Price Discrimination and Focal Points for Tacit Collusion: Evidence from the Airline Industry*

Diego Escobari† Nicholas G. Rupp‡ Joseph Meskey§

November 27, 2018

Abstract

We use unique data sets with round-the-clock posted fares and a regression discontinuity design to identify price discrimination in advance-purchase discounts. Price discrimination increases fares by 7.6% at 14 days to departure, and by 14% at 7 days to departure. While competition reduces price discrimination, it is unaffected by product variety for a multiproduct monopolist. The results show that the arbitrary thresholds of 7 and 14 days-in-advance serve as focal points for tacit collusion and to implement price discrimination in competitive markets. For round-trip tickets price discrimination depends on the days-in-advance for both the outbound and inbound flights.

Keywords: Price discrimination, Market structure, Focal points, Multiproduct monopolist, Advance purchases, Regression discontinuity, Airlines

JEL Classifications: C23, D40, L93, R41

*We thank Marco Alderighi, Alexei Alexandrov, Sebastian Calonico, Gary Fournier, Alberto Gaggero, Mark Hoekstra, Claudio Piga, and seminar participants at Texas A&M, LAMES in Medellín, SEA in Washington DC, IIOC in Philadelphia, EEA in Lisbon, and EARIE in Maastricht.

†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, Email: joemeskey@gmail.com

1 Introduction

Across a variety of markets (e.g., hotels, airlines, cruise lines, rental cars, registration fees, music performances) sellers commonly offer discounts to consumers who buy early. While there is important theoretical work explaining this phenomena, the empirical identification of price discrimination in advance purchase discounts has been elusive. After holding product characteristics fixed, the difficulty arises when trying to distinguish between cost-based differentials and discriminatory differentials. As explained in Stole (2007, p. 2225), this is particularly challenging in advance-purchase discounts where sales occur in a setting of costly capacity and aggregate demand uncertainty with the shadow cost changing across buyers and over time depending on aggregate demand expectations. The focus of this paper is to provide price discrimination estimates in advance-purchase discounts in the U.S. airline industry.

Large U.S. airlines face the complex task of pricing airline tickets for thousands of daily flights involving hundreds of thousands of passengers. According to the Bureau of Transportation Statistics in 2015, U.S. carriers offered 842 million seats for just their domestic flights. Since airlines post fares up to 332 days in advance of departure, on any given day U.S. carriers are maintaining an inventory of 766 million passenger fares for their domestic flights. Carriers rely on historical sales patterns and existing inventory levels when setting and adjusting airfares. Previous theoretical work by Dana (1998, 1999b,a) cites demand uncertainty and capacity costs as reasons why airlines offer advance purchase discounts. The purpose of this paper is to offer an alternative explanation of price discrimination that is independent of both demand uncertainty and capacity costs. We are able to control for inventory levels by examining high frequency posted fare data when estimating price discrimination in advance purchase discounts. The second objective of this paper is to determine if any focal points exist where carriers in unison increase their posted fares. We find that both one week and two weeks prior to departure serve as focal points for U.S. carriers to raise posted airfares. Our third objective is to examine the role of market concentration (both Herfindahl-Hirschman Index and number of carriers serving the market) in price discrimination. Finally, we examine whether Armstrong's (1996, 1999) theory of price discrimination by a multiproduct is consistent with empirical evidence in the U.S.

airline industry.

This paper uses unique proprietary data sets with information on posted airline ticket prices each hour prior to departure. The empirical approach uses a regression discontinuity (RD) design to allow us control for unobserved costs and identify price discrimination.[1] The assignment to the price discrimination treatment is based on whether time-to-departure exceeds a known cutoff (e.g., 7 or 14 days-in-advance requirement) and the idea behind the research design is that the cost of a ticket just below the cutoff (that received the discount) is a good comparison to a ticket just above the cutoff (that did not receive the discount).

The regression discontinuity design results provide strong evidence of both statistically and economically significant price discrimination in advance-purchase discounts of airline tickets. Using the methods in Imbens and Kalyanaraman (2012) and Calonico et al. (2014b) the point estimates indicate that for one-way economy-class tickets price discrimination increases fares by 14.0% (i.e., about \$53.91) at 7 days to departure, and by 7.6% (i.e., about \$29.07) at 14 days prior to departure. The results are consistent with higher valuation consumers purchasing closer to departure at higher prices while sellers use the timing of purchase as a mechanism to separate between consumer types. The observed discontinuities are robust across various kernel type specifications, bandwidth selection procedures, and orders of the local polynomial and bias.

When studying the role of market structure, we find that for the 7 days-in-advance threshold price discrimination is greater in more concentrated markets, consistent with the theoretical work in Stole (2007) and Möller and Watanabe (2016). Moreover, our results support the claim that the arbitrary thresholds of 7 and 14 days-in-advance serve as focal points for tacit collusion and facilitate price discrimination in competitive markets. We show that airlines have solved the coordination problem in competitive markets (see Armstrong, 1996, and Armstrong and Vickers, 2001) by using focal points and simultaneously increasing fares from competing flights at exactly the same day and hour prior to departure—immediately after midnight (Pacific Standard Time) of the 7^{th} and 14^{th} days to

[1] Hahn et al. (2001) formally show that RD designs require seemingly mild assumptions compared to those needed for other non-experimental approaches, while Lee (2008) provide a theoretical justification that causal inferences from RD designs are potentially more credible than the "natural experiment" strategies (e.g., difference-in-difference or instrumental variables).

departure. This result is additionally interesting because unlike previous literature on focal points, we provide empirical evidence of the importance of timing to solve the problem of multiple equilibria in coordination games.

Moreover, the richness of our data allows us to study the behavior of multiproduct monopolists, (monopoly carriers offering multiple flights at differentiated departure times) along with the role of refundable tickets and round-trip itineraries. We find that price discrimination is unaffected by product variety (as captured by the number of flights). This result is consistent with Gale and Holmes (1992, 1993) where a monopolist offers advance-purchase discounts to divert demand from peak to off-peak flights. For refundable tickets, we find that no evidence of price discrimination through advance-purchase discounts. Finally, for round-trip tickets price discrimination exists for 7 and 14 days-in-advance of both departure and return dates of the ticket.

Advance-purchase discount fares were introduced in the U.S. following the Airline Deregulation Act of 1978, subsequently there have been numerous studies examining advance purchase discounts. Ata and Dana (2015) explain price discrimination based on booking time is feasible in airlines because consumers learn about their demand at different times. Advance-purchase discounts as a price-discrimination devise can promote efficiency, for example, by leading to output expansion in markets with elastic demand (Schmalensee, 1981). In addition, advance-purchase discounts might be the only way to cover large fixed costs (Frank, 1983) and can serve to allocate limited capacity on peak flights (Gale and Holmes, 1992). In wholesale contracts, they can help to reduce the risk of holding excess inventories (Cachon, 2004). Möller and Watanabe (2010) present an advance selling model to explain why some goods are cheaper when bought earlier while some offer discounts to those who buy late. Su (2007) shows that when consumers with higher valuations have higher waiting costs or are more impatient declining prices are also optimal. Stokey (1979) shows that if consumers vary only in their valuations, then firms cannot use time to discriminate between them. Courty (2003) considers a monopolist which offers advance selling if the unit cost of production is below a certain threshold and conducts spot selling otherwise. In Nocke et al. (2011) a monopolist offers an advance-purchase discount to discriminate between consumers on the basis of their expected valuation. In Dana (1998, 1999a), advance-purchase discounts are the optimal pricing policy given fixed capacity and uncertain aggregate demand.

There are a considerable number of empirical studies of price discrimination in a variety of settings beyond the airline industry. For example, Shepard (1991) identifies price discrimination in gas stations, Ivaldi and Martimort (1994) for electric utilities, Leslie (2004) considers nonlinear pricing for Broadway theater, McManus (2007) on coffee shops, Cohen (2008) in paper towels, Clerides (2002) on books, Crawford and Shum (2007) for cable television, and Busse and Rysman (2005) in Yellow Pages directories. All of these prior studies, however, use product quality as a screening device. In our approach product quality is homogeneous and the screening device is the timing of the purchase.

Time variation in prices as the flight date nears has recently received considerable attention from airline researchers. Gaggero and Piga (2011) study market power, while Escobari (2012), Williams (2017) and Alderighi et al. (2015) examine the role of inventories. Bilotkach and Rupp (2012) analyze price-offer curves and Hernandez and Wiggins (2014) consider nonlinear pricing strategies. More recently Bilotkach et al. (2015) consider how active yield management affects capacity utilization, Cattaneo et al. (2016) look at low-cost carrier price discrimination, and Alderighi et al. (2016) examines changes in the distribution of prices.

The organization of the paper is as follows. Section 2 describes the data and the collection procedure, while Section 3 explains the price discrimination identification strategy using a regression discontinuity design. The price discrimination estimates are presented in Section 4 for monopoly markets, multiproduct monopolists, and for various specifications of competitive markets. This section also explains the existence of discontinuities as focal points that facilitate tacit collusion and price discrimination in competitive markets. Section 5 reports additional estimates on the refundability and round-trip tickets. Section 6 concludes.

2 Data

This paper takes advantage of two unique airline pricing data sets collected from online travel agencies to track hourly price changes as the departure date nears. The first data set of one-way tickets allows us to obtain measures of market structure and to assess the role of refundability, while the second data set of round-trip tickets includes a much larger combination of flights.

2.1 Refundable and Non-refundable One-way Fares

The first data set contains 1,908,683 price quotes for non-refundable (989,101) and refundable (919,582) economy class domestic one-way fares. We observe 1,665 different flights across 158 routes, where a route is defined as a directional pair of departure and arrival airports (e.g., Charlotte Douglas International Airport (CLT) to Chicago O'Hare International Airport (ORD) and ORD to CLT are treated as separate routes). There are nine carriers in the sample, American, Alaska, JetBlue, Delta, Frontier, AirTran, United, US Airways, and Virgin America.[2] While our sample fare data includes nine carriers, the bulk of the observations come from just four carriers, United, American, Delta, and US Airways, which reflects their dominant position in the U.S. market during the months the data was collected.

The high frequency posted fare data (24 observations per day per flight) enables us to observe the behavior of sellers whether or not a transaction occurs. Our data includes thousands of observations around the cutoff points of 21, 14, and 7 days-in-advance of departure. This is a key advantage over transaction data where observations are only recorded if a transaction occurs.[3] Moreover, these are equilibrium prices and because sellers take into account the optimal behavior of buyers when posting fares, we can also use these prices to draw inferences about the behavior of buyers as well.

The drawback of recording fares twenty-four times per day is that the sample size grows very quickly. Hence, we restrict the collection strategy in a way that also helps us control for sources of price variation that we do not want to study. To control for systematic price differences across departure dates (Gale and Holmes, 1993; and Escobari, 2009) we focus on departures for a single day, Thursday June 22, 2012. Within the same flight and at every hour prior to departure we record both refundable and non-refundable fares for one-way

[2]The only major U.S. carrier excluded from the sample is Southwest, whose fares only appear on Southwest.com. To control for potential effects of Southwest on pricing we selected airport pairs where Southwest does not offer non-stop flights.

[3]The widely used DB1B transaction data from the Bureau of Transportation Statistics (e.g., Berry and Jia, 2010) does not record the date of purchase.

economy class tickets. This allows us to control for price variation associated with frequent flier miles tickets and different fare classes (e.g., first class). In our first data set using one-way tickets simplifies the analysis since it enables us to control for price variation associated with Saturday-night stay over, and minimum and maximum stay requirements. Moreover, selecting non-stop flights also provides a cleaner comparison of flight quality and hence avoids issues such as considering connecting flights and more sophisticated itineraries.

[Table 1 (Summary Statistics), about here]

The summary statistics appear in panel A, Table 1. Fare is the price (in dollars) while Time is the number of days prior to departure. In addition to the carrier dummies, to gain some insights on the competition within routes the panel also reports summary statistics for the number of flights in a route, number of carriers in a route, number of own flights in a route, as well as market share and the Herfindahl-Hirschman Index (HHI).

2.2 Nonrefundable Round-trip Fares

With our second data set we extend the analysis to include round-trip fares. As with the first data set, we have a panel that keeps track of fares every hour for 28 days as the departure date nears. The cross-sectional unit is a pair of outbound and inbound flights, where each pair of flights connects the same two airports. Moreover, we use the same list of airports as in the first data set and to control for systematic price differences across departure dates we focus on a single outbound date and a single inbound date.[4] We gathered round trip fares of all possible combinations of all outbound flights that departed on Thursday December 3, 2015 and all inbound flights that returned on Monday December 7, 2015.

Overall this process gives us a panel with over 21 million observed prices. We restrict the sample to have only outbound and inbound flight pairs that belong to the same carrier, and to count each outbound and inbound flight only once we match flights based on departure times. For example, note that if there are 5 outbound and 5 inbound flights for a given airport pair, then we would have a total of 25 price quotes every hour. Based on departure

[4]Note that this also helps to avoid any "course of dimensionality" as the alternative outbound and inbound combinations grows exponentially.

time we match the first (the one that departs the earliest) outbound with the first inbound flight such that each flight is counted only once in the computation of the hourly round-trip fares for the airport pair. Panel B on Table 1 shows the summary statistics for Fare with the resulting sample of 1,514,833 observations. Breaking down the summary statistics of Fare for the last four weeks to departure we observe how average prices increase as the departure date nears.

3 Identifying Price Discrimination

To identify price discrimination we need to control for both product quality (e.g., ticket refundability, cabin class) and cost differences. Stigler's (1987) definition of price discrimination states that a firm price discriminates when the ratio of prices is different from the ratio of marginal costs of two goods offered by a firm. Stole (2007) explains that this definition requires a careful calculation of marginal costs to include all relevant shadow costs which is particularly true when capacity is costly and aggregate demand is uncertain as in airlines. Cost differences at different points before departure are difficult to control for because they depend on both seat inventories and demand expectations. For example, if aggregate demand is expected to be high, then the opportunity cost of selling the next available seat is also high. If expected aggregate demand is low, with the flight being likely to depart with empty seats, then the opportunity costs of the next ticket is low. Hence a measure of the price markups over marginal costs is challenging. Borenstein and Rose (1994) explain that disentangling the different sources of price dispersion is difficult due to product heterogeneities (e.g., refundability, advance-purchase discounts, etc.) which provide a basis for self-selective price discrimination and also affect costs.

To illustrate the observed price dispersion as the flight date nears, Figure 1 presents the one-way nonrefundable fares at every hour prior to departure for the American Airlines flight 1152 between Seattle-Tacoma International Airport (SEA) and Chicago O'Hare International Airport (ORD). The right-hand side of the figure zooms in to the 7 day-in-advance threshold to illustrate how fares jump immediately after midnight. For this flight the highest priced ticket is about 120% more expensive than the lowest priced ticket. The observed time variation in prices as the flight date nears is not necessarily price discriminatory. While we already control for ticket quality (e.g., refundability, cabin), the observed price variation in Figure 1 can be the result of differences in costs that depend on seat inventories, demand expectations, and time to departure.[5] In this section we explain how we use a RD design to identify the price discrimination component within this observed price dispersion.

[Figure 1 (AA Flight), about here]

Airlines use advance purchase discounts to price discriminate and separate between consumers who have different valuations for a ticket. Valuations have a positive correlation with travelers learning about their demand—higher-valuation consumers learn their demand closer to departure. Airlines can then use time to departure as a screening device to separate between consumer types. Consumers who show up early in the selling season are expected to have lower valuations and can receive the discount, while consumers who appear later face higher prices since the advance-purchase discounts have already expired. Advance purchase discounts are implemented by airlines jointly with more sophisticated pricing strategies that take into account capacity costs and aggregate demand uncertainty.

The intuition behind our RD design price discrimination identification strategy is as follows. We want to estimate the price discrimination treatment effect where the observed "assignment" variable (or "running" variable) is time-to-departure. When time-to-departure exceeds a known cutoff, e.g. 14 days, the ticket receives a discount. We know that the costs associated with tickets sold at the 12 and 16 days-to-departure are most likely different. The idea behind our research design is that the cost of a ticket just below the cutoff (that received the discount) is a good comparison to the cost of a ticket just above the cutoff (that did not receive the discount).

Hahn et al. (2001) formally show that RD designs require seemingly mild assumptions compared to those needed for other non-experimental approaches. Moreover, Lee (2008) provides a theoretical justification that causal inferences from RD designs are potentially more credible than the "natural experiment" strategies (e.g., difference-in-difference or instrumental variables) since RD design isolates the treatment variation as a consequence

[5] Escobari (2012) finds price increases as inventory decreases, and decreases as there is less time to sell. Moreover, airlines adjust prices as they learn about aggregate demand.

of agents' inability to precisely control the assignment variable near the known cutoff. This means that the approach does not need to assume that the RD design is "as good as randomized". This works for our price discrimination identification strategy because buyers do not have precise control over when they buy. Travelers who arrive late cannot go back in time to benefit from an advance-purchase discount. Moreover, increasing prices over time means that potential travelers who arrive early have little incentive to wait before making an airline ticket purchase.

Formally, for each observation i in the data let the random variable $\textsc{Fare}_i$ denote our outcome of interest.[6] The scalar regressor $\textsc{Time}_i$ (time-to-departure) is the running variable that determines the treatment assignment based on a known cutoff. Following the framework in Heckman and Vytlacil (2007) and Imbens and Wooldridge (2009), let $\{(\textsc{Fare}_i(0), \textsc{Fare}_i(1), \textsc{Time}_i)' : i = 1, 2, \ldots, n\}$ be a random sample from $(\textsc{Fare}(0), \textsc{Fare}(1), \textsc{Time})'$, with $\textsc{Fare}(0)$ and $\textsc{Fare}(1)$ being the outcomes without and with the price discrimination treatment. $\textsc{Fare}_i$ is assigned to the price discrimination treatment condition if $\textsc{Time}_i < \overline{\text{T}}$ and is assigned to the control (no price discrimination) condition if $\textsc{Time}_i \geq \overline{\text{T}}$ for a specific and known fixed value $\overline{\text{T}}$. For the one-way prices data we explore three potential known cutoffs, the seven-, fourteen-, and twenty-one-days-in-advance purchase restrictions (i.e., $\overline{\text{T}} = 7, 14, 21$), while for the round-trip data we explore the same cutoffs but counting the days-in-advance for both the outbound and inbound (return) flights.

The observed outcome is

$$\textsc{Fare}_i = \begin{cases} \textsc{Fare}_i(0) & \text{if} \quad \textsc{Time}_i \geq \overline{\text{T}} \\ \textsc{Fare}_i(1) & \text{if} \quad \textsc{Time}_i < \overline{\text{T}}. \end{cases} \tag{1}$$

We identify price discrimination (PD) as the sharp average treatment effect at the threshold $\overline{\text{T}}$ and it is given by

$$\text{PD} = \mathbb{E}[\textsc{Fare}_i(1) - \textsc{Fare}_i(0) | \textsc{Time}_i = \overline{\text{T}}]. \tag{2}$$

We can estimate PD nonparametrically following the regression-discontinuity design literature under mild continuity conditions. In particular

$$\text{PD} = \mu_+ - \mu_-, \tag{3}$$

[6]While our data has a panel structure, taking into account the panel data dimension is unnecessary for identification in an RD design setting (Lee and Lemieux, 2010).

where

$$\mu_+ = \lim_{\text{TIME}\downarrow\bar{\text{T}}} \mathbb{E}[\text{FARE}_i|\text{TIME}_i = \text{T}] \qquad \mu_- = \lim_{\text{TIME}\uparrow\bar{\text{T}}} \mathbb{E}[\text{FARE}_i|\text{TIME}_i = \text{T}].$$

Using kernel-based local polynomials on either side of the threshold we can estimate PD following Hahn et al. (2001) and Porter (2003).

4 Price Discrimination Estimates

Stole (2007) explains that the methodology of monopoly price discrimination is both useful and misleading when analyzing the effects of price discrimination in imperfect competition. Useful because one can solve for the best-response function in a Cournot quantity game by deriving the residual demand curve and modeling the response of a firm as monopolistic in this residual demand. Price discrimination, however, entails more than one price and monopoly models can be misleading because we want to obtain best-response functions in equilibrium rather than a single optimal pricing strategy. In imperfect competition, firms' profits depend on whether the additional surplus extracted by implementing price discrimination (that would be the standard monopoly result) is greater than the competitive externality created if price discrimination increases the intensity of price competition. Hence, it is not clear whether price discrimination can be sustained in imperfect competition. Moreover, it is not clear how firms manage to cope with potentially very complicated best response functions in equilibrium when price discrimination exists. In this section we first focus on estimating price discrimination in monopoly markets and then analyze price discrimination for multiproduct monopolists. Later we assess the role of market concentration and imperfect competition on price discrimination.

4.1 Monopoly Markets

We now present the empirical results from our regression-discontinuity approach to identify price discrimination in monopoly markets using non-stop one-way economy fares. Following a stricter definition than Borenstein and Rose (1994), we define monopoly markets as a single carrier offering non-stop service between two airports. Figure 2 has the number of days before departure in the horizontal axis and the logarithm of FARE (LOGFARE) on the vertical axis. The figure plots the mean LogFare collapsed into bins along with fourth order global polynomials estimated separately on each side of the 7-days-in-advance cutoff. This figure suggests that LogFare increases significantly and discontinuously once the days-in-advance crosses the threshold. The vertical distance between the points close to the discontinuity in analogous to the estimate of PD in equation (2). Figure 3 presents the analog to these results for the 14-days-in-advance cutoff, $\overline{\text{T}} = 14$. The 14-days-in-advance price discrimination treatment appears to also show a significant discontinuity at the cutoff.

[Figure 2 (RD Plot: 7 Days), about here]

[Figure 3 (RD Plot: 14 Days), about here]

Table 2 presents the sharp regression-discontinuity design estimates of price discrimination as suggested in equation (3). The dependent variable is the logarithm of the non-refundable fares (LogFare) and the running variable is Time. The cutoffs of 7-, 14-, and 21-days-in-advance ($\overline{\text{T}} = 7, 14, 21$) one-way fares are presented in separate panels. The RD estimates in this table and in the rest of the document use the second-generation bias-corrected bandwidth selection approach proposed in Calonico et al. (2014b) using the procedures in Calonico et al. (2014a). As explained in Calonico et al. (2014b)—henceforth, CCT—available bandwidth selectors typically yield a "large" bandwidth. These bandwidth selectors lead to a non-negligible bias in the distributional approximation of the estimator which in our case implies that conventional confidence intervals may substantially over-reject the null hypothesis of no price discrimination treatment effect.[7] The robust 95% confidence intervals and the robust p-values we report are based on this bias-corrected RD estimator and the corresponding consistent standard error estimator. Different columns present robustness checks for different kernel types, bandwidth selectors, the choice of the weighted first or second order ($p = 1, 2$) polynomial regressions for both sides of the cut-offs, and the order of the local polynomial bias estimator ($q = 2, 3$). The bandwidth (h) is measured in minutes and it is selected via Cross Validation (CV), the procedure suggested in Imbens and Kalyanaraman (2012)—henceforth, IK—or the CCT procedure.

[7] The procedure first bias-corrects the RD estimator to account for the effects of a "large" bandwidth choice. Then it rescales the standard error formula to account for the additional variability introduced by the estimated bias.

The estimates in the first column of Table 2, panel A are consistent with the observed discontinuity in the cutoff presented in Figure 2. The findings indicate that buying after the 7-days-in-advance cutoff has passed results in statistically significant higher fares—about 14.0%, (equivalent to \$53.91 or about 0.27 standard deviations in fares). The advantage of using the RD design is that the costs associated with tickets just above and just below the cutoffs are expected to be the same; hence, we can interpret the discontinuity as price discrimination. In the process of balancing the goal of focusing on observations close to the cutoffs and using enough observations to obtain precise estimates we employed windows with various sizes around the cutoff and arrived at similar conclusions. For this first column of panel A, with a triangular kernel along with $p = 1$ and $q = 2$, CCT suggest a relatively stringent bandwidth of $h = 52.47$ minutes. The bandwidths suggested by IK (column 2) and by CV (column 3) are less restrictive for our data—$h = 110.3$ and $h = 166$ minutes respectively. Moreover, in columns 4 and 5 we additionally experiment with different kernel types and orders for the local polynomials and bias. Across all of these specifications our price discrimination estimates at $\overline{\mathrm{T}} = 7$ are all highly significant and robust to both the kernel type selection: p, q, and bandwidth selection procedure.[8]

[Table 2 (Monopolies), about here]

Panel B provides the price-discrimination regression-discontinuity estimates at the 14-days-in-advance cutoff. The results are consistent with the discrete jump illustrated in Figure 3. The point estimate of 0.076 in the first column suggests that buying a ticket prior to the 14-days-in-advance cutoff results in a 7.6% lower fare (equivalent to \$29.07 or 0.15 standard deviations in fares). The observed discontinuities are always statistically significant across various kernel type specifications, bandwidth selection procedures, and orders of the local polynomial and bias.

When comparing panels A and B we observe that the price discrimination point estimates at $\overline{\mathrm{T}} = 14$ are about 33% to 49% smaller than the price discrimination point estimates at $\overline{\mathrm{T}} = 7$. This can be explained by the heterogeneity of consumers at different points prior to departure. For example, the differences in valuations of an airline ticket between con-

[8] The number of price quotes used in Panel A is 511,252, the same as in Figure 2. The differences in the reported number of observations across columns arises because price quotes were already grouped into bins.

sumer types that make a ticket purchase before and after $\overline{\mathrm{T}} = 14$ might not be as big as the differences in valuations between consumers that purchase tickets before and after the $\overline{\mathrm{T}} = 7$ cutoff.[9]

A common feature in monopoly price discrimination models (e.g., Mussa and Rosen, 1978 and Escobari and Jindapon, 2014) is the ability of the seller to screen consumers into different groups depending on the heterogeneity of consumers. In our case, even if the monopoly carrier sells to heterogeneous consumers if the difference between types is not big enough or if the market is dominated by a large proportion of a single common type then the carrier might decide to pool heterogeneous types into a single group and charge the same price. This appears to be the case for the 21-days-in-advance cutoff where the price discrimination estimates reported in panel C show no differences in prices around the cutoff. Across all specifications, the RD price discrimination estimates show the same result—RD estimates are not statistically significant.

4.2 Baseline Covariates

Lee and Lemieux (2010) suggest that one can use baseline covariates to help establish the validity of the RD design. The idea is that the inclusion of baseline covariates—no matter how highly correlated they are with LogFare—should not affect the estimated discontinuity. Taking advantage of the panel structure of the data, one covariate we can employ is fixed effects. While including fixed effects is unnecessary for identification in an RD design, we used flight fixed effects as baseline and "residualized" LogFare as explained in Lee and Lemieux (2010). To help validate the our results we conduct a RD analysis on the residuals and obtain nearly the same results as previously reported in Table 2.

Including other covariates is also unnecessary for identification, however, the RD approach requires that there are no discrete changes at the cutoffs in variables that affect pricing. Within the same flight in addition to time-to-departure (Time), other characteristics that potentially affect pricing through costs are demand expectations and seat inventories. It is reasonable to believe that there are no jumps in demand immediately after 12:00 a.m. (midnight Pacific Standard Time) for the 7-, 14-, and 21-days-in-advance

[9] Williams (2017) shows that the expectation of increasing prices over time provides little incentive for consumers to delay airline ticket purchases.

cutoffs. Moreover, it is also reasonable to assume that demand expectations do not systematically change upwards at exactly the cutoff points.

4.3 Multiproduct Monopolist

Many of the monopoly markets discussed above can be viewed as multiproduct monopoly markets. An interesting feature in airlines is that within the same market (i.e., directional nonstop service between an airport pair) a monopoly carrier might be offering multiple flights departing at different times of the day. Different departure times can be viewed as differentiated products depending on the heterogeneity of consumers' tastes for departure times. For example, consumers might be heterogeneous in their demand uncertainty in advance of their desired departure time, and may also vary in their disutility of flying in their less preferred departure times (see, Gale and Holmes, 1993).

Armstrong (1999) explains that determining the optimal selling strategy for a multiproduct firm facing consumers with unknown tastes is a difficult task. In airline markets consumers have unit demands so the multiproduct nonlinear pricing of Armstrong (1996) and the price discrimination by a multiproduct firm model of Armstrong (1999) might not help much to understand our results. Models that are closer to airline pricing include Gale and Holmes (1993) where a monopolist that offers tickets on two flights uses advance-purchase discounts to divert demand from the peak to the off-peak period. Moreover, Gale and Holmes (1992) show that advance-purchase discounts can assist in attaining an efficient allocation of capacity. They have a monopoly airline that offers two flights and the timing of the peak demand is uncertain. Then the monopolist will offer advance-purchase discounts for both flights to smooth out demand fluctuations. Dana (1999b) presents a model in a setting with price rigidities, costly capacity and stochastic demand in which demand shifting between two flights occurs even when the peak flight is unknown.

[Table 3 (Multiproduct Monopolists), about here]

Table 3 presents the price discrimination estimates in monopoly markets with various daily flights offered. The first column looks at markets with less than four daily flights, the second column considers markets with at least four flights but less than seven daily flights, while the third column includes all markets with at least seven flights.[10] As before, panels A and B report the price discrimination estimates at $\tau = 7$ and at $\tau = 14$, respectively.[11] The price discrimination point estimates are nearly the same across columns—about 14% at the $\tau = 7$ cutoff and about 8% at the $\tau = 14$ cutoff—implying that the number of products offered by multiproduct monopolist has no effect on price discrimination.[12]

The results indicate that monopoly airlines implement price discrimination in advance-purchase discounts simultaneously at $\tau = 14$ and at $\tau = 7$ on all flights in a market. These results are consistent with Gale and Holmes (1992, 1993) in which advance-purchase discounts are offered in both flights (they only consider two flights in their models). Note that this result is not obvious because as the flight date nears and airlines learn about aggregate demand, a monopoly seller might have an incentive to keep low prices in the off-peak flight to promote demand shifting and increase capacity utilization by distributing the remaining aggregate demand more evenly across flights.

4.4 Many Sellers and Tacit Collusion

In monopoly markets, not only is the underlying theory of price discrimination well understood (Stole, 2007), but there is also the expectation that price discrimination exists when reselling the product is difficult and consumers can be separated (Shepard, 1991). In the previous section, we focus on monopoly markets and show that price discrimination in advance-purchase discounts exists and estimate its magnitude. We did not, however, assess the role of competition on price discrimination. This research question is interesting in advance-purchase discounts because in a simple model of competitive markets price equals marginal costs, yet under price discrimination we know that at least one price deviates from marginal cost (Varian, 1989). Hence, various authors suggest that price discrimina-

[10]Splitting the sample into groups is a common approach in RD design (see, e.g., Pop-Eleches and Urquiola, 2013).

[11]As in Table 2, the price discrimination estimates at $\tau = 21$ are statistically insignificant.

[12]The estimates throughout Table 3 use a local-linear ($p = 1$) regression with a quadratic ($q = 2$) bias-correction estimate and a triangular kernel. The bandwidth (h) is measured in minutes and its selection procedure is the one proposed by CCT. The robust variance estimators are the ones proposed by CCT, computed with 3 nearest-neighbors. We find robust results for the different kernel types, bandwidth selectors, and the choice of the weighted first or second order polynomial regressions.

tion should exist only in the presence of market power (see, e.g., Stole, 2007). On the other hand, in a model closely related to airline pricing Dana (1998) shows that price discrimination can exist without the market power assumption.

As a benchmark the first column in Table 4 presents the price discrimination estimates for the whole sample (i.e., including non-monopoly routes). We find that the sharp regression-discontinuity estimates of price discrimination are about 11.5% at $\overline{\mathrm{T}} = 7$ and about 9.3% at $\overline{\mathrm{T}} = 14$. As in the monopoly specifications, we find that there is no statistically significant price discrimination at $\overline{\mathrm{T}} = 21$. Columns 2 through 7 aim to capture the role of competition by presenting the price discrimination results at different ranges of the HHI. With HHI_{pc} defined as the percentile pc of the HHI, columns 2 through 4 report the results for more competitive routes (HHI below its 33^{th}, 40^{th}, and 50^{th} percentiles) while columns 5 through 7 report the results where market power is greater (HHI above its 67^{th}, 60^{th}, and 50^{th} percentiles).

[Table 4 (Market Structure), about here]

Consistent with Stole (2007), we find that price discrimination is more prevalent in highly concentrated markets. For example, comparing the top versus the bottom 33^{th} percentile (columns 2 and 5) we observe that at the 7 day cutoff for more competitive routes price discrimination is about 5.5% versus 14.1% in more concentrated routes.[13] The same is true when comparing the top and bottom 40^{th} and 50^{th} percentiles. Moreover, this result of higher price discrimination in more concentrated markets holds for the 14 day cutoff as well.

[Table 5 (Number of Carriers), about here]

[Table 6 (Number of Flights), about here]

As a robustness check to further explore the link between competition and price discrimination, Tables 5 and 6 examine how price discrimination changes with both the number of carriers serving the market and flight frequency. Column 1 from Table 5 replicates the

[13] Note that the upper 95% robust confidence interval (0.08) for competitive routes is still below than the lower 95% robust confidence interval (0.12) for more concentrated routes.

monopoly estimates, while columns 2 through 5 report the estimates for two, three, four, and six carriers, respectively (there are no routes in the sample with exactly five carriers). At the 7-days-in-advance cutoff price discrimination decreases as the number of carriers increase—the point estimate is the lowest at 4.5% when there are six carriers. As reported in panel B, however, a similar result does not occur for the 14-days-in-advance cutoff. The estimates in panel A of Table 6 show that at $\overline{\text{T}} = 7$ price discrimination steadily decreases as the number of flights increases, dropping to essentially zero on routes with 40 or more flights. At $\overline{\text{T}} = 14$ (panel B) discrimination appears to be greater for markets with 20 to 30 flights and it has about the same magnitude (about 8%) for the remaining markets.

Overall, the results from Tables 4 and 5 show that price discrimination decreases with competition as captured by the HHI. At $\overline{\text{T}} = 14$ price discrimination, however, does not appear to decrease when looking at the number of carriers in the market. We interpret the persistence of price discrimination even in the most competitive markets—below the $33th$ percentile of the HHI and with six sellers—as evidence of tacit collusion.[14] Panel A in Table 6 suggests that as the number of flights increases cooperation becomes more difficult because there is a wider set of alternative options for the travelers to choose. Panel B shows that at two weeks to departure increased number of alternatives does not limit price discrimination opportunities. Both results are consistent with Holmes (1989), who shows that cross-price elasticities become important once competition is introduced in a price discrimination model. Under collusion, however, discrimination cross-price elasticities are no longer important as pricing depends only on the industry-demand elasticity.

4.5 Tacit Collusion and Discontinuities as Focal Points

While the location of the thresholds at 7-, 14-, and 21-days-in-advance appears intuitive as they signal one, two, and three weeks to departure, it is not clear why advance-purchase discounts should expire at particular points prior to departure. After all it is reasonable to argue that consumers' heterogeneity, aggregate demand learning, capacity cost, and other components that affect pricing should be changing smoothly as the flight date nears.

[14]This result at $\overline{\text{T}} = 14$ along with the findings at $\overline{\text{T}} = 7$ are consistent with Plott (1982), who in his review of the literature uses experiments to find that "slight" changes in the underlying structure can switch a market from "competitive" to "collusive" and vice versa.

Hence, prices should also adjust smoothly as heterogeneous consumers arrive at different points during advance sales. Theoretical work on the pricing of inventories with uncertain demand over a finite horizon helps support the absence of price jumps at thresholds (see, e.g., Gallego and van Ryzin, 1994; Zhao and Zheng, 2000; and Deneckere and Peck, 2012). On the other hand price jumps at given thresholds during advance sales might be the result of assuming a finite number of prices, periods, or consumer types. For example, with only two consumer types, a single price jump exists when in one separating equilibrium one of the types buys at a discount while the second type pays full price. Likewise it is also relatively simple to motivate the existence of a price jump at a threshold when there are only two periods or only two prices.

We argue that the discontinuities that we document right after midnight Pacific Standard Time for $\overline{\mathrm{T}} = 7, 14$ serve as focal points where sellers coordinate and jointly increase prices. These focal points help coordinate tacit collusion and implement price discrimination in advance-purchase discounts when there are many sellers. The idea of focal points was first introduced by Thomas Schelling (1960), who shows that agents are sometimes able to coordinate their behavior, to their mutual advantage, by drawing on shared perceptions that particular ways of coordinating are 'prominent' or 'salient'. In Schelling's example, he asked each of a group of respondents to imagine that he was one of two individuals, unable to communicate, trying to meet one another. Each individual had to choose some place in New York City with the hope of meeting the other. Given the large number of places to meet, this appears to be a tremendously difficult task. Interestingly, most of Schelling's respondents chose the same place, Grand Central Station. This meeting point serves as a 'focal point for each person's expectation of what the other expects him to expect to be expected to do' (Schelling, 1960, p. 57).

From Schelling's investigation, we know that players from pure coordination games make some systematic use of labels. Here sellers use the labels of 7 and 14 days to departure to their mutual benefit.[15] Note that in the absence of these labels coordination among airlines

[15] Sugden (1995) provides a theory of how labels can influence decisions in games, showing how rational players make use of information provided by labels. However, labels and incentives to collude do not necessarily mean that sellers will use them as focal points. For example, Engelmann and Müller (2011) find results that fail to support the focal-point hypothesis in markets with price ceilings.

would be a tremendously difficult task given all the multiple points in time in which airlines can increase ticket prices.[16] As explained in Holmes (1989), airlines might be willing to offer discriminatory discounts to attract consumers from rival firms. An interesting feature of our focal points is that they characterize coordination equilibria based on 'when' agents meet, rather than 'where' as previous work on focal points has placed emphasis on the meeting location (see, e.g., Mehta et al., 1994a,b; Bacharach and Bernasconi, 1997; and Knittel and Stango, 2003).

Cooperation in a market with many sellers can be sustained under fairly general conditions. The Folk Theorem states that with sufficiently patient players, almost any set of payoffs may be sustained as the outcome of a repeated game. The Folk Theorem provides the conditions in which sellers can keep supercompetitive prices in repeated interactions with strategies that sustain current cooperation under the threat of future punishment if any firm deviates from cooperation. Tacit collusion at the thresholds is sustainable due to the repeated interaction among sellers which occurs over multiple departure dates across multiple markets.

5 Further Results

5.1 Refundability of Tickets

During advance sales airlines offer both nonrefundable and refundable tickets. They offer these two ticket types because some of the potential travelers that are contemplating buying in advance might still be uncertain about their valuations to travel. While a monopolist can wait until all travelers learn their valuations and charge the monopoly price in the spot market, Courty and Li (2000) show that more consumer surplus can be extracted by offering refundable and nonrefundable tickets that force travelers to reveal their private information sequentially. Akan et al. (2015) have a continuum of periods in which consumers learn their valuations instantaneously at multiple times, while in Ata and Dana (2015) consumers learn their valuations gradually. Escobari and Jindapon (2014) present a model to explain the gap between refundable and nonrefundable tickets and show how consumers learn about

[16] To rule out the existence of other focal points we tested for discontinuities at every hour, but found no statistically significant price discrimination beyond the already known $\overline{\mathrm{T}} = 7, 14$.

their individual demand.

[Table 7 (Refundable Tickets), about here]

Refundablility of a ticket can be viewed as insurance in case consumers learn they do not need to fly and want a refund. Hence, as Table 1 indicates refundable tickets are more expensive than nonrefundable tickets because consumers are paying a premium for the opportunity of canceling their trip. Table 7 presents the sharp regression-discontinuity estimates of price discrimination to assess the role of refundability—the dependent variable is the logarithm of the refundable fare and the running variable is TIME. We observe that across all specifications for each of the three cutoffs (T = 21, 14, 7), there is no statically significant price-discrimination through advance-purchase discounts for refundable tickets.

We explain the absence of price discrimination in advance-purchase discounts for refundable tickets by comparing key differences between refundable and nonrefundable tickets. For nonrefundable tickets price discrimination in advance-purchase discounts exist because consumers with different valuations buy at different times to departure (i.e., higher valuation consumers typically purchase closer to departure). Essentially advance-purchase discounts force low valuation consumers to buy nonrefundable tickets earlier at lower prices because high valuation consumers who learn about their valuations closer to departure push prices up (see Dana, 1998). This mechanism that forces consumers to buy earlier at lower nonrefundable prices because of the existence of high valuation consumers buying later no longer works for refundable tickets. First, consumers of refundable tickets who arrive later might not necessarily have higher valuations. Second, consumers of refundable tickets who arrive early in the season can always buy and request a refund later if they decide not to fly. Hence, their decision to buy does not hinge on the characteristics of consumers who arrive later.

5.2 Round-trip Tickets

A round-trip ticket is essentially the combination of two one-way tickets. While prior research that uses airline prices assumes that a one-way ticket price is half the round-trip price (see, e.g., Borenstein and Rose, 1994; and Gerardi and Shapiro, 2009), this assumption might not be reasonable when our intention is to identify price discrimination through advance-purchase discounts. The reason is simple, a round-trip ticket contains the combination of two different travel dates, and hence the price can potentially be affected by two different number of days-in-advance. In this section we test if price discrimination through advance-purchase discounts also exists in round-trip tickets and we assess what role if any the two different departure dates have on ticket prices.

The collection strategy of the round-trip fares data set was designed such that there are four days between the departure date of the outbound flight and the departure date of the inbound flight. Hence the cutoff points of 21, 14, and 7 days-in-advance are different for inbound and outbound flights. For example, when buying a ticket 5 days prior to the departure of the outbound corresponds to 9 days in advance of the inbound flight (due to the four days between flights). Hence this ticket falls into different sides of the 7 days-in-advance cutoff point. The idea is to test for the existence of discontinuities at 21, 17, 14, 10, 7, 3 days-in-advance of outbound flights where the cutoffs of 17, 10, and 3 correspond to 21, 14, and 7 days-in-advance of inbound flights.[17]

[Figure 4 (RD Plot: Round-trip Tickets), about here]

Figure 4 presents the mean of the logarithm of fare (LogFare) collapsed into bins along with the fourth order global polynomials estimated for each of the sub-samples separated by cutoffs. This figure suggests the existence of discontinuities in most of the cutoffs, with stronger evidence for the cutoffs that are closer to the departure date. The sharp regression-discontinuity estimates of price discrimination are presented in Table 8. The dependent variable is the logarithm of the nonrefundable round-trip fare and the running variable is Time. The results are consistent with the discrete jumps in Figure 4 showing statistically significant price discrimination for the cutoffs T = 17, 14, 10, 7, and 3.[18] Testing for potential discontinuities beyond the 21 days-in-advance and within the other known cutoffs yield no additional statistically significant discontinuities.

[17]While we focus on the 28 days before departure, we collected data for up to 60 days to departure. We found no evidence of discontinuities beyond 21 days-in-advance.

[18]The estimates in this table use a local linear ($p = 1$) regression with a quadratic ($q = 2$) bias-correction estimate and a triangular kernel. The bandwidth selection procedure and the robust variance estimators (computed with 3 nearest-neighbors) follow CTT. These results are robust to the kernel type selection and the bandwidth selection procedure.

[Table 8 (Round-trip Tickets), about here]

A simple sum of the point estimates of the discontinuities at different cutoffs would predict about a 69% increase in fares during the last three weeks to departure. This figure can be compared to the 85% increase obtained by simply calculating the differences in average prices during the same period. This simple calculation shows that most of the price increase during the last month to departure can be attributed to price discrimination.[19]

6 Conclusion

In this paper we identify price discrimination in advance purchase discounts. Our identification strategy uses high frequency posted prices and a regression discontinuity design which compares hourly prices just before and after the 21-, 14-, and 7-days-in-advance cutoffs. This empirical approach controls for both existing inventory levels and capacity costs. The article takes advantage of original data sets that contain one-way, round-trip, refundable and nonrefundable prices for economy-class tickets at each hour prior to the departure date with thousands of observations surrounding the cutoff points. The difficulty in identifying price discrimination in advance purchase discounts arises because the observed price dispersion as the flight date nears can be affected by cost changes that depend on demand expectations, time to departure and seat availability. We find evidence of statistically significant price discrimination with fares increasing by 7.6% at 14 days to departure, and by 14% at 7 days to departure.

The richness of the data allows us to address various questions related to price discrimination. We find that market structure significantly affects price discrimination becoming more prevalent in highly concentrated markets (i.e., higher HHI). In addition, the 7 days-in-advance cutoff estimates show that an increase in either the number of carriers or the number of flights serving the market will reduce price discrimination. On the other hand, for the 14 days-in-advance cutoff, price discrimination did not decrease with a larger number of carriers serving the market or with a higher volume of flights. This finding is consistent

[19] A final additional result involves looking at the identity of the carrier. Table A1 in the Appendix presents estimates using nonrefundable one-way fares. The results show that the effects are widespread across carriers.

with the behavior of sellers who are tacitly colluding. Further evidence also supports the hypothesis that sellers tacitly collude and jointly increase prices immediately after midnight (Pacific Standard Time) of the 14-, and 7-days-in-advance cutoffs. These one and two week cutoffs serve as focal points to help sellers coordinate and implement price discrimination in competitive markets.

The analysis of multiproduct monopolists shows that the magnitude of price discrimination is unaffected by an increase of product variety, as captured by monopoly sellers offering more flights in a market. We also examine refundable tickets and find no evidence of price discrimination for these more expensive tickets. Finally, we show that for round-trip tickets are also subject to price discrimination with pricing adjusting based on the number of days prior to departure for both outbound and inbound flights.

References

Akan, M., Ata, B., and Dana, Jr., J. D. (2015). Revenue management by sequential screening. *Journal of Economic Theory*, 159(Part B):728–774.

Alderighi, M., Gaggero, A. A., and Piga, C. A. (2016). The hidden side of dynamic pricing in airline markets. Working Paper. Presented at the International Industrial Organization Conference.

Alderighi, M., Nicolini, M., and Piga, C. (2015). Combined effects of capacity and time on fares: Insights from the yield management of a low-cost airline. *Review of Economics and Statistics*, 97(4):900–915.

Armstrong, M. (1996). Multiproduct nonlinear pricing. *Econometrica*, 64(1):51–75.

Armstrong, M. (1999). Price discrimination by a many-product firm. *Review of Economic Studies*, 66(1):151–168.

Armstrong, M. and Vickers, J. (2001). Competitive price discrimination. *Rand Journal of Economics*, 32(4):1–27.

Ata, B. and Dana, Jr., J. D. (2015). Price discrimination on booking time. *International Journal of Industrial Organization*, 43:175181.

Bacharach, M. and Bernasconi, M. (1997). The variable frame theory of focal points: An experimental study. *Games and Economic Behavior*, 19(1):1–45.

Berry, S. and Jia, P. (2010). Tracing the woes: An empirical analysis of the airline industry. *American Economic Journal: Microeconomics*, 2(3):1–43.

Bilotkach, V., Gaggero, A. A., and Piga, C. A. (2015). Airline pricing under different market conditions: Evidence from european low-cost carriers. *Tourism Management*, 47:152–163.

Bilotkach, V. and Rupp, N. (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, 2012, ed. James Peoples.

Borenstein, S. and Rose, N. L. (1994). Competition and price dispersion in the U.S. airline industry. *Journal of Political Economy*, 102(4):653–683.

Busse, M. and Rysman, M. (2005). Competition and price discrimination in Yellow Pages advertisement. *Rand Journal of Economics*, 36(2):378–390.

Cachon, G. P. (2004). The allocation of inventory risk in a supply chain: Push, pull, and advance-purchase discount contracts. *Management Science*, 50(2):222–238.

Calonico, S., Cattaneo, M. D., and Titunik, R. (2014a). Robust data-driven inference in the regression-discontinuity design. *Stata Journal*, 14(4):909–946.

Calonico, S., Cattaneo, M. D., and Titunik, R. (2014b). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.

Cattaneo, M., Malighetti, P., Morlotti, C., and Redondi, R. (2016). Quantity price discrimination in the air transport industry: The easyJet case. *Journal of Air Transport Management*, 54:1–8.

Clerides, S. K. (2002). Book value: Intertemporal pricing and quality discrimination in the us market for books. *International Journal of Industrial Organization*, 20(10):1385–1408.

Cohen, A. (2008). Package size and price discrimination in the paper towel market. *International Journal of Industrial Organization*, 26(2):502–516.

Courty, P. (2003). Ticket pricing under demand uncertainty. *Journal of Law and Economics*, 46(2):627–652.

Courty, P. and Li, H. (2000). Sequential screening. *Review of Economic Studies*, 67(4):697–717.

Crawford, G. and Shum, M. (2007). Monopoly quality degradation and regulation in cable television. *Journal of Law and Economics*, 50(1):181–219.

Dana, Jr., J. D. (1998). Advance-purchase discounts and price discrimination in competitive markets. *Journal of Political Economy*, 106(2):395–422.

Dana, Jr., J. D. (1999a). Equilibrium price dispersion under demand uncertainty: The roles of costly capacity and market structure. *Rand Journal of Economics*, 30(4):632–660.

Dana, Jr., J. D. (1999b). Using yield management to shift demand when the peak time is unknown. *Rand Journal of Economics*, 30(3):456–474.

Deneckere, R. and Peck, J. (2012). Dynamic competition with random demand and costless search: A theory of price posting. *Econometrica*, 80(3):1185–1247.

Engelmann, D. and Müller, W. (2011). Collusion through price ceilings? In search of a focal-point effect. *Journal of Economic Behavior & Organization*, 79:291–302.

Escobari, D. (2009). Systematic peak-load pricing, congestion premia and demand diverting: Empirical evidence. *Economics Letters*, 103(1):59–61.

Escobari, D. (2012). Dynamic pricing, advance sales, and aggregate demand learning in airlines. *Journal of Industrial Economics*, 60(4):697–724.

Escobari, D. and Jindapon, P. (2014). Price discrimination through refund contracts in airlines. *International Journal of Industrial Organization*, 34(3):1–8.

Frank, R. H. (1983). When are price differential discriminatory? *Journal of Policy Analysis and Management*, 2(2):238–255.

Gaggero, A. A. and Piga, C. A. (2011). Airline market power and intertemporal price dispersion. *Journal of Industrial Economics*, 59(4):552–577.

Gale, I. L. and Holmes, T. J. (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. L. and Holmes, T. J. (1993). Advance-purchase discounts and monopoly allocation of capacity. *American Economic Review*, 83(1):135–146.

Gallego, G. and van Ryzin, G. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8):999–1020.

Gerardi, K. and Shapiro, A. (2009). Does competition reduce price dispersion? New evidence from the airline industry. *Journal of Political Economy*, 117(1):1–37.

Hahn, J., Todd, P., and Van der Klaauw, W. (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69(1):201–209.

Heckman, J. J. and Vytlacil, E. J. (2007). *Econometric evaluation of social programs, Part I: Causal models, structural models and econometric policy evaluation,* ed. J. Heckman and E. Leamer, volume VI of *Handbook of Econometrics*, pages 4780–4874. Elsevier Science B.V.

Hernandez, M. and Wiggins, S. (2014). Nonlinear pricing strategies and competitive conditions in the airline industry. *Economic Inquiry*, 52(2):539–561.

Holmes, T. J. (1989). The effects of third-degree price discrimination in oligopoly. *American Economic Review*, 79(1):244–250.

Imbens, G. W. and Kalyanaraman, K. (2012). Optimal bandwidth choice for the regression discontinuity estimator. *Review of Economic Studies*, 79(3):933–959.

Imbens, G. W. and Wooldridge, J. M. (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, 47(1):5–86.

Ivaldi, M. and Martimort, D. (1994). Competition under nonlinear pricing. *Annales d'Économie et de Statistique*, 34:71–114.

Knittel, C. R. and Stango, V. (2003). Price ceilings as focal points for tacit collusion: Evidence from credit cards. *American Economic Review*, 93(5):1703–1729.

Lee, D. S. (2008). Randomized experiments from non-random selection in U.S. House elections. *Journal of Econometrics*, 142(2):675–697.

Lee, D. S. and Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of Economic Literature*, 48(2):281–355.

Leslie, P. (2004). Price discrimination in Broadway theater. *Rand Journal of Economics*, 35(3):520–541.

McManus, B. (2007). Nonlinear pricing in an oligopoly market: The case of specialty coffee. *Rand Journal of Economics*, 38(2):512–532.

Mehta, J., Starmer, C., and Sugden, R. (1994a). Focal points in pure coordination games: An experimental investigation. *Theory and Decision*, 36(2):163–185.

Mehta, J., Starmer, C., and Sugden, R. (1994b). The nature of salience: An experimental investigation of pure coordination games. *American Economic Review*, 84(3):658–673.

Möller, M. and Watanabe, M. (2010). Advance purchase discounts versus clearance sales. *Economic Journal*, 120(547):1125–1148.

Möller, M. and Watanabe, M. (2016). Market structure and advance selling. Tinbergen Institute Discussion Paper, TI 2016-020/VII.

Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, 18(2):301–317.

Nocke, V., Peitz, M., and Rosar, F. (2011). Advance-purchase discounts as a price discrimination device. *Journal of Economic Theory*, 146(1):141–162.

Plott, C. R. (1982). Industrial organization theory and experimental economics. *Journal of Economic Literature*, 20(4):1485–1527.

Pop-Eleches, C. and Urquiola, M. (2013). Going to a better school: Effects and behavioral responses. *American Economic Review*, 103(4):1289–1324.

Porter, J. (2003). Estimation in the regression discontinuity model. Working Paper. Harvard University.

Schelling, T. (1960). *The strategy of conflict.* Harvard University Press, Cambridge, MA.

Schmalensee, R. (1981). Output and welfare implications on monopolistic third-degree price discrimination. *American Economic Review*, 71(2):242–247.

Shepard, A. (1991). Price discrimination and retail configuration. *Journal of Political Economy*, 99(1):30–53.

Stigler, G. J. (1987). *The theory of price.* Macmillan, New York, NY.

Stokey, N. L. (1979). Intertemporal price discrimination. *The Quarterly Journal of Economics*, 93(3):355–371.

Stole, L. A. (2007). Price discrimination and competition. In *Handbook of Industrial Organization*, volume 3. Elsevier.

Su, X. (2007). Intertemporal pricing with strategic customer behavior. *Management Science*, 53(5):726–741.

Sugden, R. (1995). A theory of focal points. *Economic Journal*, 105(430):533–550.

Varian, H. R. (1989). Price discrimination and imperfect competition. In *Handbook of Industrial Organization.* North-Holland.

Williams, K. R. (2017). Dynamic airline pricing and seat availability. Cowles Foundation Discussion Paper No. 3003. Yale University.

Zhao, W. and Zheng, Y.-S. (2000). Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management Science*, 46(3):375–388.

Figure 1: American Airlines, SEA-ORD, Flight 1152, Boeing 737-800

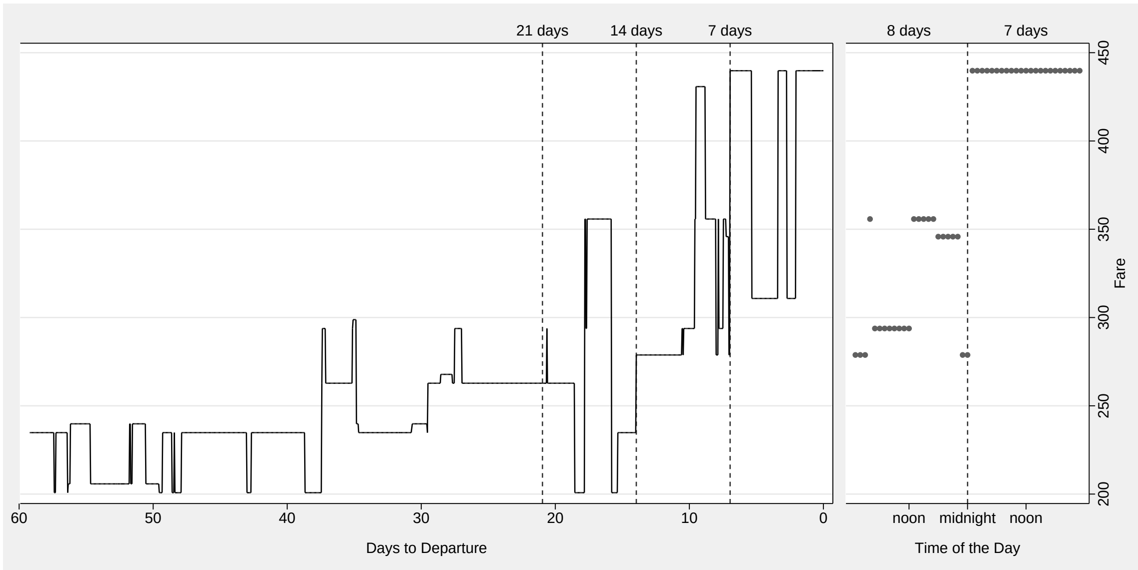

Notes: This figure shows the path of hourly prices at different times to departure for the American Airlines flight 1152 between the Seattle-Tacoma International Airport (SEA) and the Chicago O'Hare International Airport (ORD). The right-hand side of the figure zooms in to the 7 day-in-advance threshold to illustrate how fares jump right after midnight (Pacific Standard Time).

Figure 2: Regression Discontinuity Plot: 7 Days to Departure

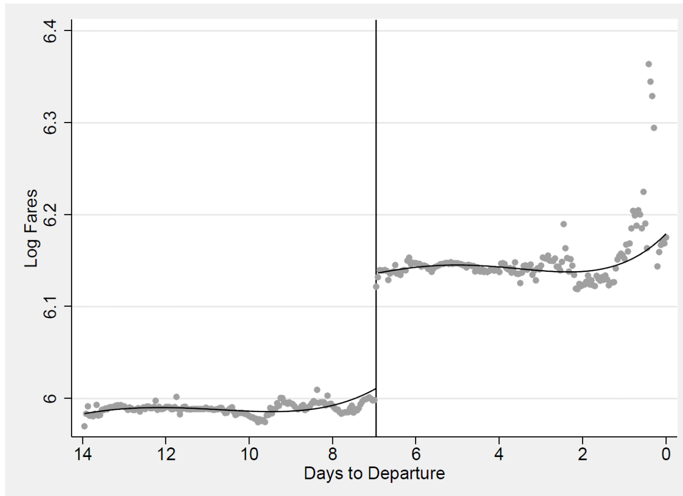

Notes: The figure shows the sample average within bin along the fourth order global polynomial estimated separately on each side of the cutoff of 7-days-in-advance. The dependent variable is logarithm of non-stop one-way domestic economy fares (LogFare). 511,252 observations.

Figure 3: Regression Discontinuity Plot: 14 Days to Departure

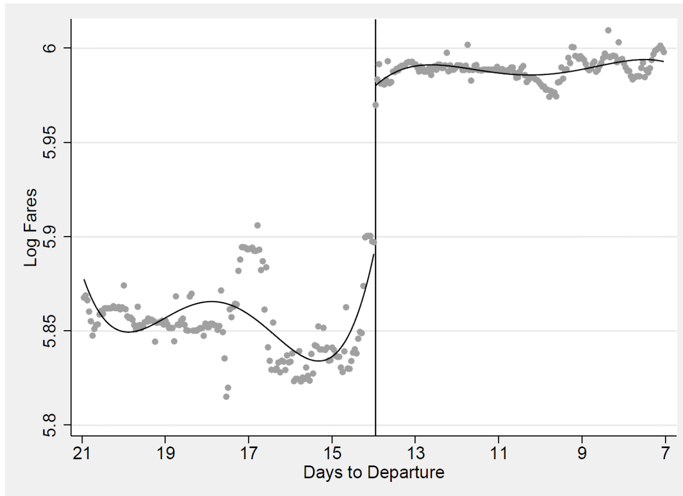

Notes: The figure shows the sample average within bin along the fourth order global polynomial estimated separately on each side of the cutoff of 14-days-in-advance. The dependent variable is logarithm of non-stop one-way domestic economy fares (LogFare). 513,695 observations.

Figure 4: Regression Discontinuity Plot: Round Trip Tickets

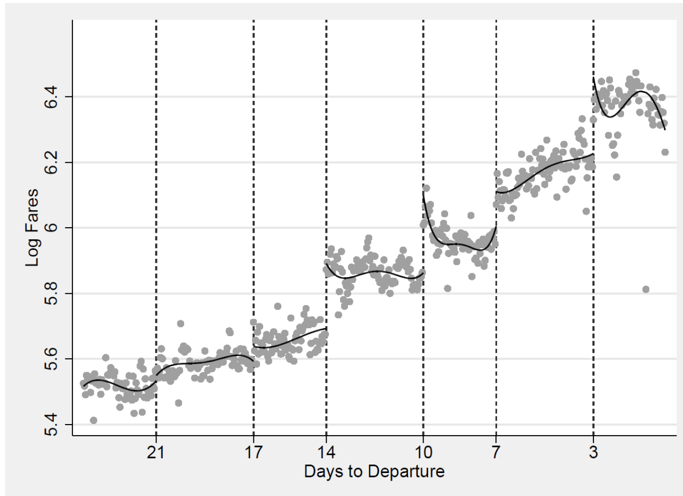

Notes: The figure shows the sample average within bin along the fourth order global polynomial estimated separately between the different cutoffs. The dependent variable is logarithm of non-refundable roundtrip domestic economy fares (LogFare). 1,514,833 observations.

Table 1: Summary Statistics

VARIABLES	mean	sd	min	max	obs
	(1)	(2)	(3)	(4)	(5)
Panel A. Refundable and Non-refundable One-way Fares:					
Fare (nonrefundable):	385.1	200.0	58.30	3,019	989,101
$\text{Time} \leq 7$	449.4	194.9	96.80	3,019	253,547
$7 < \text{Time} \leq 14$	396.3	189.9	94.80	1,619	257,705
$14 < \text{Time} \leq 21$	350.9	197.1	59.80	3,019	255,990
$21 < \text{Time} \leq 28$	352.2	202.1	58.30	1,181	221,859
Fare (refundable)	557.8	208.1	180.8	1,696	919,582
Time	14.90	7.838	0	28	989,101
American	0.254	0.435	0	1	989,101
Alska	0.0168	0.129	0	1	989,101
JetBlue	0.0239	0.153	0	1	989,101
Delta	0.189	0.392	0	1	989,101
Forntier	0.00113	0.0336	0	1	989,101
AirTran	0.0370	0.189	0	1	989,101
United	0.324	0.468	0	1	989,101
US Airways	0.143	0.350	0	1	989,101
Virgin Amer.	0.0106	0.103	0	1	989,101
# Flights in a Route	19.11	11.83	1	46	989,101
# Carriers in a Route	2.277	1.050	1	6	989,101
# Own Flights in a Route	8.954	4.778	1	23	989,101
Share Carrier in a Route	0.580	0.279	0.100	1	989,101
HHI	0.582	0.256	0.179	1	989,101
Panel B. Nonrefundable Round-trip Fares:					
Fare (nonrefundable):	337.2	179.6	38.00	2,070	1,514,833
$\text{Time} \leq 7$	505.5	241.4	58.00	2,070	371,970
$7 < \text{Time} \leq 14$	352.8	135.9	58.00	1,163	399,747
$14 < \text{Time} \leq 21$	267.4	98.00	58.00	1,525	356,927
$21 < \text{Time} \leq 28$	247.1	80.54	38.00	964.2	386,189

Notes: The sample contains 1,908,683 one-way economy-class tickets (919,582 refundable and 989,101 non-refundable) and 1,514,833 round-trip non-refundable tickets. There are 1,665 domestic flights across 158 domestic routes.

Table 2: RD Estimates: Monopoly Routes

BW Type:	CCT	IK	CV	CCT	CCT
VARIABLES	(1)	(2)	(3)	(4)	(5)
Panel A. Price Discrimination at 7 Days to Departure ($\overline{\text{T}}$=7):					
$\widehat{\text{PD}}_7$	0.140*	0.145*	0.148*	0.142*	0.139*
Robust 95% CI	[.11 ; .16]	[.11 ; .16]	[.13 ; .16]	[.12 ; .17]	[.11 ; .16]
Robust p-value	0	0	0	0	0
BW Loc. Poly. (h)	52.47	110.3	166	42.28	81.77
BW Bias (b)	80.98	94.31	166	72.76	109.9
Observations	36,421	76,066	108,626	29,482	56,478
Panel B. Price Discrimination at 14 Days to Departure ($\overline{\text{T}}$=14):					
$\widehat{\text{PD}}_{14}$	0.0755*	0.0920*	0.0868*	0.0759*	0.0783*
Robust 95% CI	[.03 ; .11]	[.05 ; .11]	[.01 ; .1]	[.03 ; .11]	[.04 ; .11]
Robust p-value	0.000464	1.69e-06	0.00915	0.000647	0.000162
BW Loc. Poly. (h)	25.76	37.34	33.40	20.20	54.16
BW Bias (b)	49.43	59.09	33.40	41.96	74.89
Observations	17,812	26,179	23,388	14,303	38,029
Panel C. Price Discrimination at 21 Days to Departure ($\overline{\text{T}}$=21):					
$\widehat{\text{PD}}_{21}$	0.00206	-0.00879	-0.00815	0.00391	0.00336
Robust 95% CI	[-.03 ; .04]	[-.02 ; .04]	[-.03 ; .01]	[-.03 ; .04]	[-.04 ; .05]
Robust p-value	0.817	0.536	0.543	0.731	0.851
BW Loc. Poly. (h)	42.49	138.7	167	31.17	60.12
BW Bias (b)	64.71	112.3	167	54.96	79.28
Observations	29,717	96,524	116,078	22,026	42,310
Kernel Type	Triangular	Triangular	Triangular	Uniform	Triangular
Order Loc. Poly. (p)	1	1	1	1	2
Order Bias (q)	2	2	2	2	3

Notes: $\widehat{\text{PD}}_{\overline{\text{T}}}$ for $\overline{\text{T}} = 7, 14, 21$ are sharp regression-discontinuity estimates of price discrimination for one-way economy class tickets with Time measured in days as the running variable. The estimates use a local-polynomial ($p = 1, 2$) regression with a quadratic or cubic ($q = 2, 3$) bias-correction estimate and a uniform or a triangular kernel. The bandwidth (h) is measured in minutes and its selection procedure is Cross validation (CV) or the one proposed by Calonico, Cattaneo and Titiunik (CCT) or Imbens and Kalyanaraman (IK). The robust variance estimators are the ones proposed by CCT, computed with 3 nearest-neighbors. * significant at the 1 percent level; † significant at the 5 percent level; ‡ significant at the 10 percent level.

Table 3: RD Estimates: Multiproduct Monopolists

Number of Flights:	Flights< 4	4 ≤Flights< 7	7 ≤Flights
VARIABLES	(1)	(2)	(3)
Panel A. Price Discrimination at 7 Days to Departure ($\overline{\text{T}}$=7):			
$\widehat{\text{PD}}_7$	0.141*	0.149*	0.135*
Robust 95% CI	[.09 ; .2]	[.1 ; .19]	[.1 ; .17]
Robust p-value	3.55e-07	1.81e-10	0
BW Loc. Poly. (h)	64.33	55.97	47.33
BW Bias (b)	97.32	87.36	71.44
Observations	9,092	13,476	14,718
Panel B. Price Discrimination at 14 Days to Departure ($\overline{\text{T}}$=14):			
$\widehat{\text{PD}}_{14}$	0.0824‡	0.0934*	0.0765*
Robust 95% CI	[-.01 ; .15]	[.04 ; .14]	[.02 ; .11]
Robust p-value	0.0923	0.000157	0.00467
BW Loc. Poly. (h)	38.81	68.96	24.51
BW Bias (b)	66.16	107.3	48.91
Observations	5,410	16,851	7,626

Notes: $\widehat{\text{PD}}_{\overline{\text{T}}}$ for $\overline{\text{T}} = 7, 14$ are sharp regression-discontinuity estimates of price discrimination with Time measured in days as the running variable. The estimates use a local-linear ($p = 1$) regression with a quadratic ($q = 2$) bias-correction estimate and a triangular kernel. The bandwidth (h) is measured in minutes and its selection procedure is the one proposed by CCT. The robust variance estimators are the ones proposed by CCT, computed with 3 nearest-neighbors. * significant at the 1 percent level; † significant at the 5 percent level; ‡ significant at the 10 percent level.

Table 4: RD Estimates: Market Structure

	All Routes	Below Percentile (HHI < HHI_{pc})			Above Percentile (HHI > HHI_{pc})		
Above/Below Percentile:		$pc = 33^{th}$	$pc = 40^{th}$	$pc = 50^{th}$	$pc = 67^{th}$	$pc = 60^{th}$	$pc = 50^{th}$
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Price Discrimination at 7 Days to Departure ($\overline{\text{T}}$=7):							
$\widehat{\text{PD}}_7$	0.115*	0.0552*	0.0722*	0.0897*	0.141*	0.128*	0.140*
Robust 95% CI	[.09 ; .13]	[.03 ; .08]	[.04 ; .09]	[.06 ; .11]	[.12 ; .16]	[.1 ; .15]	[.11 ; .16]
Robust p-value	0	9.16e-05	6.93e-08	0	0	0	0
BW Loc. Poly. (h)	16.96	19.38	18.77	18.92	55.44	30.01	25.38
BW Bias (b)	34.06	32.17	32.76	33.58	83.91	53.49	49.53
Observations	49,104	19,477	21,717	26,364	57,317	38,162	39,606
Panel B. Price Discrimination at 14 Days to Departure ($\overline{\text{T}}$=14):							
$\widehat{\text{PD}}_{14}$	0.0926*	0.0761*	0.0747*	0.0724*	0.133*	0.125*	0.114*
Robust 95% CI	[.07 ; .11]	[.05 ; .1]	[.04 ; .09]	[.04 ; .09]	[.09 ; .16]	[.09 ; .15]	[.08 ; .14]
Robust p-value	0	5.76e-08	7.43e-08	9.01e-09	0	0	0
BW Loc. Poly. (h)	20.21	23.13	21.49	22.19	24.02	25.67	25.59
BW Bias (b)	43.36	47.02	45.02	45.40	47.55	48.83	47.38
Observations	62,879	25,085	26,847	33,922	25,470	32,130	39,796

Notes: HHI is the Herfindahl-Hirschman Index with HHI_{pc} being percentile pc of the HHI. $\widehat{\text{PD}}_{\overline{\text{T}}}$ for $\overline{\text{T}} = 7, 14$ are sharp regression-discontinuity estimates of price discrimination with Time measured in days as the running variable. The estimates use a local-linear ($p = 1$) regression with a quadratic ($q = 2$) bias-correction estimate and a triangular kernel. The bandwidth (h) is measured in minutes and its selection procedure is the one proposed by CCT. The robust variance estimators are the ones proposed by CCT, computed with 3 nearest-neighbors. * significant at the 1 percent level; † significant at the 5 percent level; ‡ significant at the 10 percent level.

Table 5: RD Estimates: Number of Carriers in a Route

Number of Carriers:	One	Two	Three	Four	Six
VARIABLES	(1)	(2)	(3)	(4)	(5)
Panel A. Price Discrimination at 7 Days to Departure ($\overline{\text{T}}$=7):					
$\widehat{\text{PD}}_7$	0.140*	0.156*	0.0622*	0.0532*	0.0450*
Robust 95% CI	[.11 ; .16]	[.12 ; .18]	[.03 ; .09]	[.03 ; .08]	[.03 ; .05]
Robust p-value	0	0	0.000261	1.82e-05	0
BW Loc. Poly. (h)	52.47	23.08	23.36	20.44	13.91
BW Bias (b)	80.98	46.23	40.90	31.03	30.07
Observations	36,421	29,295	16,977	5,097	918
Panel B. Price Discrimination at 14 Days to Departure ($\overline{\text{T}}$=14):					
$\widehat{\text{PD}}_{14}$	0.0755*	0.119*	0.0873*	0.0953*	0.115*
Robust 95% CI	[.03 ; .11]	[.08 ; .14]	[.06 ; .11]	[.05 ; .12]	[.07 ; .16]
Robust p-value	0.000464	0	0	5.72e-07	7.68e-07
BW Loc. Poly. (h)	25.76	30.17	49.64	18.28	54.07
BW Bias (b)	49.43	51.14	83.26	42.62	80.16
Observations	17,812	38,509	38,809	4,718	3,853

Notes: $\widehat{\text{PD}}_{\overline{\text{T}}}$ for $\overline{\text{T}} = 7, 14$ are sharp regression-discontinuity estimates of price discrimination with Time measured in days as the running variable. The estimates use a local-linear ($p = 1$) regression with a quadratic ($q = 2$) bias-correction estimate and a triangular kernel. The bandwidth (h) is measured in minutes and its selection procedure is the one proposed by CCT. The robust variance estimators are the ones proposed by CCT, computed with 3 nearest-neighbors. * significant at the 1 percent level; † significant at the 5 percent level; ‡ significant at the 10 percent level.

Table 6: RD Estimates: Number of Flights in a Route

Number of Flights:	Flights< 10	10 ≤Flights< 20	20 ≤Flights< 30	30 ≤Flights< 40	40 ≤Flights
VARIABLES	(1)	(2)	(3)	(4)	(5)
Panel A. Price Discrimination at 7 Days to Departure ($\overline{\mathrm{T}}$=7):					
$\widehat{\mathrm{PD}}_7$	0.156*	0.145*	0.130*	0.0648*	0.000990
Robust 95% CI	[.13 ; .17]	[.12 ; .17]	[.1 ; .15]	[.02 ; .1]	[-.04 ; .04]
Robust p-value	0	0	0	0.00744	0.918
BW Loc. Poly. (h)	54.93	47.17	18.54	24.68	22.38
BW Bias (b)	91.77	73.13	34.95	47.84	34.53
Observations	44,550	39,472	13,035	9,996	4,896
Panel B. Price Discrimination at 14 Days to Departure ($\overline{\mathrm{T}}$=14):					
$\widehat{\mathrm{PD}}_{14}$	0.0887*	0.0775*	0.145*	0.0799*	0.0811*
Robust 95% CI	[.05 ; .12]	[.04 ; .11]	[.11 ; .17]	[.04 ; .11]	[.04 ; .11]
Robust p-value	3.38e-06	0.000106	0	8.98e-06	0.000113
BW Loc. Poly. (h)	26.27	29.20	30.30	52.70	20.82
BW Bias (b)	48.96	51.98	55.41	79.85	44.48
Observations	21,872	24,909	21,885	22,082	5,346

Notes: $\widehat{\mathrm{PD}}_{\overline{\mathrm{T}}}$ for $\overline{\mathrm{T}} = 7, 14$ are sharp regression-discontinuity estimates of price discrimination with Time measured in days as the running variable. The estimates use a local-linear ($p = 1$) regression with a quadratic ($q = 2$) bias-correction estimate and a triangular kernel. The bandwidth (h) is measured in minutes and its selection procedure is the one proposed by CCT. The robust variance estimators are the ones proposed by CCT, computed with 3 nearest-neighbors. * significant at the 1 percent level; † significant at the 5 percent level; ‡ significant at the 10 percent level.

Table 7: RD Estimates: Monopoly Routes (Refundable Tickets)

BW Type:	CCT	IK	CV	CCT	CCT
VARIABLES	(1)	(2)	(3)	(4)	(5)
Panel A. Price Discrimination at 7 Days to Departure ($\overline{\text{T}}$=7):					
$\widehat{\text{PD}}_7$	0.00937	0.00983	0.00870	0.00899	0.00848
Robust 95% CI	[-.01 ; .03]	[-.01 ; .03]	[0 ; .03]	[-.01 ; .03]	[-.02 ; .03]
Robust p-value	0.378	0.221	0.158	0.378	0.538
BW Loc. Poly. (h)	57.66	105.1	132.8	47.08	66.18
BW Bias (b)	87.23	106.6	132.8	80.33	87.53
Observations	36,636	66,800	83,315	30,285	42,385
Panel B. Price Discrimination at 14 Days to Departure ($\overline{\text{T}}$=14):					
$\widehat{\text{PD}}_{14}$	0.00505	0.00327	0.000707	0.00524	0.00402
Robust 95% CI	[-.01 ; .03]	[-.01 ; .03]	[-.01 ; .02]	[-.01 ; .03]	[-.02 ; .03]
Robust p-value	0.485	0.442	0.644	0.421	0.861
BW Loc. Poly. (h)	64.33	98.51	167	48.66	72.11
BW Bias (b)	100.5	101.3	167	91.24	97.22
Observations	41,190	62,684	105,907	30,960	46,280
Panel C. Price Discrimination at 21 Days to Departure ($\overline{\text{T}}$=21):					
$\widehat{\text{PD}}_{21}$	-0.000818	-0.000666	-0.00130	-0.000182	6.57e-05
Robust 95% CI	[-.02 ; .02]	[-.02 ; .02]	[-.01 ; .01]	[-.02 ; .02]	[-.02 ; .02]
Robust p-value	0.931	0.955	0.923	0.972	0.967
BW Loc. Poly. (h)	70.48	102.4	167	56.84	82.62
BW Bias (b)	106.1	97.33	167	97.90	109.1
Observations	44,691	64,587	105,302	35,850	52,149
Kernel Type	Triangular	Triangular	Triangular	Uniform	Triangular
Order Loc. Poly. (p)	1	1	1	1	2
Order Bias (q)	2	2	2	2	3

Notes: $\widehat{\text{PD}}_{\overline{\text{T}}}$ for $\overline{\text{T}} = 7, 14, 21$ are sharp regression-discontinuity estimates of price discrimination with Time measured in days as the running variable. The estimates use a local-polynomial ($p = 1, 2$) regression with a quadratic or cubic ($q = 2, 3$) bias-correction estimate and a uniform or a triangular kernel. The bandwidth (h) is measured in minutes and its selection procedure is Cross validation (CV) or the one proposed by Calonico, Cattaneo and Titiunik (CCT) or Imbens and Kalyanaraman (IK). The robust variance estimators are the ones proposed by CCT, computed with 3 nearest-neighbors. * significant at the 1 percent level; † significant at the 5 percent level; ‡ significant at the 10 percent level.

Table 8: RD Estimates: Round Trip Non-refundable Economy Class Tickets

	$\overline{\text{T}} = 3$	$\overline{\text{T}} = 7$	$\overline{\text{T}} = 10$	$\overline{\text{T}} = 14$	$\overline{\text{T}} = 17$	$\overline{\text{T}} = 21$
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
$\widehat{\text{PD}}_{\text{T}}$	0.213*	0.106*	0.127*	0.175*	0.0692*	-0.00484
Robust 95% CI	[.2 ; .24]	[.09 ; .13]	[.09 ; .15]	[.15 ; .2]	[.06 ; .09]	[-.03 ; .01]
Robust p-value	0	0	0	0	0	0.236
BW Loc. Poly. (h)	20.60	15.09	3.624	6.251	8.165	7.508
BW Bias (b)	30.62	30.08	10.37	10.31	17.37	21.02
Observations	80,727	58,292	15,310	29,156	36,957	30,569

Notes: $\widehat{\text{PD}}_{\overline{\text{T}}}$ are sharp regression-discontinuity estimates of price discrimination with Time measured in days as the running variable. The estimates use a local-linear ($p = 1$) regression with a quadratic ($q = 2$) bias-correction estimate and a triangular kernel. The bandwidth (h) is measured in minutes and its selection procedure is the one proposed by Calonico, Cattaneo and Titiunik (CCT). The robust variance estimators are the ones proposed by CCT, computed with 3 nearest-neighbors. * significant at the 1 percent level; † significant at the 5 percent level; ‡ significant at the 10 percent level.

Table A1: RD Estimates: Carrier Identity

Carrier:	American	Alaska	JetBlue	Delta	Frontier	AirTran	United	US Airways	Virgin Amer.
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A. Price Discrimination at 7 Days to Departure ($\overline{\mathrm{T}}$=7):									
$\widehat{\mathrm{PD}}_7$	0.135*	0.209*	0.249*	0.170*	0.195*	0.148*	0.106*	0.0473†	0.111*
Robust 95% CI	[.11 ; .15]	[.1 ; .29]	[.18 ; .3]	[.13 ; .2]	[.04 ; .3]	[.1 ; .21]	[.08 ; .13]	[.01 ; .08]	[.08 ; .13]
Robust p-value	0	9.29e-05	0	0	0.00872	1.87e-08	0	0.0183	0
BW Loc. Poly. (h)	38.48	28.58	31.58	20.07	28.93	30.12	56.34	23.16	15.13
BW Bias (b)	66.24	50.74	57.15	31.79	56.49	51.50	85.11	39.32	31.10
Observations	29,151	1,421	2,267	11,783	80	3,477	54,843	9,572	496
Panel B. Price Discrimination at 14 Days to Departure ($\overline{\mathrm{T}}$=14):									
$\widehat{\mathrm{PD}}_{14}$	0.0932*	0.0300	0.130*	0.187*	0.127*	0.326*	0.0480*	0.0282	0.174*
Robust 95% CI	[.05 ; .12]	[-.09 ; .17]	[.05 ; .18]	[.15 ; .22]	[.04 ; .19]	[.29 ; .37]	[.01 ; .07]	[-.01 ; .05]	[.14 ; .2]
Robust p-value	1.56e-05	0.512	0.000395	0	0.00419	0	0.00545	0.164	0
BW Loc. Poly. (h)	19.96	46.09	42.94	39.02	55.48	47.78	30.81	33.73	23.13
BW Bias (b)	42.33	72.74	74.27	65.91	86.42	82.93	53.51	58.21	49.04
Observations	15,240	2,363	3,052	22,703	222	5,399	30,209	15,029	750
Panel C. Price Discrimination at 21 Days to Departure ($\overline{\mathrm{T}}$=21):									
$\widehat{\mathrm{PD}}_{21}$	0.0138	0.0106	0.0394	0.00400	-0.00842	0.120*	0.00788	-0.00578	0.0517*
Robust 95% CI	[-.02 ; .04]	[-.09 ; .11]	[-.04 ; .1]	[-.04 ; .04]	[-.05 ; .03]	[.07 ; .18]	[-.02 ; .04]	[-.04 ; .02]	[.01 ; .08]
Robust p-value	0.466	0.803	0.369	0.994	0.558	6.26e-06	0.667	0.468	0.00514
BW Loc. Poly. (h)	44.98	64.32	62.86	30.60	55.40	24.81	43.80	38.41	45.91
BW Bias (b)	68.68	96.48	101.3	47.81	92	40.47	66.56	66.87	80.72
Observations	35,099	3,336	4,483	17,318	222	2,787	43,238	17,638	1,456

Notes: $\widehat{\mathrm{PD}}_{\overline{\mathrm{T}}}$ for $\overline{\mathrm{T}} = 7, 14, 21$ are sharp regression-discontinuity estimates of price discrimination with Time measured in days as the running variable. The estimates use a local-linear ($p = 1$) regression with a quadratic ($q = 2$) bias-correction estimate and a triangular kernel. The bandwidth (h) is measured in minutes and its selection procedure is the one proposed by CCT. The robust variance estimators are the ones proposed by CCT, computed with 3 nearest-neighbors. * significant at the 1 percent level; † significant at the 5 percent level; ‡ significant at the 10 percent level.