Randomization and price discrimination: The profitability of a mixed pricing strategy for airfares

Master Essay

Author: Jelal Younes

Supervisor: Håkan Jerker Holm

Seminar date: June 4, 2010

Lund University, Department of Economics

Abstract

In the airline industry, it is critical for carriers to vary prices offered to different customer groups in order to extract maximum willingness-to-pay from each consumer. This essay investigates a dynamic form of third-degree price discrimination in which prices are strategically adjusted as departure date approaches. It is hypothesized that by including a stochastic element in their pricing schemes, airlines can induce customers to self-select based upon their reservation prices, improving their profitability. Specifically, a mixed, randomized strategy in which expected price decreases over time may prompt risk-averse consumers with high reservation prices to purchase before customers with lesser reservation prices are offered a lower market price. Following this theoretical inquiry, an empirical examination reveals the possibility that such a pricing strategy is incorporated into pricing of airfares in the current market. Time series of flight data are studied, with a decomposition of fare prices into their systematic components revealing that an additional stochastic element may indeed be present.

Keywords: Price discrimination; airline pricing; mixed pricing strategy; inter-temporal pricing

Table of Contents:

Figures:

Tables [Appendix]:

I) Introduction

Yield management is the process of managing sales of a fixed resource to consumers with heterogeneous demand in order to maximize profits. Units of the resource in question are offered strategically such that consumers purchase specific commodities tailored to their demand, facilitating the extraction of maximum willingness-to-pay from each consumer (Ingold et al. 2000: 3). Yield management is particularly useful in the airline industry (indeed, the term originated in the industry), and has become common practice in the air transport industry since deregulation (Button and Vega 2007: 14). As flights have a fixed capacity and seats are "perishable"—that is, they generate no revenue if left unsold—airlines seek to maximize seats sold on each flight, while selling to as many high-demand customers as possible. Demand for airline seats is, of course, of varying elasticity, so it can be very profitable for airlines to adjust prices in order to attract, at different times, various elements of a diverse market (Ingold et al. 2000: 3-4). Anderson and Wilson (2003) note that yield (or revenue) management is of great importance in the airline industry, as even a slight increase in revenue per flight can have a significant impact upon an airline's profitability (p. 299).

It is well known that airlines practice third-degree price discrimination as a form of yield management. In particular, airlines offer student and senior discounts, business and economy class fares, non-linear one-way versus return flight pricing, and varying fares based upon place and time of sale. It is the last of these means of discrimination upon which this essay focuses: dynamic third-degree price discrimination based upon time (until departure) of sale. Specifically, this paper will investigate a form of discrimination in which airlines vary flight prices as time to departure approaches, inducing customers to self-select in their purchases based upon their reservation prices. It will be suggested that airlines must practice a mixed (randomized) pricing strategy in order for such a method of price discrimination to improve profitability.

Following a brief review of the relevant literature, this study proceeds to develop a theoretical model of airline ticket purchasing. This analytical model will reveal that a mixed pricing strategy can be profitable under the assumptions that the airline faces customers with

heterogeneous reservation prices and some degree of risk aversion. Subsequently, the analytical model will be expanded to allow for a more dynamic market in which customer purchases are based upon a "probability-of-purchase" function. It will be demonstrated, with some numerical illustrations, that the conclusions from the analytical model can be sustained under the new, generalized conditions. This second analytical model is included only in the Appendix to this essay.

Next, we move on to an empirical analysis of the conclusions generated in the preceding theoretical discussion. First, a number of short time series of flight prices are presented, with a view to the hypothesis that temporal pricing involves a random, strategic element. Afterwards, a longer time series of a single flight is subjected to a more thorough analysis. It is shown that much of the variability in flight prices as departure date approaches can be eliminated by controlling for a number of systematic factors. However, some variation remains, which it is suggested may be evidence of the randomized pricing strategy that this paper hypothesizes.

II)   Literature review

The literature is rife with studies focusing on price dispersion in the airline industry. In particular, it is common to explore the link between price dispersion in airfares and competition among carriers at a particular airport or on a particular route. Bilotkach (2005) asserts a consensus among researchers that fare dispersion depends positively on competition, with studies such as Borenstein and Rose (1991), Stavins (1996) and Dana (1999) generating this conclusion. Bilotkach (2005) suggests two general reasons for this sustenance of price dispersion: constrained capacity with uncertain demand, and price discrimination. It is the latter of these that is of particular interest to the study at hand.

Stavins (1996), in a highly cited article on the topic of price discrimination and market concentration in the airline industry, concludes that price discrimination is greater on more competitive routes. This conclusion stands in contrast to those that economic theory might suggest. Borenstein and Rose (1991) reiterate Stavins' finding, suggesting that further

theoretical work may be necessary to explain pricing behavior in monopolistically competitive markets. On a related thread, several other studies (such as Baye et al. 2004) investigate whether price dispersion in the industry is a disequilibrium phenomenon that is being corrected over time, or an equilibrium phenomenon. If the latter, the authors suggest, price dispersion should vary with market structure—which it does, in fact, according to the authors discussed above.

Unlike the aforementioned studies, this paper does not focus on market structure in the airline industry. While we use data on fares offered by various airlines, we do not contrast these in terms of the best-response functions that they might imply in a competitive environment. Rather, we are interested solely in the discrimination evident in the prices of the various airlines. These prices are presented as unilateral strategies on the part of each airline. Depending upon the degree of product differentiation in the airline market, this may or may not be an appropriate assumption. As noted, Borenstein and Rose (1991) suggest that the market is one of monopolistic competition. Button and Vega (2007), similarly, assert that the airline market falls somewhat short of perfect competition, despite considerable competition and freedom of entry and exit. Dana (1999) finds that optimal price strategies for *both* monopolist and oligopolist firms exhibit intra-firm price dispersion, possibly generalizing the conclusions of the monopoly model.

This paper differs from some of the previous research (Anderson and Wilson 2003, for example), in analyzing strictly the potential for strategic behavior on the part of the airline, without examining the advantages of such behavior from the perspective of the consumer. The theoretical analysis contained herein is closest to those studies that investigate inter-temporal pricing when consumers have varied reservation prices. Landsberger and Meilijson (1985), for example, contend that inter-temporal price variations may arise when firms attempt to discriminate among customers based upon their reservation prices. Other theoretical studies examining the effect of reservation prices on monopolistic pricing include Stokey (1979) and Besanko and Winston (1990).

In the empirical analysis presented in this paper, we focus on time series of price data for a flight on an unchanged route. Previous studies on changes in fares offered as departure date approaches include Bilotkach et al. (2006), who derive the general conclusion that fare prices increase at an accelerated rate as time to departure approaches, while also finding clear differences in pricing strategies by different airlines. In the same study, the authors point out that this type of longitudinal study is relatively uncharted in the literature: "Relatively few facts are known about airline fares at the micro level…the dynamics of changes in fares as the departure date nears has not been studied" (p. 2). Alam et al. (2001) perform time series analysis on flight data, though towards a somewhat different purpose—analyzing long-run relationships in prices between firms on various routes. Pels and Rietveld (2004) also base their study on a time series of flight prices, this time examining competitive price movements over time, and particularly the market influence of low-cost carriers. While each of the studies discussed above analyzes similar sets of time series data to those used in this essay, the research here is somewhat different. Instead of examining time trends of prices, we analyze the time series data to see how much of the observed price variation can be explained through systematic factors. Furthermore, we use data across a number of carriers not to compare fares between carriers, but to generalize our conclusions to a larger segment of the air transport industry.

To summarize: The research contained herein stems from the popular area of research on the link between price dispersion and market structure. Specifically, it focuses on price discrimination on the level of a single firm, with the assumption of sufficient product differentiation to ensure a market of monopolistic competition. In the paper's theoretical section, models are developed which differ from previous research by incorporating risk aversion among consumers, and focusing on profit maximization from the strategic perspective of the single airline. The conclusions that will be asserted regarding the viability of randomized pricing strategies are believed to be heretofore unseen in related literature. The empirical section of the paper follows uncharted territory in analyzing the dynamics of pricing as time passes with the goal of separating systematic from stochastic price variation. The manner in which data is gathered is unusual, allowing for a different perspective on these price dynamics. Finally, the paper attempts to bind together the unique aspects of the theoretical and empirical findings in its conclusion.

III) Theoretical model of airline ticket purchasing

One-stage market

We begin with a simple model, in which there are two players: a single airline and a number of potential consumers of airline tickets. The consumer's problem is simply to maximize the utility gained from both his wealth (w) and the possible acquisition of an airline ticket (x, which takes a value of one, with the purchase of a ticket, or zero):

$$\max u(w, x) \qquad (1)$$

The utility derived from wealth and travel is determined by the form of the utility function:

$$u = f(w) + g(x) \qquad (2)$$

That is, utility to the customer is composed of his utility-of-money function, *f(w)*, plus the utility gained from the purchase of an airline ticket, *g(x)*[1]. By separating the two components of the utility function, we assume there is no complementarity between wealth and air travel.

If we allow w_0 to be the consumer's initial level of wealth and p the purchase price of an airline ticket, we can state that a purchase will be made if and only if the following relationship holds:

$$u(w_0 - p, 1) = f(w_0 - p) + g(1) \geq f(w_0) + g(0) = u(w_0, 0) \qquad (3)[2]$$

If the airline is able to set customer-specific prices—and if it is aware of the willingness-to-pay of its customers (which we assume throughout this section)—it will naturally set $p = p^*$ to ensure equality in *(3)*. This conclusion requires us only to make the common and reasonable assumptions of monotonicity in the consumer's utility-of-money function[3] and of

---

[1] We frame *g(x)* as the utility *gained* from purchasing a ticket, rather than the other side of the coin—utility lost due to *not* purchasing a ticket. Thus, *g(1) > 0* and *g(0) = 0*.
[2] By allowed *(3)* to be satisfied with equality, we assume that a purchase is made if the consumer is indifferent between buying and not buying a ticket. We will continue to make similar assumptions throughout the essay—that agents will act as we wish if indifferent between a number of choices.
[3] Without the monotonicity assumption, an increased ticket price could increase the customer's utility. This assumption is trivial, of course.

positive utility gained from the acquisition of a ticket. Setting a price lower than p^* leaves consumer surplus which the airline could usurp by increasing p, while $p > p^*$ prompts the consumer to make no purchase at all. Optimizing at $p = p^*$ supposes only that the airline seeks to extract maximum revenue from each ticket sale.

Matters become more difficult for the airline upon the introduction of consumer heterogeneity. For simplicity, it is reasonable—for our purposes—to assume that the utility-of-money function, $f(w)$, takes the same form for each consumer. However, it is clearly not appropriate, given the focus of this analysis, to make the same assumption on the form of $g(x)$. Rather, considerable literature (for example, Brons et al. 2002) indicates that there is significant variation in consumer willingness-to-pay for airline tickets. In particular, those customers who are taking trips for business purposes have greater willingness-to-pay than leisure travelers. To a business customer, a trip is generally of more importance than to a leisure customer. Furthermore, the trip may be partially or fully financed by his employer, externalizing the cost of the ticket. For simplicity, we assume only two different forms of the $g(x)$ function, one for a business customer and the other for a leisure customer. We can use our assumptions on the nature of the customer's utility to ascertain p_B^* and p_L^* from *(3)*: the maximum willingness-to-pay of business and leisure customers, respectively. We will refer to these as the reservation prices of the customers, with *Rp(B)* the reservation price of a business customer and *Rp(L)* the reservation price of a leisure customer [*Rp(B) > Rp(L)*]. For convenience, let us normalize the utility gained from these purchases to $u = 0$.

Business customer: $u(w_0 - Rp(B),1) = u(w_0, 0) = 0$

Leisure customer: $u(w_0 - Rp(L),1) = u(w_0, 0) = 0$

The airline has a number of pricing options with respect to the reservation prices of its customers. We assume that the airline is unable to offer customer-specific prices (i.e. it cannot first-degree price discriminate). Instead, it must offer a single price which is available to all of its customers. Given that tickets sales primarily occur on the internet, this seems to be a reasonable assumption. There are five relevant prices/price ranges which the airline may set:

$$P(1) > Rp(B) > Rp(L)$$

$$P(2) = Rp(B) > Rp(L)$$

$$Rp(B) > P(3) > Rp(L)$$

$$Rp(B) > P(4) = Rp(L)$$

$$Rp(B) > Rp(L) > P(5)$$

Naturally, $P(1) > P(2) > P(3) > P(4) > P(5) > 0$ *(4)*

These five pricing scenario yield the following expected profit on a per customer basis[4], with δ equal to the proportion of business customers in the market (thus, *(1 – δ)* is the proportion of leisure customers):

$$\pi(1) = 0$$

$$\pi(2) = \delta P(2)$$

$$\pi(3) = \delta P(3)$$

$$\pi(4) = P(4)$$

$$\pi(5) = P(5)$$

Given the relation in *(4),* we can generate the following two further relations:

$$\pi(2) > \pi(3) > \pi(1) \qquad\qquad\qquad (5)$$

$$\pi(4) > \pi(5) > \pi(1) \qquad\qquad\qquad (6)$$

The relationship between $\pi(2)$ and $\pi(4)$ is ambiguous, as it depends upon δ. Thus, we conclude that the airline maximizes its expected profit by setting price equal to either the reservation price of its business customers, or that of its leisure customers, depending upon the proportions of the two in the population. We can ascertain the critical value of δ at which the profit-maximizing price shifts in the following manner:

$$\pi(2) = \pi(4)$$

$$\delta P(2) = P(4)$$

---

[4] Given constant pricing across the customer base, this expected profit figure can simply be multiplied by the number of customers *(N)* in the market to compute the market-wide expected profit to the airline.

$$\delta^* = P(4)/P(2)$$

Substituting the relations from above:

$$\delta^* = Rp(L)/Rp(B)$$

That is, when the proportion of business customers in the market is equal to the ratio of the reservation price of leisure customers to that of business customers, the airline will be indifferent between its two optimal pricing strategies of $P = P(2) = Rp(B)$ and $P = P(4) = Rp(L)$. If $\delta > \delta^*$, the airline optimizes by setting $P = P(2)$, while if $\delta < \delta^*$, the airline optimizes by setting $P = P(4)$. In the former case, the expected per-customer profit to the airline is equal to $\delta Rp(B)$, while in the latter case, the expected profit is $Rp(L)$. Generalizing to a market with N customers, the respective expected profits in the two cases are $\delta NRp(B)$ and $NRp(L)$.

Two-stage market

We now take a step in the direction of a more realistic market by including an additional stage in the interaction. We assume that there are two time periods in which a ticket sale may occur (say, today and tomorrow), with the airline able to set a market price in each of the two stages. All customers have the option to purchase a ticket in the first stage. As customers require just a single ticket (further tickets add nothing to customer utility), only those customers who have declined to purchase a ticket in the first stage have the opportunity to purchase in the second. We assume perfect information in the market: both current and future prices are known to customers at the game's beginning. Additionally, if the airline sets a mixed pricing strategy, the probability distribution of prices is known to the consumers throughout the game. This assumption is necessary to ensure that the game can reach an equilibrium in which neither party would like to adjust its strategy (in game theory terms, a Nash equilibrium).

The airline now has a number of pricing strategies that are worthy of examination. To reduce the number of variables in the model, let us assume that $\delta > \delta^*$: With the proportion of business customers in the market exceeding the critical value, the airline maximizes profit by

setting $P = P(2) = Rp(B)$ and reaping a profit (per customer[5]) of $\delta Rp(B)$—in the single-stage encounter. With the expansion to a second stage, this constitutes the most straightforward pricing strategy available to the airline—setting $P = Rp(B)$ in both stages. The profit to the airline will be the same in this case as in the single-stage market: All business customers purchase tickets, while all leisure customers do not. Thus, the airline gains the same profit as it had previously, of $\delta Rp(B)$.

Next, we examine a pricing strategy in which $P = P(2)$ in stage one and $P = P(4) = Rp(L)$ in stage two—or vice versa. Given perfect information in the market, all customers will choose to purchase at the lower price, thereby providing the airline a profit of $Rp(L)$. Since $\delta > \delta^*$, this is inferior to the profit gained by the previous strategy of setting $P = Rp(B)$ in both stages. As $P(2)$ and $P(4)$ are the two dominant pricing strategies available to the airline, we need not explore the viability of $P(1)$, $P(3)$ or $P(5)$. We can therefore assert that the airline's best outcome when utilizing straightforward, non-mixed pricing—given the composition of its customer base—comes from setting $P = P(2)$ in both stages and gaining a profit of $\delta Rp(B)$.

As we know, the goal of the airline in its two-stage pricing is to extract maximum revenue from its two customer segments, business and leisure travelers. Ultimately, the airline would like to sell tickets at $P = Rp(B)$ to each of its business customers and $P = Rp(L)$ to each of its leisure customers, gaining a profit of $\delta Rp(B) + (1-\delta)Rp(L)$. As the airline is unable to sell tickets separately to different customer groups, it would like to induce its customers to self-select based upon their type. To do this, the airline must first attempt to coax purchases from business customers, before turning its attention to leisure customers. Attracting leisure customers first cannot improve profitability, since business customers will take such an opportunity to purchase tickets at the lower prices intended for leisure customers.

Since the pure pricing strategies discussed above are incapable of moving the airline towards a higher level of profitability, we move on to explore more complex pricing options. It may be profitable for the airline to set $P = P(2)$ in stage one, before utilizing a probability distribution

---

[5] Throughout the remainder of this exposition, we will continue denoting profit on a per-customer basis. This is solely for the sake of simplicity of notation; it is trivial to generalize to the case of *N* customers.

of prices in stage two. Placing some second-stage weight on a price $P = P(1)$ could scare risk-averse business customers into purchasing in stage one. Allowing some probability of $P = P(4)$ as well, in the second stage, would then allow the airline to gain some expected profit from purchases by leisure customers in stage two. Unfortunately for the airline, this pricing strategy is ineffective. Since the reservation price of customers is computed to allow the customer precisely zero utility at the purchase price, the threat of a high price in stage two does not intimidate the business customers, as they can simply decline to purchase a ticket—providing them the same utility as if they had purchased a ticket at $Rp(B)$ in stage one. To illustrate this pricing strategy, let us say that the airline places a weight of α on $P = P(4)$ in the second stage, $0 < \alpha < 1$. It places the remaining probability weight $(1 - \alpha)$ on $P = P(1)$. As we know, the business customer gains utility of zero from purchasing a ticket at $P = P(2)$ in the first stage. When deciding whether to make such a purchase, he compares this with his expected utility in the second stage:

$$E(u) = u([1 - \alpha] \circ (w_0, 0) \oplus \alpha \circ (w_0 - P(4), 1)) \qquad (7)^6$$

The first portion of *(7)* indicates that the business customer declines to purchase a ticket with probability $(1 - \alpha)$, maintaining his wealth of w_0. As noted, we have normalized this scenario to $u = 0$ for convenience. With probability α, the customer purchases a ticket at $P = P(4)$ in the second stage, providing him utility from $(w_0 - P(4), 1)$. Since we know that $u(w_0 - P(2), 1) = 0$, we need only assume strict monotonicity in the utility-of-money function to ascertain that $u(w_0 - P(4), 1) > 0$. The payoff from *(7)* can thus be interpreted as a gamble between payoffs of $u = 0$ and $u > 0$. This is clearly preferable to gaining $u = 0$ with certainty, as the customer would if he chose to purchase in the first stage. Therefore, business customers will wait until the second stage to purchase tickets (regardless of the value of α). Further, this purchasing choice is independent of the risk preferences of the consumers, as waiting for stage two is no more risky in this instance than purchasing in stage one (since the worst possible stage two outcome provides identical utility to the guaranteed stage one outcome). The airline, meanwhile, gains profit of $\alpha P(4)$ from this pricing scheme. The firm

---

[6] Notation for gambles/lotteries is borrowed from Varian (1992). As Varian writes (p. 172), denoting a lottery (or gamble) as $p \circ x \oplus (1 - p) \circ y$ means "the customer receives prize *x* with probability *p* and prize *y* with probability *(1 − p)*."

has thus failed to induce customers to self-select, and gains a lesser profit than if it had simply set $P = P(2)$ in both stages, since $\delta P(2) > P(4)^{[7]} > \alpha P(4)$.

Now, we add a twist to the airline's pricing strategy. Instead of setting $P = Rp(B)$ in stage one, the airline offers tickets at $P = Rp(B) - \varepsilon$. Requiring $0 \leq \varepsilon \leq [Rp(B) - Rp(L)]$ necessitates that $[Rp(B) - \varepsilon]$ never rise above $Rp(B)$ nor fall below $Rp(L)$. Thus, ε should be seen as a small "carrot" intended to induce first-stage purchases by business customers. By making a purchase in stage one at this price, the business consumer gains utility $u = (w_0 - [Rp(B) - \varepsilon], 1)$. Given the constraints on ε, this $u \geq 0$.

In stage two, we propose that the airline proceed with a stochastic pricing strategy in which it simultaneously attempts to dissuade business customers from waiting until stage two to make a purchase, while also seeking to maximize the probability of selling to leisure customers. To this end, the airline utilizes the stage-two probability mix $P = P(4)$ with probability α, $0 \leq \alpha \leq 1$, and $P = P(2)$ with probability $(1 - \alpha)$. Note that the airline does not need to utilize the threat of $P = P(1)$—which is not a credible means of intimidation, since it works to the detriment of the airline: revenue from $P = P(1)$ is zero. Instead, $P = P(2)$ functions as a similar—but credible—threat, since it provides consumers with the same utility (zero) while also providing revenue to the airline. In examining the viability of this pricing strategy, we determine if this strategy would induce customers to self-select as the airline wishes, and if the strategy yields greater profitability than the "default" pricing strategy of $P = P(2)$ in both stages.

The airline seeks to induce business customers to purchase in stage one, at $P = Rp(B) - \varepsilon$. These customers will do as the airline wishes if the following condition is met:

$$u(w_0 - [Rp(B) - \varepsilon], 1) \geq u([1 - \alpha] \circ (w_0 - Rp(B), 1) \oplus \alpha \circ (w_0 - Rp(L), 1)) \qquad (8)$$

---

[7] Given our assumption on the distribution of customer types

That is, the utility of purchasing a ticket at $P = Rp(B) - \varepsilon$ must exceed the utility of a gamble in which a ticket is purchased at $P = Rp(B)$ with probability $(1 - \alpha)$ and $P = Rp(L)$ with probability α. Having normalized the scenario in which a ticket is purchased at $P = Rp(B)$ to $u = 0$, we can assert that the left side of (8) is positive, since ε is positive. The right side of the equation constitutes a gamble between gaining utility of zero, and gaining utility from the purchase of a ticket at $Rp(L)$. Since $\varepsilon \leq [Rp(B) - Rp(L)]$, we know that the utility from purchasing at $Rp(L)$ is at least as great as purchasing at $Rp(B) - \varepsilon$. Thus, (8) presents a gamble between two utilities—one of which is greater than that of a guaranteed stage-one purchase, and the other of which is less than this certain amount.

As we know from microeconomic theory, the utility to an agent of a gamble declines relative to the utility of the expected value of the gamble as the agent becomes more risk-averse. Thus, (8) is more likely to hold the greater the concavity of the business customer's expected-utility-of-money function. The airline is therefore increasingly likely to induce self-selection among its customers the more risk-averse its business customers are.

Secondly, we must check if this pricing strategy improves profitability to the airline. Up to this point, the airline's optimal pricing strategy had been, as discussed, to set $P = Rp(B)$ throughout the game (given the assumption that $\delta > \delta^*$). With this strategy, the airline had gained profit of $\delta Rp(B)$. By adopting the stochastic pricing strategy, the airline—if it succeeds in inducing its customers to self-select (i.e. if (8) holds)—is able to sell tickets at $P = Rp(B) - \varepsilon$ to all of its business customers, and $P = Rp(L)$ to all of its leisure customers—but to this second group only with probability α. Thus, expected profit to the airline from this stochastic pricing strategy is:

$$E(\pi) = \delta[Rp(B) - \varepsilon] + \alpha(1 - \delta)Rp(L) \qquad (9)$$

Assuming risk-neutrality on the part of the airline, the condition that prompts the airline to choose this stochastic pricing strategy is (recalling our assumption that $\delta > \delta^*$):

$$\delta[Rp(B) - \varepsilon] + \alpha(1 - \delta)Rp(L) \geq \delta Rp(B) \qquad (10)$$

$$\delta Rp(B) - \delta\varepsilon + \alpha(1 - \delta)Rp(L) \geq \delta Rp(B)$$

$$\alpha(1 - \delta)Rp(L) \geq \delta\varepsilon \qquad\qquad (11)$$

Simply, the profit lost on each business customer (ε) must not exceed the expected profit gained from sales to leisure customers.

From the above analysis, we can conclude that, under three conditions, the airline can improve upon the profitability of its optimal *pure* pricing strategy by setting prices in a stochastic manner. These requirements are:

1) The proportions of business/leisure customers in the population induce the airline to optimize by setting price equal to the reservation price of business customers, rather than that of leisure customers *($\delta \geq \delta* = Rp(L)/Rp(B)$)*—or at least prompt indifference from the airline in regard to its pricing strategy. This requires either a relatively high proportion of business customers in the population, or a high reservation price of business customers relative to that of leisure customers. Note that we could perform a similar analysis if the proportions were reversed (i.e. the airline optimized by setting *P = Rp[L]*); we would simply have to change the standard to which the mixed-strategy payoff is compared by the airline[8]. This condition is therefore the only menial one of the three requirements.

2) *(Consumer's) selection constraint:* The "discount" *(ε)* on the first-stage price is high enough, the probability of a low-price in stage two *(α)* is low enough, the reservation price of leisure consumers (*Rp[L]*) is high enough (relative to *Rp[B]*), and/or business customers have sufficient risk aversion that these customers prefer to purchase tickets in stage one rather than taking a gamble on stage two prices. That is, *(8)* is satisfied.

3) *(Airline's) profitability constraint:* The discount on the first-stage price is *low* enough, the probability of a low-price in stage two is *high* enough, the reservation price of leisure customers is high enough, and/or there are enough leisure customers in the market (*1 − δ*) to ensure that the stochastic pricing strategy is at least as profitable to the airline as its dominant pure strategy of setting *P = Rp(B)* in each stage. That is, *(11)* is satisfied.

---

[8] In this case—($\delta < \delta*$)—the airline's profitability constraint from *(10)* would become: $[\delta[Rp(B) - \varepsilon] + \alpha(1 - \delta)Rp(L)] \geq Rp(L)$

Numerical illustration (1): Risk-neutrality

Our first numerical example will examine if the price-mixing strategy can be effective if customers are risk-neutral. We begin by setting reservation prices for the two groups of customers: *Rp(B) = 200* and *Rp(L) = 100*. With these reservation prices, δ must be $\geq$ *100/200*. Let δ = *0.6*. We can enter these values into *(11)*, producing the airline's profitability constraint in this example:

$$\alpha(1 - \delta)Rp(L) \geq \delta\varepsilon$$

$$\alpha(1 - 0.6)100 \geq 0.6\varepsilon$$

$$66.667\alpha \geq \varepsilon \qquad (12)$$

From *(8)*, we can also identify the business customer's selection constraint:

$$u(w_0 - [200 - \varepsilon], 1) \geq u([1 - \alpha] \circ (w_0 - 200, 1) \oplus \alpha \circ (w_0 - 100, 1)) \qquad (13)$$

Let us continue to assume no complementarity between money and airline ticket; that is, $u(w, x) = u(w) + u(x)$--the utility function is separable. This allows us to remove the utility of a purchased ticket from the utility function, since it is purchased in every contingency (in this case). For the purposes of this example, we further assume that the business customer's expected-utility-of-money function is linear: that is, he is risk-neutral. As such, the utility gained from wealth is by definition independent of the possession of other wealth: $u(w_0 + w_1) = u(w_0) + u(w_1)$. With risk-neutrality, we can add that $u(w) = w$. As discussed, we have normalized the utility function such that $u(w_0 - Rp(B), 1) = 0$ for a business customer. With these assumptions, we can transform *(13)* into:

$$u(w_0 - 200, 1) + u(\varepsilon) \geq u([1 - \alpha] \circ (w_0 - 200, 1) \oplus \alpha \circ [(w_0 - 200, 1) + 100])$$

$$u(0) + u(\varepsilon) \geq u([1 - \alpha] \circ 0 \oplus \alpha \circ 100)$$

$$\varepsilon \geq 100\alpha \qquad (14)$$

With the pair of equations *(12)* and *(14)* being unsolvable (given $\alpha > 0$). Thus, with the chosen reservation prices and proportions of customer types, we are unable to create the

necessary conditions for a profitable mixed-pricing strategy, given risk-neutrality among consumers.

In fact, we can show mathematically that it is impossible to improve the airline's profitability with a mixed strategy, while heeding the consumer's selection constraint—given this assumption of risk-neutrality. Using the same process and assumptions that were used to generate *(14)*, we can produce the comparable consumer selection constraint for the general case, under the assumption of risk-neutrality. Beginning from *(8)*, we subtract $[w_0 - Rp(B)]$ from each payoff—which is acceptable given linear utility-of-money:

$$u(w_0 - [Rp(B) - \varepsilon], 1) \geq u([1 - \alpha] \circ (w_0 - Rp(B), 1) \oplus \alpha \circ (w_0 - Rp(L), 1))$$

$$u(\varepsilon) \geq u([1 - \alpha] \circ 0 \oplus \alpha \circ [Rp(B) - Rp(L)])$$

$$\varepsilon \geq \alpha[Rp(B) - Rp(L)] \tag{15}$$

Recalling the airline's expected profit function in *(9)*, we know that profit decreases with ε, and increases with α. Thus, when the airline maximizes expected profit, the consumer's constraint in *(15)* will be satisfied with equality. We now substitute this binding relation into the airline's expected profit function:

$$E(\pi) = \delta[Rp(B) - \varepsilon] + \alpha(1 - \delta)Rp(L)$$

$$E(\pi) = \delta[Rp(B) - \alpha[Rp(B) - Rp(L)]] + \alpha(1 - \delta)Rp(L)$$

$$E(\pi) = \delta[Rp(B) - \alpha Rp(B) + \alpha Rp(L)] + \alpha Rp(L) - \delta\alpha Rp(L)$$

$$E(\pi) = \delta Rp(B) - \delta\alpha Rp(B) + \delta\alpha Rp(L) + \alpha Rp(L) - \delta\alpha Rp(L)$$

$$E(\pi) = \delta Rp(B) - \delta\alpha Rp(B) + \alpha Rp(L)$$

$$E(\pi) = \delta Rp(B) + \alpha[Rp(L) - \delta Rp(B)] \tag{16}$$

With *(16)*, we see that expected profit to the airline is a linear function of its choice parameter, α. The optimal value of α is dependent simply on the sign of $[Rp(L) - \delta Rp(B)]$. If the expression is positive, profit is maximized by increasing α to its threshold ($\alpha = 1$), while if the expression is negative, the airline maximizes profit by minimizing α (setting $\alpha = 0$) (if the

expression equals zero, the airline's profit is independent of its strategy mix). In the former case of $\alpha = 1$, the binding relation in *(15)* tells us that $\varepsilon = Rp(B) - Rp(L)$. Thus, the first-stage price will be $Rp(B) - [Rp(B) - Rp(L)] = Rp(L)$. With $\alpha = 1$, we know that the second-stage price will also be $P = Rp(L)$ with certainty. In the latter case of $\alpha = 0$, the binding relation in *(15)* requires $\varepsilon = 0$. Thus, the airline optimizes by setting $P = Rp(B) - 0$ in the first stage, and $P = Rp(B)$ again with certainty in the second stage. Solving the expression $Rp(L) - \delta Rp(B)$ for δ, we again find that the critical value of δ at which the sign of the expression changes is $\delta^* = Rp(L)/Rp(B)$. Thus, we have shown that the airline's optimal pricing strategies, assuming consumer risk-neutrality, are simply the preferred pure strategies discussed above. Utilizing a mixed strategy cannot improve the airline's profitability under these conditions.

Numerical illustration (2): Risk-aversion

In this next example, we allow for risk-aversion among consumers: that is, consumers have concave expected-utility-of-money functions. Let us again (arbitrarily) set reservation prices for leisure and business customers; say $Rp(B) = \$800$ and $Rp(L) = \$500$. We now must set δ, the proportion of business customers, greater than *500/800*, to ensure that the baseline maximized profit occurs when $P = \$800$. Let $\delta = 0.75$. Finally, let us set the initial wealth of business customers to *$800* (this is relatively unimportant). Beginning with the airline's predicament, we use *(11)*:

$$\alpha(1 - \delta)Rp(L) \geq \delta\varepsilon$$

$$\alpha(1 - .75)500 \geq .75\varepsilon$$

$$125\alpha \geq .75\varepsilon$$

$$166.67\alpha \geq \varepsilon \qquad\qquad (17)$$

Next, we must ensure that business customers purchase tickets in the first stage. In this case—assuming, again, a separable utility function, which allows us to remove the utility of the ticket from the consumer's constraint—the consumer's selection constraint becomes:

$$u(w_0 - Rp(B) + \varepsilon) \geq u([1 - \alpha] \circ (w_0 - Rp(B)) \oplus \alpha \circ (w_0 - Rp(L)))$$

$$u(800 - 800 + \varepsilon) \geq u([1 - \alpha] \circ (800 - 800) \oplus \alpha \circ (800 - 500))$$

$$u(\varepsilon) \geq u([1 - \alpha] \circ 0 \oplus \alpha \circ 300) \qquad (18)$$

In the case of the risk-averse consumer presumed here, the utility of the gamble presented in *(18)* is less than the utility of the expected value of the gamble (Varian 2002: 177). If we assume, for example, that the customer's concave expected-utility-of-money function is represented by the square root function, *(18)* reduces to:

$$\sqrt{\varepsilon} \geq \alpha\sqrt{300}$$

$$\varepsilon \geq 300\alpha^2 \qquad (19)$$

We see from *(17)* and *(19)* that our problem can be solved: It is possible for the airline to utilize a profitable mixed strategy that prompts the consumer to select as the airline wishes. In order to examine precisely how profitable such a strategy can be, we maximize the airline's profit function with respect to α and ε, with the business customer's selection constraint in *(19)*. Note that we do not need to include the airline's profitability requirement *(17)* as a constraint in the maximization problem, since we can simply compare the computed maximized profit figure to the profitability of the optimal pure strategy.

$$\max_{\{\varepsilon,\alpha\}} \pi = \delta[Rp(B) - \varepsilon] + \alpha(1 - \delta)Rp(L)$$

$$\max_{\{\varepsilon,\alpha\}} \pi = .75(800 - \varepsilon) + \alpha(.25)500 = 600 - 0.75\varepsilon + 125\alpha \qquad (20)$$

Subject to (19)

Since *(20)* shows us that the airline's profit is decreasing in ε and increasing in α, *(19)* will bind when the airline maximizes profit. Thus, we insert $\varepsilon = 300\alpha^2$ into *(20)*:

$$\max_{\{\alpha\}} \pi = 600 - 0.75(300\alpha^2) + 125\alpha \qquad (21)$$

Differentiating *(21)* with respect to α yields the profit-maximizing choice of α:

$$\frac{\partial \pi}{\partial \alpha} = -450\alpha + 125 = 0$$

$$\alpha^* \approx 0.2778$$

Finally, inserting α^* into *(19)* (which we know binds) reveals that $\varepsilon^* \approx 23.148$. These parameters yield a per-customer profit to the airline of *617.361* [utilizing *(20)*]. This gain from a mixed strategy compares favorably to the optimal pure strategy of $P = Rp(B)$, which generates profit of *800(0.75) = 600*. It also, of course, compares favorably to the sub-optimal pure pricing strategy of $P = Rp(L)$, which produces a profit of *500* per customer. Thus, we have shown in this example that, under conditions of consumer risk-aversion, the airline is able to improve profitability by utilizing our suggested randomized pricing strategy.

We can generalize this process of finding the optimal mixed strategy given assumptions on the form of the risk aversion exhibited in the consumer's utility-of-money function. We proceed to explicate the case of $u(w) = \sqrt{w}$. Beginning with the consumer's selection constraint, we solve for ε:

$$u(w_0 - Rp(B) + \varepsilon) \geq u([1-\alpha] \circ (w_0 - Rp(B)) \oplus \alpha \circ (w_0 - Rp(L)))$$

$$\sqrt{(w_0 - Rp(B) + \varepsilon)} \geq (1-\alpha)\sqrt{(w_0 - Rp(B))} + \alpha\sqrt{(w_0 - Rp(L))}$$

$$(w_0 - Rp(B) + \varepsilon) \geq \left[(1-\alpha)\sqrt{(w_0 - Rp(B))} + \alpha\sqrt{(w_0 - Rp(L))}\right]^2$$

$$\varepsilon \geq \left[(1-\alpha)\sqrt{(w_0 - Rp(B))} + \alpha\sqrt{(w_0 - Rp(L))}\right]^2 - w_0 + Rp(B) \qquad (22)$$

We then address the expected profit of the airline. As before, *(22)* is binding, since the airline's payoff increases with α and decreases with ε. Thus, we insert *(22)* into the airline's expected profit function:

$$E(\pi) = \delta[Rp(B) - \varepsilon] + \alpha(1-\delta)Rp(L)$$

$$E(\pi) = \delta\left[Rp(B) - \left(\left[(1-\alpha)\sqrt{(w_0 - Rp(B))} + \alpha\sqrt{(w_0 - Rp(L))}\right]^2 - w_0 + Rp(B)\right)\right]$$
$$+ \alpha(1-\delta)Rp(L)$$

$$E(\pi) = \delta\left[w_0 - \left[(1-\alpha)\sqrt{(w_0 - Rp(B))} + \alpha\sqrt{(w_0 - Rp(L))}\right]^2\right] + \alpha(1-\delta)Rp(L) \quad (23)^9$$

In utilizing this generalized profit formula, we need only follow the process used above: Insert figures for $\delta, w_0, Rp(B)$ and $Rp(L)$ into *(23)*; differentiate *(23)* with respect to α and solve for this profit-maximizing strategy mix; and insert this optimal α into *(22)* and solve for ε as an equality. Finally, one must compute expected profit from *(23)*, and ensure that this compares favorably to that of the optimal pure strategy. In the case used here with $\delta > \delta*$, this requires:

$$\delta\left[w_0 - \left[(1-\alpha)\sqrt{(w_0 - Rp(B))} + \alpha\sqrt{(w_0 - Rp(L))}\right]^2\right] + \alpha(1-\delta)Rp(L) \geq \delta Rp(B)$$

The principal conclusion from the theoretical model presented above is the following: Relying upon the assumptions discussed above, airlines may find it beneficial to *vary* prices across time in a *stochastic* manner. Such mixing may function as a means of (third-degree) price-discrimination, allowing airlines to extract greater surplus from customers than pure strategies would allow. To generalize the two-stage game used here to a market of *S* stages, our findings suggest that a price schedule in which expected price decreases, but variance of price increases, as the market nears closure can be profitable. With a price-weighting function in which lower prices increase in probability as time elapses, airlines are able to first sell to customers who have a relatively high willingness-to-pay before lowering the *expected* market price in an attempt to attract customers who have lower willingness-to-pay. In this manner, risk-averse, higher-end customers are targeted before their lower-end counterparts, allowing the airline to price discriminate.

In the Appendix that follows this essay, we have included a more complex, numerical exposition in which we analyze the benefit to the airline of a stochastic pricing strategy in which consumers make purchasing decisions based upon a "probability-of-purchase" function. This discussion provides a more concrete—and perhaps more realistic—presentation of the viability of a mixed pricing strategy.

---

[9] The viability of the generalized formulas in *(22)* and *(23)* can be confirmed by inputting the numbers used in the example above into these equations, yielding the same values for α and ε.

The conclusion derived from the above theoretical analysis is mirrored by that found by Landsberger and Meilijson (1985): "The idea behind a time-decreasing price schedule is to make it profitable for consumers with high reservation prices to buy at the beginning of the price sequence when prices are highest. Those with low reservation prices wait to buy in the more remote future. For this strategy to succeed, consumers with high reservation prices must be discouraged from waiting until low prices are offered" (p. 429). In the model presented in this section, stochastic pricing takes advantage of the risk-aversion of customers in functioning as this means of "discouragement."

IV) Empirical examination of airline pricing

The theoretical model presented in the preceding sections may be applicable to a certain segment of the airline's pricing scheme. It is well known that ticket prices on traditional airlines increase dramatically as the departure date approaches (this is noted by Button and Vega 2007, for example), but at present we are concerned with the behavior of ticket prices several weeks before departure. To this end, our next step is to collect price data for a number of flights, in order to investigate the pricing strategies of various airlines as departure dates approach.

The data

The goal of data collection in this study was to obtain data which, as much as possible, isolates the strategic pricing decisions of airlines that are *not* made in accordance with known, predictable, and industry-wide trends (such as the Saturday-night stay-over discount). We therefore wish to remove all price effects that may be caused by the nature of competition at any particular airport, industry regulations, the day of departure or return, length of stay, or any demand-side variations (i.e. number of tickets sold). In all, 12 major trans-Atlantic routes were studies. All data was gathered from Orbitz.com in March/April 2010.

On each trip, the prices of four airlines were sampled: American Airlines, Continental Airlines, Delta Airlines, and United Airlines. These four were chosen because they are all American, which allows us to remove any pricing effects that may be due to regulations in the airline's home country. Within the American airline industry, these are the four of the major carriers for international travelers, and each provides a large selection of flights between the cities of interest. In each sample, we control for day (of the week) of departure (the methodology of which be explained shortly), as well as length of stay in the destination city (all flights include a return). Data for each flight was collected within a two-hour period, but the whole dataset was collected over a few weeks, thus controlling to some extent for variation that may be specific to a particular period of time in which prices were accessed. Between (but *not* within) the 12 samples, we vary not only the pair of cities for the journey, but also the days of departure/return and length of stay, thus controlling for these effects in our sample of prices. Pairs of cities were selected such that each journey is a major trans-Atlantic route between large American/European cities (with the city of departure rotating between the U.S. and Europe). As there are a greater number of European than American "hub" airports, only five American cities (Boston, Chicago, Los Angeles, Newark and New York) were used, while 10 European cities were chosen (Amsterdam, Berlin, Bucharest, Budapest, Copenhagen, Frankfurt, London, Madrid, Paris, and Rome).

For each flight on each carrier, the price of the cheapest available economy class ticket was recorded. If each airline provides a non-stop flight on the journey in question, this is the ticket whose price was recorded. If not all airlines provide a non-stop flight, only one-stop ticket prices were recorded (thus controlling for the unfair scenario in which a non-stop price is compared against a one-stop flight). For the pairs of cities used, one stop is the most that was required on the journey. In order to limit the number of variables in the data, this study does not concern itself with discount airlines (e.g. Ryanair, Southwest Airlines), nor is it concerned with business class fare prices.

Figure 1, below, provides a comparison of economy class ticket prices on the four airlines for a New York to London route, for departure dates between 24 weeks and one day away. All data was collected on the same day (in the example of Figure 1, data was collected on 29 March, 2010), while the date of the trip is allowed to vary. The X-axis lists the date of departure relative to the date when the data was collected. In this particular example, the return leg of the trip is always five days after the departure. Thus, the observation

corresponding to "next day" in Figure 1 departs on 30 March, 2010 and returns on 4 April. The trips corresponding to "two weeks" and "eight weeks" are 13 April – 18 April, and 25 May – 30 May, respectively. Time between observations was chosen such that each flight departs (and therefore returns) on the same day of the week in each observation (this trip is always Tuesday – Sunday). This allows us to control for day of the week of departure and return, which we will shortly see is quite important.

Figure 1: Economy Class Ticket Comparison (New York - London)
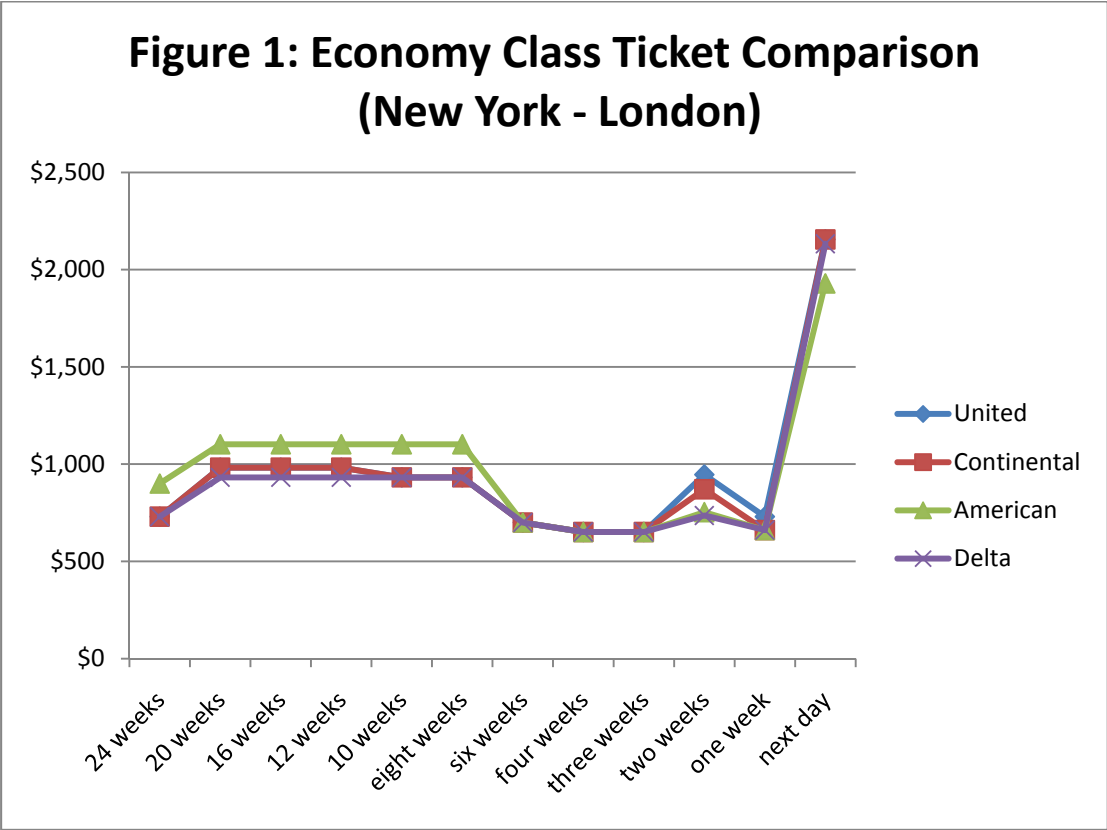
For comparison, Figure 2 shows the prices of a New York to Paris trip for the same four airlines and the same time spans to departure. Data for this trip was collected on 26 April, so the "next day" observation corresponds to 27 April, a Tuesday. On this flight, it was also chosen that five days would elapse between the original and return legs of the journey; thus, each observation again corresponds to a Tuesday – Sunday flight[10]. Finally, Figure 3 depicts prices on a Los Angeles to Berlin flight for a Wednesday – Saturday trip (so this trip does not have a Saturday-night stay-over).

---

[10] In Figure 2, United prices are hidden behind Continental's, as the two mirror each other throughout the sample.

Figure 2: Economy Class Ticket Comparison (New York - Paris)
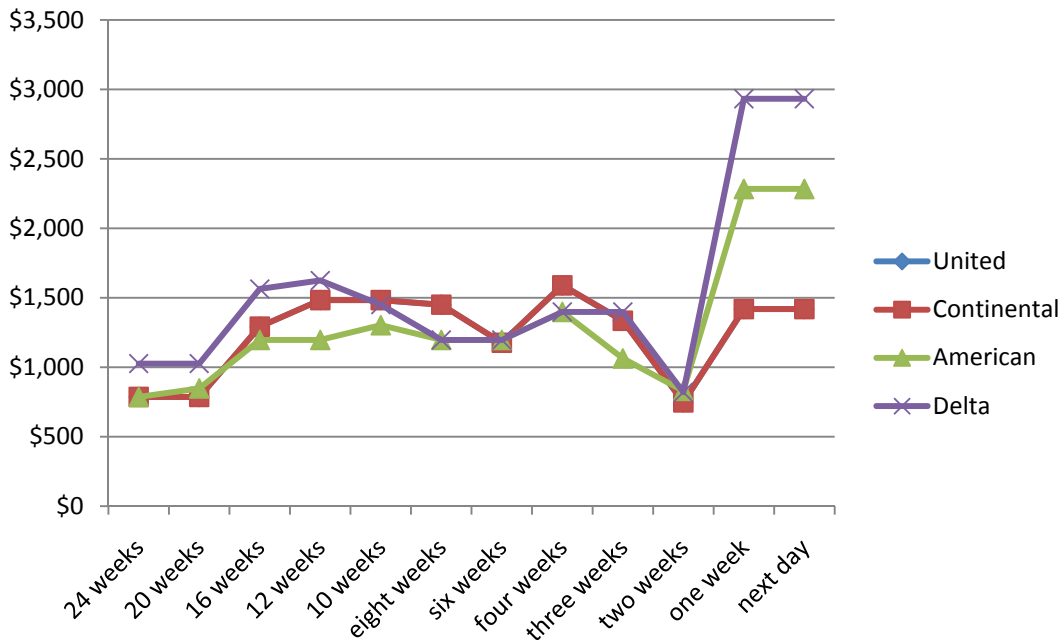
Figure 3: Economy Class Ticket Comparison (Los Angeles - Berlin)
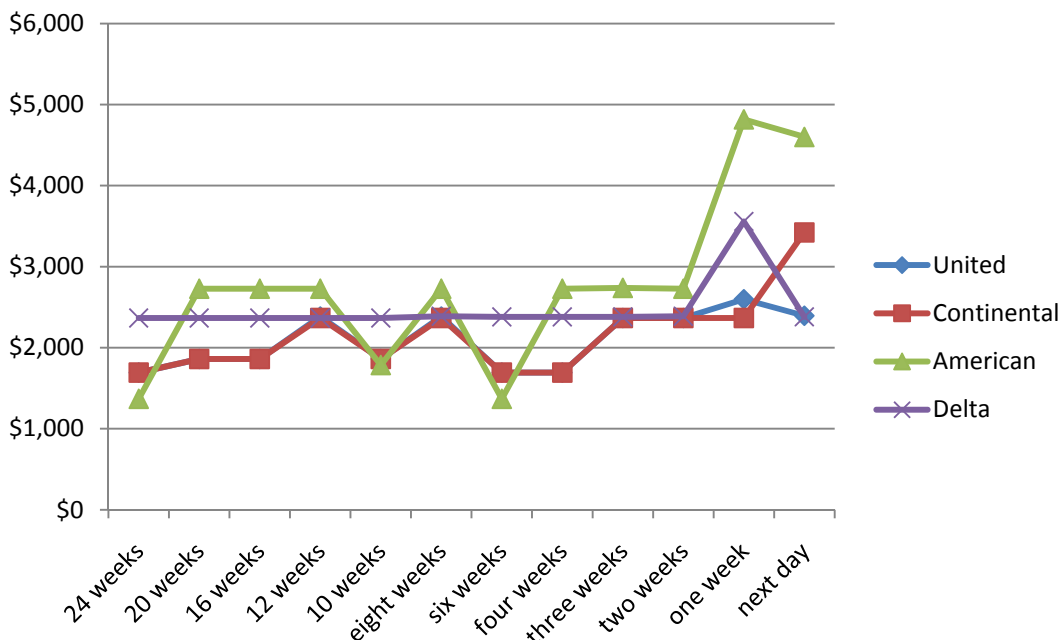
A more traditional method of tracking airline prices across time is to monitor a particular flight with fixed dates as time passes, with data collection taking place, usually daily, from a

few months before the date of departure until the day of departure (see, for example, Button and Vega 2007; Mantin and Koo 2009; Pels and Rietveld 2004). Given the time constraints for this work, however, this method is unfeasible. Thus, the method of tracking various flights with different departure dates originated. However, this method may actually be preferable to the more conventional means of data collection discussed above. As Pels and Rietveld (2004) note, the observed price quotes are dependent not only on supply-side factors, but also on demand-side factors. In particular, the authors mention, we are missing an important piece of information—the number of seats that have been sold on each flight at each point in time (p. 280). Button and Vega (2004) add that "The number of bookings, which influence the price set for subsequent passengers can fluctuate quite dramatically if, for example, a large group of people travel and book together, or if a particular flight corresponds to a unique activity such as sporting or cultural event" (p. 4). Using *different* flights arguably leaves us less vulnerable to this shortcoming. If a single flight is used, a rush of bookings on that flight will influence ticket prices from this date until the last ticket is sold. It will be unknown, unfortunately, at which date the mass sales have occurred, and therefore which prices correspond to the changed demand conditions. If, on the other hand, a number of flights are used, a congregation of bookings will only impact the price corresponding to one observation. Thus, the pattern of prices will be less impacted than it would have been if using only a single flight. The more different flights we study, the less vulnerable we will be to the effects of these demand-side fluctuations.

As noted, 12 major trans-Atlantic routes were sampled in the manner depicted in Figures 1-3. Each graph looks rather different (and, for purposes of space, they are not all included), but the point to be made here is that the flight prices over time seem to follow an unpredictable path. In general, there is a significant increase in prices when the departure is only days away, as passengers purchasing at such proximity to the time of the flight typically have low elasticity-of-demand. Before this time, however, there are some interesting and irregular variations in ticket pricing. At varying times between 24 weeks and one week until departure, ticket prices may either increase or decrease, in movements that may be difficult to explain through a conventional approach.

I propose to explain these variations with the methodology discussed in the theoretical discussion preceding this section. That is, in additional to basing market prices upon a number of systematic variables, prices also include a random component intended to prompt

customers to self-select based upon willingness-to-pay, improving the airline's profitability. To further explore the empirical viability of this proposition, we next examine a more extensive time series of flight prices.

*Extensive New York – London time series analysis*

In this section, we utilize a large dataset on the prices of a New York – London trip. In related literature, this trip is a popular one to analyze (see, for example, Bilotkach et al. 2006), as it is a popular journey with no shortage of flight options. The geography of the trip makes it fairly immune from one-stop competition, so we can assume that the non-stop ticket prices are those from which customers generally make their selections. The methodology of collection proceeds in a similar manner as in the preceding section: various flights are studied for dates relative to the collection date of 6 April, 2010. We again begin with a "next day" flight, but in this case increase the time to departure by smaller increments—two days. We always allow four days between the departure and returns legs of the journey. Thus, our first observation is for a flight from 7 April – 11 April; second observation is for 9 April – 13 April, etc. Data was again collected from the publically-available Orbitz.com, which did not allow access to price data for flights after 1 March, 2011. As this is quite an extensive time series—163 observations over 325 days—we limit the collection to only two carriers, United and Delta Airlines, selected because they showed the greatest price variation from each other in the study above.

Figures 4 and 5, below, present these time series for United and Delta flight prices, respectively. In analyzing these time series, we will create a regression which attempts to isolate any *systematic* effects on the prices of tickets. To this end, we will consult the available literature on dynamic determinants of airline prices in order to select those to be used as regressors on our data. We will also visually examine the time series, attempting to discern any systematic effects which can be reproduced in a regression. Though using the benefit of hindsight to pinpoint the determinants of prices may be seen as somewhat underhanded, it actually makes our task more difficult: The more variables that we input as potential determinants of prices, the less likely we are to end up with the random variation (i.e. residuals) that would lend evidence to the interpretation presented in this essay.

Figure 4: United Airlines prices, New York - London
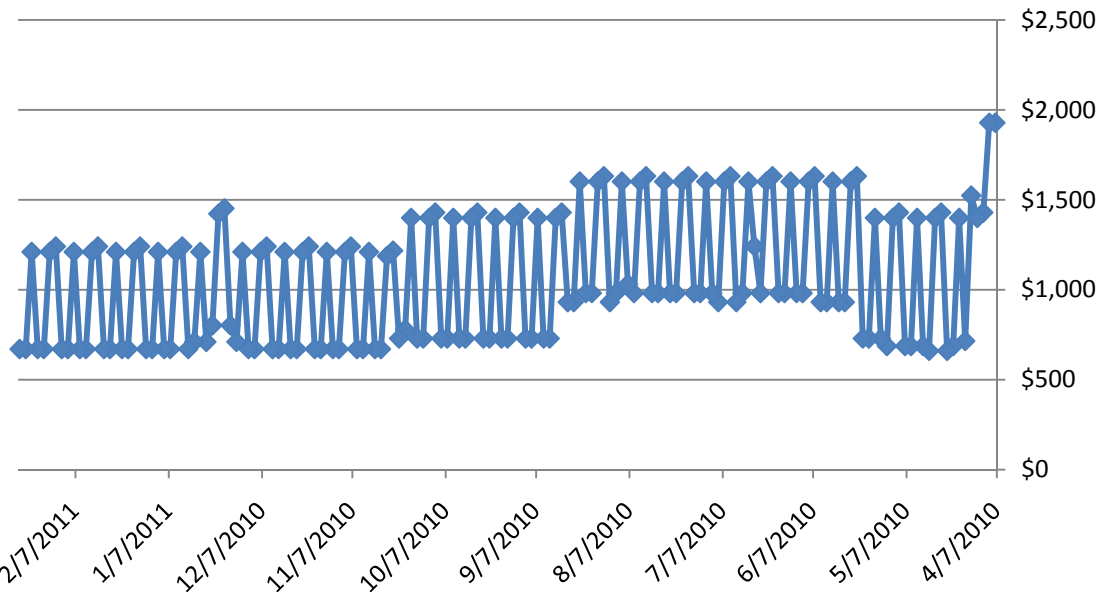
Figure 5: Delta Airlines prices, New York - London
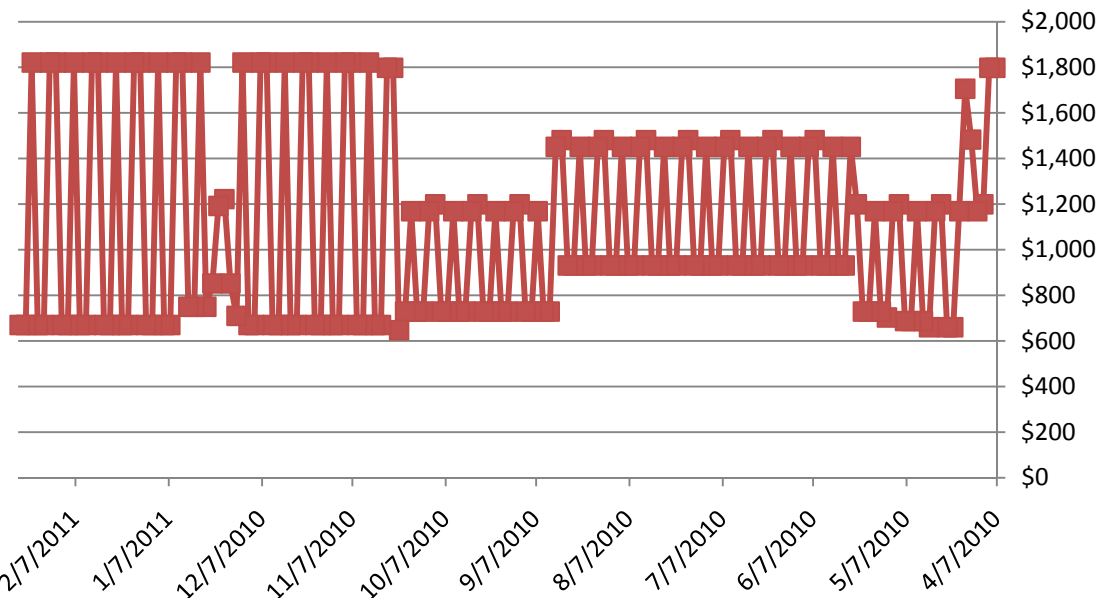
The first step in our analysis is to remove the most recent few observations from the data set. As mentioned above, these predictable price increases that take advantage of low elasticity-of-demand customers are outside the scope of this essay. In our data set, these high prices

correspond roughly to the most recent seven observations, leaving us with a remaining total of 156 observations.

We begin the model specification by adding dummy variables for day of the week on which the flight departs. The necessity of controlling for day-of-the-week is well-documented in the literature (see, for example, Sengupta and Wiggins 2006, Bilotkach et al. 2006). We include dummies for Monday, Tuesday, Wednesday, Thursday, Friday and Saturday (leaving Sunday as the control case). We regress the price data solely on these six dummy variables and a constant. Ours results indicate that the Monday and Tuesday dummies are insignificant at all confidence levels, while the Wednesday – Saturday dummy coefficients are indistinguishable from each other. The regression on Delta prices yields an adjusted R-squared of 0.781, while the regression on United data produces a comparable value of 0.802. Suspecting that a single dummy, indicating a Wednesday – Saturday departure date, may suffice, we remove the individual day-of-the-week dummies and insert this simple "WTFS" dummy in its place. While this appears to be an appropriate substitution, there appears to be some variation between Monday/Tuesday prices and Sunday prices. Thus, we include a single day-of-the-week dummy indicating a Sunday departure.

The Delta regression now produces an adjusted R-squared of 0.787, and the United regression 0.807. The addition of the WTFS dummy is thus considered an improvement. Like the individual day-of-the-week dummies, the WTFS dummy removes much of the extreme variation that is visible in Figures 4 and 5. Furthermore, utilizing this single dummy, instead of four individual day-of-the-week dummies, leaves us with several more degrees of freedom, which will be quite useful as we supplement the model with further regressors.

Promisingly, the inclusion of the WTFS dummy variable is supported by the literature: Given that we are monitoring a flight whose return leg is four days after its departure, the WTFS dummy functions as an indicator of the Saturday-night stay-over. That is, any flight departing on Wednesday, Thursday, Friday or Saturday will return following a Saturday night in the destination, while flights with other departures will not. The Saturday-night stay-over discount, Bilotkach (2005) notes, functions as a means of differentiating business from leisure customers (p. 6), since business customers are reluctant to stay in their destination through the weekend. Sengupta and Wiggins (2006) also utilize a Saturday-night stay-over dummy in their time series regression, while Piga and Bachis (2006) note that the Saturday-night stay-

over requirement is a form of price discrimination which has "dominated the [airline] industry" (p. 6) in recent years.

Next, we examine the need for dummies to control for the time of year in which the flight takes place. Button and Vega (2007) remark that there are "seasonal trends in air transportation markets that affect the aggregate demand for seats" (p. 4). Frank and Bernanke (2003) support and expand upon this view: "Seasonal price movements for airline tickets are primarily the result of seasonal variations in demand." (p. 80). As these views indicate, including seasonal variables in the regression may have the additional benefit of allowing us to control for some of the demand-side variables that we thought we would be unable follow.

While a visual examination of Figures 4 and 5 reveals a number of structural breaks, these breaks do not correspond directly with seasonal changes. In light of our preference for modeling the observed data over modeling the process which might be predicted by economic theory, we introduce a number of structural break dummy variables that correspond to the data portrayed in Figures 4 and 5. We begin with the addition of three structural breaks: One for the first 63-64 observations[11] (i.e. those departure dates that are the furthest away); the second for the next 26-28 observations; and the third for the next 49-50 observations. One can also see a fourth break following this, but this functions as the control period in our regression (recall that we eliminated the final seven observations from our dataset). Finally, a dummy variable was added to the model to indicate those flights which departed in the days surrounding major holidays. Bilotkach et al. (2006) also made this addition, yet found that indicator variables for major holidays did not significantly affect their results (pp. 6-7). Likewise, the holiday dummy variable in our regression was found to be insignificant. However, there is a large price difference in four flights departing around Christmas-time, so an additional structural break dummy was included to control for the Christmas holiday ("Break5").

Supplementing our two regressors deemed appropriate thus far (constant and WTFS dummy) with the four structural break dummies improves our regression significantly, but the residuals still exhibit variation that fluctuates between the four periods (clustered volatility). More

---

[11] The structural breaks are somewhat different for Delta and United prices, so slightly different dummy variables are used. In the case of this first period, the Delta dummy variable extends through the first 64 observations, and the United dummy through the first 63.

precisely, it appears as if the effect of the Saturday-night stay-over varies between the four periods. Our next step is therefore to interact the structural break dummies with the Saturday-night stay dummy (this interaction is facilitated by the use of the WTFS dummy, as it allows for a reasonable number of remaining degrees of freedom). We include an additional interaction term involving the fourth break, since the fluctuations in this period differ from those in the others. Since we not only have difference in periodic fluctuations, but also a shift in the intercept between the periods, we maintain the use of the break dummies alone, in addition to those interacted with the WTFS dummy. However, the dummy for the second break (without an interaction term) is insignificant at all reasonable confidence levels, so it is removed from the model. With the inclusion of the structural breaks discussed above, a simple linear (or quadratic) time trend is no longer statistically significant in the model.

The above testing process leaves us with eight regressors: a constant; WTFS dummy; four interaction terms: WTFS*Break1, WTFS*Break2, WTFS*Break3, WTFS*Break4; the un-interacted dummy variables for the first and third structural breaks; and the "Christmas-time" dummy. The adjusted R-squared values indicate that approximately 99.877 percent of the variation in Delta prices is explained with these eight regressors, while 99.377 percent of variation in the United price data is explained with the same regressors (with structural breaks slightly adjusted). Table 4, in the Appendix, provides the relevant descriptive statistics from these two regressions[12].

An examination of the available literature reveals that our model has inspected the necessity of all of the variables which previous studies indicate are important explanatory variables in ticket price changes over time on a fixed journey on a single airline. In a similar time series analysis, Pels and Rietveld (2004) used only a constant, days until departure variable, and dummy variables indicating the carrier of the flight in question. The last, of course, is not applicable to our study, while the second was deemed insignificant. In their temporal regressions, Sengupta and Wiggins (2006) utilize only the Saturday-night stay-over dummy and a set of dummy variables indicating departure and returning days of the week, both of which are used in this analysis. Bilotkach et al. (2006) use a similar set of explanatory variables to explain the price variations in question. Unfortunately, there have not been many other studies that track the dynamics of fare changes as the departure date approaches (as

---

[12] Note that we need not be concerned with some traditional econometric problems, such as multicollinearity, in our regressions, as the precision of our estimators is not of great importance to our findings.

Bilotkach et al. (2006) note (p.2); Button and Vega (2007) reiterate this view), so we do not have much previous analysis upon which to rely for selection of regressors in our model.

Figures 6 and 7[13], below, reveal that our regression lines match the observed prices to a high degree, but leave some unexplained variation in our dependent variables. Having performed a thorough search for regressors to explain the observed variation in prices, we now make the somewhat tenuous claim that the remaining variation is not systematic; instead, it is due to randomized strategic pricing on the part of the airline. That is, the residuals in the regressions constitute the random aspect of the airlines' pricing strategy. Such randomization, of course, need not be thorough across a continuum of prices: Rather, it may consist simply of randomly alternating between a few discrete prices, as demonstrated in the theoretical section of this essay. As such, there may appear to be a consistency in the observed residuals.

In order to explore the nature of each airline's mixed strategy, we examine the characteristics of the residuals in the two regressions. We perform tests for autocorrelation, heteroskedasticity and normality on the sets of residuals from both regressions, the results of which are displayed in Table 5, in the Appendix.

Figure 6: United Airlines price regression, New York – London flight

---
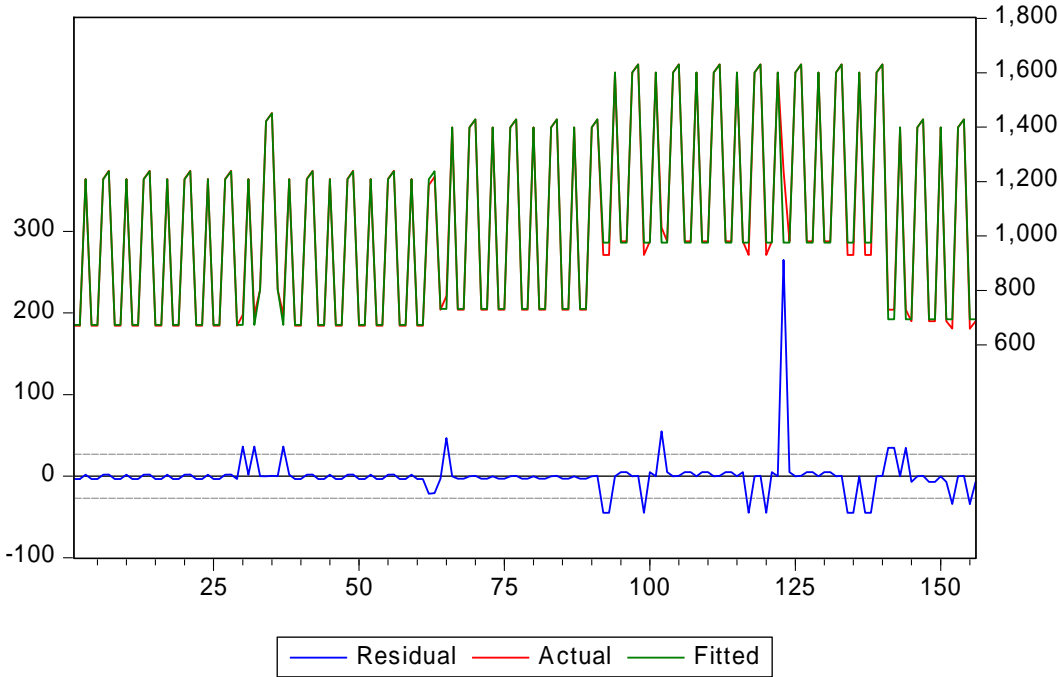
[13] The x-axes in Figures 6 and 7 (as well as those in Figures 10 and 11) mirror those in Figures 4 and 5: That is, dates closest to the origin are those furthest in the future (in this case, March 2011), and dates approach the present time as we move along the x-axis.

Figure 7: Delta Airlines price regression, New York – London flight
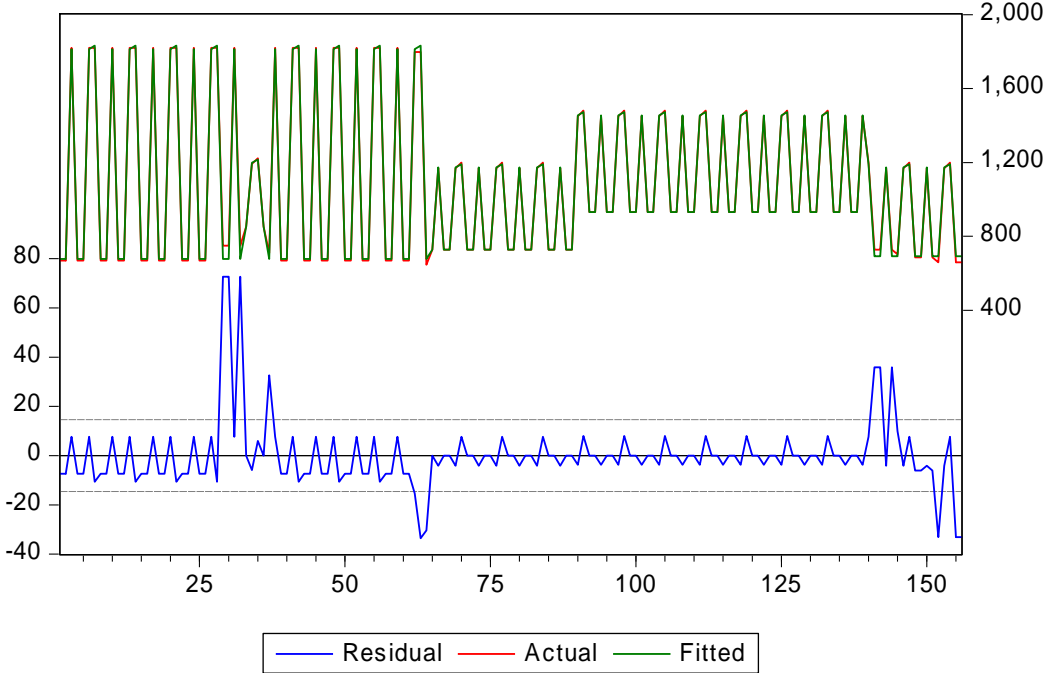
The statistics in Table 5 lead us to conclude that the residuals from the United regression are auto-correlated, while those from the Delta regression are not. Likewise, the United residuals are homoskedastic, and the Delta residuals are heteroskedastic. Our final diagnostic test reveals that neither set of residuals is normally distributed.

If, for a moment, we accept the conclusion that the residuals provide evidence of a randomized pricing strategy, we can make a number of inferences about the nature of the airlines' strategies from the distributions of the residuals. The heteroskedasticity in the Delta residuals—which can be seen clearly in Figure 7—may indicate that the airline changes the probability distribution of its mixed strategy as time passes in our sample. Early in the sample, we observe a period of high volatility in the residuals, which gives way to lower residual volatility later in the sample. Perhaps the airline uses a strategy of greater price variation during the first period before subsequently applying a strategy with less variation. Conversely, the lack of heteroskedasticity in the United price residuals may indicate a consistent mixed strategy during the period in question.

Figures 8 and 9, below, plot the residuals and their descriptive characteristics for the United and Delta regressions, respectively. Accepting our previous conclusions, these residuals

represent the particular mixed strategies that are employed by the airlines for the flight in question. That is, these residuals demonstrate the magnitude of departure from the systematic pricing scheme utilized by the airlines—due to randomized fluctuations in prices.
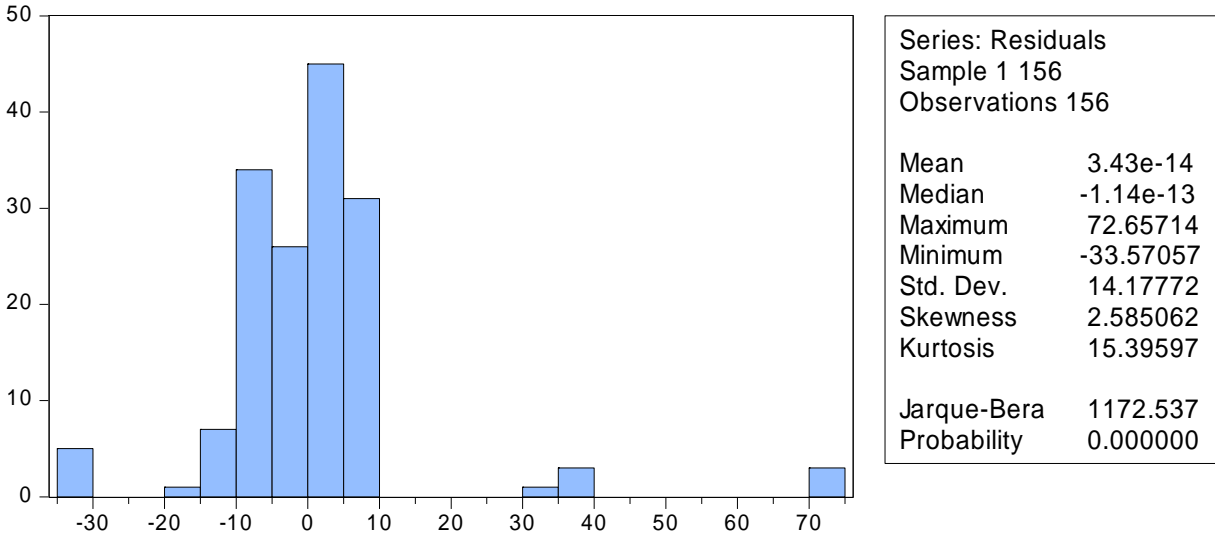
Figure 8: Residuals from United price regression

Series: Residuals	
Sample 1 156	
Observations 156	
Mean	1.90e-13
Median	-0.188955
Maximum	265.0000
Minimum	-45.00000
Std. Dev.	26.12877
Skewness	6.614725
Kurtosis	69.59001
Jarque-Bera	29960.11
Probability	0.000000
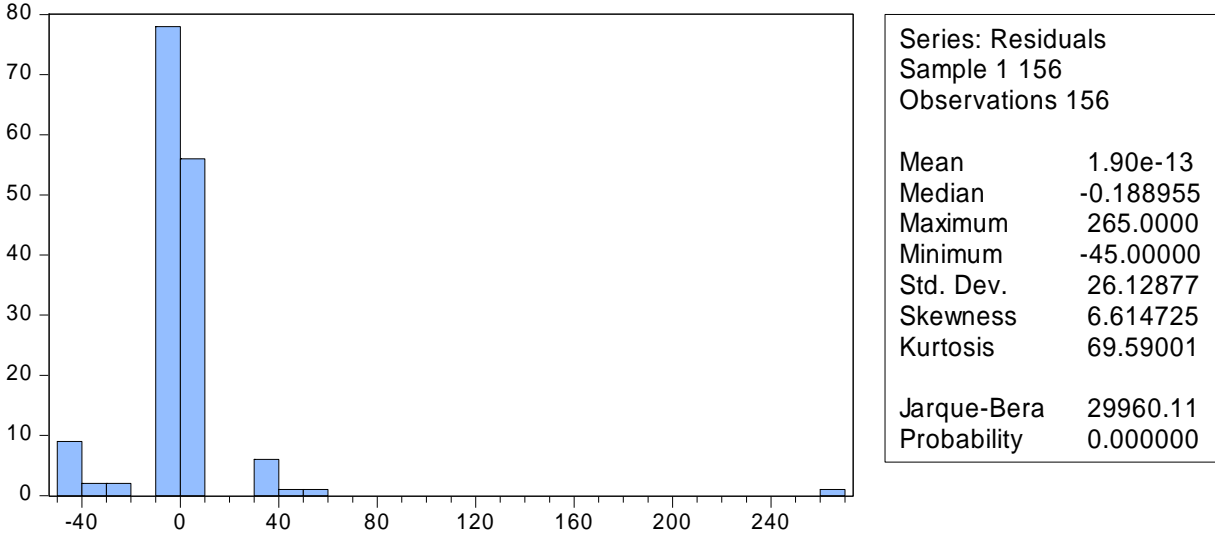
Figure 9: Residuals from Delta price regression

Series: Residuals	
Sample 1 156	
Observations 156	
Mean	3.43e-14
Median	-1.14e-13
Maximum	72.65714
Minimum	-33.57057
Std. Dev.	14.17772
Skewness	2.585062
Kurtosis	15.39597
Jarque-Bera	1172.537
Probability	0.000000

To further scrutinize the findings from the analysis of the New York – London price set, we perform a similar investigation on another set of price data. This time, we use a New York – Paris trip, with similar specifications as the New York – London trip, but now we use a six-day stay in the destination city (as discussed previously, adjusting the way in which data is sampled allows us to generalize our conclusions to a larger set of flight prices). Data was collected on 26 April 2010, three weeks after collection for the other journey, in order to

control for any pricing abnormalities that may have existed on the previous sampling date. We perform a similar analysis on the price data for this trip, experimenting with various structural break and day-of-the-week dummies, as well as constant, time trend and/or quadratic time trend, in order to find an optimal set of regressors. In the end, significance tests on the possible regressors indicate that the best set of explanatory variables includes a constant, a dummy variable for a Sunday departure (recalling that the trip is six days long, the Sunday dummy functions as an indicator of a trip which does not include a Saturday-night stay-over), a dummy variable for a single structural break, and an interaction term between these two dummies. The regression on United prices also includes dummies for Wednesday and Saturday departures, while the Delta regression includes only the latter. These sets of regressors do a somewhat poorer job of explaining variation in the prices of the two carriers than those used in the previous analysis: The United regression produces an adjusted R-squared of 0.978, while the Delta regression yields a value of 0.962. Despite an intensive search for appropriate regressors, these measures of relatively fit are low—suggesting that the airlines may be employing randomization as a larger element of their pricing scheme for this flight. Figures 10 and 11, below, plot the price data, regression lines and residuals for this New York – Paris flight, for United and Delta Airlines, respectively. Residuals are, as noted, quite large, but do not appear to be particularly systematic. This may provide further evidence of stochastic pricing.

Figure 10: United Airlines price regression, New York – Paris flight
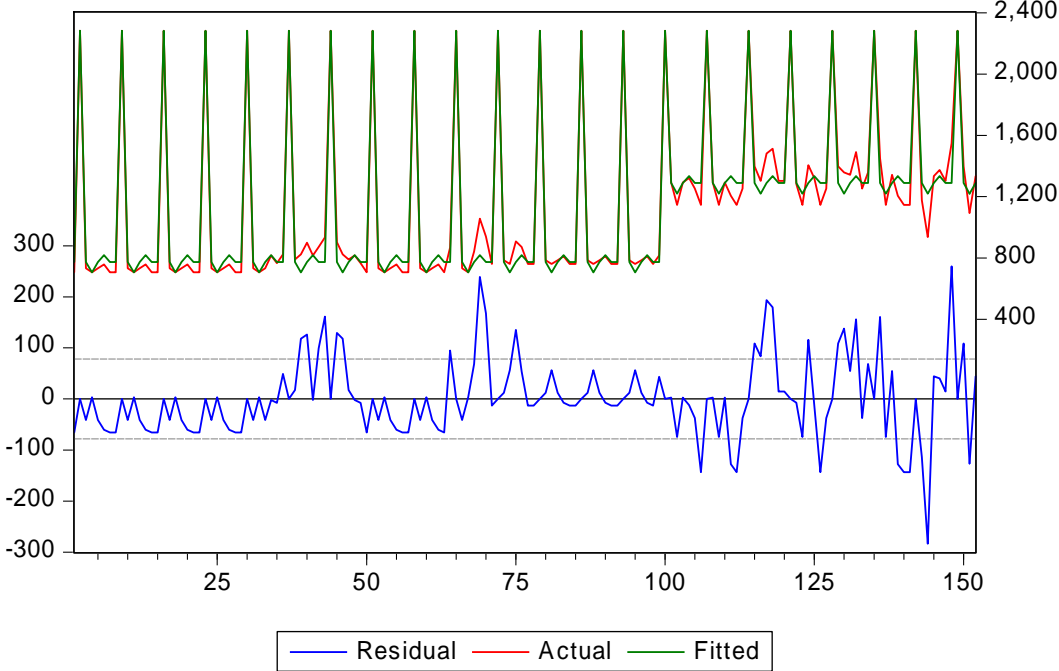
Figure 11: Delta Airlines price regression, New York – Paris flight
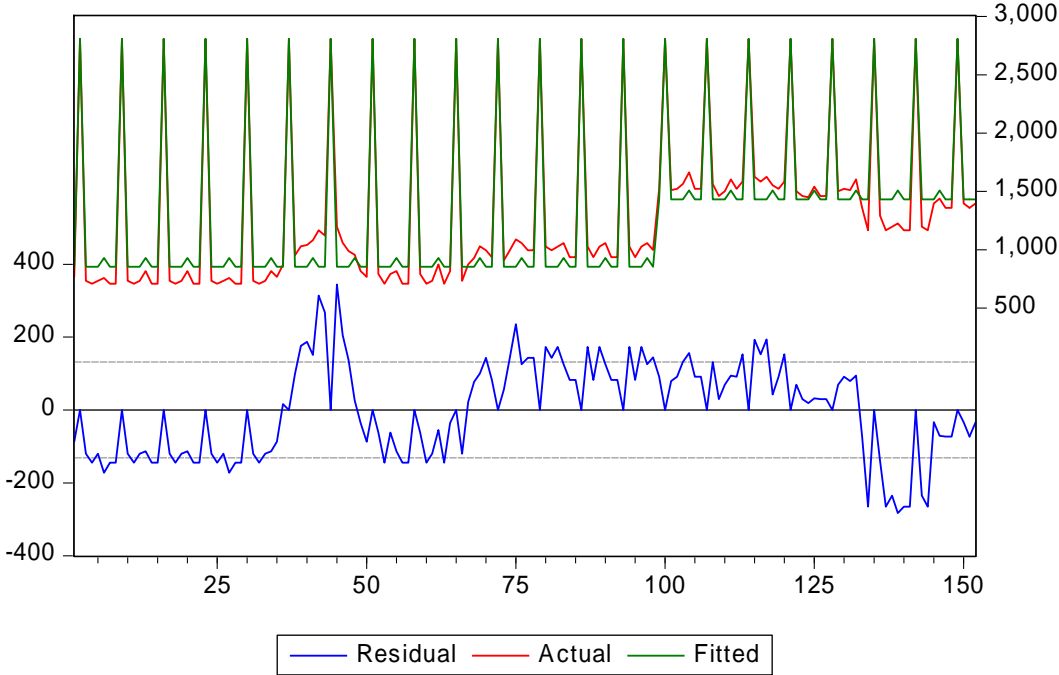
As a final note in the empirical analysis in this essay, we can assert that any price randomization is intended to take advantage of (i.e. extract surplus from) leisure passengers, as opposed to their business counterparts. As discussed above, the Saturday-night stay-over discount is a means of discriminating between business and leisure customers (this functions as an additional layer of discrimination, as airlines also offer separate business class seats intended to attract those customers with the highest willingness-to-pay[14]). The data indicate that those customers who do not stay in their destination through a Saturday night—business customers—are offered a relatively stable, high price. For the New York – Paris trip, these customers are, in fact, offered a constant price throughout the sample. Thus, all of the unexplained price variation is due to modifications of prices to those customers who do not stay over on a Saturday night. If we separate the data for the New York – London trip into those fares that have a Saturday night stay-over and those that do not, we witness a similar, though not as striking, result. Regressing this separated data on the original set of regressors (minus the interaction terms, of course), those fares pertaining to non-Saturday-night stays yield adjusted R-squared values of 0.999 for both United and Delta prices. Meanwhile, fares

---

[14] As mentioned previously, this study is not concerned with business class seats, but *is* concerned with the business/leisure distinction for economy class seats.

that *are* associated with a Saturday-night stay-over produce significantly lower adjusted R-squared values of 0.982 and 0.97 for the two carriers. If we accept the hypothesis of this essay, this difference in explained variation suggests that leisure customers have greater variance in willingness-to-pay than their business counterparts, causing the carriers to fluctuate their prices more when targeting these customers.

The above analysis provides some empirical validation of the theoretical propositions in the preceding section, but is not decisive. We cannot, of course, claim that we have performed an exhaustive search for regressors in the model, nor can we assert that any random variation is simply a result of unilateral strategic pricing by the airline. Demand-side factors, as noted, may present themselves randomly, thus causing the airlines to respond with seemingly stochastic pricing changes. In the particular case of the regression on United ticket prices for the New York – London flight, we have eliminated almost all of the variation in prices through the identified regressors. Thus, this regression provides some evidence against the propositions in this essay. However, the flip-side of this result is that we were left with significant variation in the other regressions, *using a similar set of regressors.* If we assume that systematic pricing components are largely invariant across flights (not necessarily in their magnitude, but in their presence), this lends some evidence to the view that the remaining three sets of price data exhibit stochastic variation.

Despite the qualifications noted above, it is hoped that the preceding empirical analysis suffices to *suggest* the possibility that airlines in fact employ a mixed pricing strategy, to accompany the theoretical *viability* of such a strategy discussed in the previous section. As will be noted in the conclusion, there is an opportunity for further studies to perform a deeper empirical analysis to explore the possibility of stochastic price effects in airfares.

V)   Conclusion

In its theoretical section, this essay has followed in the footsteps of studies such as Landsberger and Meilijson (1985) in examining the profitability of various inter-temporal pricing strategies when consumers have heterogeneous reservation prices for a particular good. The basic analytical model reveals that an airline can induce its customers to self-select

based upon reservation price if the preferences of these customers exhibit a sufficient degree of risk aversion. The airline can offer business customers a small, but assured, discount in the game's first stage, which may induce them to avoid a gamble on lower ticket prices in the future. The second theoretical model modifies the analytical framework and complicates the market by including a probability-of-purchase function, which implicitly assumes risk aversion among consumers. We again show that, in such a scenario, a mixed pricing strategy can function as a profitable mechanism of third-degree price discrimination. In both models, the profitability of the airline is maximized when it begins by setting a relatively high price, before decreasing the expected market price in a stochastic manner. Given the specification of the consumers' purchasing decisions, this pricing scenario induces consumers with high reservation prices to purchase before those with low reservation prices. As long as the probability of low prices in the future is high enough to allow many customers with low reservation prices to purchase tickets, but low enough to prevent high-reservation-price customers from waiting for these low prices, this method of price discrimination can be quite profitable for the airline. The theoretical study contributes to the literature with this discussion of the viability of a stochastic pricing strategy.

The empirical study in this essay charts new territory in subjecting time series of fare prices to thorough regression analysis, with the aim of stripping away all systematic pricing components. Though it cannot be concluded with certainty, it appears that a randomized pricing element may indeed be present in those time series studied. This empirical finding can be coupled with the theoretical discussion in proposing the viability of a randomized, mixed pricing strategy.

If we accept the hypothesis of this essay, we are unable to offer the consumer much guidance in his search for an ideal fare. A randomized pricing strategy like those proffered in this paper takes advantage of consumers' fear of being unable to purchase a fare at all, for a reasonable price. As long as there are customers with different reservation prices for tickets, airlines can attempt to exploit this heterogeneity through price discrimination.

With a view to the significant amount of previous research on price dispersion and competition in the airline industry, it would be interesting to perform a follow-up study on the implications of competition between carriers on the propositions contained in this essay. First, it would be necessary to ascertain to what extent fares from different carriers are substitutes for each other; a matter of some debate in the literature. With this information, one could then examine if the profitability of price randomization could be sustained in a non-monopoly market. If tickets from different airlines are, in fact, close substitutes, it would appear that a randomized pricing strategy could not be sustained in equilibrium—as long as the stochastic pricing elements are uncorrelated between carriers. If this qualification is met, randomized pricing should become less effective as competition increases, as consumers can simply wait for the inevitable fall in one of the carrier's prices. It would be pertinent to perform an analysis on time series of flight prices for a larger selection of routes, attempting to ascertain if explained variation in prices changes with the competition on a particular route.

As a secondary follow-up, some extensions to the empirical analysis in this essay could be made. As noted, the manner of data collection used here was motivated partially by time constraints. A subsequent analysis in which price data is gathered over a longer period of time may help to control for some of the flaws in the data used here. In particular, it may be beneficial to follow the more conventional approach of monitoring the price of a single flight over a period of several months.

References:

Alam, I.M.S., Ross, L.B., Sickles, R.C., 2001. Time series analysis of strategic pricing behavior in the US airline industry. Journal of Productivity Analysis 16, 49–62.

Anderson, C.K., Wilson, J.G., 2003. Wait or buy? The strategic consumer: pricing and profit implications. Journal of the Operational Research Society 54 (3), 299–306.

Baye, M.R., Morgan, J., Schloten, P., 2004. Price dispersion in the small and in the large: evidence from an internet price comparison site. Journal of Industrial Economics 52 (4), 463–496.

Besanko D.A., Winston W.L. (1990). Optimal price skimming by a monopolist facing rational consumers. Management Science 36, 555-567.

Bilotkach, V., Gorodnichenko, Y., Talavera, O., forthcoming, 2006. Are airlines' price-setting strategies different? Journal of Air Transport Management.

Bilotkach, V., forthcoming, 2005 . Understanding Price Dispersion in the Airline Industry: Capacity Constraints and Consumer Heterogeneity. Advances in airline economics, Vol. 1.

Brons, M., Pels, E., Nijkamp, P., Rietveld, P., 2002. Price elasticities of demand for passenger air travel: a meta-analysis. Journal of Air Transport Management 8, 165-175.

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

Button, K., Vega, H., 2007. The uses of the "temporal-fares-offered curve" in air transportation. Working Paper, George Mason University.

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.

Frank, R.H., Bernanke, B., 2003. Principles of Microeconomics. McGraw-Hill/Irwin.

Ingold, A., Yeoman, I., McMahon, U., 2001. Yield Management: Strategies for the Service Industries. Int. Thomson Business Press.

Landsberger, M., Meilijson, I., 1985. Intertemporal Price Discrimination and Sales Strategy under Incomplete Information. RAND Journal of Economics 16(3), 424-430.

Mantin, B., and Koo, B., 2009. Dynamic price dispersion in airline markets. Transportation Research Part E 45, 1020-1029.

Pels, E., Rietveld, P., 2004. Airline pricing behavior in the London–Paris market. Journal of Air Transport Management 10, 279–283.

Sengupta, A., Wiggins, S., 2006. Airline Pricing, Price Dispersion and Ticket Characteristics On and Off the Internet. Working Paper, Texas A&M University.

Stavins, J., 1996. Price Discrimination in the Airline Market: The Effect of Market Concentration. Federal Reserve Bank of Boston.

Stokey N., 1979. Intertemporal price discrimination. Quarterly Journal of Economics 93(3), 355-371.

Varian, H.R., 1992. Microeconomic Analysis. New York: W.W. Norton & Company, Inc.

Appendix I: Probability-of-purchase function

In the theoretical section of this essay, we established in a straightforward, analytical manner the benefit to the airline of a mixed pricing strategy. Such a strategy may prompt the airline's customers to self-select based upon willingness to pay for a ticket, allowing the airline to reap a greater portion of consumer surplus than it might be able to with a pure strategy. In further investigating the viability of a mixed pricing strategy for the airline, we introduce a more realistic response function of the consumers: a probability-of-purchasing function. In investigating the game further, it behooves us to depart somewhat from the analytical framework used in the previous section, as it becomes cumbersome when examining a greater number of customers and/or stages of market interaction. Specifically, it is difficult to analytically determine optimal purchasing and pricing actions in these conditions, so we allow Microsoft Excel (in particular, its Solver function) to compute responses for us. Using Solver allows us to easily determine profit-maximizing pricing strategies and purchasing decisions should we wish to modify reservation prices, first-stage prices, or second-stage pricing mixes. Thus, we can obtain conclusions that are applicable to a wider range of market scenarios.

Our most important modification to the model is the introduction of a probability-of-purchasing (P-of-P) function. In doing so, we recognize that the purchasing decision of a customer does not depend solely on a comparison between the market price of a ticket and his reservation price. Rather, there are a number of factors influencing a customer's purchasing decision, so his decisions are taken to be somewhat complex and, for our purposes, uncertain. We will say that the probability of him purchasing a ticket at price P is equal to the natural log of his reservation price minus the market price, divided by 10:

$$\Pr\left(R_p, P\right) = \ln\left(R_p - P\right)/10 \qquad\qquad (24)^{15}$$

This function relies upon the same assumption that is used in the simpler model above: that the purchase of a ticket is made based upon a comparison of the price offered and the internal

---

[15] If the market price is greater than or equal to the reservation price, the probability function is unidentified, but the customer, naturally, will not make a purchase (Pr = 0). This model thereby introduces the very slight modification that, if $P = R_p$, a purchase is *not* made, with certainty.

reservation price. However, for the purposes of the P-of-P function, which makes purchasing decisions stochastic, we require a transformation of the $(R_p - P)$ differential. In this model, the natural log function plays this role. This function, of course, has a negative second derivative, so the P-of-P increases at a decreasing rate as the $(R_p - P)$ differential grows. While it is difficult to justify this assumption empirically, such rationalization is actually unnecessary: Using a P-of-P function with a negative second-derivative makes it more difficult for the airline to induce its customers to self-select based upon their reservation prices, since it forces all probabilities towards the middle of the distribution. Thus, if self-selection can be demonstrated with the hindrance of a logarithmic P-of-P function, the strategy will have even greater profitability if probability-of-purchase is, for example, linear.

In constructing this second model, we (arbitrarily) assume three pricing options: $850 (High), $500 (Medium) and $300 (Low). It can be shown that the conclusions generated from this analysis are robust to changes in airline prices. The airline, as before, may not charge customer-specific prices (i.e. it cannot first-degree price discriminate), but it is of course attempting to maximize profit and may be helped in doing so if it can induce customers to self-select based upon their unique traits (i.e. it would like to third-degree price discriminate). The relevant unique customer characteristic is, as noted, reservation price. We again assume two stages in which interaction between airline and customer occurs.

For our first example, we will assume three customers, who have reservation prices of $1,000, $800 and $500, respectively. We again begin with an examination of the simplest possible strategies: setting one price for the duration of the game (in this case, in both stages). The expected profit of each simple strategy is presented in Table 1, below. It may seem counterintuitive that the customers apply a probability of purchase to the same market price of a ticket in the two stages (that is, there is a chance that they reject a ticket offer in stage one only to accept it at the same price in the subsequent stage). However, we assume here that customers have "zero memory": each ticket offer is considered anew in each stage without regard to past market interactions.

Table 1: Expected profit of simple pricing strategies				
P = High ($850)				
	Reservation price (R$_p$)	Pr: Probability of purchase in each stage: ln[Rp - p] / 10	Probability of purchase in entire game: $1 - [(1 - Pr)^2]$	Expected profit (P * Prob of purchase in entire game)
Customer 1	$1,000	0.50	0.75	$638
Customer 2	$800	0.00	0.00	$0
Customer 3	$500	0.00	0.00	$0
			Total:	**$638**
P = Medium ($500)				
Customer 1	$1,000	0.62	0.86	$428
Customer 2	$800	0.57	0.82	$408
Customer 3	$500	0.00	0.00	$0
			Total:	**$836**
P = Low ($300)				
Customer 1	$1,000	0.66	0.88	$264
Customer 2	$800	0.62	0.86	$257
Customer 3	$500	0.53	0.78	$234
			Total:	**$755**

As we can see from Table 1, there is a considerable amount of available profit that the airline is not able to capitalize upon. The maximum attainable profit is equal to the sum of the consumers' reservation prices, in this case $2300 (though the discrete pricing strategies we allow for the airline permit a maximum profit of $850 + $500 + $300 = $1650). In the scenarios described above, the airline takes only about 39%, 51% and 46% of the available profit in the high, medium and low pricing cases, respectively. As before, we proceed to analyze whether the airline can improve these payoffs by employing a mixed pricing strategy.

As in the previous model, if the airline alternates its strategies in a way that is known to customers, the customers will simply wait for their favored price and consider whether or not to purchase a ticket. For example, a pricing strategy of high in the first stage and medium in the second stage would simply encourage customers to delay their (possible) consumption until the second stage, ignoring the overpriced first-stage tickets. This would require a different P-of-P specification than the one noted above, as the first-stage probability-of-purchase would decline in favor of the second stage probability (the P-of-P function, as specified, relies on the assumption that future prices are either uncertain or non-decreasing). Thus, the airline has nothing to gain by employing a mixed, but observable and non-random, pricing strategy.

But what if the airline were to adopt a randomized mixed strategy? Because customers' probabilities of accepting a particular price are positively related to their own reservation prices, a mixed pricing strategy may prompt customers to self-select to the benefit of the airline, functioning as a mechanism of third-degree price discrimination. That is, the airline may offer a price that is fairly attractive to a customer with a high reservation price in the first stage, before lowering its price in the second stage with the expectation that customers with high reservation prices have already made their purchases. If a customer with a reservation price of $1,000 faces a market price of $800, he will be more inclined to purchase at this price if he is uncertain about which prices will be offered in the future. If, on the other hand, he is aware of future ticket prices, he will simply wait for the lowest one, as noted above. If the customer with reservation price of $1,000 in fact purchases a ticket for $800, a stochastic (and perhaps downward-trended) movement in the market price of the ticket may later allow a customer with a $500 R_p to purchase at $400, for example. Thus, the zero-marginal cost

43

airline ticket seller will have avoided his two market pit-falls: Selling to a customer significantly below his reservation price, or not selling to a customer at all.

As a simple example, let us assume that the airline utilizes a mixed strategy placing a weight of one-third on each of its three pricing strategies. We further assume that the "dice are rolled" (twice) in this randomized strategy, and the outcome is that tickets will be priced "medium" in the first stage and "low" in the second stage. At the time of purchase in stage one, the stage-two price is, of course, unknown to the customers (and indeed to the airline, in theory)—though the customers are aware of the strategy mix that the airline employs. Given this uncertainty, customers in our model will not be inclined to wait for a lower price in stage two, as there also exists a possibility of facing a *higher* price in the subsequent stage. We implicitly assume in this model that customers are uniformly risk averse: If there is any reasonable possibility of facing a higher price in future stages, they prefer to purchase—with their specified probability—a ticket at the present price.

Table 2, below, reveals the expected payoff for the exemplar mixed strategy discussed above. As we see, the expected profit of $909 compares favorably to that yielded by the pure strategies in Table 1. Of course, matters are not this simple for the airline: If it is playing a mixed strategy, there will be unfavorable outcomes as well as favorable ones. If, for example, the dice are rolled and the second-stage price offered turns out to be high instead of low, the profit to the airline ($757) will be inferior to the profit attained through the simpler strategy of a known and stable medium price throughout the game. Since there is no need to randomize in the first stage of the game—there is no past action that would have been affected by such randomization—let us allow the airline to set a price of its choice in the first stage. In the second stage, it then chooses a strategy combination: High price with probability α_1, medium price with probability α_2, and low price with probability α_3, with $0 \leq \alpha \leq 1 \ \forall \alpha$, and $\sum \alpha = 1$. To begin with, let us (arbitrarily) set $\alpha_1 = 0.2$, $\alpha_2 = 0.4$ and $\alpha_3 = 0.4$. Assume that the airline chooses to set the medium price in stage one. The expected payoff to the airline is now $850, as we see in Table 3. Table 3 also includes the expected payoff ($872) of the pricing combination denoted by $\alpha_1 = 0.1$, $\alpha_2 = 0.3$ and $\alpha_3 = 0.6$. In this particular example, we see that the expected payoff to the airline increases with α_3; that is, the airline's payoff improves as it increases the likelihood of the lowest price in the second stage (following a medium price in

the initial stage). The optimal pricing strategy is, in fact, the one illustrated by Table 2, with the low price being set with certainty in stage two (we can see this by, for example, asking Excel's Solver application to maximize the airline's profit function with respect to its relevant constraints). However, this violates the requisites of our model: To ensure a risky environment for the customer, the airline must play a randomized mix of its pricing strategies. We may wish, for instance, to mandate $0.1 \leq \alpha \; \forall \alpha$ to ensure that the consumer is kept sufficiently "in the dark" for our purposes (in this case [as one can guess], the airline optimizes by setting the high and medium prices with probability 0.1 each, and the low price with the remaining probability of 0.8, attaining an expected payoff of \$887).

Table 2: Expected profit of example mixed strategy

P(1) = Medium ($500) in stage one; P(2) = Low ($300) in stage two

	Reservation price (R_p)	Pr{1}: Probability of purchase in first stage: $\ln[R_p - 500] / 10$	Pr{2}: Probability of purchase in second stage: $\ln[R_p - 300] / 10$	Probability of purchase in entire game: $1 - [(1 - Pr\{1\})*(1 - Pr\{2\})]$	Expected profit: $(P(1) * Pr\{1\}) + (P(2) * Pr\{2\} * (1 - Pr\{1\}))$
Customer 1	$1,000	0.62	0.66	0.87	$385
Customer 2	$800	0.57	0.62	0.84	$365
Customer 3	$500	0.00	0.53	0.53	$159
				Total:	**$909**

*Note: Expected profit formula explained: Price in stage one * probability of purchase in stage one plus price in stage two * probability of second-stage purchase * probability **not** purchased in first stage.*

Table 3: Expected profit of example mixed strategies (2)

			Probability of purchasing high-priced ticket in stage 2	Probability of purchasing medium-priced ticket in stage 2	Probability of purchasing low-priced ticket in stage 2		
P(1) = Medium ($500) in stage one							
P(2): α1 = 0.2; α2 = 0.4; α3 = 0.4						Probability of purchase in entire game: $1 - [(1 - Pr\{1\})*(1 - \sum Pr\{2\})]$	Expected profit: $(P(1) * Pr\{1\}) +[(850*Pr\{2\text{-H}\} + 500*Pr\{2\text{-M}\} + 300*Pr\{2\text{-L}\})] * (1 - Pr\{1\})$
	Reservation price (R_p)	Pr{1}: $\ln[R_p - 500] / 10$	Pr{2-H}: $0.2*(\ln[R_p - 850] / 10)$	Pr{2-M}: $0.4*(\ln[R_p - 500] / 10)$	Pr{2-L}: $0.4*(\ln[R_p - 300] / 10)$		
Customer 1	$1,000	0.62	0.10	0.25	0.26	0.85	$420
Customer 2	$800	0.57	0.00	0.23	0.25	0.78	$366
Customer 3	$500	0.00	0.00	0.00	0.21	0.21	$64
						Total:	**$850**
P(1) = Medium ($500) in stage one							
P(2): α1 = 0.1; α2 = 0.3; α3 = 0.6	Pr{1}: $\ln[R_p - 500] / 10$	Pr{2-H}: $0.1*(\ln[R_p - 850] / 10)$	Pr{2-M}: $0.3*(\ln[R_p - 500] / 10)$	Pr{2-L}: $0.6*(\ln[R_p - 300] / 10)$			
Customer 1	$1,000	0.62	0.05	0.19	0.39	0.86	$407
Customer 2	$800	0.57	0.00	0.17	0.37	0.80	$370
Customer 3	$500	0.00	0.00	0.00	0.32	0.32	$95
						Total:	**$872**

Of course, the particular optimal weights and payoffs attached to the various pricing strategies are a function of the arbitrarily chosen parameters of the model: the airline's pricing choices and the customer's reservation prices (as well as the form of the probability-of-purchase function). Different values of these parameters will affect the magnitudes of the optimal pricing formulations, but *not* the general conclusions, which can be drawn logically and intuitively. The assumptions upon which these conclusions rely are as follows:

1. There exist heterogeneous consumers with varying reservation prices; i.e. consumers possess different willingness-to-pay for airline tickets.

2. Consumers' probability of purchasing an airline ticket is a positive function of their reservation price and a negative function of the market price for tickets.

3. Consumers are risk-averse: If there is uncertainty regarding future prices of tickets, consumers will make independent purchasing decisions with regard to each offer from the airline. That is, risk-averse consumers are unwilling to wait for a possible lower price, as long as there is some reasonable probability that the price will increase in the future. In this manner, the risk aversion of customers has been implicit in the model.

4. Market prices and future pricing strategies are observable to consumers. That is, if the airline sets prices in a non-stochastic manner, consumers are able to figure out the airline's pricing scheme and act accordingly.

5. The market for airline tickets takes place in multiple stages; a simplification of the continuous market which exists in reality. The market depicted above represents a market in which purchases can take place, say, at only precisely a month and a week from the departure date.

All of the above assumptions appear to be plausible, or, at worst, do not greatly impact the conclusions derived there from. Note that, for simplicity, this model neglected to take into account a number of factors which may be appropriate to include in a follow-up model with greater complexity: Varying degrees of risk aversion between customers; a more dynamic specification in which consumers react uniquely to differing probabilities and magnitudes of future prices, rather than reacting uniformly whenever there is uncertainty; and allowing airlines a continuum of prices for selection, among other possible modifications. Though the game above was presented with only two stages, for simplicity, its logic can easily be extended to a multi-stage interaction. In this case, airlines would set a probability function for

each stage (or perhaps just one function that is unchanged across stages), allowing it a greater ability to differentiate between customers with varying willingness-to-pay.

Appendix II: Tables in empirical analysis

Table 4: Regression coefficients/significance (New York - London flight)					
	United Prices		**Delta Prices**		
	Coefficient (standard error)	*Probability*	*Coefficient (standard error)*	*Probability*	
Constant	1399.189 (6.766)	0.000	1173.147 (3.791)	0.000	
WTFS*Break1	-534.812 (7.442)	0.000	-1135.977 (4.016)	0.000	
WTFS*Break2	-667.064 (9.544)	0.000	-444.1468 (5.348)	0.000	
WTFS*Break3	-625.189 (8.112)	0.000	-523.7384 (4.336)	0.000	
WTFS*Break4	-704.989 (10.874)	0.000	-480.0468 (5.976)	0.000	
WTFS*Break5	-85.47145 (27.870)	0.003	790.1023 (15.116)	0.000	
Break1	-190.8476 (8.323)	0.000	640.173 (4.594)	0.000	
Break3	202 (8.648)	0.000	281.5916 (4.719)	0.000	
Break5	213.942 (19.824)	0.000	-615.4452 (10.757)	0.000	
Sunday	29.43314 (7.044)	0.000	18.25084 (3.824)	0.000	

Table 5: Residual diagnostic tests (New York - London flight)		
	Delta price data	United price data
Autocorrelation		
Durbin-Watson test statistic	1.362	1.827
Serial correlation LM test (5 lags)		
F-statistic	8.375	0.583
Probability	0.000	0.713
Heteroskedasticity		
Breusch-Pagan-Godfrey		
F-statistic	2.163	0.812
Probability	0.028	0.606
White		
F-statistic	2.163	0.812
Probability	0.028	0.606
Normality		
Jarque-Bera statistic	1,172.537	29,960.110
Probability	0.000	0.000