

Testing between models of scarcity pricing and price discrimination has been hampered by a lack of detailed data on airline tickets. In particular, previously available data on actual transactions do not include information on either the ticket characteristics or the flights' load factors that is needed to test either set of theories. The most commonly used data to study airline pricing is the Department of Transportation's Passenger Origin and Destination Survey (Databank 1A/1B), which provides a 10% random sample of domestic U.S. tickets in a calendar quarter. These data do not include either the time of purchase or travel, or load factor data, and also lack data on ticket characteristics. As a result, individual tickets and fares cannot be linked to ticket restrictions or flight-level load factors, precluding the investigation of how prices vary in response to changes in predicted or actual load factor.

Some investigators recently have begun to gather and analyze data regarding *posted* prices gathered from online travel websites such as Orbitz (for example, see McAfee and Velde, 2006). In particular, Escobari and Gan (2007) gather and use data from posted minimum prices to test the Dana's theories.⁴ Unfortunately, we are not aware of any studies that use actual data on ticket *transactions*, or that investigate how the allocation of ticket types and price dispersion vary between peak and off-peak times.

In this paper, we use unique data on ticket transactions to test empirical implications of the leading theories of airline pricing. We directly test the relationships between fares and load factor in order to investigate how the share of high and low price tickets is affected by changes in expected and realized demand, and to assess the associated variation in price dispersion.

⁴ Escobari and Gan collected the posted minimum prices found on *Expedia.com*® for 228 flights departing on June 22, 2006. Their data are more limited than the data presented below in that they consist only of posted prices rather than transaction prices and do not include information regarding the full distribution of actual prices. A central difference between their empirical results and those presented below is that they chronicle a sharp increase in minimum fares (see their Figure 1 and Figure 2) in the last two weeks prior to departure. Our data, based on actual transactions, reveal that transactions continue to occur at very low fares as the departure date approaches. We do not know the reasons for these differences.

To carry out these tests, we use a census of all transactions through one of the major computer reservations systems (CRSs). Our data include measures of ticket restrictions, fares, and flight-level load factor at purchase and departure for each flight segment of a travel itinerary. In addition, our data include the dates of purchase and travel.

The airline pricing problem is quite important, in part because airlines are an important industry. This problem is also important because airline prices are highly dispersed and seemingly complex. In addition, a better understanding of airline pricing can perhaps lead to a better understanding of pricing in related industries, such as the hospitality industry, concerts and sporting events. These industries share a common underlying technology where capacity is costly to provide, demand on a given day or for a given event is uncertain, and capacity loses its value if it is not used. Many such industries, moreover, exhibit highly complex pricing structures that might be driven by either price discrimination or by scarcity pricing. A better understanding of airline pricing can also lead to an improved general understanding of the sources and effects of price dispersion.

The analysis below proceeds in several steps. First, we investigate the hypothesis, common to all the models, that prices are set in advance and that fares associated with particular groups of ticket restrictions do not change as demand uncertainty is realized. Then we test the central hypotheses of Dana and Gale and Holmes. In particular, we test whether there is a higher share of high price, unrestricted tickets on high demand flights, particularly in the last week prior to departure when the airline should have sold out of low priced tickets on peak flights. We also investigate more generally whether there are substantial quantity restrictions on the sale of low priced tickets on high load factor flights. In the Dana model such restrictions occur because the lowest priced tickets sell out, leaving only high priced tickets available on high demand flights. In Gale and Holmes, this variation occurs because consumers who have a low cost to taking the non-preferred flight will buy in advance, and those tickets are only available on the off-peak flight. Hence, we test the central hypotheses regarding the relative sales of high and low priced tickets on peak and off-peak flights shared by both of these theoretical models. We find only modest support for these scarcity pricing theories – the fraction of

discounted advance purchase seats is only slightly higher on off-peak flights. In addition, price dispersion is not substantially larger on high versus low demand flights, as is predicted by the scarcity pricing models. Finally, we investigate whether fares are systematically higher on flights with higher load factors. We find that tickets on flights that are unusually full do have higher fares, but the effect is relatively modest.

In contrast to finding modest support for scarcity pricing, we find that ticket characteristics that are associated with second-degree price discrimination drive much of the variation in ticket pricing. These results, taken together, suggest that scarcity pricing plays a smaller role in airline pricing than models in which ticket characteristics create fencing devices to facilitate price discrimination.

The outline of the paper is as follows. Section 2 reviews the theoretical and empirical literature on pricing and price dispersion in airlines. Section 3 describes our transaction level data. Section 4 discusses our tests of the two classes of pricing theories. Section 5 concludes.

2. Theory on Pricing and Price Dispersion in Airlines

Markets characterized by costly capacity, perishable goods, and uncertain demand often exhibit widely dispersed prices. Such variation in prices is found particularly in airlines, hotels, car rentals and other travel segments. Persistent price dispersion in a perfectly competitive market for homogenous goods was first described by Prescott (1975) and more formally developed by Eden (1990). Prescott (1975) developed a model to describe the inter- and intra-firm price dispersion that is commonly observed in the industries described above. Prescott's model posits a perishable good, such as a concert ticket or an airline seat, that entails costly capacity of λ per unit and perhaps a marginal cost, which we will ignore for now. Following Dana's presentation, it is easiest to think of two demand states, high and low, occurring with equal probability. A certain portion of seats sell out in both states, and the competitive equilibrium price for these seats is $p=\lambda$. Another set of seats sells only when demand is high, and the competitive equilibrium (zero profit) price for these seats is $p=\lambda/(1-\theta)$, where $1-\theta$ is the probability of the high demand state in which those seats would sell. The intuition is that the price must

adjust to cover the full cost of capacity, taking into account the likelihood that the seat does not sell. A substantial literature has built upon this basic Prescott model.

Dana (1999b) provides a more complete description of the model described above and extends the model to monopoly and oligopoly settings. Dana's model has a variety of important implications for pricing in industries with demand uncertain and costly capacity. One of the most relevant implications for airline pricing is that there is a pure-strategy equilibrium that generates intrafirm price dispersion without using restrictions or "fencing" devices such as advance purchase discounts or required Saturday night stays. Dana explains the intuition for the model with the example of selling tickets to an event at a stadium. Suppose a perfectly competitive seller must precommit to a schedule of prices for tickets and cannot adjust prices if anything is learned about the state of demand, e.g. tickets must be printed in advance. Demand is either "high" or "low" with equal probability. Heterogenous consumers with unit demand arrive in random order at the stadium and purchase the lowest priced ticket that is available when they arrive. Dana shows that in equilibrium the firm will offer: (a) some "low" priced tickets that will sell under either state of demand, and (b) some "high" priced tickets that only sell when demand is high. The competitive equilibrium is that the expected revenue from each ticket equals the marginal cost of capacity; thus, higher priced tickets have a lower probability of being sold. Dana also presents monopoly and oligopoly versions of this model. For example, the monopolist prices so that the expected revenue of an additional ticket equals the marginal cost of capacity plus the expected loss in revenue if the additional ticket displaces a higher priced transaction. Under all forms of market structure, firms compete in price distributions and thus there is intrafirm price dispersion.

There are several testable implications of the Dana (1999b) model on *transacted* tickets. The model predicts comparative static relationships between realized load factors and the mean and dispersion of fares. To see this, suppose the analyst observes multiple realizations of flights with the same distribution of demand. For a set of flights with the same (ex ante) distribution of demand, the set of *offered* fares is identical. But the *transacted* fares will differ. Assume, as in the Dana model, that consumers arrive and choose the lowest fare available when they arrive. In a flight with a low realized load factor, only low fare tickets are purchased. In medium load factor flights, the same low

fare tickets are purchased as well as medium fare tickets. And in high load factor flights, the low and medium fare tickets are purchased as are high fare tickets.

The Dana (1999b) model predicts four relationships between flights with different realized load factor but the same ex ante distribution of demand. First, the mean fare of transacted tickets is higher on flights with higher realized load factors. Second, there is more fare dispersion on flights with higher realized load factors. Third, the share of high-priced tickets will be larger in high demand states. Fourth, flights that have an unusually high number of tickets sold as of a given number of days before departure, will sell more high-priced tickets in the final days before departure as compared to flights that are not unusually full.⁵

The model developed by Gale and Holmes (1992, 1993) develops a similar result regarding the sales of high and low priced tickets in peak and off-peak times, but uses a different formal structure. Gale and Holmes use a mechanism design approach to model the use of advance purchase discounts in a monopoly market to divert customers with a low cost of waiting to off-peak flights. In the basic Gale and Holmes (1993) model, each consumer has a preference for either the “peak” or “offpeak” flight, but the consumer’s preferred flight is unknown until shortly before departure. Customers vary in their opportunity cost of waiting. Those customers with low waiting costs are willing to buy tickets off-peak and potentially bear the cost of flying at their less preferred time. Firms and consumers can use advance purchase discounts to contract before the uncertainty regarding preferred flights is resolved. Firms use advance purchase discounts to shift low cost-of-waiting customers to the off-peak flight. Airlines achieve this result by offering (more) advanced purchase seats on the off-peak flight. Advance purchase discounts increase output and surplus relative to the case of selling all tickets at the time of departure. Gale and Holmes (1992) allow for uncertainty in the peak period, and find that at least some advanced purchase tickets are sold on the peak flight. Dana (1998)

⁵ The model also predicts more dispersion in routes that have more competition, which is consistent with results from Borenstein and Rose (1994). However, Borenstein and Rose provide a different model yielding dispersion -- a monopolistically competitive model with certain demand. We do not test predictions regarding market structure because we seek to exploit the strength of our transaction data that include measures of flight-level load factor.

builds upon the advance purchase literature and shows that advance purchase discounts can arise in a perfectly competitive setting.⁶

The main empirical implication of Gale and Holmes models that we test is that peak flights that are expected to be full will have fewer discount/advance-purchase seats sold in equilibrium.

McAfee and Velde (2004) draw upon the yield management literature and devise results for dynamic price discrimination and the efficient allocation of seats when airlines are faced with demand uncertainty. They use data gathered from online websites to study the price paths for specific flights as departure nears. They find only weak evidence of dynamic price discrimination. Prices do not tend to fall as departure approaches despite the fact that the value of an unsold seat goes to zero at departure. Also, there is only weak evidence of the continuous adjustment of prices over time. Our data reveal similar evidence on the evolution of fares as the time of ticket purchase approaches departure.

The empirical literature on price dispersion in airlines is well-developed. Most existing studies have relied primarily on Databank 1B, and its predecessor DB1A, released by the Department of Transportation. DB1B contains information regarding route, fare, carrier, booking cabin and itinerary for a 10 percent random sample of tickets sold each quarter. As discussed above, DB1B is limited in that it does not contain information regarding the flight number, day of the week, date of purchase, load factor, or ticket characteristics such as refundability, advance purchase restrictions, and travel and stay restrictions. Such information is essential for testing the theories put forth by Dana and by Gale and Holmes.

Borenstein (1989) finds a positive relationship between a carrier's share on a particular route and the fares it charges on that route. He also finds that these higher fares do not generally spill over and raise the fares of other carriers on the route. Another strand of the literature has analyzed the effect of market structure on price dispersion. Borenstein and Rose (1994) analyze the relationship between price dispersion and market structure. They show an increase in dispersion as markets become more competitive.

⁶ Many other models argue price dispersion to be an outcome of randomization of prices by firms. Stahl (1989) and Rosenthal (1980) find decreased dispersion in more competitive markets, where the price dispersion in their markets are driven by differences in consumer search and asymmetric information.

Stavins (2002) uses a novel data set on posted prices and a subset of ticket characteristics, namely Saturday night stay-over and refundability, to find evidence consistent with both Saturday night stay and refundability being used as price discriminating instruments. Using these data, Stavins (2002) corroborates the finding of Borenstein (1989) that an increase in a carrier's share is associated with higher prices, and the finding of Borenstein and Rose (1994) that increased competition on a route is associated with higher price dispersion.

Related empirical work has studied other pricing and load factor phenomena. Sengupta and Wiggins (2006) study the effect of on-line sales on pricing. Dana and Orlov (2008) investigate whether the increased use of internet booking leads airlines to increase capacity utilization. Goolsbee and Syverson (forthcoming) investigate the effect of the threat of Southwest entry on incumbent carrier pricing. Forbes (2008) estimates the effect of delays on fares. Other research has studied the effect of airline bankruptcy or financial distress on pricing, including Borenstein and Rose (1995), Busse (2002), and Ciliberto and Schenone (2008). Berry and Jia (2008) explore a variety of demand and supply side explanations for reduced airline profitability in the last decade despite increases in both load factor and passenger miles flown.

Our contribution to this empirical literature is to test comparative static implications of the Dana and the Gale and Holmes models. We test whether the share of high-priced tickets is larger in high demand states, particularly in the period just prior to departure. More generally, we test the hypothesis that more higher-priced, unrestricted tickets will be sold during peak as compared to off-peak flights. We nest this central hypothesis in an empirical model where the baseline is the price dispersion that occurs in off-peak flights—the baseline consists of flights that have a low expected ex ante demand and a low realized ex post demand. We then examine the economic and statistical significance of whether there is an increase in the percentage of high price, unrestricted tickets on flights that have a high expected and high realized demand. Hence the model tests whether scarcity pricing of the type considered by Dana and by Gale and Holmes plays a substantial role in explaining observed levels of price dispersion as compared to a model where airlines use fencing devices as postulated in the yield management literature.

3. Data

3.A. Tickets in Our Sample

We use a census of all transactions provided by one of the major computer reservations systems (CRSs) for the fourth quarter of 2004. This CRS handles transactions for all major channels of ticket purchases: tickets purchased through travel agents, several major online travel sites, and directly from airlines, including their web-based sales. In all, these data comprise roughly one-third of all domestic U.S. ticket transactions. For each ticket sold through this CRS, the data provide information on the fare, the origin and destination, airline, flight number for each leg of the itinerary, dates of purchase, departure and return, the booking class, and whether the ticket was purchased online or offline.⁷

Following Borenstein (1989) and Borenstein and Rose (1994), we analyze the pricing of coach class itineraries with at most one stop-over in either direction. We exclude itineraries with open-jaws and circular trip tickets, and only include itineraries with four coupons or less. We analyze the prices of roundtrip itineraries; we double the fares for one-way tickets to obtain comparability. (We will control for whether tickets are one-way or roundtrip). We exclude itineraries involving travel in the first class cabin. This study includes tickets for travel on American, Delta, United, Northwest, Continental and USAir. These constituted the entire set of airlines that carry at least 5% each of U.S. domestic customers with the exception of Southwest for whom we have only limited data.⁸ We analyze tickets for travel in the fourth quarter of 2004 excluding travel on Thanksgiving weekend, Christmas, and New Years.⁹

We restrict our analysis to 90 large routes. To choose these routes, for each of the six carriers, we stratified the sample to include routes for each carrier with varied market structures.¹⁰ The routes are listed in Table 1. We include tickets by any of the six

⁷ For an analysis of online versus offline prices, see Sengupta and Wiggins (2006).

⁸ Much of Southwest's sales occur through the airline's website.

⁹ We exclude travel occurring from the Wednesday prior to Thanksgiving until the following Monday. Also, we exclude all travel beginning after December 22.

¹⁰ Routes are airport pairs. A route is a monopoly if a single carrier operates more than 90 percent of the

carriers listed above that serve any of the routes listed. One consequence of choosing large routes is that the sample consists largely of routes from airlines' hubs—though this should not pose a problem for testing the general theories of airline pricing under investigation.

3.B. Ticket Characteristics

Because we also wish to observe ticket characteristics that impact a traveler's utility (e.g. refundability, advance purchase restrictions, valid travel days or stay restrictions), we merge our transaction data to information on ticket-level restrictions.¹¹ Travel agents' computer systems can access historical data on posted prices for up to a year. We collected additional data on restrictions from a local travel agent's CRS. The historical archive contains a list of fares/restrictions where transactions occurred for travel on a specified carrier-city-pair-departure date. For each archived fare, we collected information on carrier, origin and destination, departure date from origin, fare, booking class (e.g. first class or coach), advance purchase requirement, refundability, travel restrictions (e.g. travel can only occur on Tuesday through Thursday), and minimum and maximum stay restrictions. We merged these data to the transaction data by carrier, fare, booking class and a variety of ticket characteristics.

The matching procedure is described in detail in the data appendix. Briefly, we match our transacted itineraries to the archive of fares/restrictions based upon carrier, departure date, fare, consistency between purchase date and a possible advance purchase restriction, and tickets where travel dates were consistent with the posited travel and stay restrictions. We kept matches if the tickets met these criteria and the fares were within two percent of each other. If a transaction ticket matches multiple posted fares, we took the closest match based on fare. Details are included in the appendix.

Unfortunately, some transactions did not match the data on posted prices from the travel agent's CRS.¹² Of the routes that we analyze, we were able to match 36 percent of

weekly direct flights. A route is a duopoly if it is not a monopoly route but two carriers jointly operate more than 90 percent of the flights. A route is competitive if it is neither monopoly nor duopoly.

¹¹ For confidentiality reasons, the original CRS did not provide us with the full fare basis code.

¹² The travel agent told us that the historical archive maintained by her CRS would sometimes delete some

the observed transactions. We can assess if there are systematic differences between the matched and unmatched transactions. Table 2 compares means of all transactions to those that we could successfully match to fare characteristics, and indicates only modest differences between the matched and unmatched transactions. The unmatched transactions tend to be slightly lower priced tickets – across all the carrier-routes, the matched tickets average \$424 while all tickets average \$415. The means of ticket characteristics are very similar between matched and all transactions. Matched tickets are slightly more likely to be purchased just before departure and to depart on Monday or Tuesday.

We analyze whether these unmatched tickets tend to come from a certain part of the price distribution. In Figure 1, we plot kernel density estimates of prices. Although we tend to match fares that on average are slightly higher, we are able to match fares from various parts of the fare distribution.

3.C. Measuring Realized and Expected Load Factors

The theory discussed above makes predictions that depend upon two measures of load factor: the *realized* and *expected* load factors. A central feature of our data is that we are able to estimate the load factor at various times prior to and including departure for a given airline, city-pair, and departure time (i.e. a flight). We also can estimate whether the realized demand is particularly high or low for a given flight-departure date within our sample.

To measure load factors, note that we observe all tickets sold through varied outlets by one of the major CRSs accounting for roughly one-third of all ticket transactions. This permits us to estimate total sales at the flight/day level. This estimate can then be combined with data from the Official Airline Guide, which provides the number of seats at the flight/day level, to provide an estimate of load factors.

of the posted fares, but she did not believe the deletion was systematic. We also noted that fare-ticket combinations for more recent travel, as compared to the date we accessed the data, were more complete. Except as noted below, we were unable to find a systematic pattern when comparing the more recent travel dates with the older dates where the records were less complete.

Further, while we do not know the number of tickets sold through other CRSs, we can use the available data to construct an unbiased estimate of these unobserved tickets at the airline-citypair level. In particular, the Bureau of Transportation Statistics reports monthly data on the total number of tickets sold for each city-pair by airline. Using these data we can calculate the exact share of total tickets that we observe in our CRS data for a given airline and city-pair. We then scale up the observed coupons on a particular flight by the inverse of that observed share to obtain an unbiased estimate of realized load factor for a given flight, at a given point in time.

For example, for American Flight 301 from New York La Guardia (LGA) to Chicago-O'Hare (ORD) on October 11, 2004, we measure the number of seats (129) and the number of tickets sold through the CRS that include this flight on its itinerary (26). Because American sells 36% of its tickets for direct service between LGA and ORD through our CRS, we calculate the realized load factor to be 55% ($= (26/0.36)/129$).

Of course, this load factor is measured with error, but the methodology implies that the measurement error will have zero mean at the city-pair, airline level. This procedure should also provide an unbiased estimate of the load factor at the flight level, since the CRS share is unlikely to vary systematically for particular flights or days of the week within a city-pair.¹³ Finally, note that because we observe the sequence of transactions, we also can measure the realized load factor at different dates prior to departure (e.g. the flight is half full as of 7 days before departure and two-thirds full as of 2 days before departure).

Also, we construct a measure of load factor that is systematic (or predictable) by the airline, and call it *expected* load factor. To do so, we calculate the average load factor across our sample for a particular carrier's flight for a specified day-of-the-week of travel (e.g. American flight 301 from La Guardia to O'Hare on Mondays). We have data for tickets sold for departures in a 12 week window. We calculate the average load factor for 12 departures of a given flight number-day-of-the-week, and use it to estimate the average load factor on that airline-flight-day-of-the-week.

The theories of scarcity pricing have several comparative static predictions about the characteristics of tickets sold on flights that are unusually full on peak flights and

¹³ We discuss possible attenuation bias below.

unusually empty on off-peak flights. In some of the analysis below, we separate all flights into groups based upon the expected and realized load factor. Figure 2 illustrates. The columns of the matrix divide flights based upon our measure of expected load factor into groups of expected to be “Full”, “Medium-Full”, “Medium-Empty” and “Empty”.¹⁴ For example, all of American’s flights from La Guardia to O’Hare are grouped into 4 categories of expected load factor based upon the average load factor for each FltNo-day of week. American’s flight 301 on Mondays has a relatively low average load factor (compared to other American FltNo-day of week from LGA to ORD), so all 12 of those flights in our sample are classified as expected to be “Medium-Empty”.

Next, we categorize each flight (i.e. FltNo-Departure Date) in each category of expected load factor by the realized load factor. Continuing the example above, for all American flights La Guardia to O’Hare that are expected to be “Medium-Empty”, we group each flight into 4 categories based upon realized load factor.¹⁵ American’s flight 301 on October 11 with a realized load factor of 55% is among the lowest load factor flights of those in the “Medium-Empty” expected load factor; therefore tickets on this flight are categorized as “*Expected* to be Medium-Empty and *Realized* to be Empty”.

As shown in figure 2, the top left corner consists of flights that are unusually full among the flights that are expected to be full; the bottom right corner consist of flights that are unusually empty among those that are expected to be empty.

3.D. Summary Statistics

Summary statistics of the transaction data that we include in our sample are shown in the first column of Table 2. Fares average \$415 for roundtrip travel. A stay over a Saturday night is involved in 20% of itineraries. Most tickets are purchased in the days shortly before departure; the fraction of tickets purchased within 3, 6 and 13 days before departure are 28%, 42% and 62%, respectively. The day of the week with the most initial departures is Monday and the day with the fewest departures is Saturday.

¹⁴ We create the categories “Full”, “Medium-Full”, “Medium-Empty” and “Empty” so that approximately the same number of coupons are in each category. As a result, there are more flights in the “Empty” than the “Full” category, but approximately the same number of passengers in each category.

¹⁵ We create the categories so there are approximately the same number of coupons sold for a given row of each column.

The data we analyze include 620,307 itineraries across the six carriers on these 90 routes. We measure ticket characteristics for 224,108 (or 36%) of these itineraries.

4. Testing Implications of Pricing Theories

4.1. Motivating Analysis

Prices can vary substantially as a function of days to departure. As an illustration, Figure 3 plots the prices for all round-trip tickets in our sample from Dallas-Fort Worth (DFW) to Los Angeles International Airport (LAX) on American. This figure includes both fares we could and could not match to data on ticket characteristics. Several patterns are clear.

First, for any given day in advance, there is variation in the transaction prices. However, prices on average are rising as purchase nears departure. On this route, fares do not discretely rise at 3, 7, or 14 days before departure. (On some other routes, however, we do observe such an increase).

Second, tickets appear to be sold at a discrete set of prices and these prices show up as bands of prices in the figure. These price bands can be seen in the second panel of Figure 3 which plots only fares less than \$1000. This phenomenon is consistent with work by airline pricing practitioners who write in the operations research literature – those researchers claim that airlines have fixed buckets of prices, and that yield management personnel alter the number of tickets available in each bucket.

An important phenomenon that we seek to study is the dispersion around the average prices. Although average fares rise as the purchase date approaches the departure date, we nevertheless observe some low fare tickets sold just before departure. On American's DFW-LAX route, some of the lowest fare, highly restricted coach tickets are sold up to the day of departure. Clearly, this dispersion could be caused by a variety of factors including different prices across the three months of our sample, different ticket restrictions, and different load factors on the various flights. This paper explores the

determinants of both the levels and variation in fares and how transacted fares change both as departure nears and load factors vary.

Motivating Regressions

To motivate the tests of scarcity pricing, we first analyze the association between an itinerary's fare and the ticket's restrictions and flight segment load factors. We regress the itinerary's log fare on the timing of purchase, the ticket's characteristics and restrictions, and various metrics of the load factor of the flight segments. We want to be cautious in interpreting this model as a pricing equation; there are possible explanations about the timing of purchase by different types of customers that could introduce selection concerns. Nevertheless, these regressions illustrate results consistent with our more formal tests of pricing models later in the paper.

We include several measures of ticket characteristics. *Refundable* and *Roundtrip* are indicators that the itinerary is refundable and for roundtrip travel, respectively. *TravelRestriction* is an indicator that the itinerary included a travel restriction (e.g. that all travel had to occur on Tuesday-Thursday, or that the ticket was not available on Friday or Sunday). This variable may pick up fences that separate high and low value customers. *StayRestriction* is an indicator that the ticket includes restrictions on the timing of departure and return travel (e.g. that the passenger must stay a minimum of 1 day and/or a maximum of 30 days). These restrictions are primarily minimum stay restrictions that could be used to separate customers who wish to travel and return on the same day. *SatNightStay* is an indicator for an itinerary with a stay over Saturday night; however, we do not have information on whether such a stay was required at purchase. *Advance_0_3*, *Advance_4_6*, *Advance_7_13*, and *Advance_14_21* are indicators of whether purchase occurred 0-3, 4-6, 7-13, and 14-21 days before the date of departure.¹⁶

Table 3 reports regression results of the association between fares, ticket characteristics and load factors. Each model includes fixed effects for carrier-route, the day of the week of the initial departure, and time effects (week of year). The first column

¹⁶ Purchase 21+ days in advance will serve as the excluded category. Note that our measures of advance purchase are the actual purchase dates rather than advance purchase restrictions placed on the ticket. We also have estimated the model with the advance purchase *restrictions*. In those regressions, the magnitudes of the coefficients of all other ticket characteristics and load factor are similar. Interestingly, prices are not always lower on tickets with more restrictive advance purchase restrictions.

includes only ticket characteristics as predictors of fares. Relative to travelers who purchase over 21 days in advance, passengers who purchase 14-21 days in advance pay 6% more, those who purchase 7-13 days in advance pay 18% more, those who purchase 4-6 days in advance pay 26% more, and those purchasing less than 4 days in advance pay 29% more. Passengers who purchase refundable tickets pay a 50% premium. Tickets with restrictions on the days of travel or the length of stay are sold at prices 30% and 8% lower, respectively. Passengers who stay over a Saturday night pay 13% less. It is noteworthy (and perhaps surprising) that these characteristics along with the fixed effects explain nearly 70% of the variation in fares.¹⁷

The remaining columns of Table 3 include various metrics of the actual and expected load factor of the flight segments of each itinerary. In column (2), we include the actual load factor at departure averaged over the itinerary's flight segments. Recall that this is likely to be measured with mean zero error because we "scale up" the observed tickets sold through our CRS by the CRS' share on the carrier-route; we address potential attenuation bias below. "*LF_Actual – Averaged Across Flight Segments*" is the average of each flight segments' realized load factor. We find that an increase in the actual load factor of an itinerary's flights is associated with a very modest increase in fares. A one standard deviation increase in the actual load factor averaged across flight segments (0.34) is associated with a 1.5% increase in fares (0.34*0.045).

These results suggest that load factor influences fares in a manner that is relatively small compared to ticket restrictions, and in a manner that is relatively independent of restrictions. When actual load factor is added to the model, the coefficients of ticket characteristics are very similar. Also, we find that the addition of the load factor measure does not substantially increase the fit of the model; the R^2 rises from 0.695 to 0.696.¹⁸

In column (3), we use a measure of the itinerary flight segments' *expected* load factors. As discussed above, we measure expected load factor as the average load factor for a particular carrier-flight-day of week (e.g. average load factor on American flight 301 on all Mondays in our sample). "*LF_Expected-Averaged across flight segments*" is

¹⁷ The R^2 of a regression with only the fixed effects is 0.356.

¹⁸ In unreported regressions, we include measures of load factor and fixed effects but not load factor. Adding load factor raises the R^2 from 0.356 to 0.359.

the average of each flight segment's expected load factor across all segments on an itinerary. We interpret this variable as a proxy for the component of load factor that is predictable by the airline, and could be used to set different (ex ante) price distributions on flights in response to differences in the expected distribution of demand. We find that an itinerary with flight segments that are expected to have higher load factors has slightly higher fares. A one standard deviation increase in this measure of expected load factor is associated with a 2.3% increase in fares. As with the case above using actual load factor, the association between fares and load factor is relatively small.¹⁹

The measures of actual and expected load factors are strongly positively correlated, so regressions including only one metric is likely to capture both effects. Column (4) includes measures of both actual and expected load factor. *LF_Expected* is still statistically significant – a one standard deviation increase in expected load factor is associated with a 2.1% increase in fares. However, the association between actual load factor and fares is no longer significant.

In the remaining columns, we allow for load factor to be associated with fares in a non-linear manner, and obtain similar results. It is possible that fares are high only for itineraries that involve a particularly full flight segment. In column (5), our measure of load factor is the actual load factor for the fullest flight segment (the maximum of actual load factor across all flight segments). The relationship between actual load factors and fares is very similar to the results in column (2). We include the expected load factor of the fullest flight in the model reported in column (6), and obtain results similar to those when we include the expected load factors averaged over flight segments.

Finally, we address whether the small association between fares and load factor is driven by mis-measurement of load factor. As discussed above, our load factor is measured with error because we only observe about one-third of all transactions. We 'scale up' our observed number of tickets by the inverse of our CRS' market share for the carrier-route. This 'scaled up' load factor is measured with error because individual flights will randomly have more/less than the average share of the CRS. This could lead to attenuation bias of our load factor coefficient towards zero. Under certain

¹⁹ In unreported regressions, we estimate the model using only load factor (i.e. without ticket characteristics). The coefficient estimate is 0.034 using *LF_Actual* ; it is 0.124 using *LF_Expected*.

assumptions, we can correct for the biased induced by the measurement error. If the measurement error is additive, mean zero, and independent of the true load factor, we can consistently estimate the coefficient vector using: $\hat{\beta} = (X'X - N\Omega_\epsilon)^{-1} X'y$, where Ω_ϵ is the variance matrix of the measurement error. We simulate the variance of the measurement error for different true load factors, and compute the OLS estimate of the load factor coefficient under those assumed values of Ω_ϵ . Results are reported in Table 4. The first row shows the results from column (2) of the previous table – a coefficient of LF_Actual of 0.0447. Under various assumed variances of the measurement error, the coefficient rises by a very small amount to 0.0450 to 0.0451. This suggests that measurement error is not the cause of our finding that there is a small association between fares and actual load factor.

These motivating regressions suggest that a ticket with the same characteristics but involving flights that are expected to have or actually have higher load factors, will be purchased at only a slightly higher price.

4.2. Testing for Price Rigidities

A key feature in Dana, Gale & Holmes, and the yield management literature is that prices are rigid. Airlines commit to a price schedule before demand is realized. Using a pre-determined set of prices, airlines then choose the number of seats to allocate at those prices.

In this section, we test a key assumption of these models – that prices are rigid – and find strong evidence in support of the price rigidities assumption. In order to motivate our estimation, consider the simple “stadium seating” model of Dana. The seller competes in price distributions and offers two types of seats – low-priced and high-priced seats. Consumers purchase the lowest priced seats that are available when they arrive. In the low demand state, only low priced tickets will be sold. In the high demand state, all tickets (low and high price) will be sold. The seller in Dana’s model will adjust fares *downward* in response to higher *expected* demand because seats are more likely to sell, reducing expected costs per passenger flown. However, the seller does not adjust the number of low and high priced seats for different *realized* demand states. This model

also predicts a stock-out of low fare tickets in high demand states, but the fares themselves are set in advance and do not vary with the state of demand.

All three of these theories, Dana, Gale and Holmes, and the yield management literature, in fact provide the same sharp underpinning for the empirical specification. That is, all three theories postulate that fares are set according to ticket characteristics and restrictions and do not vary across realized demand states. Gale and Holmes postulate a predetermined price schedule with more advance purchase tickets being offered on off-peak flights. The yield management literature also postulates a predetermined fare schedule where fares are set well in advance in a “pricing department” with the yield management department then allocating quantities to the various bins associated with these predetermined prices.

These models all indicate that a simple regression of fares on ticket characteristics will explain the large majority of the observed variation in prices. More specifically, suppose the analyst has data on each ticket sold and its associated characteristics. If the analyst regresses all transacted fares on dummy variables representing various ticket characteristics, the coefficients on the dummy variables representing the ticket restrictions will provide an hedonic measure of the “price” of different restrictions and in principle the R^2 will be 1.0 because ticket characteristics will perfectly explain variation in prices. Moreover, if one adds to this regression the fraction of the stadium that is full (“load factor”), the coefficient of load factor will be zero because fares are predetermined by characteristics. The reason is that prices are rigid in this model. The seller does not alter prices, but instead in the high demand state, there are more high priced tickets sold – both in absolute number and as a percentage of total sales.

We implement an analog to this empirical specification to test for price rigidities in airline pricing. We measure the “type” of ticket based upon ticket characteristics. As we show below, we believe that we accurately classify the “type” of ticket because the R^2 of estimating such a model for each route is centered around 0.84. Analogous to the thought experiment described above, when we add load factor as a dependent variable, the coefficient is not economically large, which is consistent with price rigidities.

To implement this test, we regress log fare on dummy variables for “types” of tickets. The yield management literature suggests that an airline’s planning department

creates the fare structure using a set of fare “buckets”. We do not know definitions of such buckets. However, we can create combinations of characteristics and restrictions, which we will refer to as “bins”. We use observed ticket characteristics to create a set of 72 bins. A bin is every possible combination of (a) refundability or not, (b) existence of a travel and/or stay restriction or not, (c) a Saturday night stay or not, and (d) nine categories of advance purchase restrictions: none, 1 day, 3 day, 5 day, 7 day, 10 day, 14 day, 21 day, and 30 day. Each of these types of tickets proxies a bucket to which a given number of tickets is assigned by an airline’s yield management department.

This structure of bins is consistent with all three theories described above. In the Dana model firms then stock out of low priced fares, driving customers who want to fly at a given time into higher-priced, less restricted tickets. In Gale and Holmes, airlines allocate fewer tickets to low price bins when there is an expected peak time of travel. In the yield management literature, bin prices are set in advance and tickets are then sold exhibiting the various combinations of restrictions and fares.

The empirical evidence supports this basic pricing structure. In particular, for each route we estimate log fares on dummy variables for bins. We allow each carrier to have a different fare structure for its bins by fully interacting the bin dummy variables with dummy variables for each carrier.²⁰

We assess the price rigidity assumption using the measures of R^2 for each route. Figure 4 plots the distribution of R^2 . For many routes, our bin structure explains a substantial amount of fare variation; the median R^2 is 0.84.

Now we have defined ticket “types” and found those ticket types to explain a large fraction of the total variation in fares. We are ready to test for price rigidities. We re-estimate the same regressions above except that we add a measure of actual load factor of the itinerary’s flight segments. Figure 5 plots the distribution of the coefficients of load factor. The median load factor coefficient across all routes is 0.028, and for some of the city-pairs the coefficient is statistically positive. However, the magnitude is economically small, just as in the motivating regressions above.

²⁰ We also include a dummy variable for whether the itinerary is roundtrip, which is a ticket characteristic that we do not include in the bin structure.

Although these results are consistent with airlines increasing the fares of existing fare “types”, the effect appears to be relatively small. Price rigidity appears to be a reasonable assumption. Thus, the data are relatively consistent with the assumption of airlines setting rigid prices and not adjusting prices to load factors – prices can be largely explained with groups of characteristics and prices do not increase in demand.

4.3 Testing Dana’s Predictions Regarding Price Dispersion

Next we turn to testing Dana (1999b)’s prediction regarding which flights have more fare dispersion. Flights with the same expected distribution of demand have the same *offered* fares, but the flights with higher realized load factor have more dispersion in *transacted* fares. Intuitively, passengers buy from the lowest priced fare bucket open when they purchase; so if there are more realized purchases, then more higher priced buckets observe purchases, and transacted fares are more dispersed.

We compare the average Gini coefficient for flights with the same expected demand distribution but different realized demand. We proxy for the expected distribution of demand using the expected load factor quartiles calculated above and described in Figure 2. All flights for a carrier-citypair are divided into quartiles of expected load factor based upon their average load factors for the twelve weeks for which we have data. We then divide these quartiles based on demand realizations for particular flights. Dana’s model has sharp predictions regarding the distribution of fares within this latter grouping. More specifically, for a given grouping of expected load factor (e.g. “Full”), Dana’s model predicts there will be less fare dispersion on flights with “Empty” realized load factors than those with “Full” realized load factors. In Figure 2, the Gini coefficient should rise as one moves up each column.

One empirical complication is that fares are measured for an entire itinerary that typically involves two flights that may have different expected and realized load factor quartiles. We classify an itinerary based upon the flight characteristics of the first coupon. If there is no correlation between the flight characteristics of the first and second coupon (i.e. outgoing and returning flight), this will attenuate differences in Gini

coefficients, but one still should observe higher Gini coefficients for itineraries involving flights realized to be full.

In order to measure dispersion, we calculate the Gini coefficient for each route-carrier in each category of expected-realized load factor (i.e. each cell of Figure 2), and calculate the average Gini for each cell. Because this analysis does not use data on ticket characteristics, we can use all observed transactions through our CRS.

Results are shown in the top panel of Table 5. The average within carrier-route Gini coefficient is approximately 0.28. However, there is very little variation in this metric of dispersion by either expected or realized load factor. The category with the highest dispersion has an average Gini coefficient of 0.284 while the category with the lowest dispersion has an average Gini of 0.271.²¹

This pattern of dispersion is not consistent with Dana (1999b). For flights with the same expected load factor, dispersion is *decreasing* when the realized load factor is larger. However, this decrease in dispersion is economically small with the largest change for flights that are expected to be full – the dispersion decreases from 0.284 for flights realized to be empty to 0.275 for flights realized to be full.

The bottom panel of Table 5 shows results when we restrict the sample to the tickets with matching information on ticket characteristics. For this subsample, we observe slightly less dispersion with an average within carrier-route Gini coefficient of 0.23.²² Nevertheless, we also find that dispersion does not substantially rise when realized load factor is higher, controlling for expected load factor. The average Gini coefficient is very slightly higher on flights realized Full versus Empty (e.g. from 0.243 to 0.245 on flights expected Full), however the effect is not monotonic in the realized load factor.

²¹ These Gini coefficients measure different dispersion from that of Borenstein and Rose (1994) who report an average Gini within carrier-route of 0.181 in the second quarter of 1986. Dispersion has increased since 1986 – Borenstein and Rose (2007) show that the coefficient of variation rose from below 0.4 in 1986 to between 0.5 and 0.6 in 2004. Borenstein and Rose use a larger sample of routes that we do; our average coefficient variation within carrier-route for 2004Q4 is 0.61.

²² This finding is somewhat expected given Figure 1 showing that we were unable to match some of the particularly low fare tickets.

4.4 Testing Predictions of Dana and Gale & Holmes Models Regarding Ticket Allocations in Peak Demand

We now test another key prediction posited by both Gale and Holmes and by Dana. We test Gale and Holmes' prediction that low price, advance purchase tickets will account for a larger share of tickets on off-peak flights, and correspondingly that full-fare, refundable, high price tickets will account for a larger share of tickets on peak flights. These alternative hypothesis are contrasted with the null hypothesis that tickets are simply allocated to different bins in a similar fashion on full and empty flights.

We classify tickets as “discount” by dividing itineraries into 3 groups based upon ticket characteristics. Group 1 includes refundable tickets (recall that we have already excluded first-class tickets). Group 2 includes non-refundable tickets that do not include any travel or stay restrictions. Group 3 includes non-refundable tickets involving travel and/or stay restrictions. The fraction of tickets in Groups 1-3 are 26%, 32%, and 42%, respectively. These groupings are associated with large differences in average fares, and by themselves account for a large share of the differences in observed fares. For Group 1 tickets the average fare is \$631, for Group 2 it is \$440, and for Group 3 it is \$281.²³ Hence these groupings are associated with high, medium and low fares.

In order to test these implications, we must define flights in our data that are (ex ante) high demand. To do so, for each airline on a given route, we find the average load factors for each flight number–day of week. (E.g. For American's LAX-IAD route, we calculate the average actual load factor on AA flight 76 for all 12 Mondays in our sample). Then we group each of these average load factors (at the flight-number-day of week level) into 3 tertiles – Low, Medium, and High expected load factor.²⁴ (So, for example, every AA76 on Monday is called a High Expected Load Factor flight). Next, we include in our analysis only flights that are expected to be High (Low) load factor, and actually are High (Low) load factor. This sample selection is intended to remove flights subject to unusual shocks. Finally, we count each coupon sold on these flights by

²³ A regression of fares on these groupings yields an R^2 of approximately 0.65.

²⁴ Note that we are planning to reestimate these empirical results using the quartiles in Figure 2, but we have not yet implemented this procedure.

Group and Days in Advance.²⁵ These allocations are calculated at the airline level to account for any differences across airlines. These tabulations are shown in Tables 6a-6b.

Dana's stadium model implies that low price tickets will account for a larger share of tickets sold in low demand states. Low price tickets in airlines correspond to Group 3 tickets that have more restrictions. Table 6 indicates that any such differences are economically insignificant. While all six carriers do indeed sell more Group 3 tickets in Low versus High demand states, the differences are quite small. The percentage differences between low and high demand states range from 2-5 percentage points. The differences in sales of Group 3 tickets are: United (67% vs. 65%), American (53% vs. 50%), USAir (27% vs. 23%), Delta (55% vs. 49%), Northwest (72% vs. 66%), and Continental (53% vs. 46%). Thus, there is a baseline level of allocation of high and low price seats that is driven by non-scarcity factors, and there are only modest deviations from this baseline associated with scarcity.

Finally, we test an implication of Gale and Holmes that on-peak flights will have fewer discount/advance-purchase seats sold in equilibrium. We compare the fraction of Group 3 tickets sold greater than 21 and greater than 14 days in advance. For purchases 21+ days in advance, the fractions of low price/Group 3 tickets on low and high load factor flights are American: 18% vs. 16%, Delta: 20% vs. 17%, Continental: 24% vs. 22%, United 22% vs. 18%, USAir: 11% vs. 11%, and Northwest: 28% vs. 23%. This is weak evidence in support of the Gale and Holmes prediction. If we consider all Group 3 tickets sold 14 or more days in advance, the implications are very similar. The fraction of coupons in low and high load factor flights are American: 27% vs. 27%, Delta: 32% vs. 29%, Continental: 35% vs. 31%, United: 36% vs. 32%, USAir: 17% vs. 16%, and Northwest: 44% vs. 40%. In fact, even if we consider all tickets purchased 14 or more days in advance, the fraction sold in low load factor flights are only slightly higher.

These ticket breakdowns suggest very weak evidence in support of models regarding the allocation of scarce capacity in the face of demand uncertainty. In particular, airlines have many flights that are both expected to be "empty" and are

²⁵ We include coupons for both "local" and "connecting" passengers. For the connecting passengers, we have only classified itineraries into groups for passengers traveling on the 342 largest routes, so connecting passengers with origin and destination cities from small routes are not included.

realized to be “empty”. Nevertheless, the airlines do not sell a larger fraction of lower fare tickets on these flights.

4.5 Testing for Evidence of Scarcity Pricing Using Fares on Unusually Full Flights

Dana’s model has predictions regarding the mean of transacted fares as a flight fills up approaching the departure date. To motivate this test, recall Dana’s stadium pricing example. Consider two events with the same prior distribution of demand uncertainty, so the stadium has printed the same distribution of tickets. Suppose that consumers arrive at different periods of time before the event to purchase a ticket. The stadium has printed a specific number of tickets of each price, and consumers buy the cheapest ticket available when they arrive at the ticket window. Consider the tickets purchased one hour before the event begins (call this period “T-1”). If an unusually large number of consumers have arrived and purchased greater than an hour before the event (T-2, T-3, ...), then the tickets purchased during interval T-1 will be sold at a higher price than if a “normal” number of people had arrived prior to T-1. This theory results in the prediction that when a flight is closer to capacity at a given point in time prior to departure, then its fares will be higher than fares on flights that are less full.

To test this hypothesis, we calculate at the carrier-route level, the average load factor as of a specific number of days prior to departure. Then we calculate for each ticket the deviation in load factor for this mean, conditional on the days in advance that the ticket was purchased. We then transform these deviations into percentages of the conditional mean load factor. We calculate the analogous measure of the itinerary’s fare relative to the mean, also conditional on the number of days in advance that the ticket was purchased. For tickets with multiple legs, we use the first leg, noting as above that the effects should still impact fares so long as the return is randomly distributed. This allows us to ask a question such as: for a ticket bought 7 days before departure, if the plane is 10% fuller than normal (for a plane 7 days before departure), what percent more expensive is the fare?

Figure 6 plots the kernel regression of the relationship between the percent deviation in fares and the percent deviation in load factor for each carrier. We restrict the analysis to tickets sold in the 7 days before departure. We find that tickets on flights that are unusually full do have higher fares, but the effect is relatively modest. The slope of the relationship is steeper for American than the other carriers. The relationship for American corresponds to roughly a 1.7% increase in fares for a 10% higher load factor at time of purchase (and roughly a 0.8% increase for the other carriers).²⁶ In unreported analysis, we restrict the tickets to those sold within 3 days of departure, and the results are qualitatively similar. These results suggest that the levels of fares are only modestly higher when load factors are higher.

5. Conclusions

This paper tests several of the leading models of pricing in the airline industry. First, our data on ticket prices and characteristics allow us to provide support for the assumption of price rigidities, where the “bins” of fares are well approximated with ticket characteristics. Empirical validation of this assumption is important because it is common to both sets of theories of airline pricing.

Second, we test basic implications of several of the leading models of scarcity-based pricing. Our results provide only modest evidence that pricing in the airline industry is driven by these models. Fares on flights with higher expected and realized demand are only slightly higher than flights with low demand, after controlling for ticket characteristics. In addition, fare dispersion is not significantly higher on flights with higher realized demand. And we find only weak support for the prediction from Dana and Gale and Holmes models that there will be quantity restrictions on the sale of low-priced and advance purchase tickets on high demand flights.

Taken together, there is some empirical support for scarcity-based pricing, but it appears to be relatively modest. We find much stronger evidence that certain sets of ticket characteristics drive much of the variation in ticket pricing and these ticket

²⁶ The slope is calculated using the range of Load Factor % Deviation from -0.5 to 0.5, where many itineraries are concentrated.

characteristics affect fares in a manner largely independent of load factor. While such evidence does not rule out theories of the Dana and Gale and Holmes variety, it does suggest that theories in which ticket characteristics segment customers and facilitate price discrimination may play a large role in airline pricing.

Our findings provide the foundation for further empirical investigation on the nature of airline pricing. Future research could explore how ticket characteristics are used to segment customers. Our finding that ticket characteristics are strongly associated with fares is consistent with a variety of models of second-degree price discrimination, including work in the yield management literature and Dana's (1998) analysis of advance-purchase discounts. Such models have varying implications about the choice of capacity and the efficient use of that capacity. The role of ticketing restrictions can be explored in future work using our information on ticket characteristics.

References

- Berry, Steven and Panle Jia, "Tracing the Woes: An Empirical Analysis of the Airline Industry", NBER Working Paper 14503, November 2008.
- Borenstein, Severin. "Hubs and High Fares: Dominance and Market Power in the U.S. Airline Industry", *RAND Journal of Economics*, 20(3): 344-365, Autumn 1989.
- Borenstein, Severin. "U.S. Domestic Airline Pricing, 1995-2004". Berkeley, CPC Working Paper No. CPC05-48, January 2005.
- Borenstein, Severin and Nancy Rose. "Competition and Price Dispersion in the U.S. Airline Industry," *Journal of Political Economy*, 102(4):653-683, August 1994.
- Borenstein, Severin and Nancy Rose. "Bankruptcy and Pricing Behavior in U.S. Airline Markets," *American Economic Review*, 85(2):397-402, May 1995.
- Borenstein, Severin and Nancy Rose. "How Airline Markets Work...Or Do They? Regulatory Reform in the Airline Industry", NBER Working Paper 13452, September 2007.
- Busse, Meghan. "Firm Financial Condition and Airline Price Wars," *RAND Journal of Economics*, 33(2): 298-318, Summer 2002.
- Ciliberto, Federico and Carola Schenone, "Bankruptcy and Product-Market Competition: Evidence from the Airline Industry", University of Virginia Working Paper, 2008.
- Dana, James D., Jr. 1998. "Advance-Purchase Discounts and Price Discrimination in Competitive Markets." *Journal of Political Economy*, 106 (2): 395-422.
- Dana, James D., Jr. 1999a. "Using Yield Management to Shift Demand when the Peak Time is Unknown." *RAND Journal of Economics*, 30 (Autumn): 456-74.
- Dana, James D., Jr. 1999b. "Equilibrium Price Dispersion under Demand Uncertainty: The Roles of Costly Capacity and Market Structure." *RAND Journal of Economics*, 30 (Winter): 632-60.
- Dana, James D. and Eugene Orlov, "Internet Penetration and Capacity Utilization in the US Airline Industry", mimeo, Northeastern University, February 2008.
- Eden, Benjamin. 1990. "Marginal Cost Pricing When Spot Markets are Complete." *Journal of Political Economy*, 98 (6) 1293-306.
- Escobari, Diego. 2007. "Systematic Peak-load Pricing, Congestion Premia and Demand Diverting: Empirical Evidence", mimeo, San Francisco State University.

- Escobari, Diego and Gan, Li. 2007. "Price Dispersion under Costly Capacity and Uncertain Demand". NBER Working Paper No. 13075, May 2007.
- Forbes, Silke Januszewski, "The Effect of Air Traffic Delays on Airline Prices", *International Journal of Industrial Organization*, 26(5): 1218-1232, September 2008.
- Gale, Ian, and Thomas Holmes. 1992. "The Efficiency of Advance-Purchase Discounts in the Presence of Aggregate Demand Uncertainty." *International Journal of Industrial Organization* 10 (3): 413-37.
- Gale, Ian and Thomas Holmes. 1993. "Advance-Purchase Discounts and Monopoly Allocation of Capacity." *American Economic Review* 83 (March): 135-46.
- Goolsbee, Austan and Chad Syverson, "How Do Incumbents Respond to the Threat of Entry? Evidence from the Major Airlines." *Quarterly Journal of Economics*, forthcoming.
- McAfee, R. Preston, and Vera L. te Velde. 2006. "Dynamic Pricing in the Airline Industry." In *Handbook on Economics and Information Systems*, vol. 1, edited by T.J. Hendershott. New York: Elsevier Science.
- Prescott, Edward C. 1975. "Efficiency of the Natural Rate." *Journal of Political Economy* 83 (6): 1229-36.
- Rosenthal, R. W. (1980), "A Model in Which an Increase in the Number of Sellers Leads to a Higher Price". *Econometrica*, 48(16), 1575–1579.
- Sengupta and Wiggins, 2006. "Airline Pricing, Price Dispersion and Ticket Characteristics On and Off the Internet." Texas A&M University, mimeo.
- Smith. B, Rao. B and Ratliff. R. 2001. "E-Commerce and Operations Research in Airline Planning, Marketing, and Distribution," *Interfaces*, 31(2): 37-55.
- Stahl, D. 1989. "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review*, 79(4), 1989: 700-712.
- Stahl, D. 1996. "Oligopolistic Pricing with Heterogeneous Search", October, *International Journal of Industrial Organization*, 14, 1996: 243–268.
- Stavins, Joanna. 2001. "Price Discrimination in the Airline Market: The Effect of Market Concentration." *Review of Economics and Statistics*, 83 (1): 200-02.

Data Appendix

Transactions Data

We study itineraries for travel in 2004Q4 that were purchased between June and December 2004 through the Computer Reservation System (CRS) that provided us with the data. Although we do not have data on transactions occurring prior to June (which means we miss transactions occurring 4 months before our first day of October 1, 2004), we do not expect this to substantively affect our results.

We exclude itineraries involving any international travel, more than four coupons, open jaws and circular trips, or more than one carrier. Also, we exclude itineraries with a zero fare.

We calculate a measure of flight level load factor using the tickets we observe and the CRS's share of tickets sold on a city-pair. This is described in more detail in the main text. The CRS share is calculated by finding the fraction of total coupons for non-stop travel between two cities (the "T-100 Domestic Segment" data from the Bureau of Transportation Statistics) that we observe in our transaction data. We compute these "CRS shares" at the route-carrier level.

Procedure to Merge Transaction Data to Posted Fare Data

We used the following procedure to match transactions from the CRS providing us with transaction data to posted fares from the CRS that provided archived fares.

In the first step, we matched a ticket from the transaction data to a posted fare using carrier, date of departure (but not return), booking class, and price.²⁷ In this first step, we included any fares matching within 10%.

After this first step, the resulting dataset included multiple matching posted fares for some individual transactions. This primarily included multiple matching fares with different combinations of advance purchase requirements and travel restrictions. Because

²⁷ In this first matching step, we only require fares to match within 10%. In a later step, we require fares to match much closer. In addition, we matched a transaction's date of departure to a 7-day window of days of departure in the posted fare data, and later use the match in which the dates of departure are closest.

our transaction data include no additional information to facilitate matching, we were required to make additional assumptions. In the second step of the matching procedure, we eliminate multiple matches on advance purchase. We assume that the ticket was purchased with the most restrictive advance purchase requirement for which it qualified.²⁸

For any transactions that still matched multiple posted fares, we adopted a third matching step. Prices were required to match within a 2 percent range.²⁹ Any remaining multiple matches were then screened to meet travel restrictions that involve travel on specified days of the week. For example, some posted fares required travel on a Tuesday, Wednesday or Thursday. Using the ticket's date of departure, we eliminated any multiple matches that did not satisfy the posted travel restriction. For any additional transactions with multiple matches, we assumed that any ticket meeting a travel restriction had that travel restriction. For example, a ticket matching fares with and without a travel restriction was assumed to have that travel restriction.

The final step includes the verification of minimum and maximum stay restrictions. For the minimum and maximum stay restrictions collected from the travel agent, some restrictions were explicitly given (namely 1 day, 2 days etc.). However, other posted fares were indicated to include a travel restriction but the restriction was not specifically named on the travel agent's CRS screen that we accessed. For the matches where the minimum and maximum days of stay restriction were given, we verified that the actual transactions met the specific requirements. In case of multiple matches (which comprise less than 1%), if two tickets had the same characteristics but one required a 1 day minimum stay while the other did not, and the transaction involved a 2 day stay, we match the posted fare with a 1 day minimum stay.

²⁸ For example, suppose a ticket was purchased 16 days before departure. If the first step matched both a 14 day and a 7 day advance purchase requirement, we match the transaction with the posted fare that required a 14 day advance purchase.

²⁹ We should note that the local travel agent used a different CRS than our transaction data. Since July 2004, CRSs were not required to post identical fares.

Table 1: Routes Included in Analysis

The analysis includes all carriers flying on any of these routes, where the routes are large routes for the six carriers below.

American	LAS-DFW	LAX-JFK	PHX-DFW	DFW-DEN	ORD-STL
	LAX-DFW	ORD-LGA	LAX-ORD	ORD-DFW	DFW-MCO
	SJU-MIA	STL-DFW	DFW-SNA	LGA-MIA	MIA-JFK
Delta	DFW-ATL	LAS-ATL	ATL-MIA	ATL-PHL	EWR-ATL
	MCO-ATL	LGA-ATL	TPA-ATL	ATL-FLL	BOS-ATL
	LAX-ATL	CVG-ATL	CVG-LGA	FLL-BDL	LAX-TPA
United	LAX-DEN	LAS-ORD	IAD-ORD	LAS-DEN	SEA-ORD
	LAX-ORD	DEN-ORD	ORD-SFO	SFO-LAX	ORD-LGA
	SFO-SAN	IAD-SFO	OAK-DEN	ONT-DEN	PDX-SFO
Continental	LAX-EWR	DEN-IAH	ORD-IAH	ATL-EWR	IAH-DFW
	EWR-MCO	FLL-EWR	LAS-EWR	BOS-EWR	SFO-EWR
	IAH-LAX	EWR-IAH	MSY-IAH	IAH-LAS	IAH-MCO
Northwest	MSP-PHX	MSP-LAS	DEN-MSP	DTW-LAS	PHX-DTW
	LGA-DTW	MCO-DTW	LAX-DTW	MSP-MCO	MKE-MSP
	DTW-MSP	LAX-MSP	SEA-MSP	MSP-SFO	BOS-DTW
USAir	PHL-MCO	FLL-PHL	BOS-DCA	BOS-LGA	ORD-PHL
	PHL-BOS	LGA-DCA	PHL-TPA	LAS-PHL	RDU-PHL
	MCO-CLT	CLT-PHL	LGA-CLT	CLT-BOS	PIT-PHL

Notes: These routes are large representative routes for each of the six carriers. Airport codes: ATL=Atlanta, BDL=Hartford, BOS=Boston, CLT=Charlotte, CVG=Cincinnati, DCA=Washington-Reagan, DEN=Denver, DFW=Dallas-FtWorth, DTW=Detroit, EWR=Newark, FLL=Fort Lauderdale, IAD=Washington-Dulles, IAH=Houston, JFK=NY-JFK, LAS=Las Vegas, LAX=Los Angeles Intl, LGA=NY-La Guardia, MCO=Orlando, MIA=Miami, MKE=Milwaukee, MSP=Minneapolis-St Paul, MSY=New Orleans, OAK=Oakland, ONT=Ontario, ORD=Chicago-O'Hare, PDX=Portland, PHL=Philadelphia, PHX=Phoenix, PIT=Pittsburgh, RDU=Raleigh-Durham, SAN=San Diego, SEA=Seattle, SFO=San Francisco, SJU=San Juan, SNA=Orange County, STL=St. Louis, TPA=Tampa.

Table 2: Sample Means

Variable	All Transactions	Matched Transactions
Fare (for roundtrip)	\$ 414.61	\$ 423.64
Refundable	--	0.26
Some Travel Restriction (e.g. DOW)	--	0.38
Minimum Stay Restriction	--	0.20
Maximum Stay Restriction	--	0.15
Stayed over Saturday Night	0.20	0.19
Purchased 0-3 Days in Advance	0.28	0.31
Purchased 4-6 Days in Advance	0.14	0.14
Purchased 7-13 Days in Advance	0.20	0.20
Purchased 14-21 Days in Advance	0.14	0.14
Purchased > 21 Days in Advance	0.24	0.21
Roundtrip Itinerary	0.66	0.65
Load Factor averaged across itin legs	0.91	0.90
American	0.30	0.28
Delta	0.16	0.15
United	0.14	0.15
Continental	0.16	0.18
Northwest	0.07	0.08
USAir	0.17	0.15
Monday Departure	0.19	0.20
Tuesday Departure	0.16	0.18
Wednesday Departure	0.16	0.17
Thursday Departure	0.15	0.16
Friday Departure	0.16	0.13
Saturday Departure	0.07	0.06
Sunday Departure	0.11	0.11
N	620,307	224,108

Note: Summary statistics for itineraries to travel in 2004Q4 on American, Delta, United, Northwest, Continental and USAir on the routes in our sample. The first column includes all transactions through the CRS that gave us transaction data (excluding first class tickets and itineraries involving more than four coupons, as discussed in the Data section). The second column includes only transactions we were able to match with ticket characteristics from the other CRS's archive.

Table 3: Motivating Regressions

Dependent Variable: Log(Fare)						
	(1)	(2)	(3)	(4)	(5)	(6)
	Characteristics Only	Actual LF	Expected LF	Actual & Expected LF	Max Actual LF across segments	Max Expected LF across segments
Advance_0_3	0.292 (0.011)**	0.295 (0.011)**	0.294 (0.011)**	0.294 (0.011)**	0.295 (0.011)**	0.293 (0.011)**
Advance_4_6	0.262 (0.011)**	0.265 (0.011)**	0.264 (0.011)**	0.264 (0.011)**	0.265 (0.011)**	0.264 (0.011)**
Advance_7_13	0.180 (0.008)**	0.182 (0.008)**	0.181 (0.008)**	0.181 (0.008)**	0.182 (0.008)**	0.180 (0.008)**
Advance_14_21	0.056 (0.008)**	0.058 (0.008)**	0.057 (0.008)**	0.057 (0.008)**	0.058 (0.008)**	0.057 (0.008)**
Refundable	0.497 (0.009)**	0.497 (0.009)**	0.497 (0.009)**	0.497 (0.009)**	0.497 (0.009)**	0.498 (0.009)**
Roundtrip Itinerary	-0.116 (0.004)**	-0.117 (0.004)**	-0.119 (0.004)**	-0.119 (0.004)**	-0.124 (0.004)**	-0.131 (0.005)**
Travel Restriction	-0.304 (0.004)**	-0.302 (0.004)**	-0.301 (0.004)**	-0.301 (0.004)**	-0.302 (0.004)**	-0.302 (0.004)**
Stay Restriction	-0.080 (0.005)**	-0.081 (0.005)**	-0.081 (0.005)**	-0.081 (0.005)**	-0.081 (0.005)**	-0.081 (0.005)**
Stayed Over Saturday Night	-0.131 (0.006)**	-0.126 (0.006)**	-0.123 (0.007)**	-0.123 (0.007)**	-0.126 (0.006)**	-0.121 (0.006)**
LF_Actual - Averaged across flight segments		0.045 (0.005)**		0.004 (0.007)		
LF_Expected - Averaged across flight segments			0.091 (0.008)**	0.086 (0.011)**		
LF_Actual - Maximum across flight segments					0.039 (0.004)**	
LF_Expected - Maximum across flight segments						0.081 (0.007)**
Observations	224,108	224,108	224,108	224,108	224,108	224,108
R-squared	0.695	0.696	0.696	0.696	0.696	0.696

Note: All models include fixed effects for route-carrier, day of the week of initial departure, and week of year. The \bar{R}^2 of a model with only the fixed effects is 0.356. Model estimated via least squares with robust standard errors (clustered on the calendar date of the initial departure).

** significant at 1%

Table 4: Robustness of Actual Load Factor Coefficient to Measurement Error

Model	Assumed Std Dev of Measurement Error	Coeff of LF_Actual
Original Model	0.00	0.0447
True LF = 0.55	10.76	0.0450
True LF = 0.75	12.43	0.0450
True LF = 0.95	14.00	0.0451

Table 5
Gini Coefficients by Expected and Realized Load Factors

Using All Transactions

Realized Load Factor	Expected Load Factor			
	Full	Medium-Full	Medium-Empty	Empty
Full	0.275	0.274	0.271	0.275
Medium-Full	0.279	0.271	0.273	0.282
Medium-Empty	0.283	0.280	0.276	0.280
Empty	0.284	0.277	0.274	0.283

Using Only Transactions Matched to Ticket Characteristics

Realized Load Factor	Expected Load Factor			
	Full	Medium-Full	Medium-Empty	Empty
Full	0.245	0.238	0.231	0.227
Medium-Full	0.243	0.233	0.222	0.234
Medium-Empty	0.248	0.244	0.228	0.232
Empty	0.243	0.236	0.230	0.239

Notes: Cell values are the simple average Gini coefficient for each carrier-route-load factor category. We only include a carrier-route-load factor category in the calculation if at least 100 itineraries were observed.

Table 6a: Tests of Comparative Statics of Dana and Gale & Holmes

Flights - Expected to be Low Load Factor & Are Low Load Factor						Flights - Expected to be High Load Factor & Are High Load Factor					
American - All Routes											
Groups	0 to 6	7 to 13	14 to 21	21+		Groups	0 to 6	7 to 13	14 to 21	21+	
Group 1	13%	1%	0%	0%	15%	Group 1	10%	1%	0%	0%	12%
Group 2	14%	8%	4%	6%	32%	Group 2	13%	10%	5%	9%	38%
Group 3	14%	12%	9%	18%	53%	Group 3	11%	12%	11%	16%	50%
	40%	22%	13%	25%			34%	24%	17%	25%	
Delta - All Routes											
Groups	0 to 6	7 to 13	14 to 21	21+		Groups	0 to 6	7 to 13	14 to 21	21+	
Group 1	3%	1%	1%	1%	6%	Group 1	5%	1%	0%	0%	7%
Group 2	18%	8%	6%	7%	39%	Group 2	19%	10%	7%	9%	45%
Group 3	11%	12%	12%	20%	55%	Group 3	8%	12%	12%	17%	49%
	32%	21%	19%	28%			32%	22%	20%	26%	
Continental - All Routes											
Groups	0 to 6	7 to 13	14 to 21	21+		Groups	0 to 6	7 to 13	14 to 21	21+	
Group 1	14%	3%	1%	1%	19%	Group 1	21%	5%	2%	1%	28%
Group 2	10%	7%	3%	7%	27%	Group 2	8%	6%	3%	9%	26%
Group 3	7%	11%	11%	24%	53%	Group 3	6%	8%	9%	22%	46%
	31%	21%	16%	32%			35%	19%	14%	32%	

Note: Each panel contains percentages of the total coupons on flights. There are two panels for each airline. The left panel contains flights (flight number - date of departure) that are forecasted to be low load factor and are realized to be low load factor. The right panel contains flights that are forecasted to be high load factor and are realized to be high load factor. A flight is forecasted to be high/low load factor if that flight has an average load factor in the top/bottom tertile of all flights-day of week for that carrier-route. Group 1 = Refundable tickets, Group 2 = Nonrefundable without travel or stay restrictions, 3 = Nonrefundable with travel and/or stay restrictions.

Table 6b: Tests of Comparative Statics of Dana and Gale & Holmes

Flights - Expected to be Low Load Factor & Are Low Load Factor						Flights - Expected to be High Load Factor & Are High Load Factor					
United - All Routes											
Groups	0 to 6	7 to 13	14 to 21	21+		Groups	0 to 6	7 to 13	14 to 21	21+	
Group 1	7%	1%	1%	1%	10%	Group 1	10%	1%	0%	0%	12%
Group 2	11%	7%	2%	2%	23%	Group 2	10%	8%	2%	3%	23%
Group 3	18%	14%	14%	22%	67%	Group 3	16%	17%	14%	18%	65%
	36%	22%	17%	25%			36%	26%	17%	21%	
US Airways - All Routes											
Groups	0 to 6	7 to 13	14 to 21	21+		Groups	0 to 6	7 to 13	14 to 21	21+	
Group 1	44%	8%	2%	1%	55%	Group 1	44%	10%	3%	2%	58%
Group 2	7%	6%	2%	3%	18%	Group 2	6%	5%	2%	5%	19%
Group 3	5%	5%	6%	11%	27%	Group 3	3%	5%	4%	11%	23%
	56%	20%	10%	15%			52%	20%	9%	18%	
Northwest - All Routes											
Groups	0 to 6	7 to 13	14 to 21	21+		Groups	0 to 6	7 to 13	14 to 21	21+	
Group 1	8%	2%	1%	1%	11%	Group 1	11%	2%	0%	0%	14%
Group 2	6%	3%	2%	6%	17%	Group 2	9%	4%	2%	5%	20%
Group 3	14%	14%	16%	28%	72%	Group 3	13%	13%	17%	23%	66%
	29%	18%	19%	34%			33%	19%	19%	28%	

Note: Each panel contains percentages of the total coupons on flights. There are two panels for each airline. The left panel contains flights (flight number - date of departure) that are forecasted to be low load factor and are realized to be low load factor. The right panel contains flights that are forecasted to be high load factor and are realized to be high load factor. A flight is forecasted to be high/low load factor if that flight has an average load factor in the top/bottom tertile of all flights-day of week for that carrier-route. Group 1 = Refundable tickets, Group 2 = Nonrefundable without travel or stay restrictions, 3 = Nonrefundable with travel and/or stay restrictions.

Figure 1

**Comparing the Kernel Densities of Matched and Unmatched Transactions
All Carriers and All Routes**

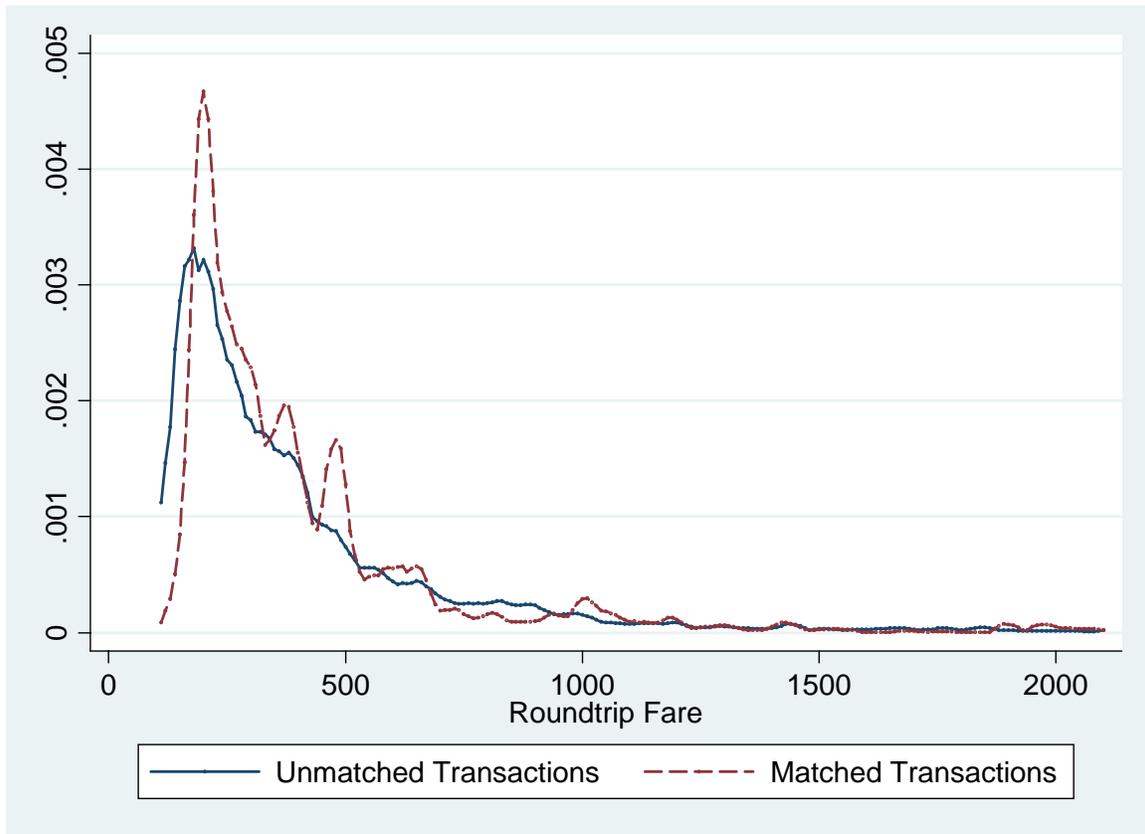


Figure 2

Dividing Sample by Expected and Realized Load Factors

		Expected Load Factors			
		Full	Medium-Full	Medium-Empty	Empty
Realized Load Factors	Full				
	Medium-Full				
	Medium-Empty				
	Empty				

This table illustrates how flights are divided to test comparative static predictions about the characteristics of tickets sold on flights that are unusually full on peak flights and unusually empty on off-peak flights. We divide flights (flight-departure date) into quartiles based upon actual load factor and expected load factor. The expected load factor is estimated as the load factor for the flight number–day of week averaged over the 12 weeks in our sample. We create the categories so there are approximately the same number of tickets in each cell. A complete description of the methodology is included in the text.

Figure 3: Transactions in 21 Days Before Departure

American: DFW-LAX

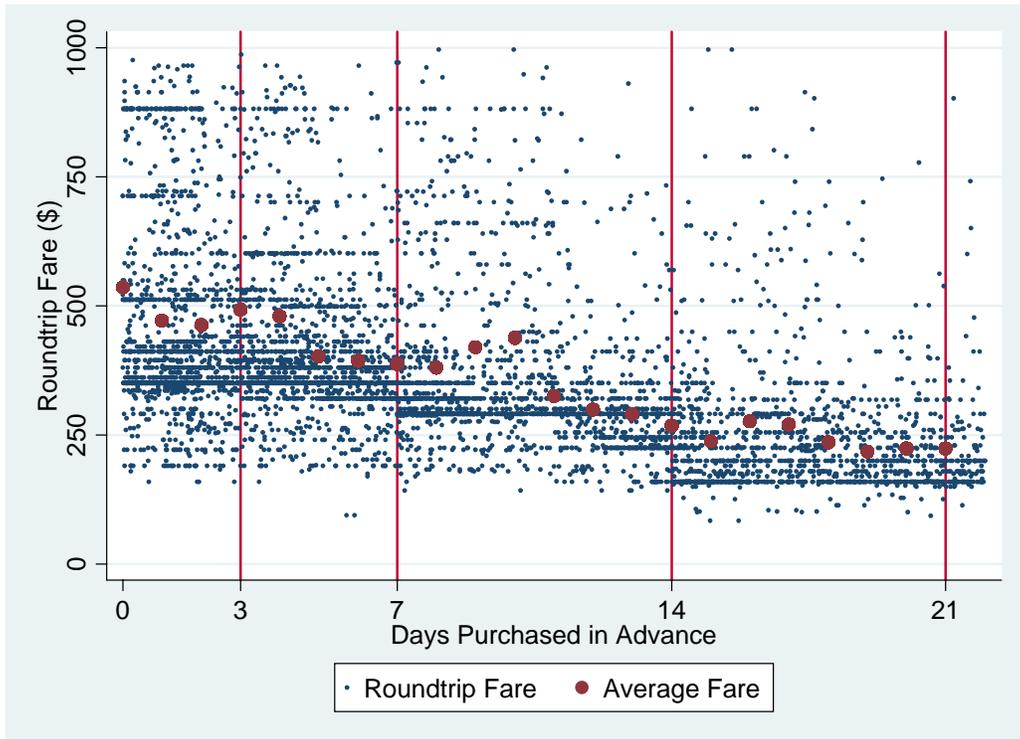
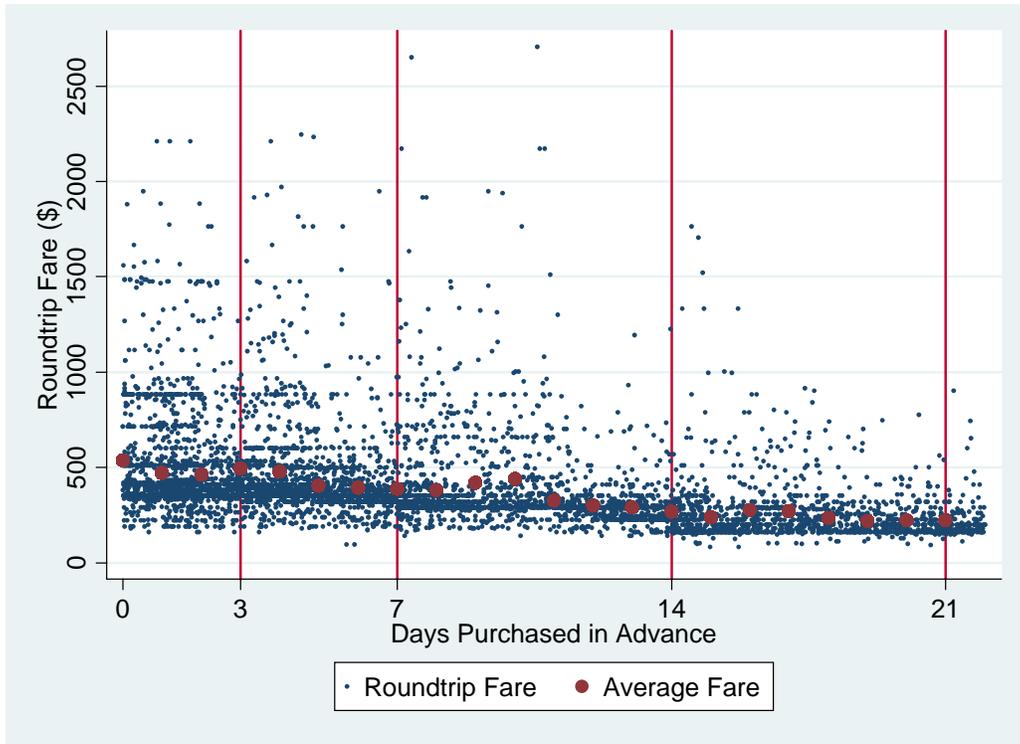


Figure 4
Histogram of R^2 for Each Route's Bins Regression

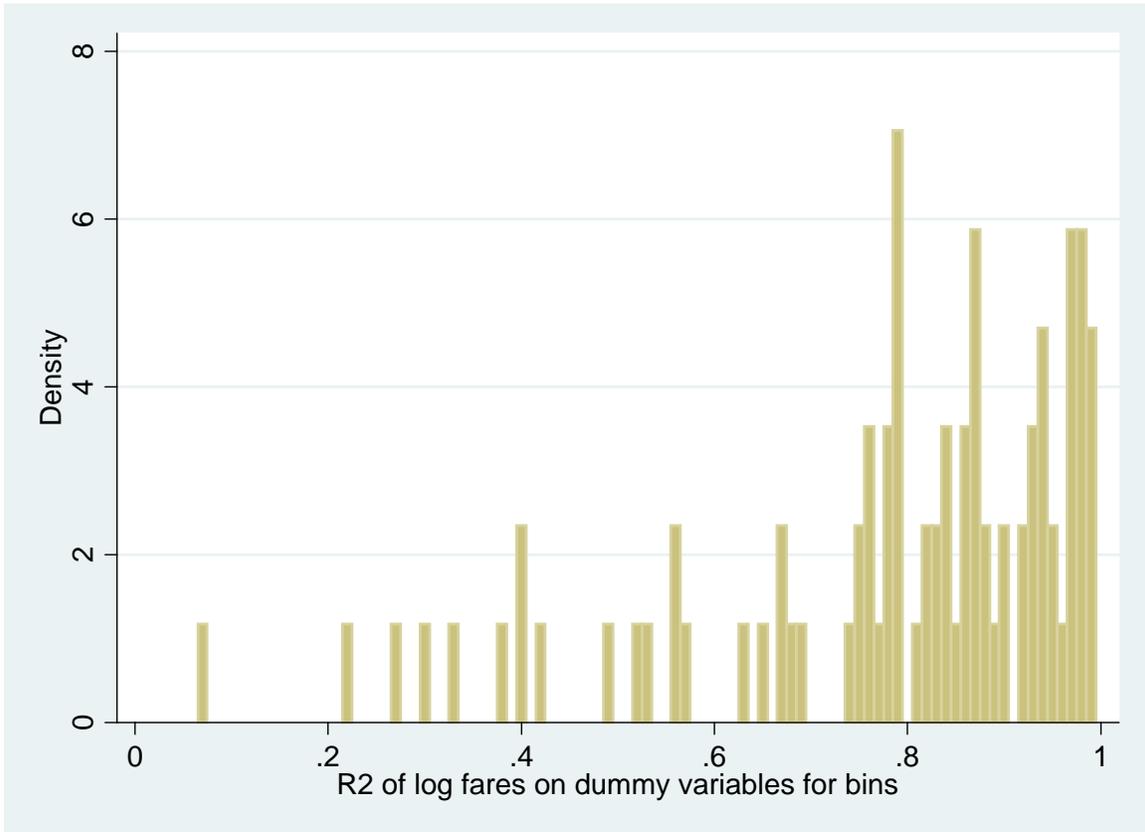


Figure 5

Histogram of Coefficient of Actual Load Factor for Each Route's Bins Regression

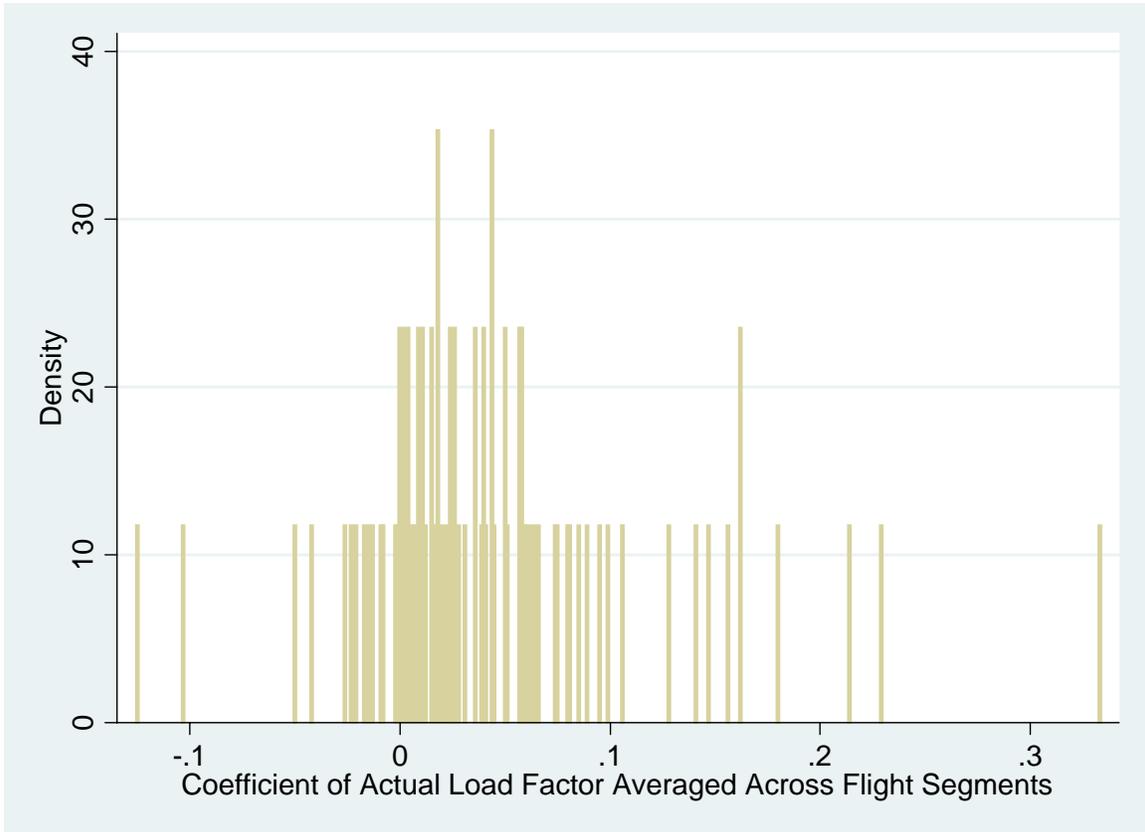
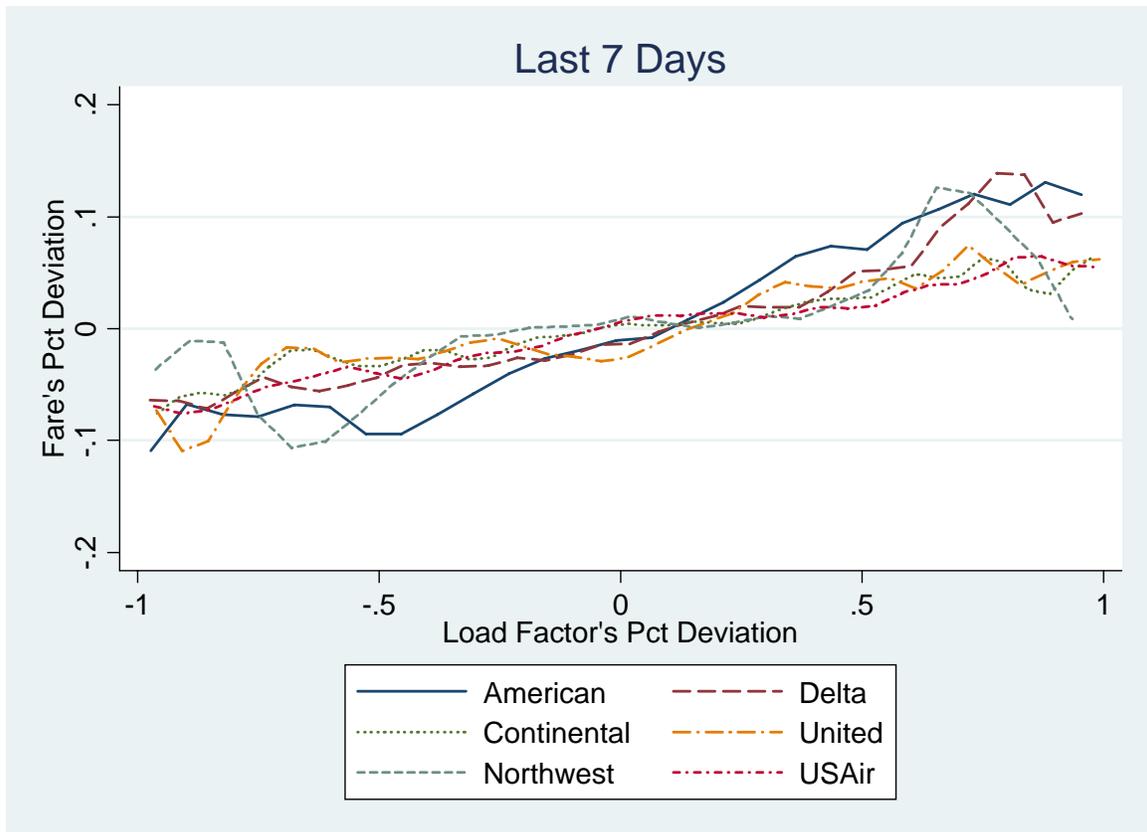


Figure 6

**Percent Deviation in Fare as a Function of
Percent Deviation in Load Factor
at Date of Purchase**



Note: Using tickets sold in 7 days before departure. All routes are included.