University of Texas Rio Grande Valley

ScholarWorks @ UTRGV

Economics and Finance Faculty Publications
and Presentations

Robert C. Vackar College of Business &
Entrepreneurship

12-14-2018

An Analysis of Dynamic Price Discrimination in Airlines

Diego Escobari

The University of Texas Rio Grande Valley

Nicholas G. Rupp

Joseph Meskey

Follow this and additional works at: https://scholarworks.utrgv.edu/ef_fac

Part of the [Finance Commons](#)

Recommended Citation

Escobari, Diego, Nicholas G. Rupp, and Joseph Meskey. "An Analysis of Dynamic Price Discrimination in Airlines." *Southern Economic Journal* 85, no. 3 (December 14, 2018): 639–62. <https://doi.org/10.1002/soej.12309>.

This Article is brought to you for free and open access by the Robert C. Vackar College of Business & Entrepreneurship at ScholarWorks @ UTRGV. It has been accepted for inclusion in Economics and Finance Faculty Publications and Presentations by an authorized administrator of ScholarWorks @ UTRGV. For more information, please contact justin.white@utrgv.edu, william.flores01@utrgv.edu.

An Analysis of Dynamic Price Discrimination in Airlines*

(*Southern Economic Journal*, Forthcoming)

Diego Escobari[†] Nicholas G. Rupp[‡] Joseph Meskey[§]

August 1, 2018

Abstract

Prices for the same flight change substantially depending on the time of purchase. This paper uses a unique dataset with round-the-clock posted fares to document significant within-day price variation. Labeling time-variation as discriminatory is difficult because the cost of an unsold airline seat changes with inventory, days before departure and aggregate demand expectations. After controlling for these factors and aggregating hourly fares to have a framework with two consumer types, we are able to identify a component that is largely consistent with dynamic price discrimination. We find higher prices during office hours (when business travelers are likely to buy) and lower prices in the evening (when leisure travelers are more likely to purchase). As the proportion of business travelers increases closer to departure, both price dispersion and price discrimination become larger. We provide an alternative explanation for the observed within-day price differentials which is related to Edgeworth price cycles.

Keywords: Pricing, Price discrimination, Price dispersion, Airlines

JEL Classifications: C23, L93

*We thank comments by Marco Alderighi, Jan Brueckner, Qihong Liu, Claudio Piga, Oz Shy, Jonathan Williams, and seminar participants at Texas A&M, Southern Economics Association Meetings, International Industrial Organization Conference, Industry Studies Association Conference, conference of the European Association for Research in Industrial Economics, and INFORMS Revenue Management and Pricing Conference.

[†]Department of Economics & Finance, The University of Texas Rio Grande Valley, Edinburg, TX 78539, Phone: (956) 665-2104, Fax: (956) 665-5020, Email: diego.escobari@utrgv.edu, URL: <http://faculty.utrgv.edu/diego.escobari>

[‡]Department of Economics, East Carolina University, Greenville, NC 27858, Phone: (252) 328-6821, Fax: (252) 328-6743, Email: ruppn@ecu.edu, URL: <http://myweb.ecu.edu/ruppn>

[§]Department of Economics, East Carolina University, Greenville, NC 27858, Phone: (252) 328-2948, Fax: (252) 328-6743, Email: meskeyj@ecu.edu

1 Introduction

The U.S. commercial aviation sector plays a large role in the U.S. economy, comprising 4.9 percent of U.S. GDP in 2010 (International Air Transportation Association). There were 812.9 million domestic and international enplanements in 2012 (Bureau of Transportation Statistics) with the fifteen largest U.S. carriers collecting \$159.5 billion in revenue in 2012. Yet, these U.S. carriers registered small profit margins of just 3.7 percent (Bureau of Transportation Statistics). Given that the composition of customers that are shopping for airline tickets changes during the day, the objective of this paper is to determine whether airlines are systematically charging different prices throughout the day to increase revenue. Dynamic price discrimination occurs when a company sells an identical good at different prices based on the time of purchase.

Dynamic price discrimination exists in a large number of industries (e.g., fashion apparel, car rentals, retail sales, hotels, e-commerce). It is well known that airlines offer substantially different fares for the same flight. Airlines use different ‘fences’ (e.g., refundability, cabin, Saturday-night-stayover, frequent-flier programs, or advance purchases) to implement multiple screening mechanisms and separate among unobserved consumer types into different groups. This has led to a surge in the literature studying airline’s price dispersion (see, e.g., Borenstein and Rose, 1994; Gerardi and Shapiro, 2009; Gaggero and Piga, 2011; and Dai et al., 2014). Labeling the observed price dispersion as price discrimination, however, is a challenging task because existing price differentials are also due to differences in quality and costs. For example, a refundable ticket is a higher quality product than a non-refundable ticket and the cost of an airline seat at different days to departure depends on seat availability and demand expectations.¹ Hence, the well-known time-variation in airline prices is not necessarily price discriminatory.

In this paper we exploit a unique panel of flights that keeps round-the-clock track of the dynamics of posted prices for two months prior to departure. We document statistically and economically significant price dispersion that occurs for the same flight and within the same day-to-departure. In addition, within this price dispersion and following a model

¹At a given day to departure and for a given inventory of seats the opportunity cost of a seat is higher if demand is expected to be larger (see, e.g., Escobari, 2012).

with two consumer types we identify a dynamic price discrimination component. On price dispersion within the same day the results show that fares peak at 3:00 a.m. when sellers typically revise prices upwards and they are the least expensive at 9:00 p.m. The estimated price difference between these two hours is on average 0.65% (\$2.44) during the two months prior to departure and is about 1.28% (\$5.34) for the last two weeks to departure.

Motivated with a two consumer types model, and after controlling for quality, costs and multiple screening mechanisms, we are able to identify a dynamic price discrimination component that arises when airlines set higher prices during office hours and lower prices later in the evening. At seven days to departure fares are 0.25% (\$1.04) more expensive during office hours.² We argue that airlines know that some consumers are more likely to buy at different times of the day and use this as a screening mechanism to separate consumers into types that are heterogeneous in their valuations. High valuation consumers (business travelers) are more likely to buy during office hours, while low valuation consumers (leisure travelers) are more likely to buy later in the evening. Business travelers who have high valuations and buy during office hours cannot shift to buy at lower prices at other times of the day as they usually buy tickets through their offices. Our results also support a higher dynamic price discrimination as the departure date nears and the proportion of business travelers increases. Moreover, the price discrimination component is more prevalent for low cost carriers when compared to legacy carriers. Because the airline pricing problem is dynamic in nature, we estimate a dynamic model that is consistent with rational expectations and with agents forming expectations about future prices. When interpreting our results with more than three consumer types, the observed within-day price dispersion is also consistent with varying intensity of competition.

There is extensive theoretical literature on discounts and price discrimination. Nevo and Hendel (2013) present an intertemporal price discrimination model for storable goods, while Chevalier and Kashyap (2011) examine price discrimination in a durable goods setting. Stokey (1979) looks at price discrimination for new products. Theoretical models that are closer to airline pricing include Dana (1998) who explains the existence of price

²While \$1.04 is a relatively small amount, given the 813 million U.S. passenger enplanements in 2012 (Bureau of Transportation Statistics), slightly higher fares during office hours can generate millions of dollars in higher airline revenue.

discrimination in the form of advance-purchase discounts when demand is uncertain, and Courty and Li (2000) who consider sequential price discrimination where a monopolist screens consumers using refund contracts.

On the empirical side, Shepard (1991) finds price gaps in gas stations that are consistent with price discrimination. Verboven (1996) looks at price discrimination as a plausible explanation behind price differences in automobiles across European countries and Leslie (2004) studies price discrimination for Broadway play tickets. More recently in the airline industry, Berry and Jia (2010) estimate the proportion of business and leisure travelers, while Escobari and Jindapon (2014) show theoretically and empirically how price discrimination through refundable tickets decline as individuals learn about their demand. Puller and Taylor (2012) under strong assumptions look at discrimination based on the day-of-the-week purchase and Siegert and Ulbricht (2014) explore how price dynamics vary with the competitive environment.

This paper also contributes to the airline price dispersion literature by shedding light on price dynamics for perishable goods under demand uncertainty. Most prior work on airline price dispersion has been empirical. Borenstein and Rose (1994), Gerardi and Shapiro (2009), and Dai et al. (2014) estimate the effect of competition, Bilotkach et al. (2010) study airlines' pricing strategies, Bilotkach and Rupp (2012) analyze the price-offer curves, and Alderighi et al. (2015) study the effects of booking time and load factors on fares. Recent theoretical studies on airline price dispersion include Czerny and Zhang (2015), who develop a framework to examine third-degree price discrimination by a monopolist at a congested airport, and Czerny and Zhang (2014), who consider airport congestion pricing when carriers can price discriminate between business and leisure passengers. On the dynamic of prices for perishable goods, Gallego and van Ryzin (1994) present a theoretical model that allows prices to adjust dynamically based on inventories and time to departure, while Deneckere and Peck (2012) develop a dynamic model of perfectly competitive price posting. Empirically, Sweeting (2012) uses Major League Baseball tickets to study the seller's pricing behavior in secondary markets, and Escobari (2012) estimates a dynamic pricing model for airline tickets and finds aggregate demand learning by airlines.

The organization of the paper is as follows. Section 2 explains the data, while the empirical strategy is presented in Section 3. The estimation results and a discussion of the

assumptions appear in Sections 4 and 5, respectively. Section 6 concludes.

2 Data

This paper uses an original data set on airline prices collected from a major U.S. Online Travel Agency (OTA) during a period when this OTA did not charge a booking fee for airline tickets. The collection strategy presents two key improvements over a similar data set used in Stavins (2001). First, we have a panel, which allows us to control for observed and unobserved heterogeneity across flights, carriers, and routes. Second, we track hourly changes in posted prices. This unique piece of information is essential to determine whether there exist within day price dispersion and whether airlines price discriminate based on the hour-of-the-day the ticket is purchased. Each cross-sectional observation in the sample is the lowest observed economy fare on a non-stop one-way flight from a carrier on a route, where a route is a pair of departure and arrival U.S. airports. Each flight is observed every hour starting at 59 days in advance until the departure date. We have 158 domestic flights in 158 different airport pairs. The carriers in the sample are American Airlines, Alaska, JetBlue, Delta, Frontier, AirTran Airways, United, US Airways, and Virgin America, with the proportion of flights of each carriers chosen to be close to its share in the U.S. market.³ Following Stavins (2001) and Escobari (2012) to control for any departure day effect on fares, all non-stop, one-way flights depart on a single day: Thursday July 12, 2012. We argue that the ex-ante high demand period of the 4th of July is distant enough from our chosen departure date. Hence, our fares are likely capture the typical demand during summer travel.⁴ Because our data involves airports located in different time zones, the time we record for each of the flights is the one associated with the originating airport.

The high frequency posted price data we utilize was specifically gathered to study the dynamics of prices within the same day to departure. The widely used and more aggregate DB1B transaction data from the Bureau of Transportation Statistics (e.g., see Borenstein and Rose, 1994; and Gerardi and Shapiro, 2009) just does not have this level of information.⁵

³The only major U.S. carrier excluded from the sample is Southwest, whose fares only appear on Southwest.com.

⁴The Memorial Day holiday, May 28, is also relatively far from our chosen departure date.

⁵A recent study with detailed transaction data at the ticket level is Hernandez and Wiggins (2014). They

In addition to much more detailed information, a less apparent advantage of using posted prices is the possibility of directly observing the behavior of the seller every hour —this is true whether a transaction occurs or not. With transaction data the econometrician is only able to observe the behavior of buyers and sellers once a sale occurs. Moreover, because we observe equilibrium prices in which the sellers take into account the optimal behavior of the buyers when posting their fares, we implicitly observe the behavior of buyers.

Borenstein and Rose (1994) explain how price dispersion in the airline industry comes from multiple sources. The data collection and the estimation approach controls for these multiple sources and allows us to isolate the price differences across purchase times within the same flight and within the same day to departure. We keep track of the price of the non-refundable economy-class ticket at every hour prior to departure. This controls for the existence of tickets sold with frequent flier miles, different fare classes, and refundable tickets (see, e.g., Escobari and Jindapon, 2014). The existence of price differences associated with Saturday-night stayover, minimum and maximum stay are controlled for by focusing on one-way flights.⁶ Moreover, selecting non-stop flights is helpful to control for connecting flights and more complicated itineraries (e.g., open jaws). Note that while we focus on one-way tickets, extending our results to round-trip tickets would need to follow the standard assumption in the literature where the price of a round-trip ticket is assumed to be two times the price of a one-way ticket (see, e.g., Borenstein and Rose, 1994, p. 677; and Gerardi and Shapiro, 2009, p. 5). Our data does not allow us to test this assumption. Generalizing our results to the U.S. airline industry would require more assumptions as the existence of different ticket characteristics make price discrimination in airlines a complex multidimensional screening problem (Armstrong, 1996). Our data on one-way fares greatly helps us to simplify the complexity of pricing strategies and is in line with the vast majority of the single-leg (i.e., one-way) theoretical models that help explain price dispersion (see, e.g., Gallego and van Ryzin, 1994; and Dana, 1998).

empirically examine the effect of competitive conditions on nonlinear pricing strategies.

⁶Under this data collection strategy if a particular flight is the back-end (return portion) of a round trip ticket, it will not enter in our sample. It only enters if it was bought as a separate one-way ticket.

3 Empirical Strategy

3.1 Pricing Equation

We estimate the following equation to capture the hour-of-day purchase effect on prices:

$$\text{LNFARE}_{ijt} = \beta \text{HOUR}_t + \sum_{k=0}^{59} \gamma_k \text{I}_{[\text{ADV}_t=k]} + v_{ij} + \varepsilon_{ijt}, \quad (1)$$

where i denotes the flight, j the airport pair, and t is time. LNFARE_{ijt} is the logarithm of the price and HOUR_t is the matrix with the hour-of-day dummies with the hour of the day corresponding to the time at the departing airport. Hence, β is the vector of coefficients of interest. $\text{I}_{[\text{ADV}_t=k]}$ is an indicator variable equal to one if the number of days prior to departure (ADV_t) is equal to k , zero otherwise. v_{ij} is the flight and airport-pair specific effect, while ε_{ijt} is the remaining disturbance.

Notice that despite its very simple structure, Equation 1 is actually very powerful. Because we have twenty-four observations within the same day prior to departure, we can estimate the model while controlling for the effect of days in advance with daily dummies. This takes into account any (linear or nonlinear) effect that ADV_t can have on fares —e.g., intertemporal price discrimination based on advance purchase requirements and any day-of-the-week purchase effect. Moreover, the fixed effect v_{ij} controls for time-invariant flight-, carrier-, and airport-pair-specific characteristics. This includes all of the controls included in price dispersion models that use either Bureau of Transportation Statistics data (e.g., Borenstein and Rose, 1994, and Gerardi and Shapiro, 2009) or posted prices data (e.g., Stavins, 2001, and Gaggero and Piga, 2011). Examples of flight-specific characteristics include aircraft size, departure time, hub effect at the flight level, and systematic peak-load pricing.⁷ Carrier-specific characteristics include managerial capacity, while airport-pair-specific characteristics include the Herfindal index, distance between airport pairs, and any hub effect that affects prices at the airport level.

3.2 Dynamic Pricing

One constraint with the estimation of Equation 1 via OLS or fixed effects is that it would assume that sellers and buyers are myopic. This means that sellers do not consider previous

⁷This one arises due to congestion known at the time the flight is scheduled (see, e.g., Escobari, 2009).

prices when posting the current price and do not have any beliefs about future prices either. Likewise, buyers that observe the current posted price must make a buying decision with no information about previous prices nor any beliefs about future prices. This is an important constraint because the airline pricing problem is dynamic in nature and agents can behave dynamically. Consumers can easily search posted prices through online travel agencies and observe part of the pricing sequence before deciding to buy. In fact, online travel metasearch engine Kayak (acquired by Priceline.com in 2012) now provides consumers with both a price trend and a buy recommendation. Beliefs are important because when consumers observe the current price they may decide to postpone their purchase decision if they believe that prices are likely to be lower in the future.

A model consistent with agents behaving dynamically, as explained in Arellano and Bond (1991) (see also Bun and Kiviet, 2006), involves the estimation of the following dynamic pricing equation:

$$\text{LNFARE}_{ijt} = \alpha \text{LNFARE}_{ij,t-1} + \beta \text{HOUR}_t + \sum_{k=0}^{59} \gamma_k \text{I}_{[\text{ADV}=k]} + v_{ij} + \varepsilon_{ijt}. \quad (2)$$

While the coefficient α is not of direct interest, the variable $\text{LNFARE}_{ij,t-1}$ is needed to incorporate dynamics. Equation 2 indicates that current prices depend on previous prices (lagged prices), but given the nature of the autoregressive term current prices depend on all previous prices and previous shocks. In the estimation of Equation 2 we treat $\text{LNFARE}_{ij,t-1}$ as endogenous. Hence $\text{LNFARE}_{ij,t-1}$ is allowed to be correlated with ε_{ijt} and earlier shocks, but not with $\varepsilon_{ij,t+1}$ and future shocks. Of course, both the hour-of-the-day and the day-in-advance dummies are treated as strictly exogenous.

While we do not model agents' beliefs and we cannot conclude that agents are forward looking, endogenous $\text{LNFARE}_{ij,t-1}$ along with a no serial correlation assumption of ε_{ijt} implies that our estimation is consistent with agents (buyers and sellers) that can behave dynamically. Sellers use Equation 2 to set prices, while buyers can use previous and current posted prices (and Equation 2 under rational expectations) to form beliefs about future prices. Heterogeneity across consumers on arrival rates and expectation formations determines whether consumers purchase the airline ticket and when they buy. Serially uncorrelated ε_{ijt} does not prevent buyers from forming expectations about future prices, rather it suggests that buyers cannot predict future price shocks based on previous infor-

mation.

To obtain consistent estimates of (α, β, γ) , equation 2 is estimated using the methods described in Arellano and Bond (1991). We take the first differences of Equation 2 to eliminate the fixed effect term v_{ij} . Then we use the resulting error term $\Delta\varepsilon_{ijt}$ to construct the moments $E(\Delta\varepsilon_{ijt}\mathbf{M})$, where we need a matrix of instruments $\mathbf{M}$ to obtain $(\hat{\alpha}, \hat{\beta}, \hat{\gamma})$ via GMM. Because HOUR_t and $\sum_{k=0}^{59} \gamma_k \mathbf{I}_{[\text{ADV}=k]}$ are treated as strictly exogenous, they serve as their own instruments. The endogeneity assumption of $\text{LNFARE}_{ij,t-1}$ along with lack of serial correlation of ε_{ijt} implies that lagged values of LNFARE_{ijt} are valid instruments for $\Delta \text{LNFARE}_{ij,t-1}$.

3.3 Dynamic Price Discrimination

Two necessary conditions must be satisfied for price discrimination to exist: (i) different valuations across consumer types, and (ii) the seller’s ability to prevent arbitrage. Airlines can restrict arbitrage by not allowing low valuation consumers to resell tickets to high valuation consumers. To completely avoid arbitrage, however, the airline also needs to prevent high valuation types from buying tickets marketed for low valuation consumers. Price discrimination models based on mechanism design call this the incentive compatibility constraint. In a simple two-type model, the pricing scheme should be such that high valuation consumers are forced to reveal their type and can only buy higher priced tickets. Airlines price discriminate in multiple dimensions and aim to separate consumer types in different categories by using ‘fences’ such as refundability of the ticket, Saturday-night stayover, or fare classes.⁸ In this paper we control for these multiple fences and focus on a single price discrimination dimension, the hour-of-day purchase of homogeneous tickets for the same flight within the same number of days before departure.

We follow McAfee and te Velde (2007) and refer to dynamic price discrimination as the practice of charging different consumer types distinct markups over marginal cost based on the time of purchase. The price dispersion within the same flight and within the same number of days before departure are captured by the vector β in Equations 1 and 2, with the 1:00 a.m. dummy serving as the omitted category.

⁸Multidimensional screening problems are difficult to characterize theoretically because of the continuum of boundary conditions (see, e.g., Armstrong, 1996).

To motivate the empirical approach and findings consider the existence of two consumer types. On the one hand, we have business travelers —or high type— who have high reservation values for a ticket and usually buy tickets through their offices representatives or travel agents during normal office hours. On the other hand, there are leisure travelers —or low type— who have lower valuations and lower costs for searching and buying tickets during the evening. Of course, leisure travelers are not restricted to buy tickets only in the evening. They could purchase airline tickets before going to work, however, mobile web browsing traffic immediately before office hours (7-8 a.m.) is only half the level compared to evening hours (9 pm).⁹ A more likely scenario is that leisure travelers search for airfares during their lunch break while at work. Given that consumers are visiting an average of 38 web sites before booking a trip,¹⁰ it is unlikely that a leisure traveler will complete their travel booking during a lunch break. Instead, we believe that leisure passengers will continue their search and complete their travel booking in the evening. The mechanism that sellers use to prevent arbitrage is simple, offer higher priced tickets during office hours and lower priced tickets during the evening. Discrimination based on hour-of-the-day purchase is considered third-degree price discrimination, since the seller knows the distribution of consumers' valuations.

4 Estimation Results

4.1 Pricing Equation

Table 1 presents the summary statistics. The initial results from the estimation of Equation 1 appear in Table 2. To gain a better understanding of the problem, the different columns of Table 2 include different sets of fixed effects and controls for days to departure and day-of-week dummies. Because office hours are only relevant Monday through Friday, we begin the analysis by restricting the sample to include only weekdays. We later include weekend data in additional specifications to provide some robustness checks on the main findings. The first three specifications do not include the days-to-departure fixed effects,

⁹See chitika.com/browsing-activity-by-hour (accessed 9 August 2017).

¹⁰See www.travelmarketreport.com/articles/Consumers-Visit-38-Sites-Before-Booking-Expedia-Says (accessed 9 August 2017).

hence we are able to identify the day-of-the-week and the day-in-advance effects on pricing. Note that because these specifications do not control for day-in-advance effects, it would be difficult to generalize the observed price differences to flights that do not depart on Thursdays.

[Table 1, about here]

Different prices over different days of the week can be interpreted as price discriminatory only under strong assumptions. First, there needs to be a mechanism to prevent high-valuation consumers from imitating the behavior of lower valuation consumers. That is, a mechanism that prevents people from easily switching their purchase days. Second, capacity costs across different days must be the comparable. This may be a concern because capacity constraints along with uncertain demand means that the opportunity cost of a seat changes depending on the expected demand (see Escobari, 2012). Expectations are updated as the departure date nears, hence inducing cost variation across different days.

[Table 2, about here]

Our main focus on Table 2 are the hour-of-day coefficients. The estimates in the last column show that when compared to 1:00 a.m. —the omitted category— fares are 1.0% more expensive at 3:00 a.m., with higher fares being statistically significant until 10:00 a.m.¹¹ The spike at 3:00 a.m. is consistent with airlines updating their fares upward after midnight when few travelers search for tickets (information we received from people in the industry). Between 11:00 a.m. and 7:00 p.m., fares are not significantly different to fares at 1:00 a.m. Finally, we find the lowest fares between 8:00 p.m. (0.3% lower) and midnight (0.1% lower). There are three plausible scenarios that could potential explain what we are observing in these hourly price changes. One scenario is that carriers follow Edgeworth pricing cycles (see, Maskin and Tirole, 1988) by adjusting prices upward in the early morning hours (at 3 a.m.) in hopes that these price increases will stick.¹² Our findings suggest that such increases do not consistently “stick” since prices typically continue to decline throughout the day reaching a bottom in the evening. A second plausible scenario which is

¹¹The interpretation takes into account that the dependent variable is $\text{LNFARE}_{ijt} \times 1000$.

¹²For more information on price leadership in the airline industry see Morrison and Winston (2010).

consistent with the observed pricing patterns is that carriers are charging higher prices during the day when price inelastic consumers are shopping for fares (business travelers) and lower fares in the evening when more price sensitive consumers are fare shopping (leisure travelers). A third scenario also exists since empirically what we find may not be consistent with a simple two type price discrimination model (business and leisure types). This third scenario introduces the possibility of a third consumer type: “early risers”. These would be highly motivated airline ticket buyers who seek to make ticket purchases prior to going to work. The hourly pricing data shows the highest prices occur at 3 a.m. hence the “early risers” pay a ticket premium compared to business and leisure travelers.

The easiest way to interpret these coefficients is as within-day and within-flight price dispersion. With constant costs and if consumers do not wait to purchase later within the same day, the observed price dispersion in Table 2 can be interpreted as price discrimination based on time-of-day purchase.¹³ This is in line with Basso et al. (2009), where agency problems’ incentives make the demand of business travelers more inelastic.¹⁴ The difference in fares between office hours and evening hours is at its largest of 1.0% when comparing prices at 9:00 a.m. versus 10:00 p.m. At the average one-way ticket price of \$375.30, this difference represents \$3.74. From the point of view of many buyers, these differences might not be economically significant. From the perspective of the sellers, however, a nearly \$4 higher ticket price has large revenue implications given the size of the market —with 813 million U.S. passenger enplanements in 2012 (Bureau of Transportation Statistics).

Note that the error in Equation 1 is unlikely to be independent across observations. Having day-in-advance fixed effects and flight fixed effects will absorb away common shocks (e.g., common cost shocks), and this is helpful if the potential within day-in-advance and within flight correlation of errors is driven by a common shock process. To further allow for correlation between errors within clusters Table 2 reports White heteroscedasticity-consistent estimates of the asymptotic standard errors, clustered at the carrier level. Cameron and Miller (2015) explain that there is no formal test to assess the clustering level. The

¹³We also control for capacity costs within the same day to departure to support the existence of a particular dynamic price discrimination component.

¹⁴In their model Frequent Flier Programs “bribe” employees to book flights at higher prices. This occurs because an employee selecting the airline will not necessarily have the right incentives to find the lowest possible price when an employer is paying the bill.

consensus described in Cameron and Miller (2015) is to be conservative by aggregating clusters when possible and that a reasonable approach is to cluster at progressively higher levels and stop when there is little change in the standard errors. Our clustering level follows this argument as we progressively clustered at the flight, route and carrier levels.

4.2 Days to Departure Dynamics and Consumers' Heterogeneity

Days to departure plays an important role in our analysis because of the heterogeneity across consumers that buy at different points prior to departure, and because consumers can behave dynamically by forming expectations and delaying their purchase decisions. To address how consumers' heterogeneity affects our results, Table 3 presents the dynamic panel estimates of Equation 2. We use different subsamples of the data across columns, broken down into four groups depending of the number of weeks prior to departure. The cutoffs $ADV = 14.8$ and 29.6 allow us to focus on the first quarter of the data ($ADV < 14.8$) and also evenly splits the sample in half (before and after $ADV = 29.6$). The differences in the estimates across columns show strong evidence that price dispersion or price discrimination (under stronger assumptions) through hour-of-day is larger the closer to departure. For example, during the last two weeks (when $0 \leq ADV < 14.8$) there are significant price differences between 6:00 a.m. and 6:00 p.m. of 0.51% (or \$2.11) in the average priced ticket (purchased within 14 days of departure), while this difference is smaller for ticket purchases made further in advance (column 3).

[Table 3, about here]

The explanation behind these results is simple and intuitive. The proportion of leisure travelers making advance purchases is larger earlier in the sales season, hence the seller has little incentive to try to separate across types and rather opts to pool traveler types together by charging nearly the same price throughout the day. As the departure date nears, however, the proportion of business travelers becomes larger, and hence it becomes advantageous for an airline to follow a separating strategy. During the last two weeks fares are more expensive early in the morning and during office hours, while they are lower in the evening (between 9:00 p.m. and 11:00 p.m.) With constant costs we interpret this

difference as price discriminatory because of the difficulty that business travelers have in purchasing tickets in the evening (outside of typical office hours).

Because the estimates in Table 3 treat $\text{LNFARE}_{ij,t-1}$ as endogenous, we need a vector M of instruments for the moments $E(\Delta\varepsilon_{ijt}M)$ that come from first differencing Equation 2. We use $\text{LNFARE}_{ij,t-2}$ and $\text{LNFARE}_{ij,t-3}$ as instruments for $\Delta\text{LNFARE}_{ij,t-2}$.¹⁵ We use two tests to address the validity of the dynamic panel specifications. Across all columns we observe large p-values associated with the serial correlation test. This provides strong evidence for a valid specification —the error term ε_{ijt} is not serially correlated. Moreover, the Sargan test for over-identifying restrictions to test the null that the instruments are not correlated with the residuals also provides strong evidence for a valid specification across all columns. For the standard errors we compared the one-step GMM robust standard errors clustered at the flight, route and carrier levels (see, Roodman, 2009) with the two-step GMM robust standard errors that include the Windmeijer finite-sample correction. Without the correction, Windmeijer (2005) explains that the standard errors tend to be severely downward biased. Following Cameron and Miller (2015), in a similar conservative approach than in the static estimation we use the Windmeijer finite-sample correction because it provides larger standard errors.

The estimates for the hour-of-day dummies in the first column (within two weeks prior to departure) show that fares are the least expensive at 9:00 p.m. followed by 6:00 p.m. Fares are the most expensive in the early morning hours, reaching a peak at 3:00 a.m. The range of prices reaches its maximum when comparing fares between 9:00 p.m. and 3:00 a.m. where the difference is about 2.30% or \$9.56 —as evaluated at the mean FARE of \$415.8 within two weeks to departure. After the peak, fares remain high during office hours, yet continue to decline during the day, with the lowest fares occurring in the evening (after 6 p.m.). As the departure date nears and the proportion of business travelers increase, we observe higher airline prices during normal office hours. Moreover, lower fares are found later in the evening when leisure travelers typically search for tickets. The estimates across different columns reveal that the pattern of price discrimination is not the same at different days prior to departure. For example, well in advance of departure (column 3), when the

¹⁵ As a robustness check we also use a different selection of instruments in the matrix M and find similar results.

proportion of business travelers is small, we find the lowest hourly ticket prices occur at 11 a.m. This result, however, is not robust when comparing models 2 and 3, and it is difficult to interpret due to the large number of dummy variables.

[Table 4, about here]

Table 4 consolidates the twenty-three hour dummies into three categories in an effort to capture the office-hours effect. The idea is to group the potentially continuum of different consumers types and focus on the difference between office and evening prices. OFFICE is a dummy variable which equals one between 9:00 a.m. and 4:59 p.m., zero otherwise. The variable EVENING equals one between 5:00 p.m. and 11:59 p.m., zero otherwise. The omitted category is MORNING (from 12:00 a.m. to 8:59 a.m.) which we will label “early risers”. All nine sets of estimates across columns pass both of the specification tests; there is no serial correlation and Sargan tests provide strong evidence that we have valid instruments.

The second half of Table 4 provides direct estimates of the price discrimination component by comparing the marginal effects of OFFICE and EVENING on fares. Column 1 focuses on the last month prior to departure and finds that fares during office hours are 0.093% (\$0.36) higher. When focusing between 29.6 and 59.3 days to departure column 4 shows no evidence of a price discrimination effect. Overall when using the entire sample of fares, the effect is small —fares are 0.037% (\$0.14) more expensive (column 7)— mainly because early in the selling season there is no effect. To further analyze the role of days to departure we include the interaction terms OFFICE $\times$ ADV and EVENING $\times$ ADV in columns 2, 5, and 8 of Table 4.¹⁶ The results show that price discrimination is greater closer to departure when the proportion of business travelers is greater. The estimates in column 2 indicate that at one week to departure (ADV= 7) fares during office hours are 0.25% (\$1.04) more expensive.¹⁷ At earlier dates the effect is smaller.

¹⁶The omitted category is MORNING $\times$ ADV. Since we control for the day in advance effect by using day dummies we allow for a very flexible nonlinear relationship between ADV and LNFARE. As a result, the coefficients on the interaction terms do not capture the marginal effect relative to the omitted category. Hence we include both OFFICE and EVENING interaction terms with ADV. The difference (evaluated at some value of ADV) captures the price discrimination effect.

¹⁷ $2.502/1000 \times 100 = 0.25\%$. $\$415.8 \times 0.25\% = \1.04 .

Note that the negative and statistically significant coefficients on OFFICE and EVENING (e.g., in columns 1, 4, and 7), mean that early morning prices are higher. There is nothing that prevents leisure travelers (that perhaps hold a regular nine-to-five job) from buying tickets early in the morning or at anytime during the day. However, given the pricing structure, leisure travelers reveal their type and buy at lower prices during the evening (when comparing to higher prices they face early in the morning or during office hours). This flexibility on when tourist travelers can buy is easy to understand in a simple price discrimination framework with two consumer types (Varian, 1989). In a two-type model we can characterize the problem with two participation constraints and two incentive-compatibility constraints.¹⁸ The seller adopts a pricing strategy such that one participation constraint and one incentive-compatibility constraint are binding. In our setting the pricing strategy makes the tourists’ incentive compatibility constraint to be binding —tourists reveal their unobserved (to the seller) type by buying at lower prices during the evening. Of course, tourists do not want to pretend to be business travelers and buy at higher prices during the office hours. The key element to be able to interpret our results as price discrimination is that business travelers cannot pretend to be tourists and benefit from lower prices. This is achieved due to the restriction that business travelers can only buy during office hours.

To determine which intraday price fluctuation scenario better fits our data, we also examine weekend fares separately. If airlines are trying to pass along price increases by hiking fares early in the morning only to see them fall throughout the day when other carriers fail to match the increase, then we should also expect to see similar price dynamics on the weekend. However, if the airlines are charging different prices based on consumer types (business or leisure traveler) then on days when a single type of traveler is searching for airline tickets (i.e., leisure travelers on the weekend), then the hourly posted prices should not change on the weekend. In sum, our argument of business and leisure travelers buying at different times of the day does not work for weekends, hence our approach is to focus on the weekday portion of the data. We will also examine if “early risers” continue to pay a premium on the weekends. Prior research has found larger price dispersion for posted prices of airfares occurs on weekends (Mantin and Koo, 2010). Therefore we use weekends

¹⁸See, e.g., page 613 on Varian (1989), page 3 on Escobari and Jindapon (2014), or Maskin and Riley (1984) for a general framework.

as an interesting exogenous factor that affects pricing dynamics and our weekend results can help us validate our findings. Based on our description of leisure and business travelers, only leisure travelers buy on weekends with no restriction on the hour of the day they can buy. Columns 3, 6, and 9 in Table 4 present the results for one-month in advance, two-months in advance and the whole sample respectively. When compared to the weekdays, the weekend estimates of OFFICE and EVENING are smaller in magnitude, suggesting that the differences between prices during MORNING the rest of the day is much smaller during the weekends. Moreover, the difference between OFFICE and EVENING is now negative and statistically significant for the last month to departure (column 3) and for the whole sample (column 9). This suggests that weekend shoppers of airline ticket prices can find slightly lower prices one month prior to departure during office hours -0.064% (\$0.25) compared to evenings. For the entire sample, once again compared to evenings lower weekend prices can also found during offices hours -0.028% (\$0.11). There is no difference two months prior to departure (column 6). We find no evidence of a “early riser” premium since the estimated coefficients for both OFFICE and EVENING are either insignificant or positive in every weekend estimation from Table 4. These weekend findings are consistent with the price discrimination scenario since we document an office hour premium occurs solely during weekdays.¹⁹

The price differentials in Table 4 are interpreted as price discrimination based on a framework with two consumer types. A two-type model is helpful to motivate the incentive compatibility constraint that prevents high valuation consumers (i.e., business travelers) from buying tickets marketed to low valuation types (i.e., tourists). However, the estimates in Tables 2 and 3 suggest that airlines continuously change prices throughout the day (to a large extent prices fall continuously within the same day). A generalization of a two-type price discrimination model to three or more types might help explain this price pattern. For example, the observed price differentials are consistent with the claim that airlines price

¹⁹Escobari (2012) and Alderighi et al. (2015) use load factor (LOAD) to measure the proportion of the seats already sold. With LOAD, a simplified version of the model would be: $\text{LNFARE} = \gamma \text{OFFICE} + \theta \text{LOAD} + u$. EVENING is the omitted category and price discrimination is directly captured by γ . If LOAD is omitted, δ is the slope when regressing LOAD on OFFICE, and $\hat{\gamma}$ is the OLS estimator, then $E(\hat{\gamma}) = \gamma + \theta \times \delta$. Escobari (2012) and Alderighi et al. (2015) find that $\theta > 0$. Moreover, $\delta < 0$ as airlines sell tickets throughout the day. Hence, $\theta \times \delta < 0$ means that γ would underestimate the price discrimination effect.

discriminate based on many traveler types with price sensitiveness that change throughout the day.²⁰ However, it is difficult to explain how airlines can effectively separate between three or more consumer types and prevent high valuation types from buying at low prices. For example, we do not know of any mechanism that prevents consumers from avoiding the high prices early in the morning and buying at lower prices during other times of the day. These higher prices are not consistent with a price discrimination model.

Our specifications in columns 3 and 4 of Table 3, as well as those in Table 4 include flight fixed effects. This means that we are controlling for observed and unobserved time invariant characteristics, which includes the Herfindahl-Hirschman Index (HHI) when constructed based on market shares. While in the next section we further explore the role of market structure, a possible explanation of the observed price differentials in Tables 2 and 3 (e.g., falling prices through the day and not significantly different prices between late morning and early evening) is that the intensity of competition can be changing throughout the day. This intensity of competition means that while holding the same number of competitors per route, it is likely that they try to keep prices high early in the morning and then they compete more aggressively by reducing prices later in the day.

4.3 Carrier’s Identity and Market Structure

To further explore the factors that can affect the observed price discrimination we run three additional specifications. In the first specification we separate legacy carriers from low-cost carriers (LCC). The second specification assesses the role of market structure, while the third explores the identity of the carriers.

[Table 5, about here]

Columns 1 and 2 of Table 5 report the estimates for the legacy carries (American, Alaska, Delta, United, and US Airways), while columns 3 and 4 represent LCCs (JetBlue, Frontier, Airtran, and Virgin). The distinction in this table is important as low-cost carriers have a lower operating cost structure and typically offer lower fares than legacy carriers.

²⁰Airlines have many years of information on prices and sales patterns for different flights and departure dates. Based on the estimates on Tables 2 and 3 that show within-day and within-flight price variation, it is reasonable to argue that airlines have good estimates of the different expected consumer price sensitiveness for purchases at different hours of the day.

Consistent with the previous estimates, price discrimination based on office-hours is positive, statistically significant, and becomes more prevalent the closer to departure. We find that price discrimination is more pronounced for LCCs than legacy carriers. Specifically, when interpreted as a percentage change the magnitude of the price discrimination point estimate at $ADV = 7$ during the month prior to departure for the LCC is about four times larger (\$6.05) than the estimate for the legacy carriers (\$1.58). From columns 1 and 3 we see that at $ADV = 7$ the legacy fare is about 0.16% (\$0.65) more expensive during office hours compared to the larger 0.6% (\$1.36) for the LCC.²¹

[Table 6, about here]

To assess the role of competition on price discrimination the specifications reported on Table 6 include the HHI. This research question is of interest given the findings by Gerardi and Shapiro (2009) that price dispersion decreases with competition, in line with the textbook treatment of price discrimination. This was further analyzed in Dai et al. (2014), who find a nonlinear relationship between price dispersion and market concentration. In our specifications HHI captures market concentration and it is constructed at the route level with market share based on the number of nonstop flights with the same departure date.

The office-hour price discrimination estimates in Table 6 are presented for three different values of the HHI: 0.52 (first quartile), 0.78 (median), and 1 (third quartile). These relatively high values for the different quartiles of HHI are evidence of a relatively high level of concentration. The price discrimination estimates during the last month to departure (column 1) show a statistically significant effect for the second and third quartiles of the HHI.²² The same is true for the third quartile of the HHI using the whole sample (column 3). While the price discrimination point estimates are greater for higher values of the HHI, the difference is not statistically significant.

²¹The dollar estimates use the average fares during the last month to departure for LCCs (\$225.65) and for legacy carriers (\$411.31).

²²Following Stavins (2001) we assume HHI is exogenous. Note that unlike Stavins (2001) we have a panel and we further control for flight fixed effects. In our specifications these fixed effects span the instruments for HHI proposed in Borenstein and Rose (1994) and used in Gerardi and Shapiro (2009).

Table 7 presents the dynamic panel estimates when running separate regressions by airline. These specifications are informative to assess whether our office-hours price discrimination findings are common to all carriers or if they are generated by the activity of only a few companies. The results show statistically significant price discrimination for American, Frontier, AirTran, and United, with noticeably larger magnitudes of price discrimination for the LCC’s Frontier (\$7.05) and AirTran (\$8.89). Perhaps, Frontier and AirTran, which have a substantial component of leisure travelers, offer lower prices in the evening to attract bookings from price sensitive leisure travelers. Overall, consistent with the findings in Bilotkach et al. (2010), we interpret these results as evidence of heterogeneity across carriers when implementing pricing strategies within the same day-to-departure. This observed heterogeneity can have important implications, for example, on the reported “hub premiums”. Lee and Luengo-Prado (2005) find that although average prices to and from hub airports tend to be higher, much of the difference can be explained by the proportion of leisure versus business travelers.

[Table 7, about here]

These estimates should be interpreted with care as the sample size for some carriers might not be large enough to allow us to obtain precise estimates. In particular, from the summary statistics in Table 1 we can see that the regressions for Frontier and Virgin are based only on a single flight each. Moreover, each specification in Table 7 has a fairly large number of days to departure dummies that are allowed to differ across columns, and there might be some underlying flight specific characteristics that change across carriers (e.g., the sample of routes differs for each carrier).

5 Discussion

In this section we further discuss four elements behind our price discrimination interpretation of the results. First, our price discrimination interpretation holds as long as business travelers buy tickets during normal business hours and leisure travelers buy tickets later in the evening. As we discuss throughout the document, it is reasonable to argue that business travelers buy during office hours at higher prices since they buy their tickets via

their workplaces. On the other hand, tourists have more flexibility on the time of the day at which they can buy; hence, is it easy to see why they would buy in the evening taking advantage of the lower prices offered at that time of the day. Notice that none of these two groups of travelers have incentives to buy early in the morning as those prices are higher. If we introduce a third consumer type termed “early riser” who purchases airline tickets on weekdays prior to going to work, then the “early riser” pays the highest ticket prices (see Table 4).

Second, airlines use historical data and have a good idea of the different price sensitivity of consumers that buy at different times of the day. This means that they know that they can charge higher prices during business hours and lower prices later in the evening. While we provide various robustness checks in the empirical analysis (e.g., costs, weekdays vs. weekends, days to departure, market concentration, legacy vs. low cost carriers), there might be other changing market conditions not controlled for in our empirical specification that we are not aware of. The ATPCO, a leader in the collection and distribution of airline fare data worldwide to online travel agents, indicates on their web site that fare distribution “occurs as often as once an hour, allowing airlines to respond to market conditions as they change”.²³

Third, we interpret the price differences as price discrimination. As argued throughout the paper, we believe this is the case because we control for costs and we focus on an homogeneous product —an economy class ticket. One potential alternative explanation that might drive price differences could be differences in quality (e.g., if business travelers were buying tickets of a higher quality than tourists). This might occur if the seller systematically changes the booking (sub)class within the same economy class at the same time we observe price changes between office hours and evening. However, we do not believe that is the case as the main differences in tickets within economy class are restrictions on the days to departure, refundability, high-season tickets, and long-haul. Our empirical approach is already controlling for days to departure, refundability, departure date and route characteristics.

Fourth, we follow the large and growing literature that uses posted prices and assume that the observed prices from the OTA are the prices at which a ticket can be purchased. Of

²³www.atpco.net/life-cycle-fare accessed 19 May 2017.

course, this also includes the (implicit) assumption common in the literature that airlines are the ones that decide prices, not the OTA. Even though the contractual arrangements between airlines and the OTA are confidential, it is reasonable to argue that airlines will not let the OTA have the power to change prices. During the time period of the data collection the major OTA we used did not charge a booking fee, which means that the price we recorded from the OTA was the same price found in the carriers' websites.²⁴

6 Conclusion

Capturing dynamic price discrimination for a perishable good that sells in advance under demand uncertainty is a challenging task because the marginal cost of a seat changes daily in advance of departure depending on demand expectations. Hence any variation in fares between days cannot be labeled as price discriminatory due to fluctuating seat inventories and demand expectations. In this paper we follow a novel approach to measure price discrimination, following the interpretation of a model with two consumer types, by using an original data set that tracks hourly posted airline prices for sixty days prior to departure. We use these data to estimate price differences based on time-of-day purchase within the same flight and within the same day to departure.

We find that fares increase by approximately \$6.50 between 2 a.m. and 3 a.m., with the highest average hourly posted fares occurring at 3 a.m. (EST) or midnight (PST). Average hourly posted fares steadily decline throughout the day reaching their lowest levels at 9 p.m. (EST). These findings suggest three plausible scenarios for these hourly price changes. One scenario is that the intensity of competition between carriers changes within the same day and consequently airline ticket prices follow an Edgeworth price cycle. Carriers initiate fare increases during the early morning hours, perhaps when competition for travelers is not too intense. However, a series of price cuts throughout the day suggests that the intensity of competition might be increasing as carriers try to attract potential travelers. Under this scenario price discrimination is unlikely to be a cause because with multiple fares throughout the day targeted at various consumer groups with different price sensitiveness,

²⁴For evidence that airlines and OTA have the same prices, see, e.g., S. Pascarella, *USA Today*, 27 Aug 2009 — “Face-off: Airline websites vs. online travel agencies.”

carriers cannot prevent high valuation consumers from buying at times when fares are lower. An alternative scenario is that these results are in line with the theoretical prediction of textbook price discrimination models where sellers separate between consumer types based on the time of day that they purchase the ticket. Within price discrimination we consider two possible scenarios of having two consumer types (business and leisure) or three consumer types (early riser, business, and leisure). On weekdays, we find that the highest prices occur prior to 8 a.m. for the early risers, while the lowest weekday prices occur after 5 p.m. for leisure travelers. On weekends, however, we find no evidence of price differences for early risers. This suggests that the most general and hence suitable scenario to explain our findings regardless of the day of the week is a simple two consumer types price discrimination model. We find higher prices during office hours (between 9 a.m. and 5 p.m.), while lower prices are found later in the evening (between 5pm and midnight). This result is consistent with airfares being higher for business travelers which make ticket purchases during working hours and leisure travelers with lower valuations who purchase tickets in the evening.

We also note that this price discrimination effect appears larger for tickets purchased closer to departure when the proportion of business travelers is higher. Specifically, we find that fares at seven days to departure during office hours are 0.25% (\$1.04) higher compared to evening posted fares. Hence for a consumer making a round-trip purchase, doing so during office hours (rather than in the evening) adds about \$2.08 to the ticket price.²⁵ Even if business travelers are forward looking and have beliefs of lower prices later in the evening, dynamic price discrimination on the part of the carriers is effective because business travelers will likely continue to buy airline tickets during office hours through their office representatives. While we have provided two plausible scenarios above to describe the observed hourly fluctuations in posted prices, there is one additional piece of evidence which favors the price discrimination scenario instead of airlines unsuccessfully attempting to raise prices early in the morning. Examining weekend posted prices data, when leisure travelers have no restrictions on when they can browse the Internet for airfares, we find no

²⁵This last calculation follows the standard assumption in the airline literature that the round-trip price is two times the one-way price. See for example Borenstein and Rose (1994, p. 677), and Gerardi and Shapiro (2009, p. 5).

difference between office hours and evening posted fares. Hence the office hours premium is solely a weekday phenomena and the price dispersion estimates are much larger within two weeks of departure, e.g., fares for the same flight are \$9.56 more expensive (\$19.12 for round-trip tickets) at 3:00 a.m. than at 9:00 p.m. of the same day. In sum, our findings are consistent with the claim that carriers are price discriminating as we find slightly higher posted prices during office hours and lower posted prices in the evenings during weekdays.

References

- Alderighi, M., M. Nicolini, and C.A. Piga, 2015. Combined effects of load factors and booking time on fares: Insights from the yield management of a low-cost airline. *The Review of Economics and Statistics*, 97, pp. 900-915.
- Arellano, M. and S. Bond, 1991. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies*, 58, pp. 277-297.
- Armstrong, M., 1996. Multiproduct nonlinear pricing. *Econometrica*, 64, pp. 51-75.
- Basso, L.J., M.T. Clements, and T.W. Ross, 2009. Moral hazard and customer loyalty programs. *American Economic Journal: Microeconomics*, 1, pp. 101-123.
- Berry, S. and P. Jia, 2010. Tracing the woes: An empirical analysis of the airline industry. *American Economic Journal: Microeconomics*, 2, pp. 1-43.
- Bilotkach, V., Y. Gorodnichenko, and O. Talavera, 2010. Are airlines price-setting strategies different? *Journal of Air Transport Management*, 16, pp. 1-6.
- Bilotkach, V. and N. Rupp, 2012. A guide to booking airline tickets online. In *Advances in Airline Economics: Pricing Behavior and Non-Price Characteristics in the Airline Industry*, Vol. 3, ed. James Peoples.
- Borenstein, S. and N.L. Rose, 1994. Competition and price dispersion in the U.S. airline industry. *Journal of Political Economy*, 102, pp. 653-683.
- Bun, M.J.G. and Kiviet, J.F., 2006. The effects of dynamic feedbacks on LS and MM estimator accuracy in panel data models. *Journal of Econometrics*, 132, pp. 409-444.
- Cameron A.C. and D.L. Miller, 2015. A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50, pp. 317-372.
- Chevalier, J.A. and A.K. Kashyap, 2011. Best prices. National Bureau of Economics working paper 16680. Cambridge, MA.

- Courty, P. and H. Li, 2000. Sequential screening. *Review of Economic Studies*, 67, pp. 697-717.
- Czerny, A.I. and A. Zhang, 2014. Airport congestion pricing when airlines price discriminate. *Transportation Research Part B*, 65, pp. 77-89.
- Czerny, A.I. and A. Zhang, 2015. Third-degree price discrimination in the presence of congestion externality. *Canadian Journal of Economics*, 48, pp. 1430-1455.
- Dai, M., Q. Liu, and K. Serfes, 2014. Is the effect of competition on price dispersion non-monotonic? Evidence from the U.S. airline industry. *The Review of Economics and Statistics*, 96, pp. 161-170.
- Dana, Jr. J., 1998. Advance-purchase discounts and price discrimination in competitive markets. *Journal of Political Economy*, 106, pp. 395-422.
- Deneckere, R. and J. Peck, 2012. Dynamic competition with random demand and costless search: A theory of price posting. *Econometrica*, 80, pp. 1185-1247.
- Escobari, D., 2009. Systematic peak-load pricing, congestion premia and demand diverting: Empirical evidence. *Economics Letters*, 103, pp. 59-61.
- Escobari, D., 2012. Dynamic pricing, advance sales, and aggregate demand learning in airlines. *Journal of Industrial Economics*, 60, pp. 697-724.
- Escobari, D., 2017. Airport, airline and departure time choice and substitution patterns: An empirical analysis. *Transportation Research Part A: Policy and Practice*, 103, pp. 198-210.
- Escobari, D. and P. Jindapon, 2014. Price discrimination through refund contracts in airlines. *International Journal of Industrial Organization*, 34, pp. 1-8.
- Gaggero, A.A. and C.A. Piga, 2011. Airline market power and intertemporal price dispersion. *Journal of Industrial Economics*, 59, pp. 552-577.
- Gallego, G. and G.J. van Ryzin, 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40, pp. 999-1020.

- Gerardi, K.S. and A.H. Shapiro, 2009. Does competition reduce price dispersion? New evidence from the airline industry. *Journal of Political Economy*, 117, pp. 1-37.
- Hernandez, M.A. and S.N. Wiggins, 2014. Nonlinear pricing strategies and competitive conditions in the airline industry. *Economic Inquiry*, 52, pp. 539-561.
- Lee, D. and M.J. Luengo-Prado, 2005. The impact of passenger mix on reported “hub premiums” in the U.S. airline industry. *Southern Economic Journal*, 72, pp. 372-394.
- Leslie, P., 2004. Price discrimination in Broadway theater. *Rand Journal of Economics*, 35, pp. 520-541.
- Mantin, B. and B. Koo, 2010. Weekend effect in airline pricing. *Journal of Air Transport Management*, 16, pp. 48-50.
- Maskin, E. and J. Riley, 1984. Monopoly with incomplete information. *Rand Journal of Economics*, 15, pp. 171-196.
- Maskin, E. and J. Tirole, 1988. A theory of dynamic oligopoly II: Price competition, kinked demand curves, and Edgeworth cycles. *Econometrica*, 56, pp. 571-599.
- McAfee, R.P. and V. te Velde, 2007. Dynamic pricing in the airline industry. In *Handbook of Economics and Information Systems*, vol. 1, edited by T.J. Hendershott. New York: Elsevier Sci.
- Morrison, S. and C. Winston, 2010. *The Evolution of the Airline Industry*, Brookings Institution Press. 1 Dec 2010.
- Nevo, A. and I. Hendel, 2013. Intertemporal price discrimination in storable goods markets. *American Economic Review*, 103, pp. 2722-2751.
- Puller, S.L. and L.M. Taylor, 2012. Price discrimination by day-of-week of purchase: Evidence from the U.S. airline industry. *Journal of Economic Behavior & Organization*, 84, pp. 801-812.
- Roodman, D., 2009. How to do xtabond2: An introduction to difference and system GMM in Stata. *Stata Journal*, 9, pp. 86-136.

- Shepard, A., 1991. Price discrimination and retail configuration. *Journal of Political Economy*, 99, pp. 30-53.
- Siegert, C. and R. Ulbricht, 2014. Dynamic oligopoly pricing: Evidence from the airline industry. Working Paper.
- Stavins, J., 2001. Price discrimination in the airline market: The effect of market concentration. *The Review of Economics and Statistics*, 83, pp. 200-202.
- Stokey, N.L., 1979. Intertemporal price discrimination. *The Quarterly Journal of Economics*, 93, pp. 355-371.
- Sweeting, A., 2012. Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets. *Journal of Political Economy*, 120, pp. 1133-1172.
- Varian, H.R., 1989. Price discrimination. R. Schmalensee, R. Willig, eds. *Handbook of Industrial Organization*, Vol. 1. North-Holland, Amsterdam, The Netherlands, 597654.
- Verboven, K., 1996. Intertemporal price discrimination in the European can market. *Rand Journal of Economics*, 27, pp. 240-268.
- Windmeijer, F. 2005. A finite sample correction for the variance of linear efficient two-step GMM estimators. *Journal of Econometrics*, 126, pp. 25-51.

Table 1: Summary Statistics

Variables	(1) mean	(2) sd	(3) min	(4) max
FARE _{ijt} :				
Whole sample	375.3	223.0	58	1,288
$0 \leq \text{ADV}_t < 14.8$	415.8	212.3	86.80	1,288
$14.8 \leq \text{ADV}_t < 29.6$	358.2	226.0	58	968
$29.6 \leq \text{ADV}_t < 44.4$	364.0	223.9	58	973
$44.4 \leq \text{ADV}_t < 59.3$	370.1	223.7	58	968
ADV _t	29.91	16.64	0	59.3
HHI _{ij}	0.780	0.259	0.2	1
Low Cost _{ij}	0.125	0.331	0	1
Days of week:				
MONDAY _t	0.184	0.387	0	1
TUESDAY _t	0.209	0.407	0	1
WEDNESDAY _t	0.206	0.405	0	1
THURSDAY _t	0.205	0.404	0	1
FRIDAY _t	0.196	0.397	0	1
Carriers:				
American _{ij}	0.299	0.458	0	1
Alaska _{ij}	0.018	0.134	0	1
JetBlue _{ij}	0.066	0.249	0	1
Delta _{ij}	0.117	0.322	0	1
Frontier _{ij}	0.006	0.078	0	1
AirTran _{ij}	0.046	0.209	0	1
United _{ij}	0.348	0.476	0	1
US Airways _{ij}	0.091	0.288	0	1
Virgin _{ij}	0.007	0.083	0	1

Notes: The sample size is 145,458.

Table 2: Pricing Equation

Variables	(1)		(2)		(3)		(4)	
	Coef	StErr	Coef	StErr	Coef	StErr	Coef	StErr
2:00 a.m.	3.341*	(0.949)	3.319*	(0.971)	3.333*	(0.955)	3.589*	(1.008)
3:00 a.m.	9.775*	(2.938)	9.819*	(2.937)	9.768†	(2.925)	10.06*	(2.874)
4:00 a.m.	9.265*	(3.275)	9.360*	(3.274)	9.260†	(3.263)	9.792†	(3.278)
5:00 a.m.	8.392†	(3.296)	8.443†	(3.301)	8.381†	(3.282)	8.555†	(3.223)
6:00 a.m.	7.993*	(2.939)	8.054*	(2.921)	7.978†	(2.929)	7.803†	(2.847)
7:00 a.m.	8.437*	(2.855)	8.478*	(2.849)	8.415†	(2.847)	7.825†	(2.938)
8:00 a.m.	8.284*	(2.662)	8.307*	(2.653)	8.280†	(2.653)	7.223†	(2.788)
9:00 a.m.	7.101*	(2.665)	7.173*	(2.653)	7.093†	(2.653)	5.964‡	(2.803)
10:00 a.m.	5.428†	(2.118)	5.485*	(2.124)	5.424†	(2.106)	4.217‡	(2.142)
11:00 a.m.	3.278†	(1.344)	3.338†	(1.349)	3.268†	(1.335)	1.890	(1.364)
12:00 p.m.	3.376†	(1.474)	3.458†	(1.478)	3.361‡	(1.463)	1.706	(1.350)
1:00 p.m.	2.935†	(1.396)	2.959†	(1.392)	2.889‡	(1.389)	1.137	(1.156)
2:00 p.m.	1.749	(1.566)	1.748	(1.572)	1.750	(1.553)	0.130	(1.260)
3:00 p.m.	0.118	(1.808)	0.147	(1.803)	0.126	(1.799)	-1.549	(1.454)
4:00 p.m.	0.782	(1.688)	0.776	(1.692)	0.794	(1.677)	-1.060	(1.511)
5:00 p.m.	0.885	(1.590)	0.907	(1.584)	0.883	(1.579)	-1.218	(1.362)
6:00 p.m.	0.514	(1.461)	0.543	(1.466)	0.512	(1.453)	-2.136	(1.419)
7:00 p.m.	0.534	(1.448)	0.562	(1.447)	0.530	(1.438)	-2.403	(1.381)
8:00 p.m.	0.367	(1.198)	0.364	(1.202)	0.357	(1.191)	-2.992‡	(1.448)
9:00 p.m.	0.0819	(0.787)	0.106	(0.805)	0.0952	(0.791)	-3.717†	(1.359)
10:00 p.m.	-0.584	(0.957)	-0.578	(0.970)	-0.583	(0.961)	-3.997†	(1.670)
11:00 p.m.	-1.583*	(0.520)	-1.581*	(0.535)	-1.578†	(0.519)	-3.797*	(1.124)
12:00 a.m.	-0.901*	(0.279)	-0.891*	(0.276)	-0.893†	(0.279)	-1.183*	(0.264)
MONDAY _t	-9.826*	(3.655)	-9.815*	(3.667)	-9.809†	(3.652)		
TUESDAY _t	-18.02*	(3.525)	-18.00*	(3.539)	-18.02*	(3.530)		
WEDNESDAY _t	-34.31*	(6.441)	-34.32*	(6.471)	-34.31*	(6.451)		
THURSDAY _t	-34.51*	(8.410)	-34.51*	(8.426)	-34.55*	(8.381)		
ADV _t	-45.04*	(9.152)	-45.02*	(9.158)	-45.02*	(9.176)		
ADV _t ²	1.249*	(0.255)	1.248*	(0.255)	1.248*	(0.255)		
ADV _t ³	-0.0106*	(0.00229)	-0.0106*	(0.00229)	-0.0106*	(0.00229)		
ROUTE FE	No		Yes		Yes		Yes	
FLIGHT FE	No		No		Yes		Yes	
ADV FE	No		No		No		Yes	
Within R ²					0.308		0.321	

Notes: The dependent variable is $\text{LNFARE}_{ijt} \times 1000$ and the total number of observations is 145,458. Figures in parentheses are White heteroskedasticity-consistent estimates of the asymptotic standard errors, clustered by carrier. ‡ significant at 10%; † significant at 5%; * significant at 1%.

Table 3: Dynamic Pricing Equation.

Sample:	(1)		(2)		(3)		(4)	
	$0 \leq \text{Adv} < 14.8$		$0 \leq \text{Adv} < 29.6$		$29.6 \leq \text{Adv} < 59.3$		$0 \leq \text{Adv} < 59.3$	
Variables	Coef	StErr	Coef	StErr	Coef	StErr	Coef	StErr
2:00 a.m.	5.201*	(1.235)	3.458*	(0.758)	0.693*	(0.253)	1.936*	(0.379)
3:00 a.m.	19.50*	(1.264)	10.46*	(0.770)	0.979*	(0.255)	5.095*	(0.383)
4:00 a.m.	3.757*	(1.293)	3.534*	(0.780)	0.689*	(0.256)	1.788*	(0.386)
5:00 a.m.	2.466‡	(1.291)	2.680*	(0.783)	0.528†	(0.258)	1.311*	(0.389)
6:00 a.m.	2.766†	(1.273)	2.745*	(0.784)	0.571†	(0.261)	1.373*	(0.391)
7:00 a.m.	2.652†	(1.284)	2.531*	(0.791)	0.399	(0.264)	1.180*	(0.395)
8:00 a.m.	1.847	(1.291)	2.101*	(0.796)	0.795*	(0.268)	1.208*	(0.399)
9:00 a.m.	1.262	(1.289)	0.808	(0.798)	0.734*	(0.273)	0.529	(0.403)
10:00 a.m.	1.332	(1.292)	0.328	(0.796)	-0.454‡	(0.274)	-0.348	(0.404)
11:00 a.m.	0.616	(1.294)	-0.476	(0.798)	-1.155*	(0.275)	-0.937†	(0.405)
12:00 p.m.	0.236	(1.293)	0.179	(0.800)	0.175	(0.275)	0.199	(0.406)
1:00 p.m.	0.491	(1.298)	-0.481	(0.803)	-0.906*	(0.276)	-0.807†	(0.407)
2:00 p.m.	-0.00449	(1.301)	-0.826	(0.802)	-1.237*	(0.277)	-1.069*	(0.407)
3:00 p.m.	0.197	(1.305)	-0.925	(0.801)	-0.843*	(0.277)	-0.833†	(0.407)
4:00 p.m.	-0.538	(1.303)	-0.610	(0.800)	-1.202*	(0.276)	-0.901†	(0.406)
5:00 p.m.	-0.362	(1.295)	-0.841	(0.795)	-1.481*	(0.274)	-1.121*	(0.404)
6:00 p.m.	-2.302‡	(1.287)	-1.670†	(0.789)	-1.285*	(0.272)	-1.365*	(0.401)
7:00 p.m.	-1.414	(1.283)	-1.234	(0.785)	-1.099*	(0.272)	-1.083*	(0.400)
8:00 p.m.	-0.826	(1.276)	-0.797	(0.780)	-0.900*	(0.270)	-0.815†	(0.397)
9:00 p.m.	-3.491*	(1.278)	-2.377*	(0.781)	-0.525‡	(0.270)	-1.423*	(0.398)
10:00 p.m.	-1.832	(1.244)	-1.589†	(0.767)	-0.326	(0.264)	-0.964†	(0.390)
11:00 p.m.	-2.005‡	(1.202)	-1.846†	(0.746)	-0.192	(0.257)	-1.034*	(0.379)
12:00 a.m.	2.309‡	(1.213)	-0.00598	(0.255)	0.435	(0.380)		
LNFARE _{ij,t-1}	0.789*	(0.0115)	0.630*	(0.00856)	0.392*	(0.0128)	0.625*	(0.00608)
Specification tests:								
Serial correlation	1.167		0.0211		0.0297		0.0900	
	[0.243]		[0.983]		[0.976]		[0.928]	
Sargan	322.5		370.5		263.2		443.8	
	[0.999]		[1]		[1]		[1]	

Notes: The dependent variable is $\text{LNFARE}_{ijt} \times 1000$. Figures in parentheses are the Windmeijer finite-sample corrected standard errors of the GMM two-step estimates. Figures in brackets are p-values. ‡ significant at 10%; † significant at 5%; * significant at 1%. ^a The null hypothesis is that the errors in the first-difference regression exhibit no second-order serial correlation (valid specification). ^b The null hypothesis is that the instruments are not correlated with the residuals (valid specification).

Table 4: Dynamic Price Discrimination. Weekdays vs. Weekends.

Sample:	$0 \leq \text{ADV} < 29.6$			$29.6 \leq \text{ADV} < 59.3$			$0 \leq \text{ADV} < 59.3$		
	Weekdays		Weekends	Weekdays		Weekends	Weekdays		Weekends
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Variables:									
OFFICE_t^a	-3.509*	-3.237*	0.0213	-1.001*	-0.836	0.138‡	-2.117*	-4.172*	0.189‡
	(0.340)	(0.708)	(0.220)	(0.120)	(0.639)	(0.0785)	(0.174)	(0.360)	(0.115)
EVENING_t	-4.439*	-7.054*	0.662*	-1.006*	-0.990‡	0.210*	-2.488*	-6.125*	0.474*
	(0.338)	(0.708)	(0.220)	(0.118)	(0.584)	(0.0789)	(0.172)	(0.359)	(0.116)
$\text{OFFICE}_t \cdot \text{ADV}_t^b$		-0.0190			-0.00387			0.0698*	
		(0.0421)			(0.0147)			(0.0108)	
$\text{EVENING}_t \cdot \text{ADV}_t$		0.169*			-0.000397			0.122*	
		(0.0409)			(0.0134)			(0.0106)	
LNFARE_{ijt-1}	0.629*	0.625*	0.743*	0.404*	0.404*	0.778*	0.625*	0.616*	0.794*
	(0.00840)	(0.00843)	(0.0372)	(0.0126)	(0.0126)	(0.0452)	(0.00602)	(0.00605)	(0.0231)
Price discrimination:									
$\beta_{\text{OFFICE}} - \beta_{\text{EVENING}}$	0.931*		-0.640*	0.00477		-0.0719	0.371†		-0.284†
	[0.00519]		[0.00368]	[0.967]		[0.331]	[0.0290]		[0.0114]
At $\text{ADV}_t = 7^c$		2.502*						1.587*	
		[1.48e-07]						[6.65e-08]	
At $\text{ADV}_t = 21$		-0.128						0.857*	
		[0.753]						[9.11e-06]	
At $\text{ADV}_t = 35$					0.0329			0.126	
					[0.838]			[0.478]	
Specification tests:									
Serial correlation ^d	-0.0511	-0.0831	-0.868	0.0194	0.0198	-0.0687	-0.0438	-0.0976	-0.826
	[0.959]	[0.934]	[0.386]	[0.985]	[0.984]	[0.945]	[0.965]	[0.922]	[0.409]
Sargan ^e	331.4	325.2	232.8	280.5	270.9	174.2	380.4	387.5	328.4
	[1]	[1]	[0.903]	[1]	[1]	[0.998]	[1]	[1]	[1]

Notes: The dependent variable is $\text{LNFARE}_{ijt} \times 1000$. Figures in parentheses are the Windmeijer finite-sample corrected standard errors of the GMM two-step estimates. Figures in brackets are p-values with the null that the corresponding coefficient is equal to zero. ‡ significant at 10%; † significant at 5%; * significant at 1%. ^a Omitted category is MORNING. ^b Omitted category is MORNING·ADV with ADV controlled with day dummies. ^c $\beta_{\text{OFFICE}} - \beta_{\text{EVENING}} + \beta_{\text{OFFICE} \cdot \text{ADV}} \cdot \text{ADV}_t - \beta_{\text{EVENING} \cdot \text{ADV}} \cdot \text{ADV}_t$. ^d The null hypothesis is that the errors in the first-difference regression exhibit no second-order serial correlation (valid specification). ^e The null hypothesis is that the instruments are not correlated with the residuals (valid specification).

Table 5: Dynamic Price Discrimination. Legacy vs Low Cost Carriers.

Sample:	Legacy		Low Cost	
	$0 \leq \text{ADV} < 29.6$	$0 \leq \text{ADV} < 59.3$	$0 \leq \text{ADV} < 29.6$	$0 \leq \text{ADV} < 59.3$
	(1)	(2)	(3)	(4)
Variables:				
OFFICE _t ^a	-2.688*	-4.077*	-5.977†	-5.326*
	(0.720)	(0.364)	(2.428)	(1.321)
EVENING _t	-5.274*	-5.090*	-14.53*	-10.88*
	(0.718)	(0.361)	(2.409)	(1.305)
OFFICE _t · ADV _t ^b	-0.0519	0.0656*	0.0861	0.106*
	(0.0426)	(0.0109)	(0.151)	(0.0397)
EVENING _t · ADV _t	0.0927†	0.0966*	0.445*	0.237*
	(0.0416)	(0.0106)	(0.143)	(0.0391)
LNFARE _{ijt-1}	0.591*	0.586*	0.749*	0.750*
	(0.00893)	(0.00639)	(0.0141)	(0.0102)
Price discrimination:				
At ADV _t = 7 ^c	1.575*	0.796*	6.046*	4.633*
	[0.00109]	[0.00725]	[0.000235]	[1.37e-05]
At ADV _t = 21	-0.449	0.362‡	1.025	2.798*
	[0.272]	[0.0623]	[0.503]	[8.61e-05]
At ADV _t = 35		-0.0715		0.963
		[0.689]		[0.158]
Specification tests:				
Serial correlation ^d	0.0115	0.0351	-0.864	-1.245
	[0.991]	[0.972]	[0.388]	[0.213]
Sargan ^e	257.3	353.6	34.18	37.07
	[1]	[1]	[1]	[1]

Notes: The dependent variable is $\text{LNFARE}_{ijt} \times 1000$. Figures in parentheses are the Windmeijer finite-sample corrected standard errors of the GMM two-step estimates. Figures in brackets are p-values with the null that the corresponding coefficient is equal to zero. ‡ significant at 10%; † significant at 5%; * significant at 1%. ^a Omitted category is MORNING. ^b Omitted category is MORNING · ADV with ADV controlled with day dummies. ^c $\beta_{\text{OFFICE}} - \beta_{\text{EVENING}} + \beta_{\text{OFFICE} \cdot \text{ADV}} \cdot \text{ADV}_t - \beta_{\text{EVENING} \cdot \text{ADV}} \cdot \text{ADV}_t$. ^d The null hypothesis is that the errors in the first-difference regression exhibit no second-order serial correlation (valid specification). ^e The null hypothesis is that the instruments are not correlated with the residuals (valid specification).

Table 6: Dynamic Price Discrimination. Market Structure.

Sample:	$0 \leq \text{Adv} < 29.6$	$29.6 \leq \text{Adv} < 59.3$	$0 \leq \text{Adv} < 59.3$
	(1)	(2)	(3)
Variables:			
OFFICE _t ^a	-2.314 (1.961)	-0.600 (0.679)	-1.620 (0.991)
EVENING _t	-1.425 (1.926)	-0.0852 (0.665)	-1.084 (0.973)
OFFICE _t · HHI _{ij} ^b	-1.252 (2.460)	-0.491 (0.844)	-0.634 (1.238)
EVENING _t · HHI _{ij}	-3.341 (2.419)	-1.177 (0.829)	-1.782 (1.216)
LNFARE _{ij,t-1}	0.647* (0.00822)	0.418* (0.0125)	0.625* (0.00602)
Price discrimination:			
At HHI _{ij} = 0.52 ^c	0.197 [0.787]	-0.158 [0.540]	0.0610 [0.870]
At HHI _{ij} = 0.78	0.740† [0.0261]	0.0198 [0.866]	0.359† [0.0344]
At HHI _{ij} = 1.00	1.200‡ [0.0543]	0.171 [0.420]	0.612‡ [0.0497]
Specification tests:			
Serial correlation ^d	-0.0338 [0.973]	0.0794 [0.937]	-0.0652 [0.948]
Sargan ^e	339.5 [1]	285.9 [1]	394.9 [1]

Notes: The dependent variable is $\text{LNFARE}_{ijt} \times 1000$. Figures in parentheses are the Windmeijer finite-sample corrected standard errors of the GMM two-step estimates. Figures in brackets are p-values with the null that the corresponding coefficient is equal to zero. ‡ significant at 10%; † significant at 5%; * significant at 1%. ^a Omitted category is MORNING. ^b Omitted category is MORNING·HHI with HHI controlled with flight fixed effects. ^c $\beta_{\text{OFFICE}} - \beta_{\text{EVENING}} + \beta_{\text{OFFICE} \cdot \text{HHI}} \cdot \text{HHI}_j - \beta_{\text{EVENING} \cdot \text{HHI}} \cdot \text{HHI}_j$. ^d The null hypothesis is that the errors in the first-difference regression exhibit no second-order serial correlation (valid specification). ^e The null hypothesis is that the instruments are not correlated with the residuals (valid specification).

Table 7: Dynamic Price Discrimination. By Carrier.

Carrier:	American	Alaska	JetBlue	Delta	Frontier	AirTran	United	US Airways	Virgin
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Variables:									
OFFICE _t ^a	-4.460*	-0.896	-8.193*	-0.859	-1.266	-4.531‡	-2.029*	-2.783	0.274
	(1.701)	(2.434)	(2.620)	(2.609)	(2.023)	(2.423)	(0.641)	(1.788)	(1.369)
EVENING _t	-7.827*	-0.678	-11.50*	-2.149	-8.313‡	-13.42*	-3.295*	-5.635*	-0.172
	(1.644)	(2.599)	(2.624)	(2.636)	(3.563)	(2.374)	(0.635)	(1.770)	(1.191)
LNFARE _{ijt-1}	0.686*	0.617*	0.808*	0.749*	0.895*	0.827*	0.474*	0.775*	0.470*
	(0.0123)	(0.0268)	(0.0144)	(0.0223)	(0.0179)	(0.0142)	(0.0131)	(0.0197)	(0.0367)
Price discrimination:									
$\beta_{\text{OFFICE}} - \beta_{\text{EVENING}}$	3.367‡	-0.218	3.302	1.290	7.047‡	8.888*	1.266‡	2.851	0.446
	[0.0482]	[0.933]	[0.207]	[0.627]	[0.0425]	[0.000168]	[0.0455]	[0.125]	[0.705]
Specification tests:									
Serial correlation ^b	-1.121	0.541	-1.060	-0.982	-1.579	0.425	1.119	1.051	-0.147
	[0.262]	[0.588]	[0.289]	[0.326]	[0.114]	[0.671]	[0.263]	[0.293]	[0.883]
Sargan ^c	55.18	5.026	13.80	19.09	3.58e-10	15.97	89.07	13.20	1.320
	[1]	[1]	[1]	[1]	[1]	[1]	[1]	[1]	[1]

Notes: The dependent variable is $\text{LNFARE}_{ijt} \times 1000$. Figures in parentheses are the Windmeijer finite-sample corrected standard errors of the GMM two-step estimates. Figures in brackets are p-values with the null that the corresponding coefficient is equal to zero. ‡ significant at 10%; † significant at 5%; * significant at 1%. ^a Omitted category is MORNING. ^b The null hypothesis is that the errors in the first-difference regression exhibit no second-order serial correlation (valid specification). ^c The null hypothesis is that the instruments are not correlated with the residuals (valid specification).