

NBER WORKING PAPER SERIES

TESTING THEORIES OF SCARCITY PRICING IN THE AIRLINE INDUSTRY

Steven L. Puller  
Anirban Sengupta  
Steven N. Wiggins

Working Paper 15555  
<http://www.nber.org/papers/w15555>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
December 2009

We thank Volodymyr Bilokach, Severin Borenstein, Jim Dana, Diego Escobari, Li Gan, Jim Griffin, Mike Mazzeo, Nancy Rose, Adam Shapiro, and Manuelita Ureta for helpful discussions, and seminar participants at Yale, MIT, the International Industrial Organization Conference, FTC Microeconomics Conference, Tennessee, 2009 AEA Meetings, Chicago Booth, UC Irvine, UC Davis, Baylor, and Carnegie Mellon. Manuel Hernandez provided excellent research assistance. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2009 by Steven L. Puller, Anirban Sengupta, and Steven N. Wiggins. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Testing Theories of Scarcity Pricing in the Airline Industry  
Steven L. Puller, Anirban Sengupta, and Steven N. Wiggins  
NBER Working Paper No. 15555  
December 2009  
JEL No. L1,L93

### **ABSTRACT**

This paper investigates why passengers pay substantially different fares for travel on the same airline between the same two airports. We investigate questions that are fundamentally different from those in the existing literature on airline price dispersion. We use a unique new dataset 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 predicts that variation in ticket prices is driven by differences between high demand and low demand periods. The second group of theories is that airlines practice price discrimination by using ticketing restrictions to segment customers by willingness to pay. We use a unique dataset, a census of ticket transactions from one of the major computer reservation systems, to study the relationships between fares, ticket characteristics, and flight load factors. The central advantage of our dataset is that it contains variables not previously available that permit a test of these theories. We find only mixed support for the scarcity pricing theories. Flights during high demand periods have slightly higher fares but exhibit no more fare dispersion than flights where demand is low. Moreover, the fraction of discounted advance purchase seats is only slightly higher on off-peak flights. However, ticket characteristics that are associated with second-degree price discrimination drive much of the variation in ticket pricing.

Steven L. Puller  
Department of Economics  
Texas A&M University  
3046 Allen  
College Station, TX 77843-4228  
and NBER  
puller@econmail.tamu.edu

Steven N. Wiggins  
Texas A&M University  
swiggins@tamu.edu

Anirban Sengupta  
Analysis Group  
San Francisco, CA  
asengupta@analysisgroup.com

## 1. Introduction

This paper investigates airline fare dispersion and fare differences using a fundamentally different approach than found in the extant literature. It is well-known that airline prices exhibit substantial price dispersion. Borenstein and Rose's seminal paper (1994) analyzes dispersion in 1986, and subsequent research has shown that dispersion has continued to grow in the last two decades (Borenstein and Rose (2007)). Existing empirical analyses of this dispersion, beginning with Borenstein and Rose and continuing with Stavins (2001) and Gerardi and Shapiro (2009), have focused on explaining the variation in dispersion across routes. These studies do not consider flight level differences or the differences in prices paid based on scarcity. Instead, these empirical analyses try to determine whether price dispersion rises or falls with changes in competition across routes and time. This previous research focuses on differences in price dispersion *across* routes.

In contrast to these papers, our paper investigates the empirically larger, and arguably more fundamental issue regarding the causes of price differences and price dispersion *within* a route. We study differences in the fares paid by individual passengers—both those traveling on the same flight, and those traveling on different flights on the same route. There exists substantial price dispersion for all legacy carriers on most routes, and the amount of dispersion on those routes with the least dispersion, whether monopolistic or competitive, is in fact quite substantial. This baseline level of dispersion found on all routes is quite large.

There is an ongoing debate as to the cause of price differences and associated fare dispersion for travel between the same cities on the same airline. One class of models posits classical price discrimination. These models argue that because business and leisure flyers have different demand elasticities, airlines use ticketing restrictions such as refundability or Saturday night stay requirements to price discriminate between groups of flyers. These models hinge on using restrictions to create fencing devices between different types of customers.

Another class of models postulates that prices are linked to the scarcity value of an empty seat. This class of models posits that fare dispersion arises from stochastic

peakload pricing in the face of uncertain demand for a perishable product. In these models, airlines set prices in order 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. In such models, price dispersion emerges because uncertain demand creates differing shadow costs of seats, leading to price heterogeneities. Dana (1999a,1999b) shows that when demand is uncertain, a perishable inventory can generate price differences linked to the probability of sale.<sup>2</sup> Gale and Holmes (1993) show that when there are peak and off-peak flights, perishable demand can lead to more discount fares being offered off-peak because advance purchase discounts are used to encourage leisure travelers to fly in off-peak periods.

Assessing the relative importance of these two classes of models – price discrimination and scarcity pricing – in generating price dispersion has been difficult due to data limitations. This paper uses a unique new dataset to investigate recent theories of stochastic peakload pricing and to assess their impact on price dispersion. The standard dataset on ticket transactions does not include ticket-level information regarding the time of sale and travel, a ticket’s characteristics, or individual flight load factors required to test either theories of scarcity pricing or price discrimination.<sup>3</sup> As a result, the standard dataset cannot be used to study how the levels or dispersion of fares vary in patterns predicted by either of the two major classes of pricing theories.

This paper uses new data that overcome these limitations, permitting the development of an extensive set of empirical tests of the central hypotheses provided by these major theories. We use transactions through one of the major computer reservation systems and information on ticket restrictions. Among the most important variables, our data include measures of ticket restrictions, fares, and flight-level load factor for each flight segment of a travel itinerary. In addition, our data include the dates of purchase and travel so that we can calculate flight-level measures of expected and realized load

---

<sup>2</sup> Dana’s model builds on the pioneering analyses of Prescott (1975) and Eden (1990).

<sup>3</sup> The most commonly used data to study airline pricing is the Department of Transportation’s Passenger Origin and Destination Survey (Databank 1B). DB1B reports data quarterly but does not include either ticket characteristics or the date of purchase and departure.

factor at the time of purchase and departure. This information is used to test theories regarding the relationships among ticket characteristics, load factors, and fares.

The rich theoretical literature on scarcity pricing offers straightforward comparative static predictions for equilibrium relationships between a flight's transacted fares and its load factor. In Dana (1999b), firms facing uncertain demand will offer a schedule of tickets with different prices, and the price is tied to the probability of sale. The realization of demand will determine which tickets are purchased from the ticket schedule and, thus, demand drives the dispersion in transacted fares. In Gale and Holmes (1993), airlines with both peak and off-peak flights will offer discount, advance purchase seats on the off-peak flight to induce customers with a low-time cost of waiting to travel in the off-peak time. The peak flight will have virtually no advance purchase, discount seats. As a result, peak flights have higher average fares and less fare dispersion than off-peak flights. This paper uses our detailed data on ticket purchases to test an array of predictions arising from the scarcity pricing theories.

Our results generally offer only limited support for the scarcity pricing theories. These theories do help explain some differences in mean fares. They do not, however, explain dispersion. The scarcity pricing prediction with the largest effects, moreover, explains substantially less of the variation in fares than do ticket restrictions.

More specifically, the scarcity pricing models predict that fare differences and fare dispersion are associated with either expected load factor (Gale and Holmes) or realized load factor (Dana). In our data fare dispersion is not systematically related to either measure of load factor and is actually quite similar for both full and empty flights. The Gini coefficient does not vary materially based on either expected or realized load factor. Likewise, there is little support for the prediction that off-peak flights will have a higher share of low priced and advance purchase tickets -- the fraction of discounted advance purchase seats is only slightly higher on off-peak flights. More generally, we do not find strong evidence of substantial quantity restrictions on the sale of low priced tickets on high load factor flights.

There is some support, however, for the prediction that fare levels vary with the expected load factor. We find robust evidence that flights with higher expected load factors exhibit higher average fares, which is consistent with Gale and Holmes'

predictions regarding peak and off-peak pricing. In particular, flights in the highest quartile of expected load factor have average fares that are 7.4% higher than flights in the lowest quartile. Also, we find some mixed support for Dana's prediction that average fares are higher on flights with higher realized load factors, holding expected load factor constant. Under one specification, fares on flights in the top quartile of realized load factor are 1% to 6% higher than fares on flights in the lowest quartile of realized load factor, with the effect larger on flights that are expected to be closer to capacity. (Under another specification, there is no systematic relationship between mean fares and realized load factor). Finally, we find that tickets sold in the last few days before departure are only slightly higher priced when a flight is becoming unusually full.

In contrast to finding modest support for scarcity pricing, we find that ticket characteristics that may be associated with second-degree price discrimination drive much of the variation in ticket pricing. For example, tickets with travel and stay restrictions on average sell for about two-thirds the price of comparable tickets without such restrictions. These differences persist, moreover, when one accounts for differences in load factors. These results, taken together, suggest that scarcity pricing plays a smaller role in airline pricing than models in which ticket restrictions are used as fencing devices to facilitate price discrimination.

The airline pricing problem is quite important, in part because airlines are an important industry and airline prices are highly dispersed and seemingly complex. Airlines use complex "revenue management" (or "yield management") systems to allocate seats. These systems undoubtedly pursue both goals of attempting to segment customers and attempting to implement stochastic peakload pricing. This paper assesses which of these goals play a larger role in generating the observed price dispersion.

In addition, a better understanding of airline pricing can perhaps lead to a better understanding of pricing in related hospitality industries. Many of the revenue management approaches developed in the airline industry are now widely used in other hospitality industries such as hotels, rental cars, cruise lines, sporting events, and trains. These industries share a common feature in that capacity is costly and it loses value if unused. Demand is also highly uncertain, and customers are heterogeneous in their willingness to pay. Many of these industries feature highly complex pricing structures

that might be driven either by price discrimination or by scarcity pricing. A better understanding of airline pricing can lead to an improved understanding of price dispersion in these other industries.

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 Airline Pricing**

Price dispersion in airlines and other hospitality industries may be driven by several factors as addressed by distinct classes of theoretical models. One set of models is classical price discrimination. In these models airlines use ticketing restrictions to induce customers to self-select into groups (e.g. business and leisure travelers) that have different demand elasticities.

Another set of models posits that pricing is tied to the scarcity value of an empty seat. Airline seats are costly but perishable. Markets characterized by perishable goods, costly capacity, and uncertain demand often exhibit highly dispersed prices. 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 explaining inter- and intra-firm price dispersion commonly found in these industries. Dana (1999b) provides a more complete description of this model and extends the model to monopoly and oligopoly settings.

Prescott's model posits a perishable good, such as a concert ticket or an airline seat, that entails costly capacity of  $\lambda$  per unit and perhaps a marginal cost, which we will ignore for now. Suppose a perfectly competitive seller must precommit to a schedule of prices for tickets and cannot adjust prices based on realizations of demand, e.g. tickets must be printed in advance. Heterogenous consumers with unit demand arrive in random order and purchase the lowest priced ticket available when they arrive. Following Dana's presentation, it is easiest to think of two demand states, high and low. 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 wherein those seats sell.

Hence 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, those tickets with a lower probability of being sold will be higher priced. In Dana’s monopoly version of this model, 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.

Dana’s model has a variety of important implications for pricing in industries with uncertain demand and costly capacity. One of the most significant 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’s (1999b) model provides several testable implications for *transacted* tickets. The model predicts comparative static relationships between realized load factors and the mean and dispersion of fares. To see this, suppose there are 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 Dana, that consumers arrive and choose the lowest fare available. In a flight with a low realized load factor, only low fare tickets are purchased. In medium load factor flights, low and medium fare tickets are purchased. In high load factor flights, low, medium, and high fare tickets are purchased.

The Dana model predicts four relationships between flights with different realized load factor *but the same ex ante distribution of demand*, *Prediction 1*: The mean fare of transacted tickets will be higher on flights with high realized load factors; *Prediction 2*: There will be more fare dispersion on flights with high realized load factors; *Prediction 3*: The share of high-priced tickets will be larger on flights with high realized load

factors; and *Prediction 4*: When bookings are unusually high for a given number of days before departure, there will be more high-priced tickets sold in the remaining days before departure as compared to flights where bookings are average or below.<sup>4</sup> We test each of these predictions and investigate their ability to organize airline pricing data and explain fare dispersion.

Gale and Holmes (1992, 1993) use a different formal structure to develop similar results regarding the sales of high and low priced tickets for peak and off-peak flights. Gale and Holmes use a mechanism design approach to model advance purchase discounts in a monopoly market. Such discounts are used to divert customers with a low opportunity 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 will only learn their preferred flight time shortly before departure. Customers also vary in their opportunity cost of waiting. In this setup, firms will offer discounted advance purchase tickets for sale in the period before the uncertainty regarding preferred flight times is resolved. These discounts are priced to entice low cost-of-waiting customers to self-select in favor of the off-peak flight. In equilibrium (more) discounted fares are offered on the off-peak flight. In a related paper, 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. These models are fundamentally about scarcity pricing because the lower fares on the off-peak flight reflect the lower opportunity cost of service during those periods. Dana (1998) builds upon the advance purchase literature and shows that advance purchase discounts can arise in a perfectly competitive setting.<sup>5</sup>

The Gale and Holmes models provide three additional empirical implications that we add to our list of predictions from scarcity pricing models. *Prediction 5*: There will be fewer discount/advance-purchase seats sold in equilibrium on peak flights (i.e. flights that are expected to be full), and that average fares will be higher on those flights;

---

<sup>4</sup> 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 understand the causes of *within* carrier-route dispersion.

<sup>5</sup> 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.

*Prediction 6:* Peak flights will have higher average fares than off-peak flights; *Prediction 7:* Tickets sold on off-peak flights will exhibit greater fare dispersion because those flights will have a greater number of discounted seats.

Finally, note that one could envision a hybrid model involving scarcity pricing within a price discriminating mechanism. Although we do not formally develop such a model, we can make some assessment of whether our findings are consistent with such a hybrid model. This issue is discussed in Section 4.3.1.

The analysis below is made possible because we have data not previously available for study in the well-developed literature on airline pricing (see below). 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 or stay restrictions. Such information is essential for testing the theories put forth by Dana and by Gale and Holmes.

Our investigation differs substantially from other empirical work on airline pricing. One 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 using a cross-section of markets. They show an increase in dispersion as markets become more competitive. Gerardi and Shapiro (2009) investigate the effect of competition on dispersion using a panel data set, and in contrast to Borenstein and Rose find that increases in the competitiveness of a route reduce dispersion. They interpret their findings to suggest that more competition reduces the ability of incumbent carriers to implement price discrimination. Stavins (2002) uses a novel data set on posted prices and a subset of ticket characteristics to find evidence consistent with both Saturday night stay and refundability being used as price discriminating devices. Using these data, Stavins 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.

Another branch of the literature uses posted fares to study the evolution of fares over time. McAfee and Velde (2004), for example, 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. Escobari and Gan (2007) gather and use data from posted minimum prices to test Dana's theories.

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 (2008) 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.

Our contribution to this empirical literature is to better understand the extent to which the pricing on a carrier-route is driven by scarcity pricing or by price discrimination. We test the seven comparative static implications of the Dana and the Gale and Holmes models that we list above. We test whether the levels and dispersion of fares are systematically different on flights that have a higher expected or realized demand. We nest this central hypothesis in an empirical model where the baseline is the pricing that occurs in off-peak flights with low expected and realized demand. We compare how transacted fares differ when load factor is higher. Hence the paper 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 differences and price dispersion as compared to a model where airlines use ticket restrictions as fencing devices as postulated in the price discrimination literature.

### 3. Data

#### 3.A. Tickets in Our Sample

We use a census of all transactions conducted via 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, origin and destination, airline, flight number for each leg of the itinerary, dates of purchase, departure and return, booking class, and whether the ticket was purchased online or offline.<sup>6</sup>

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 fares of roundtrip itineraries; we double the fares for one-way tickets to obtain comparability. We exclude itineraries involving first class travel. This study includes tickets for travel on American, Delta, United, Northwest, Continental and USAir. These constitute the entire set of airlines that individually carried at least 5% of U.S. domestic customers with the exception of Southwest for whom we have only limited data.<sup>7</sup> We analyze tickets for travel in the fourth quarter of 2004 excluding travel on Thanksgiving weekend, Christmas, and New Years.<sup>8</sup>

We restrict our analysis to large routes. To choose these routes, for each of the six carriers, we stratified the sample to include the largest routes for each carrier with varied market structures.<sup>9</sup> The routes are listed in Table 1. We include tickets by any of

---

<sup>6</sup> For an analysis of online versus offline prices, see Sengupta and Wiggins (2006).

<sup>7</sup> Much of Southwest's sales occur through the airline's website.

<sup>8</sup> We exclude travel occurring from the Wednesday prior to Thanksgiving until the following Monday. Also, we exclude all travel beginning after December 22.

<sup>9</sup> Routes are airport pairs. A route is a monopoly if a single carrier operates more than 90 percent of the 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.

the six carriers listed above that serve any of the routes listed.<sup>10</sup> One consequence of choosing large routes is that the sample consists largely of routes from airlines' hubs—though this should not pose an issue for testing the general theories of airline pricing under investigation.

### 3.B. Ticket Characteristics

Because we 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 a separate data set that contains ticket-level restrictions.<sup>11</sup> Travel agents' computer systems can access historical data on fares 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 for travel on a specified carrier-route-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.

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 Appendix A.

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

---

<sup>10</sup> Our final sample studies 86 routes. Three of the 90 routes listed on Table 1 are large routes for two different carriers, and one route (San Juan, PR to Miami) does not have complete data.

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

<sup>12</sup> The travel agent told us that the historical archive maintained by her CRS would sometimes delete some 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.

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. The data in the table indicate 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 particular segments of the fare distribution. In Figure 1, we plot kernel density estimates of fares. 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 theories of scarcity pricing discussed above make predictions that depend upon two measures of load factor -- *realized* and *expected* load factors. 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 transactions. This information can 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.

A central feature of our data is that we are able to estimate the realized load factor at various times prior to and including departure for a given airline, route, departure date and flight number. As a result, we also can calculate average load factors over various sets of flights and use this average to proxy for the *expected* load factor that can be systematically predicted by the airlines. We can also determine the deviations between realized and expected demand by measuring deviations from mean load factors for individual flights and for particular intervals prior to departure.

---

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-route level. In particular, the Bureau of Transportation Statistics reports monthly data on the total number of tickets sold for each route 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 route. 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.

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 percent of its tickets for direct service between LGA and ORD through our CRS, we calculate the realized load factor to be 55 percent ( $=(26/0.36)/129$ ).

This load factor is, of course, measured with error, but the methodology implies the measurement error will have zero mean at the airline-route 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 route.<sup>13</sup> Finally, note that because we observe the sequence of transactions, we can also measure the realized load factor at different dates prior to departure (e.g. the flight is half full as of seven days before departure and two-thirds full as of two days before departure).

This basic methodology also can be used to calculate a carrier's average load factor on a flight, day of the week, or route. We use this average to measure the *expected* load factor that can be systematically predicted by the airlines. To do so, we calculate the average load factor across our sample for a particular carrier's flight number 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.

---

<sup>13</sup> We discuss possible attenuation bias below.

The theories of scarcity pricing provide comparative static predictions about the characteristics of tickets sold on unusually full flights during peak periods and on unusually empty off-peak flights. In some of the analysis below, we separate all flights into groups based upon the expected and realized load factor. For each carrier-route, we separate flights into quartiles, as illustrated in Figure 2. We first separate flight number/day of the week observations into quartiles based upon expected load factors, as illustrated by the columns in Figure 2. We denote these groups as “Full”, “Medium-Full”, “Medium-Empty” and “Empty”.<sup>14</sup> 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 flight number/day of the week. American’s flight 301 on Mondays has a relatively low average load factor (compared to other American flight numbers/day-of-the-week averages from LGA to ORD), so all 12 of those flights in our sample are classified as expected to be “Medium-Empty”.

Within each of these quartiles of expected load factor, we then categorize individual flight *dates* based upon *realized* load factors for those particular departure dates. In other words, among those flights expected to be full, flights are grouped into quartiles based upon the realized load factor for each date. Continuing the example above, for all American flights La Guardia to O’Hare that are expected to be “Medium-Empty”, we categorize each realization of actual load factor into one of four categories: “Full”, “Medium-Full”, “Medium Empty”, and “Empty”.<sup>15</sup> American’s flight 301 on October 11 with a realized load factor of fifty-five percent 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 realized to be unusually full among the flights that are expected to be full; the bottom right corner consists of flights that are unusually empty among those that are expected to be empty.

In order to test robustness, we perform many of the tests using an alternative measure of *expected* load factor. Rather than use the sample average load factor of each

---

<sup>14</sup> 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.

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

carrier-route-flightnumber-day-of-week to define expected load factor, we use an alternative measure based upon variation in the time of day and season of the flight. This alternative definition is intended to account for the fact that airlines may be able to predict that certain times of the day or certain weeks of the year may have systematically different demand. (However, recall that we have already excluded dates surrounding Thanksgiving and Christmas). We create seven timeslots that have systematically different average load factors in our sample: weekdays 1-5am, weekdays 6-9am, weekdays 10am-1pm, weekdays 2pm-7pm, weekdays 8pm-midnight, Saturdays, and Sundays. Likewise, we divide flights by the calendar week of the year. Then we measure a flight's *expected* load factor as the sample average load factor for the carrier-route-timeslot-week-of-year. We use this measure to categorize flights into quartiles of expected load factor, and sort flights as described above into cells of Figure 2.

### 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 percent of itineraries. Most tickets are purchased in the days shortly before departure; the fraction of tickets purchased within three, six and thirteen 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.

The data we analyze include 620,307 itineraries across the six carriers on these 86 large routes. We measure ticket characteristics for 224,108 (or 36%) of these itineraries.

## 4. Testing Implications of Pricing Theories

### 4.1. Motivating Analysis

In order to illustrate our data, 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 number of days in advance, there is substantial variation in transaction prices. Prices, however, rise on average as the date of purchase nears the departure date. Second, tickets appear to be sold at discrete sets of prices and these prices show up as bands of prices in the figure. These price bands can better 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 fare buckets, and that yield management personnel choose 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 for example, some very low priced tickets are sold up to the day of departure. Clearly, this dispersion could be caused by a variety of factors. For example, higher priced tickets could be those with few restrictions on refundability; or the high priced tickets could be for travel on flights with higher expected or actual load factors. 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 a ticket's restrictions as well as its 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 itineraries' flight segments. We want to be cautious in interpreting this model as a pricing equation; these regressions only characterize equilibrium relationships in our data. Nevertheless, these regressions provide suggestive results consistent with our more formal investigation below.

We include several measures of specific 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.<sup>16</sup> The omitted category is more than twenty-one days.

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

---

<sup>16</sup> 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.

noteworthy that these characteristics along with the fixed effects explain nearly 70% of the variation in fares.<sup>17</sup> It is also worth noting that the magnitude of these coefficients changes little across the columns as we introduce measures of load factor.

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. "*LF\_Actual – Averaged Across Flight Segments*" is the average of realized load factor across all flight segments of the itinerary. 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 percent increase in fares (0.34\*0.045).

These results indicate a much smaller influence of load factor on fares as compared to ticket restrictions. When actual load factor is added to the model, the coefficients of the ticket characteristics are largely unchanged. We also find that the addition of the load factor measure does not substantially increase the fit of the model; the R<sup>2</sup> only rises from 0.695 to 0.696.<sup>18</sup>

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 route-carrier-flightnumber-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 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 a slightly higher fare. A one standard deviation increase in this measure of expected load factor is associated with a 2.3 percent increase in fare. As with the above

---

<sup>17</sup> The R<sup>2</sup> of a regression with only the fixed effects is 0.356.

<sup>18</sup> In unreported regressions, we include only fixed effects and then separately estimate adding load factor. Adding load factor raises the R<sup>2</sup> from 0.356 to 0.359.

case using realized load factor, the association between higher fares and load factor is relatively small.<sup>19</sup>

The measures of realized and expected load factors are strongly positively correlated, so regressions including only one of these metrics are likely to capture both effects. Column (4) includes measures of both realized and expected load factor. *LF\_Expected* remains statistically significant – a one standard deviation increase in expected load factor is associated with a 2.1 percent increase in fares. The association between realized load factor and fares, however, is no longer significant.

In the remaining columns, we allow for a non-linear association between load factors and fares, 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 realized load factor for the fullest flight segment (the maximum of realized load factor across all flight segments). The relationship between realized 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.

One possible explanation for the small association between fares and load factor is that our measure of load factor is unbiased but measured with error. As discussed above, our load factor is measured with error because we only observe about one-third of all transactions. Such a form of measurement error could induce attenuation bias and affect our estimated coefficient of load factor. In Appendix B, we implement a measurement error correction. Those results suggest that measurement error is *not* the cause of our finding that there is a small association between fares and load factor.

These motivating regressions suggest that an itinerary with certain ticket characteristics, such as refundability or no travel restrictions, will tend to have a substantially higher fare. Moreover, this association between ticket characteristics and fares appears to be relatively independent of whether the flight segments are high or low demand. As compared to ticket characteristics, load factor is not associated with substantially higher fares.

---

<sup>19</sup> 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*.

Given this motivating analysis, next we test the seven predictions of the scarcity pricing models that are described in section 2.

#### 4.2 Testing Scarcity Pricing Predictions Regarding Pricing

We now turn to a direct test of Dana (1999b)'s central prediction regarding the relationships between load factors, fares, and fare dispersion. According to that theory, flights with the same expected distribution of demand have the same *offered* fares, but those flights with higher realized load factors will have higher mean *transacted* fares and more fare dispersion. Intuitively, passengers buy from the lowest priced fare bucket that is open when they purchase; so if there are more realized purchases, then there will be more transactions in higher priced buckets leading to higher average fares and more dispersion.

We test if flights with the same expected load factor but higher realized load factors have both higher mean fares (Prediction 1) and more dispersion in fares (Prediction 2). 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-route 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 individual flight realizations within these quartiles based on the demand realizations for individual flights. Dana's model provides 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 has two predictions: (1) mean transacted fares will be higher on flights with "Full" realized load factors as compared to flights with "Empty" realized load factors, and (2) there will be more fare dispersion on flights with "Full" realized load factors as compared to those with "Empty" realized load factors.<sup>20</sup> These predictions

---

<sup>20</sup> 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 means/Gini coefficients, but one still should observe higher means/Gini coefficients for itineraries involving flights realized to be full.

imply that as one moves up any column of expected load factor in Figure 2, both mean fares and dispersion should rise.

The results for mean fares are reported in Table 4 as deviations from the mean fare on a route. More specifically, we measure the mean fare for each carrier-route, and then calculate the percentage by which mean fares for each category of load factor deviates from that overall mean for the carrier-route.<sup>21</sup> To summarize these deviations across all carrier-routes, Table 4 reports the passenger-weighted average deviation within each of the sixteen load factor cells. These averages reflect the central tendency of how fares respond to load factors across carriers and routes.

Prediction 1 for Dana’s model is that mean fares should rise as one moves up each column in Table 4. As one goes up any column in this table, realized load factor increases holding expected load factor constant. The results for this prediction differ depending on how expected load factor is measured. The top portion of Table 4 uses expected load factor measured at the flight number-day-of-week level. The top panel shows no rise in average fares as realized load factors rise, conditional on expected load factor. In fact, average fares fall slightly as realized load increases.

In contrast, the bottom panel of Table 4 defines expected load factor using particular groups of time slots and calendar weeks (see section 3.C). Using this definition, as realized load factor rises, so do average fares, holding expected load factor constant (i.e. moving up the columns). This effect is strongest for flights that are expected to be closer to capacity, that is flights expected to be “Full”. For the expected to be “Full” flights, the flights realized to be “Empty” have average fares 1.1% above the mean for the carrier-route, while the flights realized to be “Full” have average fares 7.5% above the mean. A similar but smaller effect is observed on flights expected to be “Medium-Full”. Almost no effect is observed for flights expected to be “Medium-Empty” and “Empty”. Hence at least according to the second measure of expected load factor, there is some support for Prediction 1 of Dana’s theory regarding mean fares and realized load factors on flights expected to be closer to capacity.

Table 4 also provides evidence regarding Prediction 6 from the Gale and Holmes theory on mean fares. That theory indicates that average fares should be higher on peak

---

<sup>21</sup> We use within carrier-route deviations to control for carrier and route differences in average fares.

flights -- that is, flights expected to be full. The evidence in Table 4 is highly consistent with this prediction. In particular, fares on flights expected to be full average about five to twelve percent higher than fares on flights expected to be empty. This result holds for both panels in Table 4 which show results for both measures of expected load factor. Hence there is strong support for Gale and Holmes Prediction 6 that average fares will be higher on peak versus off-peak flights.

Next, we test the scarcity pricing models' predictions regarding the dispersion of fares, Predictions 2 and 7. In order to measure dispersion, we calculate the Gini coefficient for each carrier-route in each category of expected-realized load factor (i.e. each cell of Figure 2), and calculate the passenger-weighted average Gini for each cell. Because this analysis does not use data on ticket characteristics, we can use all transactions observed through our CRS.

Results are shown in Table 5. The average within carrier-route Gini coefficient is approximately 0.28.<sup>22</sup> However, there is almost no variation in this metric of dispersion associated with changes in expected or realized load factor.

This pattern of dispersion is not consistent with Prediction 2 from Dana (1999b), which predicts increased dispersion for flights with higher realized load factor, conditional on expected load factor. As with the analysis above, we show results using two different measures of expected load factor. The top panel of Table 5 shows that for flights with the same expected load factor, dispersion is *decreasing* when the realized load factor is larger. This decrease in dispersion, however, 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 using the alternative measure of expected load factor shows very similar results. Remarkably, fare dispersion is very similar on all flights, independent of the actual and expected load factor.

The patterns of dispersion in Table 5 also do not support Prediction 7 of the Gale and Holmes model that off-peak flights will have greater fare dispersion because those

---

<sup>22</sup> 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.

flights will have a greater number of discounted seats. Our proxy for off-peak flights are the flights both expected and realized to be “Empty”; likewise peak flights are those both expected and realized to be “Full”. Using the first definition of expected load factor in the top panel of Table 5, “Empty-Empty” flights have a Gini coefficient of 0.283 while “Full-Full” flights have a Gini of 0.275. Using the alternative definition of expected load factor in the bottom panel, the Gini coefficients are 0.288 and 0.282 for “Empty-Empty” and “Full-Full”, respectively. The magnitude of the difference in dispersion is economically inconsequential. These Gini coefficients have a straightforward interpretation: the expected fare difference of two randomly selected tickets on a peak flight is 56.4% of the mean fare, while on the off-peak flight it is 57.7%.<sup>23</sup> These results indicate that neither theoretical model of scarcity pricing helps explain the factors leading to fare dispersion.

#### 4.3 Testing Scarcity Pricing Predictions Regarding Ticket Allocations in Peak Demand

We now test two key predictions posited by both Gale and Holmes and by Dana. We test Gale and Holmes’ Prediction 5 that discount, advance purchase tickets will account for a larger share of tickets on off-peak flights. We also test Dana’s similar Prediction 3 that full-fare, high price tickets will account for a larger share of tickets on peak flights. These hypotheses 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. For Group 1 tickets the mean is \$631, for Group 2 it is \$440, and for Group 3 it is \$281. Hence these groupings are associated with high, medium and low fares. The fraction of tickets sold in Groups 1-3 are 26%, 32%, and 42%, respectively.

---

<sup>23</sup> The Gini coefficient has the property that twice the Gini is the expected difference of two randomly chosen fares as a proportion of the mean fare.

These groupings are associated with large differences in average fares, and by themselves account for a large share of the differences in observed fares.<sup>24</sup>

In order to test the theoretical implications, we must define flights in our data that are peak and off-peak flights. We define peak (off-peak) flights as those that are expected to be “Full” (“Empty”) and realized to be “Full” (“Empty”), using the same classification as described above in Section 3 and illustrated in Figure 2. We use the metric of expected load factor based on the flight number-day of week, but the appendix contains results based on the timeslot metric of expected load factor. The analysis focuses on the two corner cells from the diagonal of Figure 2. Full flights are those that are in the upper quartile of expected load factor and the upper quartile of realized load factor. Empty flights are in the lower quartile of both. We use these groups as the extremes of full and empty flights where there should be the largest and most significant differences in ticket allocation systems associated with scarcity. Next we break down tickets sold on these flights into groupings based on the ticket groups described above (Groups 1-3) and according to the number of days in advance a ticket was purchased. The latter categorization is to determine whether the sale of certain categories of tickets is restricted in certain periods (a la Gale and Holmes) or whether they sell out (a la Dana).

We count each coupon sold on these peak and off-peak flights by Group and Days in Advance.<sup>25</sup> These allocations are calculated at the airline level to account for any differences across airlines. These tabulations are shown in Tables 6a-6b.

First, we test Prediction 5 of Gale and Holmes that off-peak flights will have more discount/advance-purchase seats sold in equilibrium. We compare the fraction of Group 3 tickets sold greater than twenty-one and greater than fourteen days in advance. For purchases more than twenty-one days in advance, the fractions of low price/Group 3 tickets sold on low versus high load factor flights for various airlines include: American: 18% vs. 18%; Delta: 20% vs. 17%; Continental: 23% vs. 25%; United 23% vs. 20%; USAir: 11% vs. 14%; and Northwest: 28% vs. 26%. Thus, two of the six airlines sell

---

<sup>24</sup> A regression of fares on these groupings yields an  $R^2$  of approximately 0.65.

<sup>25</sup> 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.

*fewer* discount/advance-purchase seats on off-peak flights. And for the airlines that sell more such tickets on off-peak flights, the differences are small.

If we consider Group 3 tickets sold fourteen or more days in advance, the differences are similar. The fraction of Group 3 tickets sold on low versus high load factor flights are American: 28% vs. 31%, Delta: 32% vs. 28%, Continental: 35% vs. 35%, United: 35% vs. 34%, USAir: 17% vs. 19%, and Northwest: 43% vs. 43%. The comparisons using the alternative metric of expected load factor yield results even less supportive of the theoretical predictions (see appendix table 1).

Finally, we test Dana's Prediction 3 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 differences in tickets sold on peak versus off-peak flights are economically small. 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 vary across airlines. The differences in sales of Group 3 tickets are: American (53% vs. 52%), Delta (54% vs. 45%), Continental (53% vs. 48%), United (67% vs. 66%), USAir (28% vs. 28%), and Northwest (71% vs. 67%). In fact, off-peak flights do indeed have a higher fraction of "discount" tickets but the effect is only 0% to 9% percent, depending upon the airline. In the appendix tables, we use the alternative metric of expected load factor and find the differences to be even smaller, ranging from -2% to 4%. Thus, the allocation in the off-peak flights illustrates that there is a baseline allocation of high and low price seats that is driven by non-scarcity factors. Apparently, there are only modest deviations from this baseline associated with scarcity.

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 realized to be "Empty". Nevertheless, the airlines do not sell a substantially larger fraction of low fare tickets on these flights.

### 4.3.1 A Hybrid Model

Finally, we address whether our empirical tests of scarcity pricing are robust to a more general hybrid model of pricing in a setting with stochastic demand and customer heterogeneity in demand elasticity. As noted above, one could imagine a hybrid model involving scarcity pricing within a price discriminating mechanism. Suppose that airlines are able to segment customers into business and tourist groups using ticketing characteristics. After segmenting customers into such groups, the airlines then employ a Dana-like scarcity pricing model within each group. Suppose that high value customers also have a higher variance of demand, and therefore are higher cost to serve from an uncertainty point of view. In this setting, higher fares for refundable tickets could reflect at least in part a more uncertain demand for business customers rather than merely reflecting customer heterogeneity in demand elasticities.

We can assess the assumption of this hybrid model that business customers have a higher variance of arrivals. More specifically, Group 1 customers pay higher average fares and purchase tickets that are refundable, given our definition of Group 1 above. Group 3 customers pay lower average fares and have restricted tickets. Our data permit us to examine the standard deviation of arrivals of both groups. In our data, the standard deviation of Group 3 customers at the flight level is higher than the standard deviation of Group 1 customers. This suggests that higher prices for unrestricted tickets are not driven by differences in costs of service associated with demand uncertainty.

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

Prediction 4 from Dana's model is that mean transacted fares will increase when a flight is becoming unusually full. To motivate this test, recall Dana's example of pricing for events such as concerts. 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 at each price, and consumers buy the cheapest ticket available when they arrive at the ticket window. Consider tickets

purchased in the last week before the event. If an unusually large number of consumers have arrived and purchased tickets prior to that last week, then the tickets purchased during the last week will be sold at a higher price than if the “normal” number of people had purchased prior to the last week. This element of the model results in the prediction that when a flight is closer to capacity at a given point in time prior to departure, then going forward fares on that flight 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 and fare as of each specific day before departure. We then calculate deviations from this mean – calculating deviations in both mean fares and load factor for each day prior to departure. We then transform these deviations into percentages of the conditional mean fare and load factor. For example, if the average load factor was 50 percent for American from La Guardia to Chicago six days in advance, then a ticket sold in the last six days on a flight that was 55 percent full would have a ten percent positive deviation in load factor. We calculate the analogous measure of the itinerary’s deviation in fare relative to the mean, based on the number of days in advance that the ticket was purchased. For tickets with multiple legs, we use the first leg. These calculations permit us to determine whether fares are on average higher near departure when a flight is nearer capacity, as predicted by Dana.

Figure 4 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 somewhat higher fares, but the effect is 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 when load factors are 10% higher at the time of purchase (and roughly a 0.8% increase for the other carriers).<sup>26</sup> In unreported analysis, we restrict the tickets to those sold in the last 3 days before departure, and find qualitatively similar results. These results suggest that the levels of fares are only modestly higher when load factors are higher.

---

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

#### 4.5 Summary of Key Findings Regarding Scarcity Pricing

Before evaluating the relative importance of scarcity pricing and price discrimination, it is useful to summarize the above results. There is no support for the predictions regarding fare dispersion. Flights expected to be peak do not have more dispersion than those expected to be off-peak (Prediction 7 from Gale and Holmes). Also, flights realized to be Full do not have more dispersion than those realized to be Empty, conditional on expected load factor (Prediction 2 from Dana).

Dana's predictions regarding the levels of fares have some support (Predictions 1 and 4). Under one of the two metrics of expected load factor, higher realized load factor conditional on expected load factor is associated with higher fares. This effect arises primarily on flights above the median expected load factor. Under the other metric, the data do not support the model's predictions. Also, we find that flights unusually full for a given number of days before departure, exhibit modestly higher fares in the few days before departure. This effect is largest on American but is generally modest.

Tests about the allocation of different types of tickets across peak and off-peak flights yield mixed results (Predictions 3 and 5). For some carriers, the sizes of the relative allocations are opposite that predicted by theory. When the relative percent allocations are consistent with the models' predictions, the percent differences are relatively small and suggest a small number of seats reallocated.

However, one prediction is borne out fairly robustly in the data. Gale/Holmes predict that flights with high expected load factors will have higher fares than those with lower expected load factors (Prediction 6). We find that flights with higher expected load factors do have highest average fares. Relative to flights expected to be in the bottom quartile of expected load factor ("Empty"), flights in the top quartile of expected load factor ("Full") exhibit fares that are about 7.4% higher. We interpret these findings to indicate that some portion of the variation in fares on a carrier-route is driven by factors associated with scarcity.

#### 4.6 Relative Impact of Scarcity and Ticket Restrictions

In this section, we show that even the scarcity pricing prediction with the *most* empirical support explains substantially less variation in fares than ticket restrictions. More specifically, we compare the extent to which fare variation is driven by scarcity pricing versus the impact of ticket restrictions which may be associated with price discrimination. This comparison reveals that even the scarcity pricing prediction with the strongest empirical support – Gales and Holmes Prediction 6 – explains substantially less of the variation in fares than do ticket restrictions.

This difference can be seen in Table 7, which shows percentage deviations in fares from the mean for the carrier-route, similar to Table 4. In contrast to the previous table, Table 7 shows the percentage deviations broken down by expected load factor (the predictor based on Gale and Holmes) as compared to ticket characteristics. We classify tickets characteristics in the same manner as in section 4.3. Group 1 includes only highly flexible, refundable tickets; Group 2 includes tickets that are not refundable but ones that do not include restrictions on days of travel or stay; and Group 3 includes nonrefundable tickets with restrictions. We believe that these grouping of characteristics are likely to capture ticketing restrictions that are used to segment customers and implement price discrimination.

Table 7 shows that the magnitude of variation in mean fares explained by ticket restrictions is substantially larger than the variation explained by expected load factor. Group 1 tickets average 47% *above* the mean, Group 2 tickets average 10% *above* the mean, and Group 3 tickets average 29% *below* the mean. Thus, tickets with travel and stay restrictions on average sell for substantially less than comparable tickets without such restrictions. In contrast, the effect of expected load factor is just over 7%, as we note above. The relative magnitude of these effects is not surprising given the results of the motivating regressions that found the coefficients of ticket characteristics to be much larger than those of expected or realized load factor. Table 7 also shows that to the extent that expected load factor is associated with higher fares, the primary differences are found in the fares of restricted Groups 1 and 2 tickets.

## 5. Conclusions

This paper tests basic implications of several of the leading models of scarcity-based pricing. Our results provide modest evidence that fare levels in the airline industry are driven by scarcity pricing models, and no evidence that fare dispersion is explained by these theories. When expected load factor is above the median, flights with higher realized demand do feature higher fares than flights with low demand using one of our measures of expected load factors. Hence there is some support for the scarcity pricing theories.

In contrast, there is no support for the ability of the scarcity pricing models to explain fare dispersion. Fare dispersion is not significantly higher on flights with higher realized demand. We also 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 that support is modest. We find much stronger evidence that certain sets of ticket characteristics drive much of the variation in ticket pricing, and these ticket 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 play a larger 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.

Gerardi, Kristopher S. and Adam Hale Shapiro. 2009. "Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry", *Journal of Political Economy*, 117(1): 1-37, February 2009.

Goolsbee, Austan and Chad Syverson, "How Do Incumbents Respond to the Threat of Entry? Evidence from the Major Airlines." *Quarterly Journal of Economics*, 123(4):1611-1633., November 2008.

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.

## Appendix A: Data

### *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.<sup>27</sup> 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

---

<sup>27</sup> 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.<sup>28</sup>

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.<sup>29</sup> 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.

---

<sup>28</sup> 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.

<sup>29</sup> 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.

## Appendix B: Robustness of Load Factor Coefficient to Measurement Error

We address whether the small association between fares and load factor (section 4.1 and Table 3) is driven by measurement error in load factor. As discussed in the main text, 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 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  $\Omega_\epsilon$  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  $\Omega_\epsilon$ . Results are below. 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.

### 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 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.

<b>American</b>	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
<b>Delta</b>	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
<b>United</b>	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
<b>Continental</b>	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
<b>Northwest</b>	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
<b>USAir</b>	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
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)**
Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
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  $R^2$  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**

## Fare Deviations from Carrier-Route Mean by Expected and Realized Load Factors

*Expected Load Factor Defined using Flight Number and Day of Week*

<u>Realized Load Factor</u>	<u>Expected Load Factor</u>			
	Full	Medium-Full	Medium-Empty	Empty
Full	5.8%	-1.4% †	-2.8%	-5.7%
Medium-Full	6.3% *	0.6% †	-2.7%	-5.8% †
Medium-Empty	4.3% †	2.1%	-2.3% †	-3.9% †
Empty	6.6%	1.8%	-0.4%	-2.4%

\* indicates that cell is statistically different from the one below it at 5% level, sign correct

† indicates that cell is statistically different from the one below it at 5% level, sign not correct

*Expected Load Factor Defined using Timeslot and Week of Year*

<u>Realized Load Factor</u>	<u>Expected Load Factor</u>			
	Full	Medium-Full	Medium-Empty	Empty
Full	7.5% *	4.1% *	0.1% *	-3.9% *
Medium-Full	4.3% *	2.6% *	-1.3% *	-5.3%
Medium-Empty	1.8%	1.6% *	-2.3% †	-4.6%
Empty	1.1%	-0.1%	-1.2%	-3.7%

\* indicates that cell is statistically different from the one below it at 5% level, sign correct

† indicates that cell is statistically different from the one below it at 5% level, sign not correct

## Notes:

(1) For each carrier-route, we calculate the percentage deviation in mean fare for each category of actual and expected load factor from the mean fare for the entire carrier-route. This table reports the passenger-weighted average percentage deviation across all carrier-routes. We only include a carrier-route if at least 1600 itineraries were observed.

(2) The construction of both metrics of expected load factor is described in the main text. Briefly, the first panel estimates the expected load factor using the average load factor over the sample for each route-carrier-flightnumber-day of week. The second panel estimates the expected load factor using the average load factor over the sample for each route-carrier-week of year-timeslot (the 7 timeslots are: weekdays 1-5am, weekdays 6-9am, weekdays 10am-1pm, weekdays 2pm-7pm, weekdays 8pm-midnight, Saturdays, and Sunday).

**Table 5**  
Gini Coefficients by Expected and Realized Load Factors

*Expected Load Factor Defined using Flight Number and Day of Week*

<u>Realized Load Factor</u>	<u>Expected Load Factor</u>			
	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

*Expected Load Factor Defined using Timeslot and Week of Year*

<u>Realized Load Factor</u>	<u>Expected Load Factor</u>			
	Full	Medium-Full	Medium-Empty	Empty
Full	0.282	0.280	0.269	0.267
Medium-Full	0.285	0.278	0.266	0.281
Medium-Empty	0.281	0.279	0.271	0.279
Empty	0.278	0.280	0.266	0.288

Notes:

(1) Cell values are passenger-weighted averages of the Gini coefficient for each carrier-route-load factor category. We only include a carrier-route if at least 1600 itineraries were observed.

(2) The construction of both metrics of expected load factor are contained in the main text. Briefly, the first panel estimates the expected load factor using the average load factor over the sample for each route-carrier-flightnumber-day of week. The second panel estimates the expected load factor using the average load factor over the sample for each route-carrier-week of year-timeslot (the 7 timeslots are: weekdays 1-5am, weekdays 6-9am, weekdays 10am-1pm, weekdays 2pm-7pm, weekdays 8pm-midnight, Saturdays, and Sunday).

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

<b>Flights - Expected to be Low Load Factor &amp; Realized Low Load Factor</b>						<b>Flights - Expected to be High Load Factor &amp; Realized High Load Factor</b>					
<b>American - All Routes</b>											
<b>Groups</b>	<b>0 to 6</b>	<b>7 to 13</b>	<b>14 to 21</b>	<b>&gt;21</b>		<b>Groups</b>	<b>0 to 6</b>	<b>7 to 13</b>	<b>14 to 21</b>	<b>&gt;21</b>	
Group 1	12%	1%	0%	0%	14%	Group 1	9%	1%	0%	0%	10%
Group 2	13%	7%	5%	8%	33%	Group 2	12%	10%	5%	9%	38%
Group 3	13%	12%	10%	18%	53%	Group 3	10%	11%	13%	18%	52%
	39%	20%	15%	26%			31%	23%	19%	28%	
<b>Delta - All Routes</b>											
<b>Groups</b>	<b>0 to 6</b>	<b>7 to 13</b>	<b>14 to 21</b>	<b>&gt;21</b>		<b>Groups</b>	<b>0 to 6</b>	<b>7 to 13</b>	<b>14 to 21</b>	<b>&gt;21</b>	
Group 1	4%	1%	1%	2%	8%	Group 1	4%	1%	0%	0%	6%
Group 2	17%	8%	6%	7%	38%	Group 2	19%	10%	8%	12%	49%
Group 3	11%	11%	12%	20%	54%	Group 3	7%	11%	11%	17%	45%
	32%	20%	19%	29%			30%	21%	20%	29%	
<b>Continental - All Routes</b>											
<b>Groups</b>	<b>0 to 6</b>	<b>7 to 13</b>	<b>14 to 21</b>	<b>&gt;21</b>		<b>Groups</b>	<b>0 to 6</b>	<b>7 to 13</b>	<b>14 to 21</b>	<b>&gt;21</b>	
Group 1	15%	3%	1%	1%	20%	Group 1	17%	4%	2%	1%	24%
Group 2	10%	7%	3%	7%	26%	Group 2	8%	6%	4%	10%	27%
Group 3	8%	10%	12%	23%	53%	Group 3	5%	8%	10%	25%	48%
	33%	21%	16%	31%			30%	19%	15%	36%	

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 classified as having a high/low *Expected* load factor if that flight has an average load factor in the top/bottom quartile of all flights-day of week for that carrier-route. A flight is classified as having a high/low *Realized* load factor by grouping all flights in each expected load factor quartile into quartiles based on the flight's realized load factor, as described in the text. 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**

<b>Flights - Expected to be Low Load Factor &amp; Realized Low Load Factor</b>						<b>Flights - Expected to be High Load Factor &amp; Realized High Load Factor</b>					
<b>United - All Routes</b>											
<b>Groups</b>	<b>0 to 6</b>	<b>7 to 13</b>	<b>14 to 21</b>	<b>&gt;21</b>		<b>Groups</b>	<b>0 to 6</b>	<b>7 to 13</b>	<b>14 to 21</b>	<b>&gt;21</b>	
Group 1	8%	1%	1%	2%	11%	Group 1	8%	1%	0%	1%	10%
Group 2	11%	7%	2%	2%	22%	Group 2	11%	8%	2%	3%	24%
Group 3	18%	14%	13%	23%	67%	Group 3	15%	16%	14%	20%	66%
	36%	22%	16%	26%			34%	26%	17%	24%	
<b>US Airways - All Routes</b>											
<b>Groups</b>	<b>0 to 6</b>	<b>7 to 13</b>	<b>14 to 21</b>	<b>&gt;21</b>		<b>Groups</b>	<b>0 to 6</b>	<b>7 to 13</b>	<b>14 to 21</b>	<b>&gt;21</b>	
Group 1	43%	8%	2%	1%	54%	Group 1	37%	11%	3%	2%	53%
Group 2	7%	6%	1%	3%	18%	Group 2	5%	5%	3%	7%	19%
Group 3	5%	6%	6%	11%	28%	Group 3	3%	6%	5%	14%	28%
	56%	20%	9%	15%			45%	21%	11%	23%	
<b>Northwest - All Routes</b>											
<b>Groups</b>	<b>0 to 6</b>	<b>7 to 13</b>	<b>14 to 21</b>	<b>&gt;21</b>		<b>Groups</b>	<b>0 to 6</b>	<b>7 to 13</b>	<b>14 to 21</b>	<b>&gt;21</b>	
Group 1	8%	2%	0%	1%	11%	Group 1	10%	2%	0%	0%	13%
Group 2	6%	3%	2%	6%	17%	Group 2	8%	4%	2%	5%	20%
Group 3	14%	14%	15%	28%	71%	Group 3	10%	14%	17%	26%	67%
	29%	19%	18%	34%			28%	20%	20%	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 classified as having a high/low *Expected* load factor if that flight has an average load factor in the top/bottom quartile of all flights-day of week for that carrier-route. A flight is classified as having a high/low *Realized* load factor by grouping all flights in each expected load factor quartile into quartiles based on the flight's realized load factor, as described in the text. Group 1 = Refundable tickets, Group 2 = Nonrefundable without travel or stay restrictions, 3 = Nonrefundable with travel and/or stay restrictions.

**Table 7**

Fare Deviations from Carrier-Route Mean by *Expected* Load Factor  
and Ticket Restrictions

*Expected Load Factor Defined using Flight Number and Day of Week*

Expected Load Factor	Group (capturing ticket characteristics)			ALL
	1	2	3	
Full	56.8%	13.4%	-28.1%	8.6%
Medium-Full	48.0%	10.2%	-28.7%	4.0%
Medium-Empty	40.8%	9.5%	-29.7%	1.6%
Empty	44.4%	9.1%	-30.9%	1.1%
ALL	47.3%	10.5%	-29.4%	

*Expected Load Factor Defined using Timeslot and Week of Year*

Expected Load Factor	Group (capturing ticket characteristics)			ALL
	1	2	3	
Full	62.3%	12.4%	-28.5%	7.5%
Medium-Full	49.6%	10.8%	-29.2%	3.8%
Medium-Empty	40.9%	10.0%	-28.9%	3.3%
Empty	39.3%	8.6%	-30.9%	0.3%
ALL	47.3%	10.5%	-29.4%	

Notes:

(1) For each carrier-route, we calculate the percentage deviation in mean fare for each category of expected load factor and Group from the mean fare for the entire carrier-route. This table reports the passenger-weighted average percentage deviation across all carrier-routes. We only include a carrier-route if at least 1600 itineraries were observed.

(2) Groups are defined in the main text. Group 1 includes refundable tickets. Group 2 includes non-refundable tickets that have neither travel nor stay restrictions. Group 3 includes non-refundable tickets with travel and/or stay restrictions.

(3) The construction of both metrics of expected load factor is described in the main text. Briefly, the first panel estimates the expected load factor using the average load factor over the sample for each route-carrier-flightnumber-day of week. The second panel estimates the expected load factor using the average load factor over the sample for each route-carrier-week of year-timeslot (the 7 timeslots are: weekdays 1-5am, weekdays 6-9am, weekdays 10am-1pm, weekdays 2pm-7pm, weekdays 8pm-midnight, Saturdays, and Sunday).

(4) The row means are positive for each category of expected load factor ("ALL") because the distribution of percentage deviations is right skewed.

## Appendix Table 1A: Robustness Tests of Comparative Statics of Dana and Gale & Holmes

### Testing Robustness to Definition of Expected Load Factor

In these tables, Expected Load Factor is defined as the mean realized load factor averaged by carrier-origin-destination-timeslot-week of year. The timeslots are: weekdays 1-5am, weekdays 6-9am, weekdays 10am-1pm, weekdays 2pm-7pm, weekdays 8pm-midnight, Saturdays, and Sunday.

<b>Flights - Expected to be Low Load Factor &amp; Realized Low Load Factor</b>
--

<b>Flights - Expected to be High Load Factor &amp; Realized High Load Factor</b>
--

#### 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	9%	2%	0%	0%	11%
Group 2	12%	7%	5%	8%	32%	Group 2	13%	11%	5%	9%	38%
Group 3	13%	11%	10%	19%	53%	Group 3	10%	11%	13%	17%	51%
	39%	20%	15%	27%			32%	24%	18%	27%	

#### 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	4%	1%	1%	1%	7%	Group 1	3%	1%	0%	0%	5%
Group 2	19%	9%	6%	7%	41%	Group 2	18%	10%	8%	11%	47%
Group 3	9%	11%	11%	20%	52%	Group 3	6%	12%	11%	19%	48%
	33%	20%	19%	29%			28%	22%	20%	30%	

#### 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	15%	4%	1%	1%	21%	Group 1	18%	4%	2%	1%	24%
Group 2	9%	7%	3%	8%	27%	Group 2	8%	6%	3%	10%	27%
Group 3	8%	10%	11%	24%	53%	Group 3	6%	8%	10%	25%	49%
	33%	21%	15%	32%			31%	19%	14%	36%	

## Appendix Table 1B: Robustness Tests of Comparative Statics of Dana and Gale & Holmes

### Testing Robustness to Definition of Expected Load Factor

In these tables, Expected Load Factor is defined as the mean realized load factor averaged by carrier-origin-destination-timeslot-week of year. The timeslots are: weekdays 1-5am, weekdays 6-9am, weekdays 10am-1pm, weekdays 2pm-7pm, weekdays 8pm-midnight, Saturdays, and Sunday.

<b>Flights - Expected to be Low Load Factor &amp; Realized Low Load Factor</b>
--

<b>Flights - Expected to be High Load Factor &amp; Realized High Load Factor</b>
--

#### 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	9%	1%	1%	1%	12%	Group 1	8%	1%	1%	0%	10%
Group 2	11%	7%	2%	3%	23%	Group 2	10%	7%	3%	3%	23%
Group 3	19%	14%	12%	20%	65%	Group 3	17%	17%	14%	20%	67%
	39%	22%	15%	24%			35%	25%	17%	23%	

#### 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	38%	10%	3%	2%	54%
Group 2	7%	6%	1%	3%	18%	Group 2	5%	6%	3%	7%	20%
Group 3	5%	6%	5%	10%	27%	Group 3	2%	5%	5%	14%	26%
	57%	20%	8%	14%			46%	21%	10%	23%	

#### 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%	0%	1%	12%	Group 1	10%	2%	0%	1%	13%
Group 2	8%	4%	3%	6%	20%	Group 2	8%	3%	2%	4%	17%
Group 3	15%	14%	15%	25%	69%	Group 3	13%	14%	17%	26%	70%
	31%	19%	18%	32%			31%	19%	20%	30%	

**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**

		<b>Expected Load Factors</b>			
		Full	Medium-Full	Medium-Empty	Empty
<b>Realized Load Factors</b>	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 expected load factor and then actual load factor. 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

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.