Price Discrimination through Refund Contracts in Airlines

Escobari, Diego and Jindapon, Paan

The University of Texas - Pan American, University of Alabama

10 February 2014

Price Discrimination through Refund Contracts in Airlines*

(*International Journal of Industrial Organization*, Forthcoming)

Diego Escobari† Paan Jindapon‡

February 10, 2014

Abstract

This paper shows how an airline monopoly uses refundable and non-refundable tickets to screen consumers who are uncertain about their travel. Our theoretical model predicts that the difference between these two fares diminishes as individual demand uncertainty is resolved. Using an original data set from U.S. airline markets, we find strong evidence supporting our model. Price discrimination opportunities through refund contracts decline as the departure date nears and individuals learn about their demand.

Keywords: Price discrimination; Refund contracts; Airlines; Individual demand learning

JEL Classifications: C23; D42; D82; L93

*We thank James Dana, Li Gan, James Haag, Monica Hartmann, Tim Hazledine, Junsoo Lee, Sang-Ho Lee, Bradford Patterson, Harris Schlesinger, Roger Sherman, Joseph Walsh, and seminar participants at CSU Fullerton, East Carolina University, University of Alabama, University of Chile, University of Louisville, University of Texas-Pan American, IIOC, SEA, AEA, and WEAI for helpful comments. Thanks in particular to the editor and three anonymous referees, whose comments helped improve the paper. Jindapon thanks the Culverhouse College of Commerce and Business Administration for financial support and Escobari thanks the Bradley Foundation for financial support. The data set collection was funded by the Private Enterprise Research Center at Texas A&M. Stephanie C. Reynolds provided excellent research assistance.

†Department of Economics and Finance, University of Texas-Pan American, escobarida@utpa.edu

‡Department of Economics, Finance and Legal Studies, University of Alabama, pjindapo@cba.ua.edu

1 Introduction

Consider a potential traveler who is planning to buy a plane ticket in advance. However, at this moment, he is not certain whether he will travel. Let his valuation for traveling be v. If the airline offers him a refundable ticket now, he will be willing to pay v for that ticket. But if the airline offers him a non-refundable ticket, his willingness to pay should be less than v.[1] Once he knows with certainty whether he wants to travel, there is no reason why he should be willing to pay more for a refundable ticket than for a non-refundable ticket.

This paper presents a theory that explains how a monopolist can use refundable and non-refundable tickets to screen consumers and extract more surplus when consumers have to select a contract before knowing their demand with certainty.[2] In the simple two-period model, the airline will offer both ticket types in advance to consumers who are uncertain about their demand and have different willingness to pay. We derive optimal refundable and non-refundable fares that depend on each consumer's willingness to pay and the probability of travel. Consumers with high willingness to pay buy refundable tickets, and consumers with low willingness to pay buy non-refundable tickets. Furthermore, we find that the difference between these two fares consists of a quality component—the refundability value—and a price-discrimination component. A comparative-statics analysis provides an empirical implication of the model: the gap between the two fares diminishes as the date of departure approaches and consumers become more certain about their individual demand. Therefore, the airline's ability to separate consumer types and to price discriminate vanishes.

In the empirical section we test the main empirical implication from the theory. We collected from the online travel agency Expedia.com an original panel data set of refundable fares, non-refundable fares, and seat inventories across 96 U.S. domestic monopoly routes at various days prior to the departure date. The data collection focuses on posted one-way economy-class fares to control for price differentials associated with other ticket

[1] This ex-ante willingness to pay for a non-refundable ticket is analogous to the option price in Cicchetti and Freeman (1971). Schmalensee (1972) and Graham (1981) show that the option price may be greater or smaller than the expected willingness to pay.

[2] Even though we focus on airlines, the results can also be applied to other industries where goods are sold in advance with a refundable option, such as cruises, car rentals, and lodging.

restrictions; for example, Saturday-night stayover, minimum and maximum stay, first-class travel, and connecting legs. These restrictions are commonly used as 'fences' to implement other forms of price discrimination and congestion pricing and to deal with aggregate demand uncertainty. The panel structure of the data controls for unobservable time-invariant carrier-, route-, and flight-specific characteristics. Moreover, the fact that both fares were collected at the same time for the same seat allows us to control for unobserved time-variant seat-specific characteristics. The estimation method, which takes into account the dynamic adjustment between the difference in fares and seat inventories, shows strong evidence that price discrimination through refund contracts vanishes as the departure date nears. In addition, a nonparametric specification indicates that most of the individual demand uncertainty, as implied by the carriers' pricing strategy, is resolved during the last two weeks prior to the departure date.[3]

In the literature, Gale and Holmes (1993) and Dana (1998) use advanced-purchase discounts as a means of price discrimination to improve capacity utilization in monopolistic and competitive markets respectively. In contrast, we show that an airline monopoly can use a refundability option to screen consumers and increase the airline's expected profit. Courty and Li (2000) suggest a theoretical model for a monopolist that price discriminates via refund contracts consisting of a price a buyer has to pay in advance and a refund the buyer can receive after he learns his valuation of the good. While Courty and Li's purpose is to find an optimal refund contract consisting of an advance payment and a refundable amount, our goal is to find an optimal contract consisting of a non-refundable price and a refundable price for each type. In a related work, Akan et al. (2011) present a generalization of Courty and Li (2000) with consumers who learn their valuations gradually and with a seller that can vary the length of time during which the tickets are refundable. Bilotkach (2009) presents a model explaining refund contracts under costly capacity and demand uncertainty.

The rest of the paper is structured as follows. Section 2 presents the theoretical model,

[3]There are many ways that uncertainty could be resolved. In this paper, by resolved uncertainty, we mean an increase in the probability that a ticket holder wants to travel on a particular date as the trip date nears. These results, in addition, help explain the large price dispersion in airlines documented in Borenstein and Rose (1994) and more recently in Gerardi and Shapiro (2009).

including the consumer's problem and the airline's problem, and describes the airline's price menu in equilibrium. Section 3 presents the empirical analysis by first describing the data, then setting the empirical model, and closing with the results. Section 4 concludes.

2 Theoretical Analysis

Consider a monopolistic airline that sells homogeneous seats on a flight to consumers whose type, high (H) or low (L), is not observable to the airline. In airline markets, it is common to see consumers buy tickets in advance despite uncertainty in their travel plans. In our two-period model, consumers have unit demands and have to decide whether to buy a ticket in period 1. However, they learn in period 2 whether they want to fly or not (i.e., demand equals 1 or 0). Let travel and no travel be mutually exclusive states of nature in which a consumer wants and does not want to travel respectively. The risk in the state of nature that each consumer faces is an individual risk that is independent from those of other consumers. For $\theta = H, L$, we let π_θ denote the probability that type θ consumer wants to travel and hence $1 - \pi_\theta$ the probability that the consumer does not want to travel. Both types have a positive valuation of traveling v_θ and share the same utility function u with $u' > 0$ and $u'' \leq 0$. We normalize u so that $u(0) = 0$.

2.1 The Consumer's Problem

In period 1, the airline offers refundable and non-refundable tickets to all consumers. If a type θ consumer buys a refundable ticket at price p in period 1 and learns that he wants to travel in period 2, then he will use the ticket and his utility will be $u(v_\theta - p)$. If he learns that he does not want to travel, he will request a refund and his utility will be his status quo, $u(0)$, which is equal to zero. In contrast, if the consumer buys a non-refundable ticket at price p, his utility will be $u(v_\theta - p)$ if he wants to travel and $u(-p)$ otherwise. Under expected utility theory, type θ consumer's expected utility from buying a refundable ticket at price p is denoted by

$$U_\theta^r(p) = \pi_\theta u(v_\theta - p) \tag{1}$$

and type θ consumer's expected utility from buying a non-refundable ticket at price p is denoted by

$$U_\theta^{nr}(p) = \pi_\theta u(v_\theta - p) + (1 - \pi_\theta)u(-p). \tag{2}$$

Note that there is no time value of money. If a consumer does not buy a ticket, his utility in period 2 will be zero in both states.[4] Then type θ's reservation price for a refundable ticket is v_θ and type θ's reservation price for a non-refundable ticket is c_θ such that $U_\theta^{nr}(c_\theta) = 0$; i.e.,

$$\pi_\theta u(v_\theta - c_\theta) + (1 - \pi_\theta)u(-c_\theta) = 0. \tag{3}$$

Note that c_θ is an increasing continuous function of π_θ from $[0, 1]$ onto $[0, v_\theta]$.[5]

Now we explain how a consumer decides which type of ticket to buy when the airline simultaneously offers non-refundable and refundable tickets in period 1. Formally, let the airline offer a price menu (p^{nr}, p^r) in which p^{nr} represents a non-refundable price and p^r represents a refundable price. Each consumer's action set includes buy a refundable ticket, buy a non-refundable ticket, and not buy a ticket. We find that type θ consumers best response is given by:

(i) buy a refundable ticket if $U_\theta^r(p^r) \geq U_\theta^{nr}(p^{nr})$ and $U_\theta^r(p^r) \geq 0$,

(ii) buy a non-refundable ticket if $U_\theta^{nr}(p^{nr}) > U_\theta^r(p^r)$ and $U_\theta^{nr}(p^{nr}) \geq 0$,

(iii) buy no ticket if $U_\theta^{nr}(p^{nr}) < 0$ and $U_\theta^r(p^r) < 0$.

2.2 The Airline's Problem

We now turn to the airline's pricing problem. In particular, we are interested in a separating equilibrium where type H consumers buy refundable tickets and type L consumers buy non-refundable tickets. We solve for the optimal separating price menu and show that it constitutes an equilibrium under reasonable conditions.

Let the numbers of type H and type L consumers in period 1 be N_H and N_L and the expected numbers of type H and type L consumers that want to travel in period 2 be

[4] It is possible to extend the model by imposing a cost on the consumer who wants to travel but does not have a ticket so that his utility is lower than zero. However, the differences in the results are immaterial.

[5] For example, if $u(x) = \ln(1 + x/1000)$, $v_L = 500$, and $\pi_L = 0.6$, then we find that $c_L = 268$.

$n_H = \pi_H N_H$ and $n_L = \pi_L N_L$ respectively. The airline, which has zero marginal cost and a capacity at least $n_H + n_L$, announces p^{nr} and p^r at the beginning of period 1. Since the airline does not know each consumer's type, we let the airline derive its belief about each consumer type from N_H and N_L. After observing the prices, consumers' strategies could be either pooling (i.e., both types choose the same action) or separating (i.e., each type chooses a different action). We define an equilibrium as a combination of the airline's beliefs and strategy (p^{nr}, p^r) and each consumer's strategy given θ and (p^{nr}, p^r) so that the airline's expected profit and each consumer's expected utility are maximized.

We are interested in separating equilibria in which each consumer type buys a different type of ticket. However, we cannot identify which consumer type buys which type of ticket without further assumptions about v_θ and π_θ for $\theta = H, L$. Thus we assume negative correlation between valuation and certainty about travel (i.e., $v_H > v_L$ and $\pi_H < \pi_L$). This is a realistic assumption because most business travelers have a higher willingness to pay for a ticket and are less certain about their travel than pleasure travelers. As a result, type H consumers buy refundable tickets, type L consumers buy non-refundable tickets, and the airline's expected profit is equal to $N_L p^{nr} + n_H p^r$. Given these consumers' strategies, we solve for the airline's optimal price menu below and present the solution in Proposition 1. Nonetheless, this separating price menu does not always constitute an equilibrium because it may be possible for the airline to increase expected profit by using an alternative price menu inducing other responses from the consumers. We provide necessary and sufficient conditions for existence of this separating equilibrium in Proposition 2.

The airline's optimization problem can be written as:

$$\max_{p^{nr}, p^r} N_L p^{nr} + n_H p^r \tag{4}$$

s.t.

$$U_H^r(p^r) \geq U_H^{nr}(p^{nr}) \tag{5}$$

$$U_L^{nr}(p^{nr}) \geq U_L^{nr}(p^r) \tag{6}$$

$$U_H^r(p^r) \geq 0 \tag{7}$$

$$U_L^{nr}(p^{nr}) \geq 0. \tag{8}$$

The incentive-compatibility constraints, (5) and (6), are required for type H consumers to prefer buying a refundable ticket and type L consumers to prefer buying a non-refundable ticket, while the participation constraints, (7) and (8), are required for both consumer types to buy a ticket. There are two possibilities: either $c_L \geq c_H$ or $c_L < c_H$. If $c_L \geq c_H$, type L consumers are willing to pay more for a non-refundable ticket while type H consumers are willing to pay more for a refundable ticket. It follows that the airline can separate the consumers by simply setting $(p^{nr}, p^r) = (c_L, v_H)$.[6]

Now consider the case where type L consumers are willing to pay less for a non-refundable ticket than type H consumers (i.e., $c_L < c_H$). If the airline sets $(p^{nr}, p^r) = (c_L, v_H)$, type H consumers will find non-refundable tickets more attractive than refundable tickets because $U_H^r(v_H) < U_H^{nr}(c_L)$. To obtain the separating equilibrium, the airline must lower the refundable price to be no larger than m so that $U_H^r(m) = U_H^{nr}(c_L)$. That is,

$$\pi_H u(v_H - m) = \pi_H u(v_H - c_L) + (1 - \pi_H) u(-c_L). \tag{9}$$

We say that the difference between v_H and m is the information rent taken by type H consumers. The airline offers this discount to prevent type H consumers from buying a non-refundable ticket. In this case the solution is $(p^{nr}, p^r) = (c_L, m)$.[7] We summarize these results in Proposition 1.

Proposition 1 *Given the following assumptions:*

(i) $v_H > v_L > 0$.

(ii) $0 < \pi_H < \pi_L < 1$.

(iii) $u' > 0$ *and* $u'' \leq 0$.

[6] For example, let $u(x) = \ln(1 + x/1000)$. For type L consumers, $v_L = 500$ and $\pi_L = 0.6$, so $c_L = 268$. For type H consumers, $v_H = 800$ and $\pi_H = 0.3$, so $c_H = 185$. Since $c_L \geq c_H$, the airline sets $(p^{nr}, p^r) = (268, 800)$.

[7] Let $u(x) = \ln(1 + x/1000)$. For type L consumers, $v_L = 500$ and $\pi_L = 0.6$, so $c_L = 268$. For type H consumers, $v_H = 800$ and $\pi_H = 0.5$, so $c_H = 323$. We find that $m = 678$. Since $c_L < c_H$, the airline sets $(p^{nr}, p^r) = (268, 678)$.

The airline's optimal price menu so that type L consumers buy non-refundable tickets and type H consumers buy refundable tickets is given by

$$(p^{nr}, p^{r}) = \begin{cases} (c_L, v_H) & \text{if } c_L \geq c_H \\ (c_L, m) & \text{if } c_L < c_H \end{cases} \tag{10}$$

where c_θ and m are defined in (3) and (9) respectively.

Proof See the above argument.

Given the separating response in which type L consumers buy non-refundable tickets and type H consumers buy refundable tickets, the airline maximizes its profit by always selling non-refundable tickets at c_L. If a type L consumer is willing to pay more for a non-refundable ticket than a type H consumer (i.e., $c_L \geq c_H$) then the type H consumer will find a refundable ticket priced at v_H more attractive than a non-refundable ticket. Despite the fact that a non-refundable ticket costs less, it gives a type H consumer lower expected utility than a refundable ticket. Hence the optimal pricing menu is $(p^{nr}, p^{r}) = (c_L, v_H)$. In contrast, if a type H consumer is willing to pay more for a non-refundable ticket than a type L consumer (i.e., $c_H > c_L$) he will find a non-refundable ticket priced at c_L more attractive than a refundable ticket priced at v_H. So the airline has to lower the refundable price to m so that the expected utilities from both ticket types are equal, and by our tie-breaking rule, he will buy a refundable ticket.

Even though the airline sells tickets to the two consumer types at different prices, we cannot claim that the difference between the two prices, $p^{r} - p^{nr}$, is solely due to price discrimination because the two ticket types have different qualities. We decompose the fare difference into two components: $p^{r} - v_L$ and $v_L - p^{nr}$. The second component, which is equal to $v_L - c_L$, is the refundability value to type L consumers because it is the difference in willingness to pay for a refundable ticket versus a non-refundable ticket. If $c_L \geq c_H$, the first component equals the difference in willingness to pay for a refundable ticket between the two consumer types. Since the airline can sell refundable tickets to type H consumers at a higher price than what type L consumers are willing to pay, we say that this component is the airline's price discrimination. If $c_L < c_H$, the price discrimination component is smaller because the airline has to lower the refundable price to separate consumers.

2.3 Equilibrium Prices

We have shown in Proposition 1 that the airline's optimal price menu conditional on obtaining a separating equilibrium in which type H consumers buy refundable tickets and type L consumers buy non-refundable tickets is $(p^{nr}, p^r) = (c_L, v_H)$ if $c_L \geq c_H$ and $(p^{nr}, p^r) = (c_L, m)$ if $c_L < c_H$. However, this separating response may not be best to the airline and hence the described price menu may not constitute an equilibrium. For instance, if there are only a few type H consumers, a pooling response in which all consumers buy refundable tickets may be more profitable. We provide necessary and sufficient conditions for the airline to find the separating response most profitable in the following proposition.

Proposition 2 *Given the following assumptions:*

(i) $v_H > v_L > 0$.

(ii) $0 < \pi_H < \pi_L < 1$.

(iii) $u' > 0$ *and* $u'' \leq 0$.

Necessary and sufficient conditions for existence of an equilibrium where type L *consumers buy non-refundable tickets and type* H *consumers buy refundable tickets are*

$$\frac{N_H}{N_L} \geq \frac{\pi_L v_L - c_L}{\pi_H (v_H - v_L)} \tag{11}$$

if $c_L \geq c_H$ *and*

$$\frac{\pi_L v_L - c_L}{\pi_H (m - v_L)} \leq \frac{N_H}{N_L} \leq \frac{c_L}{\pi_H (v_H - m)} \tag{12}$$

if $c_L < c_H$, *where* c_θ *and* m *are defined in (3) and (9) respectively.*

Proof See Appendix.

To assure that the separating response where the airline sells refundable tickets to type H consumers and non-refundable tickets to type L consumers constitutes an equilibrium, we must show that the airline's maximum profit given this response is not lower than its maximum profit from each of other consumers' responses.[8] The inequality in (11) and the

[8] In the proof of Proposition 2, we compare the airline's profit from the optimal price menu in (10) to profits from other optimal price menus given other consumers' responses.

left inequality in (12) guarantee that the separating price menu from (10) yields a higher profit than the optimal price menu conditional on obtaining a pooling equilibrium with both types buying refundable tickets. Both inequalities suggest that N_H/N_L must be large enough because the airline makes more profit from type H and less profit from type L with the separating price menu. To maximize its profit while making both type H and type L consumers buy refundable tickets, the airline must set p^{nr} really high so that no one would want to buy a non-refundable ticket. With this pooling price, the airline makes less profit from type H consumers because it has to lower p^r to v_L so that refundable tickets attract all consumers. However, the airline makes more profit from type L consumers by selling them refundable instead of non-refundable tickets because of risk aversion. We show in the appendix that $c_L \leq \pi_L v_L$ whenever $u'' \leq 0$ and hence the profit from selling non-refundable tickets to type L consumers, $N_L c_L$, is smaller than $N_L \pi_L v_L$ which is the airline's profit when selling them refundable tickets.[9]

In addition to pooling-on-refundable response described above, we also consider other consumers' responses in the proof of Proposition 2 and find that the separating price menu dominates all other responses regardless of N_H and N_L when $c_L \geq c_H$. However, when $c_L < c_H$, another candidate of equilibrium arises because the airline might prefer selling only refundable tickets to type H consumers and no tickets at all to type L consumers. Remember that the optimal separating price menu given $c_L < c_H$ is $(p^{nr}, p^r) = (c_L, m)$. Since the airline sells refundable tickets at a discounted price, m, it may be more profitable to the airline, when N_H is large enough, to set p^{nr} very high so that a non-refundable ticket is unattractive to everyone and increase p^r to v_H.[10]

In short, the separating price menu in (10) constitute an equilibrium when the proportion of type H to type L population is neither too large nor too small. If there are relatively

[9] Consider the example in Footnote 5. If $N_H = N_L = 50$, the separating price menu yields $(15 \times 800) + (50 \times 268) = 25,400$ in profit which is greater than profit from the pooling-on-refundable price menu, $(15 \times 500) + (30 \times 500) = 22,500$. If the N_H/N_L ratio was not large enough, say $N_H = 10$ and $N_L = 90$, the airline would try to sell refundable tickets to all consumers and make $(3 \times 500) + (54 \times 500) = 28,500$ in profit instead of $(3 \times 800) + (90 \times 268) = 26,520$ from separating.

[10] Following with the example in Footnote 8, with $N_H = 90$ and $N_L = 10$. The separating price menu yields $(45 \times 678) + (10 \times 268) = 33,190$ in profit which is smaller than profit from selling only refundable tickets to type H, $45 \times 800 = 36,000$.

many type L consumers, the airline will prefer selling refundable tickets to both types at v_L. On the other hand, if there are relatively many type H consumers and the willingness to pay for a non-refundable ticket of type L is less than that of type H, the airline will prefer selling only refundable tickets to type H at v_H and no tickets at all to type L.

The restrictions (11) and (12) can be simplified when consumers are risk neutral because, as we see from (3), $c_L = \pi_L v_L$ when $u'' = 0$, and as a result, the inequality in (11) and the left inequality in (12) are automatically satisfied. So we can say that when consumers are risk neutral and $c_L \geq c_H$, the airline's separating price menu given in Proposition 1 always constitutes an equilibrium. In contrast, if $c_L < c_H$, we still need to impose the right inequality of (12) even when $u'' = 0$. Nonetheless, we can say that it is very likely to see the separating price menu in equilibrium if consumers are slightly risk averse because both restrictions on N_H/N_L are easier to satisfy.

2.4 Empirical Implication

We now provide an interesting empirical implication from this theoretical model: the gap between refundable and non-refundable prices diminishes as the flight date nears and the consumers become more certain about their travel plans. Let π_H and π_L be continuous functions of τ, the time length between period 1 and period 2 in the theoretical model (i.e., the number of days between when agents consider buying a ticket and the departure date). We assume the following:

(i) $\pi_L(\tau)$ and $\pi_H(\tau)$ are non-increasing in τ.

(ii) $\pi_L(\tau) \geq \pi_H(\tau)$ for all τ.

(iii) $\pi_L(0) = \pi_H(0) = 1$.

The possibility that each passenger becomes more certain about traveling because he may have more information as the departure date approaches can justify (i). Conditions (ii) and (iii) mean that type L travelers are more certain about traveling than type H travelers on any day prior to departure and that there is no uncertainty about traveling on the departure date. These are practical conditions because most business travelers are more uncertain about their travel than pleasure travelers when they buy tickets in advance, and all agents know with certainty whether they are flying on the departure date.

We learn from (3) that c_θ is a continuous function of π_θ and that $dc_\theta/d\pi_\theta > 0$. It follows from the above assumptions that when τ is small (i.e., close to departure date) both π_L and π_H will be close to 1 and, as a result, $c_L < c_H$. In contrast, when τ is large, it is possible that the difference between π_L and π_H is so large that $c_L \geq c_H$. We show below that regardless of whether $c_L \geq c_H$ or $c_L < c_H$ the gap between refundable and non-refundable prices are smaller as τ decreases.

If $c_L \geq c_H$, the airline sets $(p^{nr}, p^r) = (c_L, v_H)$. Since π_L increases as τ decreases, we find from (3) that c_L also increases and the gap between p^r and p^{nr} becomes smaller. On days close to departure, π_H will be large enough that $c_H > c_L$ and the airline sets $(p^{nr}, p^r) = (c_L, m)$. From (9) we find that, as π_H approaches 1, m converges to c_L. Therefore, the difference between the refundable and non-refundable fares diminishes as the departure date approaches (i.e., $\tau \to 0$). When $\tau = 0$, the airline sets $(p^{nr}, p^r) = (c_L, c_L) = (v_L, v_L)$.[11] There is only one price on the departure date because there is no benefit from paying more for a refundable ticket when individual demand uncertainty is fully resolved. The airline's ability to screen consumers vanishes, making price discrimination opportunities through refund contracts disappear.

3 Empirical Analysis

3.1 Data

We collected from the online travel agency Expedia.com the lowest posted refundable and non-refundable one-way economy-class fares for 96 flights that departed on June 22, 2006. Following Stavins (2001) we focused on a single day, Thursday, to control for price differentials associated with systematic peak-load pricing over days of the week. The data form a panel with 96 cross-sectional units observed over 28 periods. Each cross-section corresponds to a specific carrier's non-stop flight between a city pair. Fares were recorded every three days, from 82 days prior to departure to one day prior to departure, i.e., $\tau = 1, 4, 7, \ldots, 82$. The carriers considered are American, Alaska, Continental, Delta, United, and US Airways.

[11] As the date of travel approaches the model allows for the possibility that we move from the separating equilibrium to other equilibria. This will not be an issue if N_H and N_L satisfy the conditions in Proposition 2 for all τ.

The share of each carriers' flights in the data set was chosen to be close to the carrier's share of the U.S. market.

A monopoly route, as defined by Borenstein and Rose (1994), is a route on which a single carrier operates more than 90 percent of the weekly direct flights. Following a similar but stricter criterion, each of our 96 monopoly routes is operated by a carrier who is the sole supplier of non-stop service between the city pair. Tickets with one or more stops and first-class travel tickets are considered to be of significantly different quality. By picking non-stop flights and one-way fares we control for price differences associated with different ticket restrictions, cost differences associated with round-trip tickets, and price variations related to more sophisticated itineraries. Dealing with one-way tickets is particularly important in our model because the traveler may also be uncertain about the return portion in a round-trip ticket.

[Insert Figure 1 here.]

Figure 1 shows the average refundable and non-refundable fares across all 96 monopoly routes at different points prior to the departure date. The dashed lines are 95% confidence intervals for the means and are calculated separately for every point in time using the t-distribution. The figure shows that there is a strong tendency for non-refundable fares to increase faster than refundable fares, suggesting that as the flight date nears consumers resolve their individual demand uncertainty.

3.2 Main Result

The panel structure of the data allows us to control for time-invariant flight-, carrier-, and route-specific characteristics. This includes all time-invariant characteristics in Stavins (2001), who used a cross section of tickets. However, there are time-variant cost components that arise at seat-level. Stochastic peak-load pricing, as explained in Borenstein and Rose (1994), depends on the degree of price adjustment as demand is revealed over time. Prices on the same flight can vary with the purchase date and with the probability, at the time the ticket is sold, that demand will exceed capacity. In models of aggregate demand uncertainty where sellers commit to a price schedule, Dana (1998, 1999) considers the existence of an effective cost of capacity, which also changes across seats.

To control for these time-variant seat-specific characteristics, we take advantage of the fact that both refundable and non-refundable fares were obtained at the same time for the same seat. Hence, taking the difference between these two fares will wipe out these characteristics.[12] The basic logarithmic specification of the reduced-form model that we estimate is given by

$$\ln(p^{r}_{ijt} - p^{nr}_{ijt}) = \sum_{s=4,7,\ldots,82} \beta_s I_t^{[\tau_t=s]} + X'_{ijt}\delta + \nu_{ij} + \varepsilon_{ijt}, \tag{13}$$

where the subscript i refers to flight, j to route, and t to time. Following the theoretical section, p^r and p^{nr} are the refundable and non-refundable ticket prices, and τ is the number of days prior to departure when the two prices were recorded. The indicator variables $I_t^{[\tau_t=s]}$ are each equal to one at different number of days to departure, i.e., $\tau = 4, 7, \ldots,$ or 82. Because $I_t^{[\tau=1]}$ is the omitted category, β_τ captures the logarithm of the difference in the price gap between τ days and 1 day to departure. X'_{ijt} is a vector of controls that includes the lagged dependent variable and the load factor ($LOAD$), defined as the ratio of occupied seats to total seats.[13] Finally, ν_{ij} denotes the unobservable flight-specific effect and ε_{ijt} denotes the remaining disturbance.

[Insert Table 1 here.]

From the summary statistics presented in Table 1, we find that refundable fares are, on average, 51% larger than non-refundable fares. As expected, variation in fares across flights is very close for both fare types, but non-refundable fares appear to have more within-flight variation than do refundable fares. Table 2 presents the estimates of the

[12]Controlling for these costs has limited the availability of empirical papers on airline price discrimination. Stavins (2001) looks at price differentials due to ticket restrictions; however, those ticket restrictions are understood to solve the peak-load pricing problem as well (see Courty and Li, 2000, p. 716).

[13]Given that overbookings are usually a small fraction of the total number of tickets, $LOAD$ is assumed to be proportional to bookings. $LOAD$ was obtained from the seat-availability map, where the available or prime seats reported by Expedia.com are counted as available seats. Escobari (2009), Escobari (2012) and Alderighi et al. (2012) explain the importance of controlling for seat availability.

coefficients of Equation 13. All columns are estimated using flight fixed effects to control for route-, carrier-, and flight-specific characteristics. As robustness checks, the second column controls for capacity utilization using $LOAD$, while the third column includes $LOAD$ and the lagged dependent variable. The numbers in parentheses are t-statistics based on cluster-robust standard errors, clustered by airline. Column 1 shows that at a 1% significance level, with the exception of 4 days in advance, the price gaps are larger than the gap from 1 day in advance. The decline in the gap appears to occur during the last 13 or 16 days prior to departure, with statistically nonsignificant fluctuations in the gap during earlier dates. The estimates indicate that, for example, the price difference is 211% larger when $\tau = 7$ days and 298% larger when $\tau = 16$ days.[14] Columns 2 and 3 show that the percentage gap decline is smaller when controlling for capacity utilization and even more so when additionally including the lagged dependent variable. We can see that across all specifications the main result holds—the difference between refundable and nonrefundable fares is smaller closer to departure.

[Insert Table 2 here.]

As a robustness analysis we also use the *difference* and *system* GMM methods proposed in Holtz-Eakin et al. (1988), Arellano and Bond (1991), Arellano and Bover (1995), and Blundell and Bond (1998) to estimate

$$\ln(p^r_{ijt} - p^{nr}_{ijt}) = \sum_{s=1}^{3} \gamma_s \tau_t^3 + X'_{ijt}\delta + \nu_{ij} + \varepsilon_{ijt}. \tag{14}$$

We model the lagged dependent variable and $LOAD$ in X'_{ijt} as endogenous and weakly exogenous respectively. A weakly exogenous $LOAD$ means that the number of seats sold up to time t can be affected by past realizations of the differences in fares.[15]

Table 3 presents the results. For comparison purposes, the first column shows the OLS estimates of the pooled regression model. Exploiting the panel structure of the data, the

[14] Sweeting (2012) runs a similar specification, but of prices—not a price gap—on a set of dummies to analyze the dynamics of prices in secondary markets for Major League Baseball tickets.

[15] Bilotkach et al. (2011) present a specification in which the load factor is affected by previous prices, consistent with our assumption of weak exogeneity of $LOAD$. Moreover, weak exogeneity is also consistent with Deneckere and Peck (2012) and the dynamic demand estimation in Escobari (2012).

second column presents the within-flight regression estimates. While the signs of the day-in-advance coefficients in both specifications are consistent with a fluctuating difference in earlier dates and a sharp decrease in the gap close to departure, the squared and cubic terms are not statistically significant.

The third and fourth columns report the two-step first-differenced GMM panel estimator as proposed in Holtz-Eakin et al. (1988) and Arellano and Bond (1991). The t-statistics in parentheses are based on the Windmeijer finite-sample correction for the standard errors of the two-step estimates. The GMM difference panel estimator works by taking first-differences to eliminate time-invariant characteristics and assumes that the error term, ε_{ijt}, is not serially correlated. Moreover, in this estimation the series $LOAD_{ijt}$ may be endogenous in the sense that $LOAD_{ijt}$ is correlated with ε_{ijt} and earlier shocks but uncorrelated with $\varepsilon_{ij,t+1}$ and subsequent shocks. Then lagged values of $LOAD_{ijt}$ are valid instruments in the first-differenced equation. Column 3 uses $LOAD_{ij,t-2}$ as an instrument, while column 4 uses $LOAD_{ij,t-3}$ and earlier lags as instruments.[16] To deal with the problem that the new error term, $\varepsilon_{ijt} - \varepsilon_{ij,t-1}$, is correlated with the lagged dependent variable, $\ln(p^{r}_{ij,t-2} - p^{nr}_{ij,t-2})$ is also used as an instrument. To address the validity of the specifications, columns 3 and 4 also report two tests. To test the hypothesis that the error term, ε_{ijt}, is not serially correlated, we test whether the differenced error term is second-order serially correlated. The p-values reported for the serial correlation test provide strong support for a valid specification. To test the overall validity of the instruments, we provide a Sargan test of over-identifying restrictions. The validity of lagged levels dated $t-2$ as instruments in column 3 is rejected. However, it is not rejected for lagged levels dated $t-3$ (and earlier) as instruments.

[Insert Table 3 here.]

Blundell and Bond (1998) point out that, when the explanatory variables are persistent over time, lagged levels of these variables are weak instruments for the regression equation in differences. To reduce the potential bias and imprecision of the difference estimator,

[16] τ_t, τ_t^2, and τ_t^3 are treated as strictly exogenous across all specifications.

we use the system estimator as suggested in Blundell and Bond (1998). Usual weak-instrument tests cannot be implemented in these GMM specifications, so we use the known bias in the difference GMM by comparing its sample performance with the pooled OLS and within-group estimators, which have known properties in dynamic panel data, and see if the system GMM improves the precision of the estimates. This system estimator combines the regression in differences with the regression in levels. The instruments for the regression in levels are the lagged differences of the corresponding variables. The validity of the instruments relies on an additional assumption: There is no correlation between the differences of $LOAD$ and the flight-specific effects, but there may be correlation between the levels of $LOAD$ and the flight-specific effects. Columns 5 and 6 in Table 3 report the two-step system GMM estimator, with the figures in parentheses being t-statistics based in the Windmeijer robust estimator. The serial-correlation test shows strong support for the assumption of no serial correlation, and the Sargan provides strong support for the validity of the instrument list. In addition, the difference Sargan that tests for the additional moment conditions used in the levels equations accepts their validity. We find that the coefficients on τ_t, τ_t^2, and τ_t^3 are statistically significant in all GMM specifications and that the signs are consistent with a quick decline in the gap close to departure.

[Insert Figure 2 here.]

Finally, our last robustness check is aimed at addressing nonlinearities in a very flexible way by estimating a nonparametric model with fixed effects. The estimation follows the kernel methods of Racine and Li (2004) for a mix of discrete (τ) and continuous ($LOAD$) data types.[17] This approach allows for interactions between τ and $LOAD$ and for nonlinearities in and between both. The smoothing parameters are calculated with least squares cross-validation. The results are summarized in the partial regression plot of Figure 2, which presents the multivariate regression function via a bivariate plot with $LOAD$ held constant at its median and permits a direct comparison of the parametric and nonparametric results. The black dots are the nonparametric results, with the bars representing asymptotic standard errors. The solid line is the gap profile as estimated by our preferred

[17]This estimator has better finite-sample properties than the popular local kernel estimator.

specification of column 6, Table 3. Consistent with the dummy variable regression of Table 2, there is a sharp decrease in the gap during the last two to three weeks prior to departure. The gap between p^r and p^{nr} closes as the departure date nears and consumers are more certain about their travel plans. While there does not seem to be much learning during earlier dates, most of the learning takes place during the last two weeks prior to departure.

4 Conclusions

This paper shows the importance to airlines of offering a menu of prices namely refundable and non-refundable fares. We show that a monopolistic airline can separate consumers who are uncertain about their demand for travel and have different willingness to pay. The fact that individual demand uncertainty is not fully resolved by the time the individual buys a ticket is used by the seller to price discriminate and extract more surplus. In our model, buyers can use refund contracts to insure against uncertainty in consumption. One implication from the theoretical model is that the gap between refundable and non-refundable fares is a function of the individuals' travel uncertainty. If there is no uncertainty in individual demand, there is no difference in buying a refundable or a non-refundable ticket, and hence there should be no difference between these two fares.

The empirical section looks at the dynamics of the price gap, refundable minus non-refundable, in 96 monopoly routes and tests whether the individual demand uncertainty implied by the carrier's pricing strategy is resolved as the departure date approaches. After controlling for unobserved time-invariant flight-, carrier-, and route-specific characteristics, unobserved time-variant seat-specific characteristics, and potential sources of endogeneity, the results show that the theoretical predictions are empirically supported. Second-degree price discrimination in the form of refund contracts shrinks as the departure date nears. Nonlinear parametric specifications and a nonparametric regression show that most of the individual demand uncertainty is resolved during the last two weeks, when the opportunity for price discrimination through refund contracts declines.

References

Alderighi, M., M. Nicolini, and C.A. Piga, 2012. Combined Effects of Load Factors and Booking Time on Fares: Insights from the Yield Management of a Low-Cost Airline. FEEM Working Paper.

Akan, M., B. Ata, and J.D. Dana, Jr., 2011. Revenue Management by Sequential Screening. Northwestern University Working Paper.

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, 277-297.

Arellano, M. and O. Bover, 1995. Another Look at the Instrumental Variable Estimation of Error-Components Models. *Journal of Econometrics* 68, 29-51.

Bilotkach, V., 2009. Refundable Advance Purchase Tickets and Advance Purchase Restrictions under Aggregate Demand Uncertainty. UC Irvine Working Paper.

Bilotkach, V., A. Gaggero, and C.A. Piga, 2011. Airline Pricing under Different Market Conditions: Evidence from European Low-Cost Carriers. Rimini Centre for Economic Analysis, WP 11-47.

Blundell, R. and S. Bond, 1998. Initial Conditions and Moment Restrictions in Dynamic Panel Data Models. *Journal of Econometrics* 87, 115-143.

Borenstein, S. and N.L. Rose, 1994. Competition and Price Dispersion in the U.S. Airline Industry. *Journal of Political Economy* 102, 653-683.

Cicchetti, C.J. and A.M. Freeman, III, 1971. Option Demand and Consumer Surplus: Further Comment. *Quarterly Journal of Economics* 85, 528-539.

Courty, P. and H. Li, 2000. Sequential Screening. *Review of Economic Studies* 67, 697-717.

Dana, Jr., J.D., 1998. Advance-Purchase Discounts and Price Discrimination in Competitive Markets. *Journal of Political Economy* 106, 395-422.

Dana, Jr., J.D., 1999. Equilibrium Price Dispersion under Demand Uncertainty: The Roles of Costly Capacity and Market Structure. *Rand Journal of Economics* 30, 632-660.

Deneckere, R. and J. Peck, 2012. Dynamic Competition with Random Demand and Costless Search: A Theory of Price Posting. *Econometrica* 80, 1185-1247.

Escobari, D., 2009. Systematic Peak-load Pricing, Congestion Premia and Demand Diverting: Empirical Evidence. *Economics Letters* 103, 59-61.

Escobari, D., 2012. Dynamic Pricing, Advance Sales, and Aggregate Demand Learning in Airlines. *Journal of Industrial Economics* 60, 697-724.

Escobari, D. and Gan, L., 2007, 'Price Dispersion under Costly Capacity and Demand Uncertainty', NBER working paper 13075, Cambridge, MA.

Gale, I.L. and T.J. Holmes, 1993. Advance-Purchase Discounts and Monopoly Allocation of Capacity. *American Economic Review* 83, 135-146.

Gerardi, K.S. and A.H. Shapiro, 2009. Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry. *Journal of Political Economy* 117, 1-37.

Graham, D., 1981. Cost-Benefit Analysis under Uncertainty. *American Economic Review* 71, 715-725.

Holtz-Eakin, D., W. Newey, and H. Rosen, 1988. Estimating Vector Autoregressions with Panel Data. *Econometrica* 56, 1371-1395.

Li, Q. and J.S. Racine, 2007. *Nonparametric Econometrics*. Princeton University Press, Princeton, NJ.

Racine, J.S. and Q. Li, 2004. Nonparametric Estimation of Regression Functions with both Categorical and Continuous Data. *Journal of Econometrics* 119, 99-130.

Schmalensee, R., 1972. Option Demand and Consumer's Surplus: Valuing Price Changes under Uncertainty. *American Economic Review* 62, 813-824.

Stavins, J., 2001. Price Discrimination in the Airline Market: The Effect of Market Concentration. *Review of Economics and Statistics* 83, 200-202.

Sweeting, A., 2012. Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets. *Journal of Political Economy* 120, 1133-1172.

Appendix

In this appendix we first prove two lemmas that we use to derive the result in Proposition 2 and then we prove the proposition.

Lemma 1 *If* $u'' \lesseqqgtr 0$, *then* $c_\theta \lesseqqgtr \pi_\theta v_\theta$.

Proof. We rewrite (3) as

$$\frac{-u(-c_\theta)}{u(v_\theta - c_\theta) - u(-c_\theta)} = \pi_\theta. \tag{15}$$

We multiply both sides of (15) by v_θ / c_θ to obtain

$$\frac{-u(-c_\theta)/c_\theta}{[u(v_\theta - c_\theta) - u(-c_\theta)]/v_\theta} = \frac{\pi_\theta v_\theta}{c_\theta}. \tag{16}$$

If $u'' \lesseqqgtr 0$, the left-hand side of (16) $\gtreqqless 1$ and, as a result, $c_\theta \lesseqqgtr \pi_\theta v_\theta$.

Lemma 2 *If* $u'' \lesseqqgtr 0$, *then* $m \gtreqqless \frac{c_L}{\pi_H}$.

Proof. We rewrite (9) as

$$\frac{-u(-c_L)}{u(v_H - c_L) - u(v_H - m)} = \frac{\pi_H}{1 - \pi_H}. \tag{17}$$

We multiply both sides of (17) by $(m - c_L)/c_L$ to obtain

$$\frac{-u(-c_L)/c_L}{[u(v_H - c_L) - u(v_H - m)]/(m - c_L)} = \frac{\pi_H (m - c_L)}{(1 - \pi_H) c_L}. \tag{18}$$

If $u'' \lesseqqgtr 0$, the left-hand side of (18) $\gtreqqless 1$ and, as a result, $m \gtreqqless \frac{c_L}{\pi_H}$.

Proof of Proposition 2

Since the airline cannot observe each consumer's type, the airline may try to set prices so that consumers either pool or separate their strategies. We summarize the consumers' responses to the airline's price menu as follows:

(i) Pooling 1 : The airline sells refundable tickets to both types (RB).

(ii) Pooling 2 : The airline sells non-refundable tickets to both types (NB).

(iii) Separating 1 : The airline sells only refundable tickets to type H (RH).

(iv) Separating 2 : The airline sells only non-refundable tickets to either type H or L (NL if $c_L > c_H$ or NH if $c_L < c_H$).

(v) Separating 3 : The airline sells refundable tickets to type H and non-refundable tickets to type L (RHNL).

We consider these responses under two cases: $c_L \geq c_H$ and $c_L < c_H$. For RB the airline maximizes its profit by setting $p^{nr} \geq \max\{c_H, c_L\}$ and $p^r = v_L$ in both cases. For NB the airline sets $p^{nr} = c_H$ if $c_L \geq c_H$, $p^{nr} = c_L$ if $c_L < c_H$, and $p^r > v_H$ in both cases.

For RH the airline sets $p^{nr} > c_L$ if $c_L \geq c_H$, $p^{nr} \geq c_H$ if $c_L < c_H$, and $p^r = v_H$ in both cases. For NL in case $c_L > c_H$, the airline sets $p^{nr} = c_L$ and $p^r > v_H$. For NH in case $c_L < c_H$, the airline sets $p^{nr} = c_H$ and $p^r > v_H$. Note that NL is not possible when $c_L \leq c_H$ and that NH is not possible when $c_L \geq c_H$.

For RHNL, Proposition 1 states that the airline sets $p^{nr} = c_L$ and $p^r = v_H$ if $c_L \geq c_H$ and $p^{nr} = c_L$ and $p^r = m$ if $c_L < c_H$. The airline's profits for each response by the consumers in the two cases are summarized in Table A-1.

[Insert Table A-1 here.]

Let $\Pi(\cdot)$ denote the airline's maximum profit given consumer responses to the airline's price menu. If $c_L \geq c_H$ we find that $\Pi(\text{RHNL}) \geq \Pi(\text{RH})$ and $\Pi(\text{RHNL}) \geq \Pi(\text{NL})$. Given that $u'' \leq 0$, Lemma 1 suggests $c_H \leq \pi_H v_H$. Hence $\Pi(\text{RHNL}) \geq \Pi(\text{NB})$. If $\frac{N_H}{N_L} \geq \frac{\pi_L v_L - c_L}{\pi_H(v_H - v_L)}$, then $\Pi(\text{RHNL}) \geq \Pi(\text{RB})$.

If $c_L < c_H$ $\Pi(\text{RH}) \geq \Pi(\text{NH})$ because $c_H \leq \pi_H v_H$. Given that $u'' \leq 0$, Lemma 2 suggests $\pi_H m \geq c_L$. Hence $\Pi(\text{RHNL}) \geq \Pi(\text{NB})$. If $\frac{N_H}{N_L} \geq \frac{\pi_L v_L - c_L}{\pi_H(m - v_L)}$, then $\Pi(\text{RHNL}) \geq \Pi(\text{RB})$. If $\frac{N_H}{N_L} \leq \frac{c_L}{\pi_H(v_H - m)}$, then $\Pi(\text{RHNL}) \geq \Pi(\text{RH})$.

Table 1: Summary statistics

Variables	Mean	Std. Dev.	Min.	Max.	Obs.
p^r_{ijt}					
overall	494.486	169.181	144.000	1715.310	2628
between		156.974	144.000	735.497	96
within		64.167	141.262	1474.299	27.375
p^{nr}_{ijt}					
overall	327.749	171.588	64.000	914.000	2628
between		156.654	74.107	665.786	96
within		70.204	164.642	852.249	27.375
τ_t	41.500	24.238	1.000	82.000	2688
$LOAD_{ijt}$	0.591	0.241	0.038	1.000	2688

Table 2: Regression estimates, separate day dummies

	(1)		(2)		(3)	
VARIABLES	Coef	t-stat	Coef	t-stat	Coef	t-stat
$I_t^{[\tau_t=4\text{ days}]}$	0.527	(1.575)	0.510	(1.518)	-0.212	(-0.368)
$I_t^{[\tau_t=7\text{ days}]}$	2.106***	(5.375)	2.067***	(5.361)	1.035***	(2.852)
$I_t^{[\tau_t=10\text{ days}]}$	2.614***	(5.290)	2.523***	(5.502)	1.580***	(3.638)
$I_t^{[\tau_t=13\text{ days}]}$	2.565***	(4.141)	2.451***	(4.297)	1.285**	(2.489)
$I_t^{[\tau_t=16\text{ days}]}$	2.977***	(4.306)	2.800***	(4.613)	1.653***	(2.989)
$I_t^{[\tau_t=19\text{ days}]}$	3.013***	(4.190)	2.803***	(4.505)	1.672***	(2.619)
$I_t^{[\tau_t=22\text{ days}]}$	2.999***	(5.066)	2.737***	(5.800)	1.670***	(3.748)
$I_t^{[\tau_t=25\text{ days}]}$	2.963***	(4.708)	2.674***	(5.423)	1.550***	(3.101)
$I_t^{[\tau_t=28\text{ days}]}$	3.036***	(4.751)	2.727***	(5.623)	1.565***	(3.031)
$I_t^{[\tau_t=31\text{ days}]}$	3.092***	(4.903)	2.737***	(6.072)	1.615***	(3.500)
$I_t^{[\tau_t=34\text{ days}]}$	3.124***	(5.081)	2.733***	(6.471)	1.562***	(3.425)
$I_t^{[\tau_t=37\text{ days}]}$	3.187***	(5.081)	2.757***	(6.465)	1.654***	(3.540)
$I_t^{[\tau_t=40\text{ days}]}$	3.069***	(4.980)	2.609***	(6.587)	1.336***	(2.678)
$I_t^{[\tau_t=43\text{ days}]}$	3.418***	(6.369)	2.927***	(9.515)	1.711***	(4.084)
$I_t^{[\tau_t=46\text{ days}]}$	3.325***	(6.522)	2.809***	(10.10)	1.597***	(4.079)
$I_t^{[\tau_t=49\text{ days}]}$	3.331***	(6.698)	2.796***	(10.95)	1.532***	(3.921)
$I_t^{[\tau_t=52\text{ days}]}$	3.442***	(7.401)	2.878***	(12.56)	1.615***	(4.517)
$I_t^{[\tau_t=55\text{ days}]}$	3.439***	(7.441)	2.863***	(12.93)	1.636***	(4.709)
$I_t^{[\tau_t=58\text{ days}]}$	3.392***	(7.049)	2.795***	(12.38)	1.555***	(4.545)
$I_t^{[\tau_t=61\text{ days}]}$	3.429***	(7.006)	2.818***	(11.96)	1.667***	(5.306)
$I_t^{[\tau_t=64\text{ days}]}$	3.291***	(5.722)	2.665***	(8.860)	1.529***	(4.352)
$I_t^{[\tau_t=67\text{ days}]}$	3.249***	(5.028)	2.601***	(7.173)	1.466***	(3.658)
$I_t^{[\tau_t=70\text{ days}]}$	3.257***	(4.891)	2.600***	(6.902)	1.603***	(5.278)
$I_t^{[\tau_t=73\text{ days}]}$	2.993***	(3.390)	2.315***	(3.978)	1.320***	(2.624)
$I_t^{[\tau_t=76\text{ days}]}$	3.003***	(3.425)	2.315***	(4.053)	1.396***	(2.856)
$I_t^{[\tau_t=79\text{ days}]}$	2.861***	(3.194)	2.161***	(3.686)	1.204**	(1.962)
$I_t^{[\tau_t=82\text{ days}]}$	3.177***	(4.309)	2.469***	(5.897)		
$LOAD_{ijt}$			-1.350**	(-2.138)	-0.953**	(-2.221)
$\ln(p^r_{ij,t-1} - p^{nr}_{ij,t-1})$					0.515***	(8.332)
Observations	2,628		2,628		2,519	
Within R-squared	0.312		0.319		0.492	

Notes: The dependent variable is $\ln(p^r_{ijt} - p^{nr}_{ijt})$. t-statistics in parentheses based on cluster-robust standard errors, clustered by airline; ***p-value<0.01, **p-value<0.05, *p-value<0.1. All specifications estimated with flight fixed effects.

Table 3: Regression estimates, cubic model

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	Within	GMM *Dif*	GMM *Dif*	GMM *Sys*	GMM *Sys*
VARIABLES	levels	groups	$t-2$	$t-3$	$t-2$	$t-3$
$\tau_t/10^2$	7.970**	11.578**	10.349***	12.423***	11.782***	12.043***
	(2.222)	(2.537)	(2.610)	(3.860)	(6.475)	(5.411)
$\tau_t^2/10^4$	-17.228*	-24.928*	-21.435***	-25.486***	-25.525***	-26.080***
	(-1.865)	(-1.891)	(-4.211)	(-4.782)	(-5.019)	(-5.183)
$\tau_t^3/10^6$	11.236	15.898	13.680***	16.220***	16.783***	17.152***
	(1.608)	(1.501)	(4.077)	(4.389)	(4.136)	(4.353)
$LOAD_{ijt}$	-0.434***	-0.828**	-0.068	0.102	-0.317	-0.289
	(-5.578)	(-2.156)	(-0.026)	(0.064)	(-0.129)	(-0.124)
$\ln(p_{ij,t-1}^r - p_{ij,t-1}^{nr})$	0.854***	0.530***	0.572***	0.554***	0.566***	0.560***
	(22.736)	(10.772)	(6.221)	(6.073)	(6.679)	(6.168)
Serial correlation test^a(p-value)			0.605	0.619	0.604	0.609
Sargan test^b(p-value)			0.004	0.066	0.689	0.988
Difference Sargan test^c(p-value)					1.000	1.000

Notes: The dependent variable is $\ln(p_{ijt}^r - p_{ijt}^{nr})$. Columns 2 through 6 control for carrier-, route-, and flight-specific characteristic. t-statistics in parentheses for the OLS and the Within groups based on cluster-robust standard errors, clustered by airline. t-statistics in parentheses for the two-step system GMM based on Windmeijer WC-robust estimator; ***p-value<0.01, **p-value<0.05, *p-value<0.1. [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). [c] The null hypothesis is that the additional instruments $t-3$ are not correlated with the residuals (valid specification).

Table A-1: The airline's profits given consumers' responses

Consumers' responses	$c_L \geq c_H$	$c_L < c_H$
RB	$N_H\pi_H v_L + N_L\pi_L v_L$	$N_H\pi_H v_L + N_L\pi_L v_L$
NB	$N_H c_H + N_L c_H$	$N_H c_L + N_L c_L$
RH	$N_H\pi_H v_H$	$N_H\pi_H v_H$
NL	$N_L c_L$	n.a.
NH	n.a.	$N_H c_H$
RHNL	$N_H\pi_H v_H + N_L c_L$	$N_H\pi_H m + N_L c_L$

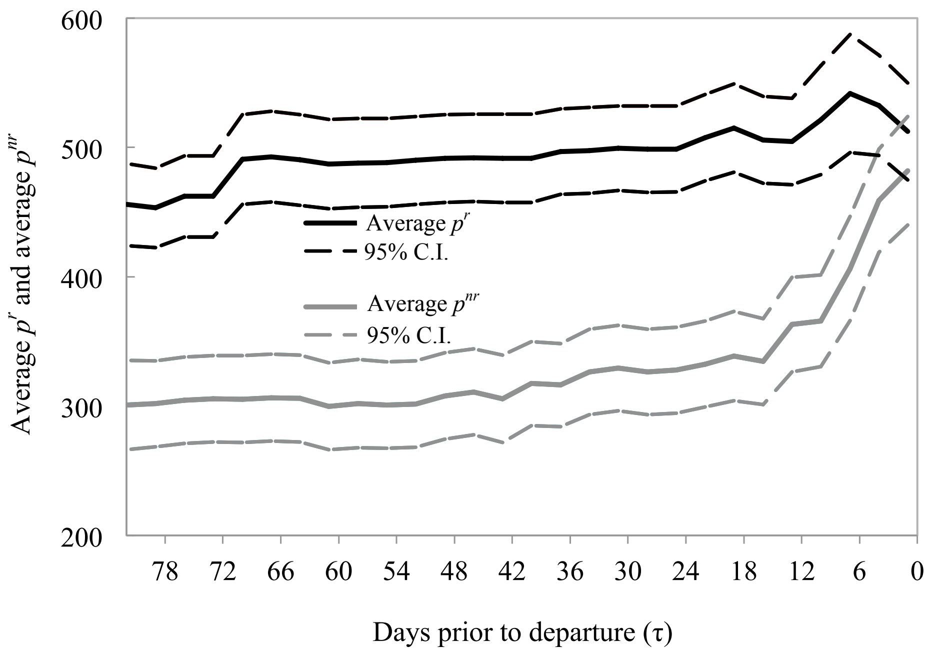

Figure 1: Average p^r and p^{nr} with 95% confidence intervals

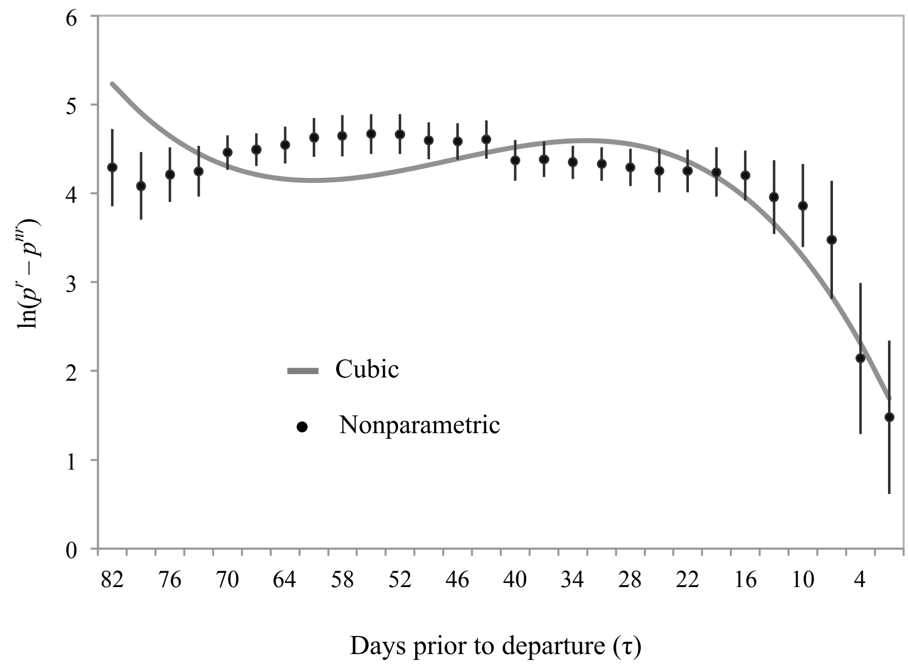

Figure 2: Nonparametric partial regression plot and cubic model