

Price Discrimination by Day-of-Week of Purchase:

Evidence from the U.S. Airline Industry

Steven L. Puller¹ and Lisa M. Taylor²

December 2011

Abstract

This paper identifies a source of price discrimination utilized by airlines -- price discrimination based on the day-of-the-week that a ticket is *purchased*. Using unique transaction data, we compare tickets that are similar in every observable aspect except the day-of-week *purchased*, that is, traveling on the same day of week on the same route on the same airline with the same ticket restrictions on flights with the same load factors and purchased the same number of days in advance. We find that fares are 5% lower when purchased on the weekend. We conjecture that this is a form of second-degree price discrimination. If airlines believe that weekend purchasers are more likely to be price-elastic leisure travelers, then they may offer lower prices on weekends when the mix of purchasing customers makes demand more price elastic. This conjecture is supported by the finding that the weekend purchase effect is distinctly larger on routes with a mixture of both business and leisure customers than on routes primarily traveled by leisure customers. These results have implications for other industries that have the ability to change prices daily based upon the types of customers who purchase on a specific day.

JEL codes: L1, L8, L9

¹ Texas A&M University and NBER. Corresponding author. Email: puller@econmail.tamu.edu

² Texas A&M University. Email: lisamari.taylor@gmail.com. We thank Alex Brown, Richard Dunn, Adi Mayer, and Steve Wiggins for very helpful comments.

1. Introduction

It is well-known that airlines use a variety of mechanisms to price discriminate between customers with different willingness to pay for travel. The existing theoretical and empirical literature has investigated several of these mechanisms including advance purchase restrictions, non-refundability, minimum stay requirements, and Saturday night stay requirements. Advance purchase restrictions can be used to segment consumers by their value of time (Gale and Holmes (1993)) and may be sold disproportionately to customers with low valuation (Dana (1998)). Tickets with Saturday night stay restrictions and other travel and refundability restrictions have lower fares, suggesting that ticket restrictions are used to price discriminate (Stavins (2001); Puller, Sengupta, and Wiggins (2009)).

However, the literature has not studied whether airlines segment customers by the day-of-week of *purchase*. In principle, this could be a valuable segmenting device. Travelers who *purchase* on the weekend (but travel any day of the week) may have different price elasticities than those who *purchase* during the week. Moreover, it would be very feasible to implement “day-of-week-of-purchase” pricing because airlines have the ability to dynamically change prices daily using sophisticated computer reservation systems. Current revenue management systems used by airlines allow revenue management analysts to reassess pricing daily during the booking process.³

In this paper, we make a simple and straightforward contribution to the literature. We find that airlines charge lower fares for observably similar tickets based on the day-

³ See Belobaba (2009) for a description of revenue management systems used by major carriers.

of-week of purchase, and that this phenomenon is consistent with price discrimination. This finding is important in its own right because airlines are increasingly using complex pricing schemes, and revenue management systems are becoming progressively more sophisticated. This analysis provides insights into the mechanisms of airline pricing.

Our finding also has implications for a variety of other industries in which sophisticated pricing schemes can be applied. For example, the revenue management systems developed for airlines are being deployed in other hospitality industries including hotels, rental cars, cruise lines, and trains. And more generally, the study of price discrimination by time of purchase could have implications for e-commerce. The dynamic pricing of online retail markets could take advantage of changing prices based on the demand elasticities of consumers likely to be purchasing on any given day or specific times of the day. Although the general topic of intertemporal price discrimination has received considerable attention in the literature, this is the first paper to our knowledge to empirically investigate price discrimination based on day-of-week of *purchase* that is independent of the actual day of *consumption*.⁴

One obstacle to identifying whether airlines price differently on specific days of the week is obtaining sufficiently detailed data in order to address various selection issues. For example, travelers purchasing on the weekend could pay less because they choose tickets with more restrictions or fly on less popular flights; such selection behavior could lead one to incorrectly conclude that airlines set different fares on

⁴ Of course, differential pricing based on day of *consumption* is widely studied in both the theoretical and empirical literature; for example, see the literature on peak-load pricing.

weekends. One would need to control for a variety of ticket characteristics to accurately assess whether airlines price differently on weekends. The most common data used in existing airline pricing research – the U.S. Department of Transportation’s Airline Origin and Destination Survey (DB1B) – do not include purchase or departure date nor ticket restrictions or load factors; thus, it is not adequate to properly control for other factors that could affect pricing. Likewise, data on posted airfares gathered via web-scraping are not sufficient to address this issue. Although many observations pertaining to a single flight can be collected over time, the restrictions and load factors associated with such fares are generally unobservable (or at least hard to obtain). Moreover, web-scraping exercises that gather the lowest available fare for particular dates of travel may inappropriately compare different flight times and different sets of ticket restrictions.

We use a unique new dataset of ticket transactions to overcome many of these obstacles. Our data include individual ticket restrictions and information on the load factors of the itinerary’s flights. After controlling for a large set of ticket restrictions and load factors, we find that tickets purchased on weekends are sold at fares that are 5% lower than fares purchased on weekdays. We interpret this finding as differential pricing of weekend purchases.

We show that this empirical regularity is consistent with second-degree price discrimination. Routes with a mix of business and leisure customers are likely to have a different composition of passengers purchasing on weekends versus weekdays, creating incentives for airlines to lower fares on weekends when the demand is more price elastic. We find that the weekend purchase effect is larger on such “mixed” routes – it is 7% on

“mixed” routes versus 2% on other routes. We argue that this is highly suggestive that airlines implement price discrimination by the day-of-week of purchase.

Our results differ from the only other paper that investigates a day of purchase effect in airlines. Mantin and Koo (2010) analyze a collection of fares from Farecast.com and find that, for a given route, average price is not affected by purchase day of the week but that price dispersion is higher Friday through Sunday. The differences between our findings and those of Mantin and Koo most likely arise from fundamental differences in the data. Mantin and Koo use posted online fares while we use transacted fares. Our data make it possible to account for factors not controlled for by Mantin and Koo, including ticket restrictions and flight-level load factors.

2. Methodology

Our empirical strategy is to test whether tickets transacted on weekends have different mean fares than tickets transacted on weekdays, after controlling for a large set of factors that could reflect weekend purchasers choosing tickets and flights that are observably different. It is critical to control for a host of possible selection issues so that any fare differences can be attributed to supply side behavior.

It is straightforward to see why selection issues could complicate inference. As we show with data below, the unconditional mean fare for weekday purchases (Monday through Friday) is \$365, compared with \$290 for tickets purchased on the weekend (Saturday and Sunday). Although the mean fare is lower for weekend purchasers, this could reflect the types of tickets purchased by weekend purchasers rather than airlines

pricing tickets differently on weekends. Customers purchasing on weekends may be more likely to purchase tickets for off-peak travel times or tickets for more restricted travel. Theoretical work has shown that airlines may have incentives to charge higher prices on flights with higher demand (Gale and Holmes (1992, 1993)) or on flights with more demand uncertainty (Dana (1999)). Likewise, weekend purchasers may buy tickets further in advance than weekday purchasers. The number of days in advance of departure that the ticket is purchased has been shown to have a significant impact on price, with ticket prices increasing as departure nears. This price increase is most dramatic in the last 7-14 days before departure.⁵ These factors could confound a weekend purchase effect if tickets purchased on weekends are cheaper only because they are for travel during periods of lower or more certain demand, or if they are purchased further in advance than tickets bought on Monday through Friday.

To address this concern, we include a rich set of variables to control for these confounding factors. First, we include controls for the number of days that the itinerary was purchased before departure, so that any inference is conditional on tickets purchased the same number of days in advance.

Second, we control for various measures of demand for the specific flights involved in an individual itinerary. We include three different metrics of the itinerary flights' "load factor" (i.e. the fraction of the flight's seats that are occupied), as we describe in more detail in the data section below. We include measures of the ex ante

⁵ See Stavins (2001); Pels and Rietveld (2004); Puller, Sengupta and Wiggins (2009); and Mantin and Koo (2010).

expected load factor for the specific flight, the ex post realized load factor, and the load factor as of the date a ticket is purchased. To control for any residual variation in demand or demand uncertainty, we control for a variety of factors for each flight segment of the itinerary: the week of the year, the day of the week of travel, and the time of day of travel.

Third, we control for ticket restrictions that may be used to segment customers using other forms of price discrimination. Airlines can use ticket restrictions to discriminate between business customers with low price elasticities and leisure customers with higher price elasticities. With most restrictions, consumers face a trade-off between price and flexibility of travel plans. As a result, one would expect airlines to offer lower fares on more restricted tickets that target customers who are more price elastic. We include indicators of whether the ticket includes some form of advance purchase requirement, travel restriction, or length of stay restriction. Also, we include whether the itinerary included a Saturday night stayover and whether the ticket was refundable. Finally, we include an indicator of whether the ticket was “full fare coach” because such tickets may allow changes in reservations or be eligible for upgrade and additional frequent flyer benefits.

The resulting baseline model can be written as:

$$(1) \text{Log}(Fare)_i = \beta_0 + \beta_1 \text{WeekendPurchase}_i + \beta_2 \text{AdvancePurchaseDays}_i + \beta_3 \text{LoadFactors}_i + \beta_4 \text{Timing}_i + \beta_5 \text{Restrictions}_i + \beta_6 \text{Carrier-RouteFixedEffects}_i + \varepsilon_i$$

where subscript i indicates an individual itinerary. The variables included in $\text{AdvancePurchaseDays}_i$, LoadFactors_i , Timing_i , and Restrictions_i are described with the results of the model estimation. WeekendPurchase_i is an indicator of whether the ticket was purchased on a Saturday or Sunday; β_1 is interpreted as the percentage by which a ticket purchased on the weekend is priced lower than an observably similar ticket purchased on a weekday.

3. Data

Our data are individual ticket transactions for travel on large domestic routes in the fourth quarter of 2004 on the major legacy carriers in the U.S. As we describe below, we have detailed information on itinerary-level ticket restrictions, flight characteristics, and fares. These data are the same as those used in Puller, Sengupta, and Wiggins (2009, hereafter PSW), so we will provide a brief description in this paper and refer the interested reader to PSW for further details. These data are much more detailed than information contained in “DB1B” – one of the standard datasets used in airline pricing research. Our data include all ticket transactions made through one of the major computer reservations systems for travel occurring in the fourth quarter of 2004. This computer reservation system (CRS) handles approximately one-third of total transactions for domestic airline tickets. For each itinerary, we have information on the date of

purchase and fare; in addition, for each flight segment of an itinerary, we have information on the date of travel, the carrier, origin and destination, flight number, and class of service.

First, we merge the census of transactions to several other data sources so that we can measure flight-level load factor on the observed itineraries. In order to measure the load factor for each flight segment of an itinerary (e.g. American Airlines flight 301 from New York La Guardia to Chicago O'Hare on October 11, 2004), we merge observed itinerary counts to information from the Official Airline Guide on flight times and aircraft capacity. These data allow us to calculate the realized load factor (total tickets / total seats) for each flight segment of an itinerary.⁶

Next, we merge each itinerary to another dataset with information on ticket restrictions using an archive of fares available from a travel agent's computer reservation system. This archive contains a list of offered 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 requirements, refundability, travel restrictions, and minimum and maximum stay restrictions. We merged these data to our transaction data. Further details are available in PSW (2009). Although we only could match restrictions for a subset of our total transactions, PSW (2009) illustrate that the matching process does not introduce substantial selection concerns.

⁶ We account for the fact that our CRS does not sell all tickets by adjusting observed sales by the route-carrier specific CRS share, which is calculated with data from the Bureau of Transportation Statistics T-100 data. See PSW for details.

Table 1: Routes in Our Sample 2004Q4

Route by Origin and Destination City (Airport Code).

Atlanta (ATL) – Boston (BOS)	Detroit (DTW) – New York-La Guardia (LGA)
Atlanta (ATL) – Cincinnati (CVG)	Detroit (DTW) – Los Angeles Intl (LAX)
Atlanta (ATL) – Dallas-Fort Worth (DFW)	Detroit (DTW) – Orlando (MCO)
Atlanta (ATL) – Newark (EWR)	Detroit (DTW) – Minneapolis-St Paul (MSP)
Atlanta (ATL) – Fort Lauderdale (FLL)	Detroit (DTW) – Phoenix (PHX)
Atlanta (ATL) – Las Vegas (LAS)	Fort Lauderdale (FLL) – Philadelphia (PHL)
Atlanta (ATL) – Los Angeles Intl (LAX)	Hartford (BDL) – Fort Lauderdale (FLL)
Atlanta (ATL) – New York-La Guardia (LGA)	Houston (IAH) – Las Vegas (LAS)
Atlanta (ATL) – Orlando (MCO)	Houston (IAH) – Los Angeles Intl (LAX)
Atlanta (ATL) – Miami (MIA)	Houston (IAH) – Orlando (MCO)
Atlanta (ATL) – Philadelphia (PHL)	Houston (IAH) – New Orleans (MSY)
Atlanta (ATL) – Tampa (TPA)	Houston (IAH) – Chicago-O’Hare (ORD)
Boston (BOS) – Charlotte (CLT)	Las Vegas (LAS) – Minneapolis-St Paul (MSP)
Boston (BOS) – Washington-Reagan (DCA)	Las Vegas (LAS) – Chicago-O’Hare (ORD)
Boston (BOS) – Detroit (DTW)	Las Vegas (LAS) – Philadelphia (PHL)
Boston (BOS) – Newark (EWR)	Los Angeles Intl (LAX) – Minneapolis-St Paul (MSP)
Boston (BOS) – New York-La Guardia (LGA)	Los Angeles Intl (LAX) – Chicago-O’Hare (ORD)
Boston (BOS) – Philadelphia (PHL)	Los Angeles Intl (LAX) – San Francisco (SFO)
Charlotte (CLT) – New York-La Guardia (LGA)	Los Angeles Intl (LAX) – Tampa (TPA)
Charlotte (CLT) – Orlando (MCO)	Milwaukee (MKE) – Minneapolis-St Paul (MSP)
Charlotte (CLT) – Philadelphia (PHL)	Minneapolis-St Paul (MSP) – Phoenix (PHX)
Chicago-O’Hare (ORD) – Philadelphia (PHL)	Minneapolis-St Paul (MSP) – Seattle (SEA)
Chicago-O’Hare (ORD) – Seattle (SEA)	Minneapolis-St Paul (MSP) – San Francisco (SFO)
Chicago-O’Hare (ORD) – San Francisco (SFO)	Newark (EWR) – Fort Lauderdale (FLL)
Chicago-O’Hare (ORD) – St. Louis (STL)	Newark (EWR) – Houston (IAH)
Cincinnati (CVG) – New York-La Guardia (LGA)	Newark (EWR) – Las Vegas (LAS)
Dallas-Fort Worth (DFW) – Houston (IAH)	Newark (EWR) – Los Angeles Intl (LAX)
Dallas-Fort Worth (DFW) – Las Vegas (LAS)	Newark (EWR) – Orlando (MCO)
Dallas-Fort Worth (DFW) – Los Angeles Intl (LAX)	Newark (EWR) – San Francisco (SFO)
Dallas-Fort Worth (DFW) – Orlando (MCO)	New York-JFK (JFK) – Los Angeles Intl (LAX)
Dallas-Fort Worth (DFW) – Chicago-O’Hare (ORD)	New York-JFK (JFK) – Miami (MIA)
Dallas-Fort Worth (DFW) – Phoenix (PHX)	New York-La Guardia (LGA) – Miami (MIA)
Dallas-Fort Worth (DFW) – Orange County (SNA)	New York-La Guardia (LGA) – Chicago-O’Hare (ORD)
Dallas-Fort Worth (DFW) – St. Louis (STL)	Orlando (MCO) – Minneapolis-St Paul (MSP)
Denver (DEN) – Dallas-Fort Worth (DFW)	Orlando (MCO) – Philadelphia (PHL)
Denver (DEN) – Houston (IAH)	Philadelphia (PHL) – Pittsburg (PIT)
Denver (DEN) – Las Vegas (LAS)	Philadelphia (PHL) – Raleigh-Durham (RDU)
Denver (DEN) – Los Angeles Intl (LAX)	Philadelphia (PHL) – Tampa (TPA)
Denver (DEN) – Minneapolis-St Paul (MSP)	San Diego (SAN) – San Francisco (SFO)
Denver (DEN) – Oakland (OAK)	Washington-Dulles (IAD) – Chicago-O’Hare (ORD)
Denver (DEN) – Ontario (ONT)	Washington-Dulles (IAD) – San Francisco (SFO)
Denver (DEN) – Chicago-O’Hare (ORD)	Washington-Reagan (DCA) – New York-La Guardia (LGA)
Detroit (DTW) – Las Vegas (LAS)	

We restrict our analysis to the six major legacy carriers in 2004 -- American, Delta, United, Continental, USAir, and Northwest. These are all the carriers that served at least 5% of domestic travelers with the exception of Southwest which is excluded due to data limitations. We focus on a set of 85 domestic routes that represent a stratified sample of the largest routes for the six carriers with varied market structures. A list of the routes is included in Table 1. We exclude tickets sold for the first class cabin. We study only round-trip itineraries that include nonstop service from the origin to destination and back.⁷ In order to avoid unusual travel periods, we exclude tickets for flights during Thanksgiving weekend (Wednesday through Monday), Christmas and New Year's (after December 22).

We test whether the weekend purchase effect is consistent with price discrimination using cross-sectional variation in route characteristics. In particular, we expect such price discrimination to be more feasible on routes with a mix of leisure and business travelers, for reasons that we describe below. Therefore, we use two different measures to classify routes as either "leisure" or "mixed". The first measure is a tourism index similar to that utilized by Borenstein and Rose (1994) and Gerardi and Shapiro (2009). The tourism index is equal to the ratio of 2004 accommodations income to total personal income for the Metropolitan Area of the destination airport (from the Bureau of Economic Analysis). Those routes with a tourism index above the 80th percentile of our routes are classified as "leisure" routes (this amounts to routes with destinations of Las

⁷ Results are unchanged qualitatively (and very similar quantitatively) when including one-way ticket purchases with one-way fares doubled to obtain analogous round-trip fares.

Vegas, Orlando, New Orleans, Miami, and Fort Lauderdale); the remaining routes are classified “mixed”.

The second measure is a business travel index of domestic airline travel to and from cities that is for “business purposes.” This index is constructed by Borenstein (2010) based on the 1995 American Travel Survey.⁸ We use this index to define a route as “leisure” if the route is in the *bottom* 20th percentile of our routes on the business travel index. Because the index distinguishes between origin and destination city, we define routes by direction for purposes of creating this categorization of itineraries. The destinations that are classified as “leisure” in our sample are Orlando, Las Vegas, Tampa, San Diego, Fort Lauderdale, Miami, Phoenix, Seattle, and Denver. The remaining routes are classified as “mixed”.

Table 2 lists summary statistics. The vast majority of tickets are sold on weekdays. The most popular ticket restrictions are non-refundability, stay restrictions, and travel restrictions. In addition, many tickets come with advance purchase requirements.

⁸ The index is described in Borenstein (2010) and the data are available on the NBER website.

Table 2: Sample Means

Variable	Sample Mean
Round-trip fare	\$361.72
Weekend purchase	0.05
American Airlines	0.28
Delta	0.15
United	0.15
Continental	0.20
Northwest	0.10
USAir	0.12
Number of days in advance of departure purchased	18.54
Sunday departure	0.11
Monday departure	0.22
Tuesday departure	0.18
Wednesday departure	0.17
Thursday departure	0.14
Friday departure	0.12
Saturday departure	0.05
Sunday return	0.13
Monday return	0.11
Tuesday return	0.14
Wednesday return	0.16
Thursday return	0.19
Friday return	0.21
Saturday return	0.07
1 day advance purchase requirement	0.01
3 day advance purchase requirement	0.08
5 day advance purchase requirement	0.0005
7 day advance purchase requirement	0.21
10 day advance purchase requirement	0.03
14 day advance purchase requirement	0.31
21 day advance purchase requirement	0.02
30 day advance purchase requirement	0.004
Refundable	0.18
Travel restriction	0.45
Stay restriction	0.31
Saturday stay included	0.29
Full fare coach (Y class)	0.05

Note: Summary statistics for round-trip two segment itineraries to travel in 2004Q4 on American, Delta, United, Northwest, Continental and USAir on the routes in our sample. All variables except fare and number of days in advance purchased are dummy variables.

4. Results

A. Evidence of a Weekend-Purchase Pricing Effect

We estimate our baseline model (equation 1) to test whether observably similar tickets that are purchased on the weekend are priced differently from tickets purchased on a weekday. As we discuss in section 2, we must control for a variety of potentially confounding factors to ensure that customers who purchase on weekends are not merely purchasing tickets with different restrictions or tickets for flights with lower load factors.

In order to illustrate these selection issues, we first estimate models that do not control for the confounding effects. Table 3 presents coefficient estimates from models in which we progressively control for more confounding factors. These estimates suggest the sign of the selection bias and demonstrate the need for a rich set of detailed data on itineraries, as we have in our data.

In column (1), we regress log fare only on a set of fixed effects for carrier-routes and an indicator for whether the ticket was purchased on a weekend. This model uses within carrier-route variation in fares and finds that tickets purchased on weekends are 12% cheaper than tickets purchased during the week if one does not condition on any ticket or flight characteristics. As we show below, this difference is smaller after one accounts for the fact that passengers buying on the weekend may purchase more restricted tickets or fly on emptier flights.

Table 3: Estimation Results

Dependent variable: log(fare)						
	(1)	(2)	(3)	(4)	(5)	(6)
	No Controls	Add controls for timing	Add controls for Ticket Characteristics	Add controls for Load Factor	Measure 1 of "Mixed"	Measure 2 of "Mixed"
Weekend purchase	-0.12*	-0.08*	-0.05*	-0.05*		
	(0.01)	(0.01)	(0.00)	(0.00)		
Weekend purchase					-0.02*	-0.02*
					(0.01)	(0.01)
Weekend purchase * "Mixed" route					-0.05*	-0.05*
					(0.01)	(0.01)
Constant	Yes	Yes	Yes	Yes	Yes	Yes
Carrier-route fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Day-in-advance fixed effects, day- timeslot fixed effects, week-of-year fixed effects	No	Yes	Yes	Yes	Yes	Yes
Dummy variables for advance purchase requirements, full coach fare status, refundability, travel restriction, stay restriction, Saturday night stay	No	No	Yes	Yes	Yes	Yes
Realized load factor at departure, load factor as of day of purchase, expected load factor	No	No	No	Yes	Yes	Yes
Observations	145425	145425	145425	145425	145425	145425
Adjusted R-squared	0.3903	0.5756	0.7373	0.7407	0.7408	0.7408

* significant at the 1% level; standard errors are given in parentheses.

Estimated using ordinary least squares with robust standard errors, clustered by departure date.

In column (2), we address many of the selection issues by controlling for the timing of travel and the number of days before departure that the ticket is purchased. In particular, we account for the fact that *travel* (as opposed to purchase) on certain days of the week and during certain times of the day may lead to higher fares for a variety of reasons, e.g. peak-load pricing. We divide every day into five timeslots and interact the day of week of travel with this timeslot fixed effect to generate 35 day-timeslot fixed effects.⁹ For example, this will allow for fares to differ when traveling on a Monday morning versus a Saturday afternoon. We allow for such travel time fixed effects for both legs of the roundtrip itinerary. In addition, we include fixed effects for the number of days in advance that the ticket was purchased, which for our data are 192 day-in-advance fixed effects. Finally, in order to allow for overall fares to vary over the calendar quarter of our sample (e.g. due to changes in fuel costs), we include fixed effects for each of the 12 weeks in our sample.¹⁰

⁹ The timeslots are 1-5am, 6-9am, 10am-1pm, 2-7pm, and 8pm-midnight.

¹⁰ It is important to recognize that for any given flight, the purchase day of the week cannot vary randomly with purchase days in advance. That is, for the same flight, it is not possible to observe both a ticket purchased 23 days in advance on a Tuesday and another ticket purchased 23 days in advance on a Friday. Thus one might be concerned that the observed weekend purchase effect is merely a mechanical anomaly resulting from two other strong empirical regularities in airline pricing – the weekly cycle in peak/off-peak travel and the “booking curve” effect that fares are higher as departure nears. We perform the following simulation exercise to determine if a weekend purchase effect could arise from these two other pricing patterns. First, we calculate the joint probability distribution of days in advance purchased and the day of week of *travel*, and we use this to generate 10,000 simulated purchases. Second, we use our data to estimate the percentage discount of buying each day in advance relative to buying 0 days in advance. Third, we use our data to calculate average fares for *travel* for each day of the week for travelers buying 0 days in advance. We use each of these three calculations to simulate 10,000 fares and purchase patterns. Then we estimate our benchmark model testing for a weekend pricing effect on these simulated data. We estimate coefficients of zero on the weekend purchase variable. Therefore, we do not believe that other mechanical patterns could generate the weekend purchase effect that we estimate.

Results in column (2) indicate that tickets purchased on weekends sell at fares 8% lower than tickets purchased on weekdays, after controlling for the timing of travel, the number of days in advance that the ticket is purchased, and the week of year of travel.¹¹ Of course, even with these controls, it is possible that tickets purchased on weekends carry different restrictions or involve flights with different load factors. We address these additional selection issues next.

In column (3), we control for various ticket characteristics that could impact consumer utility and be used as segmenting devices for other forms of price discrimination. First, we control for whether the itinerary includes advance purchase restrictions – we include dummy variables for each of the types of restrictions we observe: 1-day, 3-days, 5-days, 7-days, 10-days, 14-days, 21-days, and 30-days.¹² Second, we control for whether the fare was a “full fare coach” (Y class) fare that often includes certain forms of ticket flexibility to change the reservation. Third, we control for whether the fare basis code indicates that the ticket is refundable. Fourth, we control for whether the ticket includes restrictions on the days of week of travel or on the minimum and/or maximum days of stay. Finally, we control for whether the ticket involved a stay over a Saturday night. The Saturday night stay restriction had been one of the more powerful devices for segmenting business from leisure travelers, according to industry experts. We do not observe the presence of the restriction, but we can

¹¹ Although Table 3 does not report all coefficients in the interest of space, we find that fares increase for purchases made closer to departure, especially in the last two weeks before departure.

¹² Note that in these specifications, we continue to control for the number of days in advance that the ticket was purchased; these advance purchase restrictions capture whether an advance purchase was a requirement on the ticket.

measure whether the travel involved a Saturday night stay, and we include this metric as our proxy for Saturday night stay restrictions. We believe that these sets of observable ticket characteristics capture the vast majority of differences between tickets.

Column (3) finds that after controlling for ticket characteristics in addition to the controls from column (2), tickets purchased on weekends have fares that are 5% lower than tickets purchased on weekdays.¹³ It is not surprising that the *WeekendPurchase* coefficient estimate falls after controlling for ticket characteristics. Consumers purchasing on weekends are more likely to purchase more restrictive tickets. To see this, separately we calculate the fraction of tickets sold that include some form of restriction by whether the ticket was purchased on a weekend. Compared to weekday purchases, weekend purchases are more likely to include an advanced purchase restriction (by 72% to 66%), less likely to be full fare coach (2% vs. 5%), less likely to be refundable (10% to 18%), more likely to include a travel and/or stay restriction (64% vs. 55%), and more likely to include a Saturday night stay (55% vs. 27%).¹⁴ However, the results in column (3) suggest that even after correcting for a different composition of tickets being sold on the weekend, purchases on Saturday and Sunday occur at lower fares.

¹³ Table 3 does not report coefficients of the ticket characteristics in the interest of space. We find that *ceteris paribus*, full fare coach tickets sell at a 69% higher fare, refundable tickets sell at a 20% higher fare, travel and stay restrictions are associated with lower fares of 24% and 5%, respectively, and a stay over a Saturday night is associated with a 14% lower fare. All of these differences are statistically different from zero.

¹⁴ These patterns indicate why it is important to control for various ticket characteristics and why studying this issue using “lowest posted fare” data without such characteristics could lead to incorrect inference.

In the specification in column (4), we include various metrics of the load factor of the specific flights on the itinerary. These load factor metrics capture route and flight specific peak-load pricing that is not captured with the day-timeslot fixed effects. Although it is not a topic of this paper, these measures of load factor could capture effects of peak-load pricing and the effect of demand uncertainty (Dana, 1999). Our strategy here is to capture as many demand-side driven factors of pricing, so that we can compare the pricing for tickets that are observably the same. To that end, we include multiple metrics of load factor simultaneously. First, we include the realized load factor (at departure) averaged across both flight segments of the itinerary. Second, we include the load factor as of the day the ticket was purchased. For example, if a ticket was purchased 7 days before departure and the flight was 50% full as of 7 days before departure, we define the load factor at the date of ticket purchase to be 50%. Our measure of the load factor at purchase for an itinerary is the average of this metric across both flight segments of the itinerary. Third, we include a measure of expected load factor – the load factor that can be systematically predicted by the airlines. We create seven timeslots that have systematically different average load factors, and compute expected load factor as the sample average realized load factor for the carrier-route-timeslot-week-of-year.¹⁵ Details of this calculation can be found in PSW (2009). Because these three measures of load factor are highly correlated, we make no attempt to

¹⁵ The seven timeslots are: weekdays 1-5am, weekdays 6-9am, weekdays 10am-1pm, weekdays 2-7pm, weekdays 8pm-midnight, Saturdays, and Sundays.

disentangle or interpret the coefficients of each. Rather we include all measures simultaneously to control for systematic differences in demand that impact pricing.

Column (4) reports our preferred specification with controls for all of the factors described above – timing of actual travel, days in advance the ticket was purchased, ticket characteristics and restrictions, and load factor. Controlling for each of these factors, fares on observably similar tickets are 5% lower when purchased on the weekend versus a weekday. We interpret this empirical finding to be differential pricing on the weekends of otherwise similar tickets. Because we have controlled for a very rich set of ticket characteristics and determinants of demand, we believe it unlikely that this empirical regularity is driven by selection effects of weekend purchasers buying tickets that are unobservably more restricted or for travel on unobservably lower demand flights. In order for this finding to be driven by selection, tickets must differ in characteristics *not* captured by how many days in advance they are purchased, actual advance purchase restrictions, refundability, travel restrictions, length of stay restrictions, full fare coach status, and whether a Saturday night stay is involved.¹⁶ Alternatively, tickets must differ in demand characteristics of the flight segments that are *not* captured by 35 different weekly timeslots and the actual, expected, and purchase date load factors. Although such unobservables cannot be completely ruled out, we believe these findings to be strongly suggestive of differential pricing of weekend purchases.

¹⁶ In unreported regressions, we also control for the first digit of the fare basis code and we continue to find a weekend purchase effect.

B. Evidence that the Weekend-Purchase Pricing Effect is Consistent with Price Discrimination

Next, we present evidence that this robust empirical regularity is consistent with price discrimination in which the day-of-week of purchase is used as a fencing device. We conjecture that airlines recognize that customers purchasing on weekends have a higher price elasticity of demand and use yield management systems to lower fares on weekends.

We test this conjecture by comparing the size of the weekend purchase effect on routes in which the overall price elasticity is more likely to vary by the day-of-week of purchase. In order to motivate this test, consider the following stylized description of two different routes.

Suppose there are two types of travelers: business and leisure travelers. Assume that business travelers are less price elastic. Airlines may use a variety of fencing devices to segment customers (e.g. refundability or travel restrictions). However, these devices may not perfectly separate business from leisure travelers, creating room for additional price discrimination by day-of-week of purchase.

Consider a hypothetical route in which half of the customers are business travelers and the other half are leisure travelers. Suppose that business travelers make all purchases on weekdays. However, leisure travelers purchase more uniformly across the days of the week. There is evidence of this supposition in our data – itineraries that involve a stay over a Saturday night (a proxy for leisure travelers) exhibit purchase patterns that are less concentrated on weekdays (relative to weekends) than itineraries

that do not involve a Saturday night stay. As a result of the “mixed” nature of this route, the customer composition on weekends consists of a relatively higher fraction of leisure travelers than the composition on weekdays. Thus the overall price elasticity is higher on weekends. Airlines have an incentive to lower fares of (otherwise identical) tickets purchased on weekends. (Of course, business travelers would have an incentive to respond by shifting purchases to the weekend, but institutional constraints of the workweek may prevent this).

Consider another hypothetical route in which all travelers are leisure travelers. Because there are no business travelers who buy solely on weekdays, the price elasticity of customers purchasing on weekends is very similar to the price elasticity of those purchasing during the week. On this route, airlines have less incentive to change the fares of (otherwise identical) itineraries on the weekend.

Clearly, this stylized description does not fully characterize all of the real-world heterogeneities in consumer types and purchase patterns. However, it does suggest that price discrimination by day-of-week of purchase would lead to a larger “weekend purchase effect” on routes with a mix of business and leisure travelers than on routes with a dominant share of leisure travelers. We test this conjecture by allowing the size of the weekend purchase effect to vary by route type.

In the last two columns of Table 3, we interact the *WeekendPurchase* dummy variable with a dummy variable for whether the route is “mixed”. As we describe in section 3, we use two measures of whether a route is “leisure” or “mixed”. Results from using the tourism index are presented in column (5). The weekend purchase effect is 5%

larger on mixed routes – it is 2% on leisure routes while it is 7% on mixed routes.

Column (6) reports results using the alternative classification of leisure and mixed routes based upon the “business travel index”. We obtain identical results under this alternative definition of mixed and leisure.

Using both metrics, we find a statistically larger weekend purchase effect on mixed routes. Thus, the relative magnitudes of these estimated weekend purchase effects are consistent with the stylized model of pricing lower during the weekends when the customer mix is more price elastic.

Our finding of lower fares on weekends is notable for reasons beyond implications about price discrimination. Industry experts have conjectured that airlines *increase* fares on weekends because fewer tickets are sold on weekends, and that airlines use the low purchase weekend period to coordinate on fare increases.¹⁷ These types of analyses – based on web-scraping lowest fares – do not account for differences in ticket characteristics or the travel times of the web-scraped fares. Our finding suggests that, to the extent that such weekly pricing dynamics exist, that the price discrimination effect dominates. Alternatively, airlines may simply not have engaged in the weekly pricing dynamics during our sample period of 2004Q4. Nevertheless, the net effect of day-of-week of purchase price discrimination and any weekly pricing dynamics is that fares purchased on weekends are lower.

¹⁷ For a description of this possibility, see Borenstein (2004) or popular press such as McCartney (2011).

5. Conclusions

This paper investigates a robust empirical regularity that was not previously identified – the existence of a weekend purchase effect on airline ticket prices. Although we cannot definitively conclude that the effect reflects price discrimination, the cross-sectional variation in the size of the effect is certainly suggestive of a price discrimination mechanism.

These findings have implications beyond pricing in the airline industry. If the customer composition varies by day-of-week, such a pricing strategy could be utilized in other hospitality industries that use yield management software. Future research could test for day-of-week pricing in other hospitality industries such as hotels, car rentals, trains, and cruises. More generally, future research could test for similar pricing behavior in e-commerce in which prices can be adjusted by day of week (or even time of day) when customer composition is different.

References

- Belobaba, Peter P., 2009. Fundamentals of Pricing and Revenue Management. In:
Belobaba, P., Odoni, A., Barnhart, C. (Eds.). The Global Airline Industry. Wiley,
73-112.
- Borenstein, S., 2010. An Index of Inter-City Business Travel for Use in Domestic
Airline Competition Analysis. UC Berkeley, mimeo. Index is available at:
<http://www.nber.org/data/bti.html>
- Borenstein, S., 2004. Rapid Price Communication and Coordination: The Airline Tariff
Publishing Case (1994). In: Kwoka, J.E., White, L.J. (Eds.). The Antitrust
Revolution, 4th edition.
- Borenstein, S., Rose, N.L., 1994. Competition and price dispersion in the U.S. airline
industry. *Journal of Political Economy*, 102 (4), 653-683.
- Dana, J.D., Jr., 1998. Advance-purchase discounts and price discrimination in
competitive markets. *Journal of Political Economy*, 106 (2), 395-422.
- Dana, J.D., Jr., 1999. Equilibrium price dispersion under demand uncertainty: the roles
of costly capacity and market structure. *RAND Journal of Economics*, 30 (4),
632-660.
- Gale, I., Holmes, T., 1992. The efficiency of advance-purchase discounts in the presence
of aggregate demand uncertainty. *International Journal of Industrial
Organization*, 10 (3), 413-437.

- Gale, I., Holmes, T., 1993. Advance-purchase discounts and monopoly allocation of capacity. *American Economic Review*, 83 (1), 135-146.
- Gerardi, K., Shapiro, A.H., 2009. Does competition reduce price dispersion? New evidence from the airline industry. *Journal of Political Economy*, 117 (1), 1-37.
- McCartney, S., 2011. "Whatever You Do, Don't Buy an Airline Ticket On...", *Wall Street Journal* (WSJ.com).
- Mantin, B., Koo, B., 2010. Weekend effect in airfare pricing. *Journal of Air Transport Management*, 16 (1), 48-50.
- Pels, E., Rietveld, P., 2004. Airline pricing behaviour in the London-Paris market. *Journal of Air Transport Management*, 10 (4), 277-281.
- Puller, S., Sengupta, A., Wiggins, S., 2009. Testing theories of scarcity pricing in the airline industry. Texas A&M University, mimeo.
- Stavins, J., 2001. Price discrimination in the airline market: the effect of market concentration. *Review of Economics and Statistics*, 83 (1), 200-202.