Price Discrimination through Refund Contracts in Airlines*

Diego Escobari†, Paan Jindapon‡

October 5, 2008

Abstract

We show how a monopoly airline uses two ticket types, refundable and non-refundable, to screen consumers with different willingness to pay. Our theoretical model suggests that the difference between these two fares consists of refundability value and price discrimination, and the fare difference diminishes as risk-averse passengers learn about their individual demand uncertainty. Using an original dataset, we find that, after controlling for unobserved seat- and flight-specific characteristics, the empirical evidence from the U.S. airline industry supports our theory.

*We thank Junsoo Lee, Harris Schlesinger, Roger Sherman, and seminar participants at the IIOC 2008 for helpful comments. Escobari thanks the Bradley Foundation for financial support and Jindapon thanks the Culverhouse College of Commerce and Business Administration. The dataset collection was funded by the Private Enterprise Research Center at Texas A&M.

†Department of Economics, San Francisco State University, escobari@sfsu.edu

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

1 Introduction

Consider a consumer who is thinking about buying a plane ticket in advance. At this moment, he is not certain whether he will be able to travel on the date specified on the ticket. Assume that his valuation of a seat on the corresponding flight is \$500 and there is no time value of money. If the airline offers him a refundable ticket, he will be willing to pay up to \$500 for the ticket today. If the ticket is not refundable, his willingness to pay for such a ticket should be less than \$500. In this paper, we propose a theory to explain how a monopolist can use refund contracts to screen risk-averse consumers, and we provide empirical evidence from the U.S. airline industry that supports this theory.

Under the expected utility framework, one may use Cicchetti and Freeman's (1971) option price, not to be confused with the price of a financial option, as an *ex ante* willingness to pay for a private good. The option price is the maximum willingness to pay for the ticket when it is bought in advance and the true state (travel or not travel) is not yet known. If the probability that you can make the trip is 0.5, then the expected willingness to pay is \$250 which may be greater or smaller than the option price[1]. Considering a monopolist who sells a non-refundable ticket in advance to heterogenous buyers, Cicchetti and Freeman show that option prices exist and the monopolist who has perfect information can perfectly price-discriminate by charging each individual his option price. In this paper, we apply the option price concept to a monopoly problem with asymmetric information and explain how a firm can use refund contracts as a means of second-degree price discrimination.

[1]When the state-dependent utility function is linear, the option price is equal to the expected willingness to pay. However, Schmalensee (1972) and Graham (1981) show that, when state-dependent utility functions are concave, the option price may be greater or smaller than expected willingness to pay.

The idea in the model is simple; the firm offers two types of contracts, refundable tickets and non-refundable tickets, to consumers who have either high or low willingness to pay for a seat on the corresponding flight. The consumers privately know their type, but do not know their demand when choosing a contract. We find that an optimal price menu consists of a refundable price that attracts only the high-type consumers and a non-refundable price that attracts only the low-type consumers. We extend the model to include the case where the firm offers the two types of tickets in multiple periods prior to departure and potential passengers learn about their individual demand as the date of departure approaches. On any day before departure, having a refundable ticket is more valuable to the consumers because circumstances may change and they can claim a refund if they do not want to fly. Based on our model, we find that the difference between these two fares consists of a quality component and a price-discrimination component. Moreover, as the date of departure approaches, the consumers become more certain about their individual demand, and the gap between the two fares diminishes. Using an original dataset from 96 monopoly routes in the U.S. airline market, we find strong evidence supporting the predictions of the model. Even though we focus on the airline industry throughout the paper, by no means are results from the model restricted to this industry. This pricing behavior can be observed in other industries where goods are sold in advance with a refundability option, such as cruises and lodging.

Early papers on price discrimination under asymmetric information include Mussa and Rosen (1978), Maskin and Riley (1984), and Landsberger and Meilijson (1985). To design a contract extracting surplus from each type of consumer, Mussa and Rosen use quality discounts, while Maskin and Riley use quantity discounts, and Landsberger and Meilijson use a time-decreasing price schedule. Recent literature on screening consumers in the airline industry includes Gale and Holmes (1993), Dana (1998), Courty and Li (2000), and Akan *et at.* (2008). Under a capacity constraint, when individual demand is uncertain, Gale and Holmes (1993) use advance-purchase discounts to increase the firm's expected profit by improving efficiency in capacity utilization. Dana (1998) introduces advance-purchase discounts as a means of price discrimination and improvement in capacity utilization in competitive markets. Courty and Li (2000) suggest a theoretical model for a monopolist that price discriminates via refund contracts consisting of a price that a buyer has to pay in advance and a refund that the buyer can receive after he learns his valuation of the good. In the first period, each traveler only knows the probability distribution of his valuation for the ticket, and in the second period, he learns his actual value for the ticket. While Courty and Li's purpose is to find an optimal refund contract consisting of advance payment and refundable amount, our goal is to find an optimal contract consisting of a non-refundable price and a fully refundable price for each type. Another difference from Courty and Li is that we allow consumers to be risk averse. Akan *et al.* (2008) present a generalization of Courty and Li with consumers that learn their valuations gradually and with the seller that can vary the length of time that the tickets are refundable.

Gale and Holmes (1993) is also related to our paper regarding the existence of individual demand learning. They consider a monopolist who faces capacity constraints and wants to expand output by diverting demand from a peak-time flight to an off-peak-time flight. Consumers are uncertain about their individual demands and the monopolist benefits from this by offering a discount for the off-peak flight together with an advance purchase requirement. It is related to our paper in the sense that consumers in Gale and Holmes also learn about their demand over time to find out which flight they prefer. In our model there is a single flight and consumers have to learn about whether or not to fly. While Gale and Holmes consider the existence of a single fare at each point in time, in our model the screening mechanisms works through an additional refundability option.

The paper is structured as follows. Section 2 presents the theoretical model. The two period model with homogeneous consumers is developed in Section 2.1, where we show the conditions under which the monopolist can make more profits by selling non-refundable tickets in advance. Section 2.2 includes the existence of heterogeneous consumers and sets the mechanism design problem that allows the seller to screen consumers and price discriminate. The model's implications for airline pricing are presented in Section 2.3. Section 3 presents the empirical analysis by describing the data (3.1), setting the empirical model (3.2), and presenting the results (3.3) that provide empirical support for the predictions of the theoretical model. Finally, Section 4 concludes.

2 Theoretical Analysis

2.1 Ex-ante willingness to pay

Consider a two-period model where, in period 1, a monopolist sells homogeneous goods to risk-averse consumers whose demand for the good is either 1 or 0. Without knowing his own demand with certainty, each consumer has to make a decision whether to buy the good in period 1. Then, he learns his demand in period 2. For example, airlines sell plane tickets in advance. Each potential passenger makes a decision about buying a ticket before he knows with certainty whether he wants to fly or not on the date of departure.

Formally, let a and b be mutually exclusive states in which a consumer's demand is 1 and

0, respectively. The risk in the state of nature that each consumer faces is an individual risk which is independent from those of other consumers. Let π denote the probability that state a occurs and hence the probability that state b occurs is $1-\pi$. In other words, in period 2, each consumer will be willing to pay up to a surplus denoted by $S > 0$ with probability π and 0 with probability $1-\pi$. As a result, the expected surplus is equal to πS. Let w be a consumer's wealth and $u_s^0(w)$ and $u_s(w)$ be utility functions in state s, for $s = a, b$, when the good is unavailable and available to the consumer, respectively. For example, $u_a^0(w)$ represents his utility function when he wants to travel but he does not have a ticket. Using the definition of compensating variation, we derive S from $u_a(w - S) = u_a^0(w)$ and since the consumer does not want the good in state b, then $u_b(w) = u_b^0(w)$.

We assume that, when the good is available for consumption, consumers prefer state a to state b, i.e., $u_a(w) > u_b(w)$, marginal utility of income in state a is higher than in state b, i.e., $u_a'(w) \geq u_b'(w) > 0$, and state-dependent utility functions are concave, i.e., $u_a''(w), u_b''(w) \leq 0$ for all w. The firm sells goods in advance by offering a price menu consisting of a (discounted) non-refundable price (D) and a refundable price (R). We assume that expected utility theory holds and thus define

$$U_d(D) \equiv \pi u_a(w - D) + (1-\pi)u_b(w - D) \tag{1}$$

as expected utility from buying the good at non-refundable price D, and

$$U_r(R) \equiv \pi u_a(w - R) + (1-\pi)u_b(w) \tag{2}$$

as expected utility from buying the good at refundable price R. Note that there is no time value of money. If a consumer buys a non-refundable ticket in period 1, his final wealth in period 2 will be $w - D$ irrespective of the state of nature. If the customer buys a refundable ticket instead and observes that the state of nature is a in period 2, he will claim a seat on the plane and his final wealth will be $w - R$. If the state of nature is b, then he will claim a refund and his final wealth will be equal to his initial wealth.

Since each consumer's surplus is S, the reservation price for a refundable ticket is S and the maximum willingness to pay for a non-refundable ticket is a real number θ such that $U_d(\theta) = U_r(S)$. After the airline announces D and R at the beginning of period 1, a consumer buys a non-refundable ticket if $D \leq \theta$ and $U_d(D) > U_r(R)$, and a refundable ticket if $R \leq S$ and $U_d(D) \leq U_r(R)$.[2] If $D > \theta$ and $R > S$, then he will not buy a ticket. Given a price menu (D, R), the consumer's decision to buy can be illustrated in Figure 1. The (D, R) space is divided into three regions by $U^0 \equiv \max\{U_d(\theta), U_r(S)\}$ and the dashed curve representing all (D, R) such that $D \leq \theta$, $R \leq S$, and $U_d(D) = U_r(R)$. Along the dashed curve, $R > D$, $\frac{dR}{dD} > 0$, and $\frac{d^2R}{dD^2} < 0$ if $u'_a(w) > u'_b(w) > 0$ and $u''_i(w) < 0$. In region 1 the consumer buys a non-refundable ticket, in region 2 the consumer buys a refundable ticket, and in region 3 the consumer does not buy any ticket. Note that the region 1 includes the vertical part of U^0 except (S, θ), and the region 2 includes the horizontal part of U^0 and the dashed curve.

Let τ denote the difference between θ and the expected surplus, i.e., $\tau \equiv \theta - \pi S$. The definitions of θ and τ are similar to Cicchetti and Freeman's (1971) option price and option value, respectively. Schmalensee (1972) and Graham (1981) show that when state-dependent utility functions are concave, an option value (τ) can take either sign. Lemma 1 shows that, under the aforementioned assumptions regarding the utility functions in our setup, $\theta \geq \pi S$ and hence τ is positive.

[2]When $U_d(D) = U_r(R)$, we assume, without loss of generality, that the consumer chooses a refundable ticket over a non-refundable ticket.

Lemma 1 *If* $u_a'(w) \geq u_b'(w) > 0$ *and* $u_s''(w) \leq 0$ *for* $s = a, b$, *then* $\theta \geq \pi S$.

Proof Since $u_a'' \leq 0$, then $u_a(w - \pi S) = u_a(\pi(w-S)+(1-\pi)w) \geq \pi u_a(w-S)+(1-\pi)u_a(w)$, which is equivalent to $\pi[u_a(w - \pi S) - u_a(w - S)] \geq (1 - \pi)[u_a(w) - u_a(w - \pi S)]$. The latter part is equal to $(1 - \pi)\int_{w-\pi S}^{w} u_a'(x)dx$ which is greater than $(1 - \pi)\int_{w-\pi S}^{w} u_b'(x)dx$ because $u_a'(w) \geq u_b'(w)$ for all w. It follows that $\pi[u_a(w - \pi S) - u_a(w - S)] \geq (1 - \pi)[u_b(w) - u_b(w - \pi S)]$ which can be rearranged as $\pi u_a(w - \pi S) + (1 - \pi)u_b(w - \pi S)] \geq \pi u_a(w - S) + (1 - \pi)u_b(w)$. Hence $U_d(\pi S) \geq U_r(S) = U_d(\theta)$. Therefore, it follows that $\theta \geq \pi S$. ■

Example 1. Consider an individual with a state-dependent utility function $u_a(w) = w^{0.7}$ in state a and $u_b(w) = w^{0.6}$ in state b. We can see that $u_a'(w) > u_b'(w) > 0$ and $u_a''(w), u_b''(w) < 0$ for all $w > 0$. This individual is considering buying a ticket for a flight departing in the future. Assume that his initial wealth is 2,000 and his surplus for the flight is 500. Thus his reservation price for a refundable ticket would be 500. If, as of today and based on what he knows about his individual demand, the probability that he will demand a seat on the departure date (π) is 0.5, he will be willing to pay 360 for a non-refundable ticket because $U_r(500) = U_d(360)$. Note that his ex-ante willingness to pay is higher than his expected surplus (250). ■

What should the monopolist do when the ex-ante willingness to pay is higher than the expected surplus? If there are N customers similar to the individual in Example 1 and 50 percent of them have a positive demand on the date of departure, the airline will be able to make more profit by selling these customers non-refundable tickets than selling refundable tickets. The revenue from selling non-refundable tickets at 360 each would be $360N$ as opposed to $250N$ if refundable tickets are sold at 500 each and half of the buyers claim a full refund. Therefore, the airline should set the non-refundable price at 360 and make the refundable price greater than 500 so that everyone buys a non-refundable ticket. In the next section, which considers more than one type of consumer, Lemma 1 implies that it is profitable for the monopolist to sell non-refundable tickets to at least one type.

2.2 Price discrimination

Consider now the case of consumers with heterogeneous willingness to pay. There are two types of consumers who privately know their own type: high type (h) and low type (l). They either demand 1 unit of the good (one trip) in state a or 0 units in state b. For consumers of either type, the probability that state a occurs is π. For $i = h, l$, type i consumer has an initial wealth w_i and a state-dependent utility function with $u'_{ia}(w) \geq u'_{ib}(w) > 0$ and $u''_{is}(w) \leq 0$ for $s = a, b$. If state a occurs, type i consumer is willing to pay S_i which can be derived from $u_{ia}(w_i - S_i) = u^0_{ia}(w_i)$.[3] Without loss of generality, let $S_h > S_l$. We define

$$U_{id}(D) \equiv \pi u_{ia}(w_i - D) + (1 - \pi)u_{ib}(w_i - D), \tag{3}$$

and

$$U_{ir}(R) \equiv \pi u_{ia}(w_i - R) + (1 - \pi)u_{ib}(w_i), \tag{4}$$

to be type i consumer's expected utility if he buys a non-refundable ticket at price D and a refundable ticket at price R, respectively. The maximum willingness to pay of type i consumer for a non-refundable ticket is θ_i which can be derived from $U_{id}(\theta_i) = U_{ir}(S_i)$. We define $\tau_i \equiv \theta_i - \pi S_i$.

[3]See definition of u^0 in Section 2.1.

Let the numbers of type h and type l consumers in period 1 be N_h and N_l, respectively. There are n_h of type h and n_l of type l consumers that demand the good in period 2. Let $n_h = \pi N_h$ and $n_l = \pi N_l$. The firm announces prices D and R at the beginning of period 1 and the firm's capacity is at least $n_h + n_l$. The information about each consumer's willingness to pay is private. To price discriminate via a price menu with $R > D$, the firm has to sell non-refundable tickets to the type l consumers and refundable tickets to the type h consumers. The firm's problem can be described as follows.

$$\max_{D,R} N_l D + n_h R \tag{5}$$

s.t.

$$U_{hr}(R) \geq U_{hd}(D), \tag{6}$$

$$U_{ld}(D) \geq U_{lr}(R), \tag{7}$$

$$U_{hr}(R) \geq U_{hr}(S_h), \tag{8}$$

$$U_{ld}(D) \geq U_{ld}(\theta_l). \tag{9}$$

The incentive-compatibility constraints (6) and (7) are required for the type h consumers to prefer buying a refundable ticket and the type l consumers to prefer buying a non-refundable ticket, while the participation constraints (8) and (9) are required for all consumers with positive valuation to buy a ticket. We consider the two possibilities: 1) $\theta_l > \theta_h$, and 2) $\theta_l \leq \theta_h$.

Proposition 1 *If $\theta_l > \theta_h$ and $N_l - N_h$ is large enough, the firm sets $(D, R) = (\theta_l, S_h)$.*

Proof Constraints (8) and (9) imply $R \leq S_h$ and $D \leq \theta_l$. When the two constraints are binding, i.e., $D = \theta_l$ and $R = S_h$, $U_{ld}(\theta_l) = U_{lr}(S_l) > U_{lr}(S_h)$ and hence constraint (7) is satisfied. Since $\theta_l > \theta_h$, then $U_{hd}(\theta_l) < U_{hd}(\theta_h) = U_{hr}(S_h)$ and hence constraint (6) is satisfied. When $(D, R) = (\theta_l, S_h)$, the firm's profit is equal to $N_l\theta_l + n_h S_h$. If $D > \theta_l$ and $R = S_h$ or S_l, the corresponding profit will be $n_h S_h$ or $(n_h + n_l)S_l$. If $R > S_h$ and $D = \theta_h$ or θ_l, the corresponding profit will be $(N_h + N_l)\theta_h$ or $N_l\theta_l$. By Lemma 1, $\tau_i \geq 0$ for $i = l, h$. When $S_l < S_h$, $\tau_i \geq 0$, and $N_l > \frac{\tau_h}{\theta_l - \theta_h} N_h$, the profit from selling with $(D, R) = (\theta_l, S_h)$ is greatest. ■

We illustrate this result in the (D, R) space in the left panel of Figure 2. We define $U_i^0 \equiv \max\{U_{id}(\theta_i), U_{ir}(S_i)\}$ for $i = l, h$. Region 1 satisfies constraints (6) to (9) and the firm's profit is maximized at point A. According to Proposition 1, the firm's profit is $N_l\theta_l + n_h S_h$, while the ex-post consumer surplus is $n_l(S_l - \theta_l) - (N_l - n_l)\theta_l$. As a result, social welfare is $n_l S_l + n_h S_h$. Note that when $\theta_l > \theta_h$, the consumer surplus of type h consumers is fully extracted by the firm in the second period. This is not true for the case where $\theta_l \leq \theta_h$.

Proposition 2 *If $\theta_l \leq \theta_h$, $N_l - N_h$ is large enough, and there exists a real number δ such that $U_{hr}(\delta) = U_{hd}(\theta_l)$ and $\delta \geq \frac{\theta_l}{\pi}$, then the firm sets $(D, R) = (\theta_l, \delta)$.*

Proof If constraint (9) is binding, then $D = \theta_l$. If constraint (6) is also binding, then $R = \delta$. Since $U_{hr}(\delta) = U_{hd}(\theta_l)$, then $\delta \geq S_l$. It follows that $U_{ld}(\theta_l) = U_{lr}(S_l) \geq U_{lr}(\delta)$ and hence constraint (7) is satisfied. Since $\theta_l \leq \theta_h$, then $U_{hr}(\delta) = U_{hd}(\theta_l) > U_{hd}(\theta_h) = U_{hr}(S_h)$ and hence constraint (8) is satisfied. When $(D, R) = (\theta_l, \delta)$, the firm's profit is equal to $N_l\theta_l + n_h\delta$. If $D > \theta_h$ and $R = S_h$ or S_l, the corresponding profit will be $n_h S_h$ or $(n_h + n_l)S_l$. If $R > S_h$ and $D = \theta_h$ or θ_l, the corresponding profit will be $N_h\theta_h$ or $(N_h + N_l)\theta_l$. By Lemma 1, $\tau_i \geq 0$ for $i = l, h$. When $S_l < S_h$, $\tau_i \geq 0$, $N_l > \frac{\theta_h - \theta_l}{\theta_l} N_h$, and $\delta > \frac{\theta_l}{\pi}$, the profit from selling with $(D, R) = (\theta_l, \delta)$ is greatest. ■

We graphically derive δ in the right panel of Figure 2. Region 2 contains all the price vectors satisfying the constraints (6) to (9) and the firm maximizes profit at point B. The value of δ in Proposition 2 has to be so low that a refundable ticket is more attractive than a non-refundable ticket to type h consumers, but it has to be higher than S_l so that type l consumers find a non-refundable ticket more attractive than a refundable ticket when $D = \theta_l$. Moreover, it has to be high enough so that it is profitable to separate two groups of customers. The difference $S_h - \delta$ is a gain to type h consumers with positive surplus, because the firm has to lower the refundable price to make a refundable ticket attractive to them. When $\theta_l \leq \theta_h$, the firm's profit is $N_l\theta_l + n_h\delta$, the ex-post aggregate consumer surplus is $n_l(S_l - \theta_l) - (N_l - n_l)\theta_l + n_h(S_h - \delta)$, and hence social welfare is $n_lS_l + n_hS_h$. Note that, as an information rent, $n_h(S_h - \delta)$ is transferred from the firm to type h consumers. This information rent does not exist in the case where $\theta_l > \theta_h$. To illustrate this result, consider the following example

Example 2. Let the individual in Example 1 be a type l consumer whose surplus is 500. We already know that when $\pi = 0.5$, the optimal non-refundable price is 360. Now assume that type h consumers, whose state-dependent utility functions are $u_a(w) = 2w$ in state a and $u_b(w) = w$ in state b, have an initial income of 2,000 and a surplus of 1,000. We find that $U_{hd}(360) = U_{hr}(720)$. Hence the seller's optimal refundable price is 720. With the price menu (360,720), type l consumers buy a non-refundable ticket at 360 and type h consumers buy a refundable ticket at 720. ■

2.3 Implications for the empirical analysis

We can apply our model to airline pricing, where refundable and non-refundable tickets are available in multiple periods prior to consumption. As described in (3) and (4), if a consumer of type i buys the good in advance at a non-refundable price D, then his expected utility is $U_{id}(D)$, and if he buys the good in advance at a refundable price R, then his expected utility is $U_{ir}(R)$. The good in the market is an airline seat in a flight departing on day $t = 0$. The firm starts offering tickets T days in advance, i.e., on day $t = T$, and offers these tickets everyday after that until the flight departs at $t = 0$ or until the flight sells out, whichever comes first. The firm offers a daily price menu consisting of a non-refundable price $D(t)$ and a refundable price $R(t)$. Similarly to the previous section, we assume that there are two types of consumers: type h whose surplus is S_h and type l whose surplus is S_l, with $S_h > S_l$. There are N_h and N_l consumers of type h and type l, respectively on each day.[4] On day T, consumers of both types enter the market. Depending on the non-refundable price $D(T)$ and on the refundable price $R(T)$, each type chooses one of the following actions: buy a non-refundable ticket, buy a refundable ticket, or leave the market. On the following day, another group arrives. Each consumer in this group either buys a non-refundable ticket at price $D(T-1)$, buys a refundable ticket at price $R(T-1)$, or leaves the market. The situation repeats everyday until departure, with consumers using the decision rule described in the previous section.[5] On day t, the

[4] As seen in Proposition 2, N_h and N_l can take different values. We only need $N_l - N_h$ to be large enough for the pricing strategy to be profitable.

[5] This formulation assumes independence across time periods, where each new group of individuals in unrelated to the previous group. The empirical formulation initially considers this case, but later relaxes this assumption to allow for dynamics.

probability that individuals will fly on the date of departure is $\pi(t)$. This probability can as well be interpreted as the knowledge, at the time of purchase, that an individual has about his individual demand. In an extreme case where a consumer who buys a ticket at exactly the same moment he walks into the aircraft, i.e. at $t = 0$, he is clearly very certain about traveling, and hence $\pi(0) = 1$. Since consumers who buy in advance are less certain about their travel demand, then π is expected to be lower. Therefore, it is reasonable to assume that $\pi(t)$ is decreasing in t: as the flight date approaches, individuals are expected to know more about their demand. Because $\pi(t)$ is decreasing in t, $\theta_i(t)$ is also decreasing in t, for $i = h, l$. Since $S_h > S_l$, then $\theta_h(t) > \theta_l(t)$ for all $t = 0, \ldots, T$. The firm that can sell two types of ticket, refundable and non-refundable, will set prices according to Proposition 2, that is, on day t, $D(t) = \theta_l(t)$ and $R(t) = \delta(t)$. Consider the following example.

Example 3. There are two types of passengers. A type l passenger has a state-dependent utility function $u_a(w) = w^{0.7}$ and $u_b(w) = w^{0.6}$ with a surplus for the flight equal to 500. A type h passenger has a state-dependent utility function $u_a(w) = 2w$ and $u_b(w) = w$ with a surplus of 1,000. Both types have an initial wealth of 2,000. From Example 2, in which the probability of flying (π) is 0.5, we find that an optimal price menu is (360,720). In this example we solve for price menus given various values of π in [0.5,1] interval. This result is illustrated in Figure 3. ■

From the above example, we obtain an important empirical implication from the theory. The difference between the refundable and non-refundable fare, $\delta(t) - \theta_l(t)$ decreases as π increases. This implies that as the departure date nears and consumers learn about their individual demands, the gap between refundable and non-refundable fares closes. At the day of the flight, when individual demand uncertainty is resolved, there will only prevail a single price. There does not exist any additional benefit from buying a refundable ticket. The seller can no longer separate consumer types based on their individual demand uncertainty. Section 3 seeks to find support for this empirical prediction. Moreover, the rate at which fares get closer together gives us important information about what the carrier believes is the speed at which consumers learn about their individual demand. In addition to this prediction, Figure 4 shows that the difference in fares can be broken down in two parts: the difference in quality $S_l - \theta_l(t)$, and the price discrimination component $\delta(t) - S_l(t)$.

3 Empirical Analysis

3.1 Data

To see if actual airline fares are consistent with the theoretical model, we collected from the online travel agency *expedia.com* the lowest refundable and lowest non-refundable one-way economy-class posted fares for 96 flights that departed on June 22^{nd}, 2006. Following Stavins (2001), we focus on a single day, Thursday, to avoid having to control for price differentials associated with systematic peak load pricing. Some days are known to be more congested than others, therefore systematic peak load pricing suggests higher prices during these more congested days. The data forms a panel with 96 cross section observations and 28 observations in time. Each cross section observation corresponds to a specific carrier's non-stop flight between a city pair. Fares were recorded every three days, with the first cut in time corresponding to 82 days prior to departure and the last to one day prior to departure. The carriers considered are American, Alaska, Continental, Delta, United, and US Airways. The share of each of the carriers' flights on the dataset was chosen to be close to the share of the carrier in the U.S. market.

By picking non-stop flights and one-way fares, we control for price differences associated with fare restrictions and cost differences associated with round trip tickets (e.g., Saturday-night-stay, minimum- and maximum-stay) as well as controlling for variation in prices related to more sophisticated itineraries (e.g., travelers connecting to different cities). Dealing with non-stop tickets is particularly important in our model because for round-trip tickets, the traveler may also be uncertain about the return portion of the ticket. By picking economy-class tickets we control for other potential sources of consumer's heterogeneity, as first-class travelers may behave differently from economy-class travelers. Moreover, there also exists price dispersion associated with frequent flyer programs, as some consumers obtain tickets from miles earned from previous trips with the carrier. These rewarded tickets are excluded from the sample. The goal in this section is to use a restricted sample of tickets to find if there is support for the empirical predictions from the model. Complex itineraries, additional fare classes, round-trip tickets, or international destinations would make the results easier to generalize to the entire industry, but would impose an important burden on the empirical section. Price dispersion could then be the result of different dimensions of price discrimination, marginal cost pricing, or product differentiation. In this paper we do not attempt to explain multiple sources of price dispersion in the airline industry; we aim to explain the difference of refundable versus non-refundable fares.

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, all 96 routes in the sample are considered monopoly routes, as the carrier in the route 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 a significantly different quality.

To illustrate the data, Figure 4 shows the average of refundable and non-refundable fares across all 96 monopoly routes at different points prior to departure. Because routes with higher average fares have a greater weight on the dynamics of the average fares across flights, Figure 4 partials out the flight fixed effects from average fares. The figure shows that if carriers do not vary in their flight specific characteristics, there is a strong tendency for non-refundable fares to increase faster than refundable fares. Comparing this graph with its theoretical counterpart, Figure 3, it shows that as the flight date nears, consumers resolve their individual demand uncertainty.

3.2 Empirical model

As explained in the previous section, the construction of the dataset controls for various sources of price dispersion in airline ticket pricing that arise within the same flight. This allows us to focus on refundable and non-refundable fares for the same flight at different points in time prior to departure. In this section we present the empirical specification and explain how we control for unobserved time-invariant flight-specific characteristics and unobserved time-variant seat-specific characteristics.

The panel structure of the data allows us to estimate the model controlling for time-invariant flight-specific characteristics. Even though we chose to call them unobserved flight-specific characteristics because each cross-sectional observation in the panel is a flight, it is important to notice that each flight belongs to a carrier and to a route. Hence, we are controlling for unobserved carrier-specific characteristics and unobserved route-specific characteristics as well. This represents a control for all time-invariant characteristics included as regressors in Stavins (2001), who used a cross section of tickets in her analysis. Flight-specific characteristics include aircraft type and information about congestion that is known at the time the flight is scheduled (for example, an 8:00 a.m. flight may be priced differently from an 11:00 a.m. flight); route-specific characteristics include route distance and airport characteristics; and carrier-specific characteristics include managerial capacity and brand loyalty associated with a frequent flyer program. All these time-invariant characteristics can be associated with various sources of costs that arise at the flight-level.

There are, however, important cost components that arise at the individual seat-level that are part of the unobserved seat-specific characteristics. These change from seat to seat and could depend, among other things, on the number of days to departure, the number of available seats in the aircraft, the probability that expected demand will exceed capacity, and current and expected fuel costs. Borenstein and Rose (1994) and Dana (1998, 1999) explain the existence of different cost components that change from seat to seat and arise because of capacity constraint.

Borenstein and Rose (1994) consider two cost-based sources of price dispersion. Both are types of peak-load or congestion pricing. The first source is *systematic* peak-load pricing which reflects variations in the expected shadow cost of capacity at the time the flight is scheduled. This implies that higher fares will be set in congested periods known at the time of departure. Under the assumption that *systematic* peak load pricing affects all tickets in the aircraft in the same way, this one can be regarded as an unobserved flight-characteristic. The second cost-based source is *stochastic* peak-load pricing, and refers to uncertainty in the aggregate demand at the flight level once capacity choices are made. At any point prior to departure, higher fares will be set if the flight is expected to be peak, i.e., demand is likely to be greater than allocated capacity. The resulting variation in prices will be a consequence of price flexibility and how airlines learn about aggregate demand as the departure date nears and there are fewer seats available in the aircraft. *Stochastic* peak-load pricing changes from seat to seat and over time and is regarded as an unobserved time-variant seat-specific characteristic.

On the other hand, Dana (1998, 1999) considers the existence of two different measures of costs associated with each ticket. He assumes that carriers commit to a schedule of prices and cannot depart from this schedule as sales progress. The first measure of cost is the unit marginal costs of production, or operational marginal cost. This corresponds to the cost incurred only with the seats that are sold. The second is the unit cost of capacity, which is called effective cost of capacity when adjusted by the probability that the ticket is sold. The key idea here is that different seats within the same flight will have different selling probabilities, hence different effective costs of capacity. Both of these costs can be viewed as unobserved time-variant seat-specific characteristics.

Obtaining a measure for these alternative legitimate cost definitions would be a formidable task. It would require, at a minimum, information about each seat's selling probability and expectations about the demand each time a seat is sold. The difficulty in measuring or controlling for these costs, and any other cost that can change across seats, e.g., fuel costs, has limited the availability of empirical papers that assess price discrimination practices in airlines. Stavins (2001), who looks at the implied price differentials due to ticket restrictions, such as the Saturday night stay-over restriction and advance purchase discount restrictions, is an important contribution. However, to correctly point out that a pricing strategy is discriminatory, price differences should be net of all cost components. Stavins regards those ticket restrictions as being fully price discriminatory. The problem is that those ticket restrictions are understood to solve the peak-load pricing problem as well (see Courty and Li, 2000 p. 716). Dana (1999) states that price differences associated with ticket restrictions potentially consist of a price discrimination component, a peak-load pricing component, and a quality component, which can be used to analyze demand uncertainty for a perishable asset.

To control for all these unobserved time-variant seat-specific characteristics, we take advantage of the fact that both posted fares, refundable and non-refundable, were obtained at the same point in time for the same seat in the aircraft. Therefore, taking the difference between these two fares wipes out all unobserved seat-specific characteristics. This price difference, illustrated in Figure 3 as $\delta - \theta_l$, isolates the price discrimination and quality components. By looking at the dynamics of this price difference, we can analyze the individual demand learning implied by the airlines' pricing strategy. The logarithmic specification of the dynamic reduced form model that we will estimate is given by

$$\begin{aligned} \ln(REF)_{ijt} - \ln(NONREF)_{ijt} &= \alpha[\ln(REF)_{ij,t-1} - \ln(NONREF)_{ij,t-1}] \\ &\quad + \beta_1 DAYADV_{ijt} + \beta_2 \Delta LOAD_{ijt} + \nu_{ij} + \varepsilon_{ijt}, \quad (10) \end{aligned}$$

where the subscript i refers to flight, j to route, and t is time. Let ν_{ij} denote the unobservable flight specific effect and ε_{ijt} denote the remaining disturbance. Also $\ln(REF)$ is the natural logarithm of the price for the refundable ticket, $\ln(NONREF)$ is the natural logarithm of the price for the non-refundable ticket, $DAYADV$ is the number of days in advance prior departure the posted fares where recorded, and $LOAD$ is the number of occupied seats as a fraction of total seats in the aircraft. $\Delta LOAD$, defined as

$\Delta LOAD_{ijt} = LOAD_{ij,t} - LOAD_{ij,t-1}$, is a measure of sales in period t.

From the summary statistics presented in Table 1, we can calculate that refundable fares are, on average, 50.87% 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 refundable fares. Given that refundable fares are higher than non-refundable fares, the theoretical framework predicts that β_1 should be positive−as departure date nears and consumers learn about their individual demand, the gap between these two fares decreases. The variable $\Delta LOAD$ is included to take into account the possibility that sales may not be uniform over time. The theoretical framework assumes that the number of tickets sold in each period is constant; however, this may not be the case as more tickets are usually sold closer to departure. Figure 5 illustrates the evolution of sales by showing a nonparametric regression of daily sales as a percentage of total capacity on days prior departure using the observations across the 96 monopoly flights. The bandwidth of 1.53 was obtained by least squares cross-validation. There are few transactions 48 days or more prior to departure, and sales are larger during the last month.

The coefficient α is not of direct interest, but including a dynamic component in the equation relaxes the strict exogeneity assumption common to static panel data models and allows us to assume only weak exogeneity of the regressors. Consistent estimates of the parameters can then be obtained using the *difference* and *system* GMM estimators proposed in Holtz-Eakin *et.at.* (1988), Arellano and Bond (1991), Arellano and Bover (1995), and Blundell and Bond (1998). These estimation procedures create a more flexible empirical specification by allowing consumers to behave dynamically, while controlling for unobserved flight-specific characteristics. Weak exogeneity of $\Delta LOAD$ means that sales in period t can be affected by current and past realizations of the differences in fares, but must be uncorrelated with future realization of the error term. Consumers observe previous differences in fares, implying that the decision to buy a ticket today can be affected by previous price levels. Furthermore, weak exogeneity does not mean that consumers do not take into account expected future changes in fare differences in their decisions to buy or not to buy a ticket; it means that future unanticipated shocks in fare differences do not influence the decision to buy a ticket.

If consumers are very certain about their travel plans, then carriers will not be able to gain much from offering a menu of fares. The true benefit of screening among travelers comes from the fact that various days before departure, individuals do not know exactly their valuations for a ticket, and the seller can exploit this imperfect knowledge by making the consumer reveal information gradually. Therefore, measuring the speed at which the price gap closes reveals what the carrier's pricing strategy implies about consumers' individual demand learning. Equation (10) assumes that the speed at which consumers obtain private information about their individual demand is constant. A straight forward way to test for nonlinearities would be to include a quadratic and cubic term for $DAYADV$. However, these alternative specifications are still restrictive. To address this issue more accurately, we also estimate a more flexible nonparametric panel data regression model with fixed effects of the form

$$\ln(REF)_{ijt} - \ln(NONREF)_{ijt} = g(DAYADV_{ijt}, \Delta LOAD_{ijt}) + \nu_{ij} + \varepsilon_{ijt}, \qquad (11)$$

where $g(\cdot)$ is an unknown smooth function, and as before ν_{ij} is the flight fixed effect and ε_{ijt} is the remainder effect. The fixed effects to control for flight-specific characteristics are kept outside the smoothing function to avoid the curse of dimensionality. Even though we enjoy greater flexibility with this specification, a drawback is that we need to assume strict exogeneity. Given the nature of the regressors, $DAYADV$ is discrete and $\Delta LOAD$ is continuous, $g(\cdot)$ will be estimated using kernel methods for mixed data types as explained in Racine and Li (2004) and Li and Racine (2007). The smoothing parameters will be calculated with least squares cross-validation.

3.3 Results

The results from the estimation of Equation (10) are reported in Table 2. For comparison purposes, the first column shows the OLS estimates of the pooled regression model. The figures in parentheses are t-statistics based on heteroskedasticity-consistent estimates of the asymptotic standard errors. Here, $DAYADV$ has no significant effect on the price gap; however, this specification does not control for flight-, route- and carrier-specific characteristics. To control for these characteristics, the second column presents the within flight regression estimates. Exploiting the panel structure of the data allows control for all time-invariant regressors that affect price dispersion included in Borenstein and Rose (1994) and Stavins (2001), as well as controlling for other unobserved flight-, carrier-, and route-specific characteristics. Consistent with the theoretical predictions, the within flight estimates show that the coefficient on $DAYADV$ is positive and highly significant, meaning that the carriers' pricing strategy implies that consumers learn their preferences as the flight date nears.

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 in Equation (10) to eliminate time invariant flight and carrier specific characteristics and assumes that the error term, ε, is not serially correlated. Moreover, in this estimation the series $\Delta LOAD_{ijt}$ may be endogenous in the sense that $\Delta LOAD_{ijt}$ is correlated with ε_{ijt} and earlier shocks, but $\Delta LOAD_{ijt}$ is uncorrelated with $\varepsilon_{ij,t+1}$ and subsequent shocks. Then, lagged values of $\Delta LOAD_{ijt}$ are valid instruments in the first-differenced equation. Column three uses $\Delta LOAD_{ij,t-2}$ as instrument, while column four uses $\Delta LOAD_{ij,t-3}$ and earlier as instruments.[6] To deal with the problem that the new error term, $\varepsilon_{ijt} - \varepsilon_{ij,t-1}$, is correlated with the lagged dependent variable, $\ln(REF)_{ij,t-2} - \ln(NONREF)_{ij,t-2}$ is also used as an instrument. Consistent with the theoretical section, $DAYADV$ is positive and significant in both GMM *difference* specifications.

To address the validity of the specifications, columns 3 and 4 also report three 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 modest 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 for column 3 and of lagged levels dated $t-3$ (and earlier) as instruments for column 4 is not rejected in any of the specifications at conventional significance levels. The difference Sargan test reported in column 4 that tests the validity of the additional instruments, $t-3$, in this specification accepts its validity.

Blundell and Bond (1998) point out one statistical shortcoming with the GMM *difference*

[6] $DAYADV_{ijt}$ is considered strictly exogenous in these estimations, since it was selected by the researcher to change every three days from 82 days in advance to 1 day in advance.

estimator. 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, we use the *system* estimator as suggested in Blundell and Bond (1998). 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 the following additional assumption: There is no correlation between the differences of $\Delta LOAD$ and the flight-specific effects, but there may be correlation between the levels of $\Delta LOAD$ and the flight specific effects.

Columns 5 and 6 of Table 2 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 still shows mild support to the assumption of no serial correlation, while the Sargan and the difference Sargan provide strong support for the validity of the instrument list and the additional set of instruments used in column 6. $DAYADV$ is positive and highly significant, providing now even stronger support to the theoretical predictions.

Table 3 presents the results for fares in levels. The estimates are very similar to the ones reported in Table 2 with one important difference. Now the serial correlation test provides strong support to the assumption of no serial correlation in all four GMM specifications. Sargan and difference Sargan tests support the validity of the instruments and an additional set of instruments. As suggested by the theory, $DAYADV$ is positive and highly significant in the *system* GMM specifications. When comparing the *within* specification with the GMM specifications, it is interesting to notice that once we control for the potential endogeneity of $\Delta LOAD$, it is no longer statistically significant. None of the estimated coefficients in the GMM specifications show a significant coefficient for our proxy for sales, meaning that the exogenous component of $\Delta LOAD$ has no impact on the price gap. This in evidence that the effect of $\Delta LOAD$ on fare differences reported in the *within* specification comes just from the interaction from previous fare differences and sales.

The results above show that consumers learn about their individual demands, but says little about the speed at which consumers learn. Fares' rate of convergence is assumed to be constant because all specifications are linear on $DAYADV$. Alternative specifications could include a quadratic of cubic terms; however, we opted for a more flexible nonparametric model as represented in Equation (11). Nonparametric regression results for Equation (11) are illustrated in Figure 6. Because fares are in logarithms and days prior to departure is a time measure, the slope of the nonparametric regression can be interpreted as the speed at which carriers expect consumers to learn about their demands. While there is some learning throughout the time horizon considered, most of the learning takes place during the last two weeks prior to departure. This is when most of the individual demand uncertainty is resolved. By looking at Figure 5, we can see that the existence of two different fares is justified because the bulk of ticket sales occur during the last month prior to departure, when consumers' individual demand uncertainty is still important and the seller can use this uncertainty to extract more surplus from the consumer. In this case refunds act as an insurance against some uncertainties related to consumption.

The fact that individual demand uncertainty is not resolved by the time consumers buy a ticket contrasts some other alternative price dispersion models for airlines, including Dana (1999). He considers that all individual demand uncertainty is fully resolved at the moment of purchasing the ticket. In this case price dispersion is driven by capacity constraints, in particular the combination of costly capacity and demand uncertainty.

4 Conclusion and Discussion

The paper sets out to show the importance to airlines of offering a menu of prices, namely refundable and non-refundable fares. The theory section investigates the conditions under which a monopolist can increase its profits by offering different refund contracts in advance. We show that airlines can price discriminate when buyers with heterogeneous willingness to pay are uncertain about their demand for travel. 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, hence there should be no difference between these two fares.

The empirical section looks at the dynamics of prices in 96 monopoly routes and tests whether the individual demand uncertainty implied by the carrier's pricing strategy are resolved as the departure date approaches. After controlling for unobserved time-invariant flight-, carrier-, and route-specific characteristics, unobserved time-variant seat-specific charateristics, 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. Nonparametric regressions show that most of the individual demand uncertainty is resolved during the last two weeks, when the opportunity for price discrimination declines.

References

[1] Akan, M., B. Ata and J.D. Dana, Jr., 2008, Revenue Management by Sequential Screening, Working Paper, Kellogg School of Management, Northwestern University.

[2] 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.

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

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

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

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

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

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

[9] 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.

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

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

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

[13] Landsberger, M. and I. Meilijson, 1985, Intertemporal Price Discrimination and Sales Strategy under Incomplete Information, *Rand Journal of Economics* 16, 424-430.

[14] Li, Q. and J.S. Racine, 2007, Nonparametric Econometrics, Princeton University Press, Princeton, NJ.

[15] Maskin, E. and J. Riley, 1984, Monopoly with Incomplete Information, *Rand Journal of Economics* 15, 171-196.

[16] Mussa, M. and S. Rosen, 1978, Monopoly and Product Quality, *Journal of Economic Theory* 18, 301-317.

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

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

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

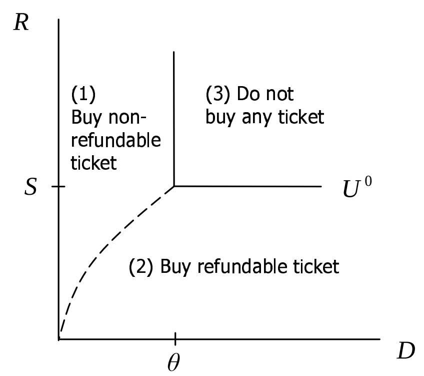

Figure 1: Type of ticket to buy in (D, R) space

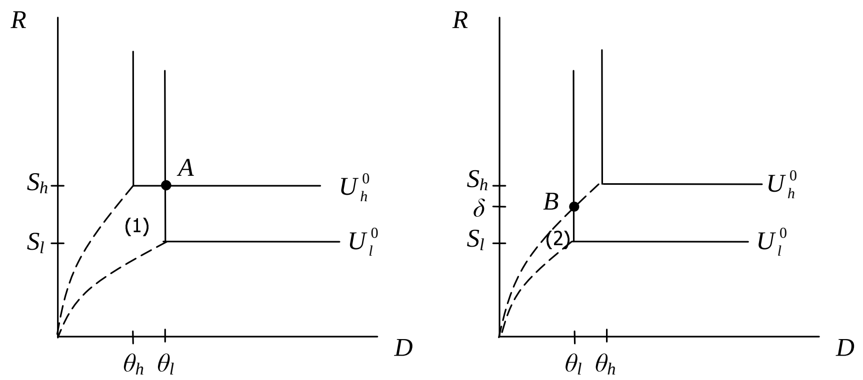

Figure 2: Type of ticket to buy in (D, R) space for heterogeneous buyers

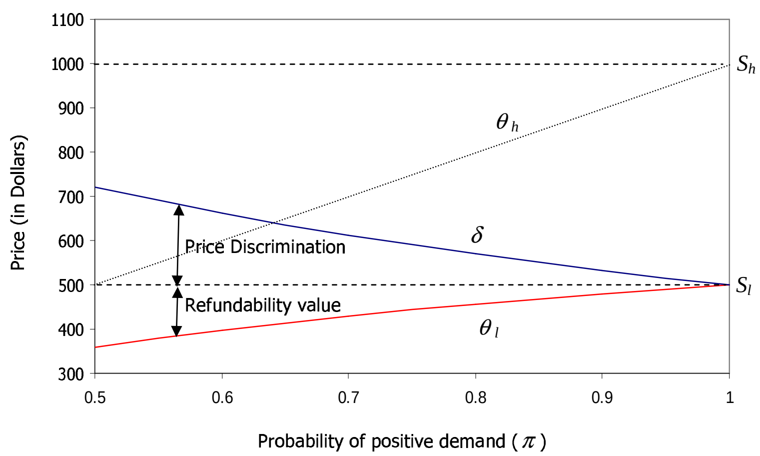

Figure 3: Optimal refundable and non-refundable prices

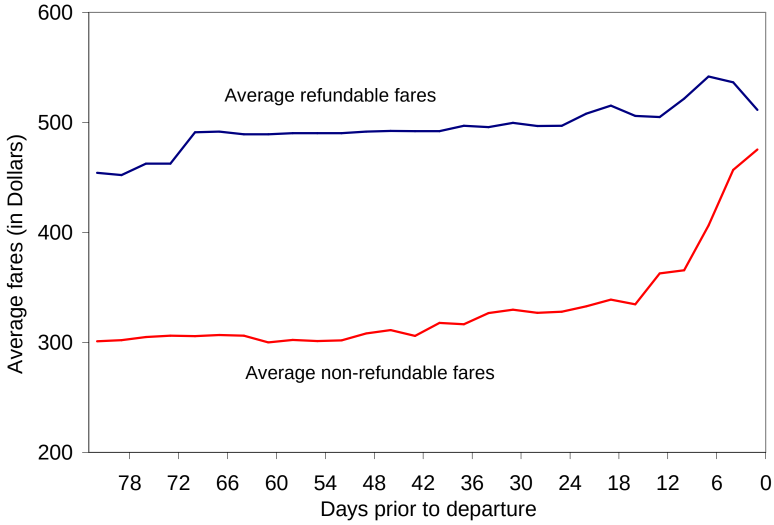

Figure 4: Average refundable and non-refundable fares at different t while controlling for flight specific characteristics

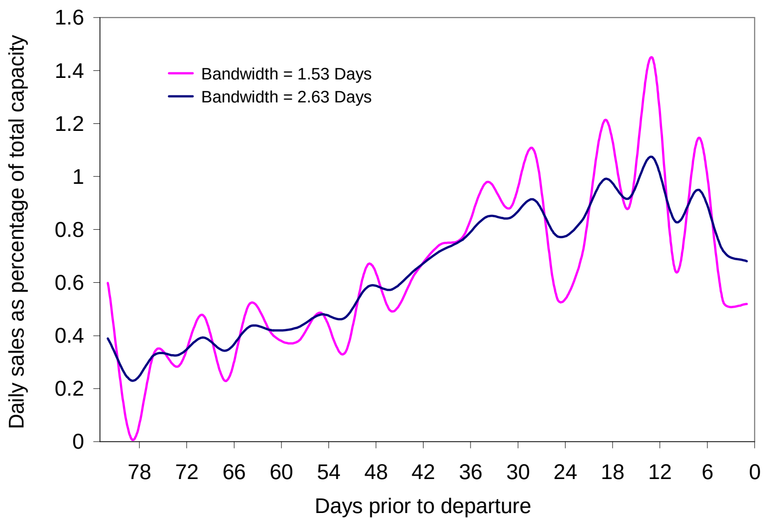

Figure 5: Nonparametric regression of daily sales on days to departure

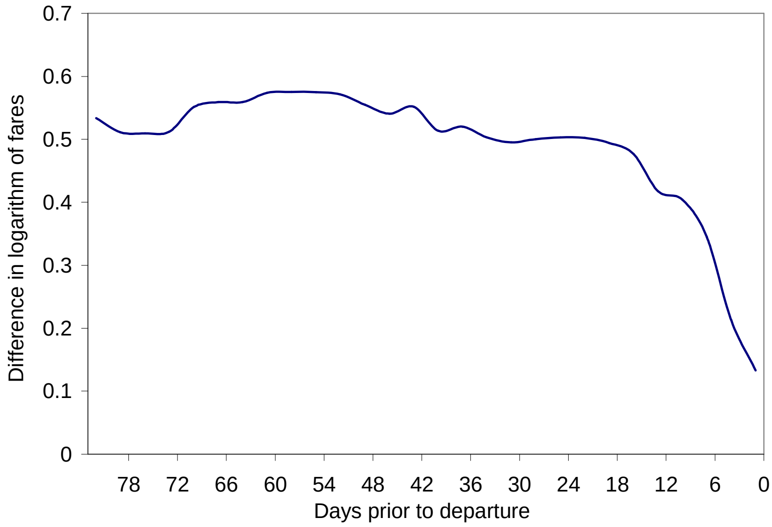

Figure 6: Nonparametric regression of difference on log of fares on days to departure

Table 1: Summary statistics

Variables	Mean	Std. Dev.	Min.	Max.	Obs.
REF					
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^a
NONREF					
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^a
DAYADV	41.500	24.238	1.000	82.000	2688
LOAD	0.591	0.241	0.038	1.000	2688

Notes: ^{a} Number of observations in time $= \bar{T}$, with one observation every three days.

Table 2: Regression estimates (prices in logs)

	OLS	Within	GMM *difference*		GMM *system*	
	levels	groups	$t-2$	$t-3$	$t-2$	$t-3$
$\ln(REF)_{ij,t-1} - \ln(NONREF)_{ij,t-1}$	0.944	0.697	0.879	0.868	0.788	0.790
	(120.459)	(19.049)	(8.904)	(9.357)	(10.663)	(11.193)
$DAYADV_{ijt}/10^3$	1.086	1.664	0.998	1.040	1.629	1.571
	(0.536)	(11.326)	(2.541)	(2.873)	(4.971)	(4.979)
$\Delta LOAD_{ijt}$	-0.001	0.180	0.232	0.138	-0.160	-0.112
	(-0.006)	(2.016)	(0.713)	(0.567)	(-0.545)	(-0.484)
Serial correlation test[a] (p-value)			0.032	0.037	0.035	0.035
Sargan test[b] (p-value)			0.133	0.178	0.691	0.988
Difference Sargan test[c] (p-value)				0.463		1.000

Notes: The dependent variable is $[\ln(REF)_{ijt} - \ln(NONREF)_{ijt}]$ and the number of observations is 2519. t-statistics in parentheses based on White heteroskedasticity robust standard errors for the first and second columns. t-statistics in parentheses based on Windmeijer WC-robust estimator for the GMM specifications. [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 3: Regression estimates (prices in levels)

	OLS	Within	GMM *difference*		GMM *system*	
	levels	groups	$t-2$	$t-3$	$t-2$	$t-3$
$REF_{ij,t-1} - NONREF_{ij,t-1}$	0.940	0.648	0.906	0.908	0.757	0.760
	(92.828)	(10.131)	(3.658)	(3.817)	(7.106)	(7.128)
$DAYADV_{ijt}$	0.387	0.488	0.290	0.311	0.479	0.456
	(5.459)	(9.361)	(1.585)	(1.695)	(5.054)	(4.881)
$\Delta LOAD_{ijt}$	17.537	67.696	195.139	124.231	7.859	11.439
	(0.581)	(2.467)	(1.840)	(1.194)	(0.086)	(0.152)
Serial correlation test[a] (p-value)			0.410	0.411	0.410	0.410
Sargan test[b] (p-value)			0.084	0.155	0.691	0.988
Difference Sargan test[c] (p-value)				0.566		1.000

Notes: The dependent variable is $[REF_{ijt} - NONREF_{ijt}]$. See notes on Table 2.