

**easyJet[®] Pricing Strategy:
Should Low-Fare Airlines Offer Last-Minute Deals?**

Oded Koenigsberg
Columbia Business School
Columbia University, New York, NY 10027
ok2018@columbia.edu

Eitan Muller
Stern School of Business
New York University, New York, NY 10012
Leon Recanati Graduate School of Business Administration
Tel Aviv University, Tel Aviv, Israel 69978
emuller@stern.nyu.edu

Naufel J. Vilcassim
London Business School
London, England NW1 4SA
nvilcassim@london.edu

October 2007

We are grateful to Stelios Haji-Ioannou, John Stephenson, and Ben Meyer at easyJet for making the data available for analysis and to Asim Ansari, Eyal Biyalogorsky, Fabio Caldieraro, Raghuram Iyengar, Don Lehmann, Olivier Toubia, Rajeev Tyagi, Garrett van Ryzin, Peter Rossi, and two anonymous reviewers for their thoughtful comments and suggestions. The authors would also like to thank Hernan Bruno for his excellent research assistance.

easyJet[®] is a registered trademark.

easyJet[®] Pricing Strategy: Should Low-Fare Airlines Offer Last-Minute Deals?

Abstract

easyJet, one of Europe's most successful low-cost short-haul airlines, has a simple pricing structure. For a given flight, all prices are quoted one-way, a single price prevails at any point, and, in general, prices are low early on and increase as the departure date approaches. We observe from these policies and from the empirical section of this paper that easyJet employs three distinct strategies: 1) it does not offer last-minute deals, 2) it offers a single class and lets price be the sole variable that controls demand, and 3) it varies the time at which tickets are first offered for sale (duration of sale). The first two policies are in stark contrast to traditional airline pricing strategies. Many airlines offer last-minute deals, either directly or via resellers. Second, the current prevailing practice is to control demand via seat allocation to various classes rather than by offering a single class and letting price be the sole variable that controls demand.

The main objective of this research is to study the conditions under which offering a last-minute deal is optimal under the single-price policy. We also learn how the duration of ticket sales is affected by consumer characteristics. We find that, for an intermediate capacity level, uncertainty with respect to the arrival of the business segment will cause the firm to offer last-minute deals and thus partially price-discriminate within the tourist segment. The same is true for uncertainty with respect to the actual behavior of the firm: if consumers are uncertain whether the firm will offer last-minute deals, then, in equilibrium, both in a one-shot game and in a repeated game, the firm will, with some probability, offer such deals. In addition, we found that for an intermediate capacity level, the larger the number of segments (that differ in price sensitivity), the longer the duration of the period in which tickets are offered for sale.

1. Introduction

Low-cost carriers have become major players in the airline industry around the world. Airlines such as easyJet and Ryanair in Europe and Southwest and JetBlue in the U.S. are forcing major changes in pricing schemes. easyJet has emerged as one of the most successful low-cost airlines in Europe since its launch in 1995. One key aspect of its marketing strategy is a simple fare structure in which all fares are quoted one-way and a single price is quoted for all seats on a given flight at any point in time and without any restrictions (such as a required Saturday-night stay). However, the price charged for a seat on a given flight changes over the period between seats on the flight being made available for booking and the date of departure. All easyJet sales are booked directly either online or by telephone. The company's website (www.easyJet.com) describes its pricing policy as being "based on supply and demand, and prices usually increase as seats are sold on every flight. So, generally speaking, the earlier you book the cheaper the fare will be."

In Figures 1a and 1b, we plot over time the number of seats sold on a given date and the price charged per seat for a flight between Liverpool, England, and Alicante, Spain, departing on a Monday in the winter of 2003 and another between Stansted (a London airport) and Edinburgh, Scotland, departing on a Monday in the summer of 2003. Certain distinct patterns are evident from Figures 1a and 1b. We note that seat sales in both cases exhibit a discrete pattern: for the Liverpool-Alicante flight, sales are spread out over time but there is a surge in sales two to three weeks after the flight becomes available for booking; for the London-Edinburgh flight, sales are sparse in the first half of the period, followed by a surge in sales on a certain date and then moderate activity toward the end of the period. Thus, there seems to be

some uncertainty in the pattern of demand. Moreover, while in both cases there is a clear upward trend in the price charged, there is a distinct difference in the magnitude of the change.

Figure 1a: Seats sold and prices paid (in British pounds) for a one-way ticket from Liverpool to Alicante (flight date January 27, 2003)

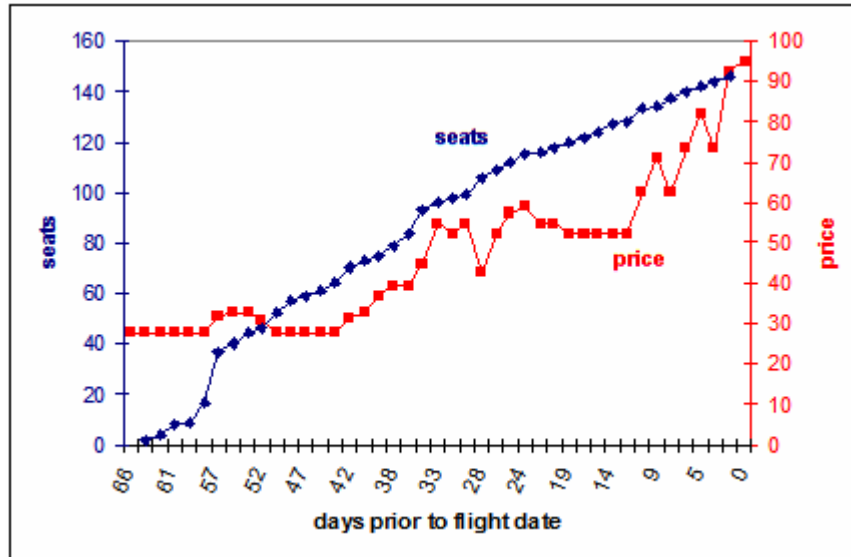
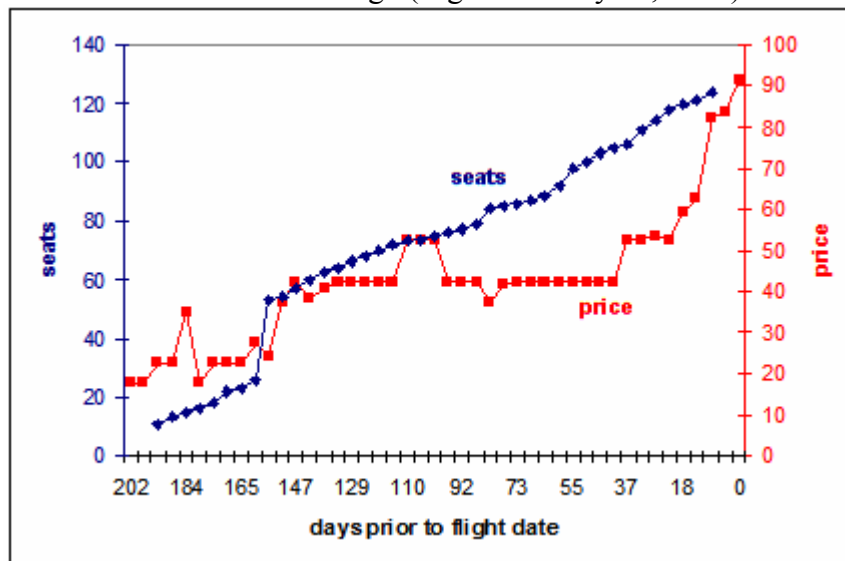


Figure 1b: Seats sold and prices paid (in British pounds) for a one-way ticket from London to Edinburgh (flight date July 21, 2003)



For the Liverpool-Alicante flight, the price increases about 2.5 times, from about £40 to £100, whereas the price increase for the London-Edinburgh flight is about 4.5 times, from around £20 to a high of £90.

We observe from these policies and from the empirical section of this paper that easyJet employs three distinct strategies: 1) it does not offer last-minute deals, 2) it offers a single class and lets price be the sole variable that controls demand, and 3) it varies the time at which tickets are first offered for sale (duration of sale). The first two policies are in stark contrast to traditional airline pricing strategies. Many airlines offer last-minute deals, either directly or via resellers. For example, in some European airports, one can buy tickets at greatly reduced prices for same-day flights. Second, as we show in the next section, the current prevailing practice is to control demand via seat allocation to various classes rather than by offering a single class and letting price be the sole variable that controls demand.

The main objective of this research is to study the conditions under which offering a last-minute deal is optimal under a single-price policy. We also learn how the duration of ticket sales affects the consumer characteristics. We find that, for an intermediate capacity level, uncertainty with respect to the arrival of the business segment will cause the firm to offer last-minute deals and thus partially price-discriminate within the tourist segment. The same is true for uncertainty with respect to the actual behavior of the firm. That is, if consumers are uncertain whether the firm will offer last-minute deals, then, in equilibrium, both in a one-shot game and in a repeated game, the firm will, with some probability, offer such deals. In addition, for an intermediate capacity level, we found that the larger the number of segments (that differ in price sensitivity), the longer the duration of the period in which tickets are offered for sale.

The rest of this paper is organized as follows; in section 2, we relate our research to the extant literature on airline pricing. In section 3, we describe our empirical study that analyzes the pattern of the data to construct the model's assumptions. In section 4, we present the structure of the model and its underlying assumptions. This presentation is followed by the

derivation of our primary analytical results in section 5. In section 6, we conclude by identifying issues for future research.

2. Airline Ticket Pricing and Yield Management

Our research on airlines' low-fare strategies draws from the literature on airline pricing in economics (Dana 1998, 1999; Kretsch 1995; Borenstein and Rose 1994; Morrison and Winston 1990) and marketing (Biyalogorsky et al. 1999, 2005 Carpenter and Hanssens 1994) and on revenue-management literature in operations research. Our modeling of consumers' choices is in line with that of Dana (1998), who assumed that there are two types of consumers—leisure travelers and business travelers—who have differing (point) valuations for the service and differing probabilities of usage. In our model, we assume that the consumer segments' valuations are derived from differing (uniform rather than point) distribution functions. Like Dana (1998), we also assume that the two segments have differing levels of uncertainty regarding flights. We assume that leisure consumers regard the flight with certainty but that business consumers have uncertainty regarding the need to take the flight and thus the value placed on it. A major difference between our model and Dana's (1998) is that, when demand exceeds capacity, Dana's model resolves excess demand by rationing while in our model the firm uses price to control demand. The resulting continuous distribution allows us to let price depend on remaining capacity and thus control remaining demand.

The analysis here also ties into operations research literature on revenue (yield) management. McGill and van Ryzin (1999) provide a detailed survey of revenue management wherein they discuss four areas: forecasting, overbooking, seat inventory control, and pricing.

Traditionally, the operations literature on airline revenue management has been concerned with seat inventory and capacity-planning problems (see Talluri and van Ryzin (2004) for an extensive review). This is in contrast to revenue management in other industries; Talluri and van Ryzin (2004) noted that “some industries use price-based RM (retailing), whereas others use quantity-based RM (airlines). Even in the same industry, firms may use a mixture of price- and quantity-based RM. For instance, many of the RM practices of the new low-cost airlines more closely resemble dynamic pricing than the quantity-based RM of the traditional carriers” (p. 176).

Recently, a growing body of revenue-management research has examined pricing decisions, including Talluri and van Ryzin (2004), You (1999), Feng and Gallego (1995), Gallego and van Ryzin (1994, 1997), and Watherford and Pfeifer (1994). The majority of this research assumes that consumers are *myopic*. In a recent survey on dynamic pricing, Elmaghraby and Keskinock (2003, p. 1298) wrote that “An important element that is largely missing, both in most of the academic literature and price optimization software, is the consideration of strategic customer behavior” and that “An interesting but equally challenging research direction would be to incorporate into the models customers’ strategic purchasing behavior in response to the firm’s pricing strategy.” For example, while Talluri and van Ryzin (2004) focus almost exclusively on myopic consumers, when they do consider strategic consumers they consider only the case in which the firm has no capacity constraint. Recent exceptions that include strategic consumers are the working papers by Jerath, Netessine and Veeraraghavan (2007), Ovchinnikov and Milner (2007), and Liu and van Ryzin (2006).

Our research thus differs from past research in several aspects. First, we assume that consumers are forward-looking and therefore behave strategically. Note that although airline tickets are a perishable good, there are some similarities to the strategic behavior of consumers

in durable goods (see Desai, Koenigsberg and Purohit (2004), and Shulman and Coughlan (2007) for recent examples). Second, under the condition in which the single-price business model is used by low-cost carriers, we investigate the conditions under which it pays to offer last-minute deals. Third, in most other yield-management models, demand exceeding capacity was resolved by rationing while in our model firms use price to control demand.

3. Empirical Analysis

In this section we utilize data from a European low-cost carrier, easyJet, first to observe its pricing and demand patterns. Using these observations, we learn about the existence of market segments that vary by price sensitivity, a central assumption of our model.

For our empirical analysis, we collected data on twenty-three easyJet flights between six different European city pairs during the year 2003, some in winter and others in summer. The flights departed on two different days of the week, Monday and Sunday. For each flight, the data include the number of seats sold per day (if any) and the price at which each seat was sold, from the first day on which seats went on sale through the date of departure. This period ranged from 63 days for a flight between Liverpool, England, and Alicante, Spain, to 211 days for a flight between East Midlands, England, and Barcelona, Spain. Additionally, we received data on the total number of seats available for each of easyJet's flights between the city pairs. Since we later assume the existence of two segments of consumers (tourists and business travelers), the first phase of analysis was to test the validity of this assumption. The data suggest that customers arrive at discrete points in time, so a latent-class Poisson regression model is an appropriate method (Wedel et al. 1993).

The negative binomial (NB) distribution also has been used in previous research to model aggregate demand for airline seats as it overcomes some well-known limitations of a Poisson model. We note that an aggregate NB model for demand can be derived by assuming that demand is Poisson at the individual level and by accounting for heterogeneity among individuals using a gamma distribution. In our analysis, we model demand at the individual level using a Poisson model but account for heterogeneity using a latent-class approach, which also can be interpreted as providing a finite approximation to any mixing distribution, such as the gamma. Therefore, the demand model we use is quite flexible. We also added a time dimension to the Poisson regression so as to capture the fact that both bookings as well as price increase towards final flight time.

Let the index j denote the relevant latent segment. If the customers' arrival rate is given by $\lambda_j = \lambda_{j_0} \exp(\beta_j p + g_j(t))$, for some increasing continuous function $g(t)$, then the probability that in any given period y customers buy tickets at a price p is

$$P(Y = y / p, \beta_j, \theta_j) = \sum_{j=1}^L \theta_j \lambda_j^y \exp(-\lambda_j) / y!, \text{ where } \sum_{j=1}^L \theta_j = 1$$

In the preceding equation, θ_j is the probability that the arriving customer belongs to the latent class, or segment $j = 1, 2, \dots, L$. Since we noted from the data that bookings seem to increase at an increasing rate with time, we have specified the function $g(t)$ as a quadratic equation of time. We estimate the parameters of the latent-class Poisson regression model using maximum-likelihood methods. Because the number of segments is unknown, we incrementally add segments until there is no improvement in the fit of the model as measured by the Bayesian information criterion (BIC); in other words, we increase the number of

segments from one to two to three and so on until the BIC value is minimized (see Chintagunta, Jain, and Vilcassim (1991)). The results of the estimation are shown in Table 1.

We note from Table 1 that the fit of the Poisson regression model improves considerably when going from one segment to two segments in all cases except for two. For all flights, the two-segment solution fits the data better than the three-segment solution. Thus, there is strong empirical support for our modeling assumption of two segments of customers.

Table 1: Bayesian information criterion for the twenty-three flights

Flight Number	BIC Values		
	One Segment	Two Segments	Three Segments
1	111.6	107.1	121.3
2	103.2	87.8	101.9
3	186.6	153	169.8
4	202.8	131.1	147.6
5	72.2	84.4	98.5
6	80.3	87.2	99.0
7	247.7	166.9	169.9
8	149.2	127.1	142.5
9	116.8	89.0	102.6
10	126.9	116.3	132.2
11	175.0	140.3	157.4
12	170.4	145.9	162.3
13	170.3	70.3	81.8
14	166.6	100.7	116.0
15	165.3	156.7	167.7
16	171.3	158.8	175.3
17	99.2	73.2	84.8
18	76.2	75.4	86.6
19	178.5	143.8	154.9
20	152.4	140.3	155.0
21	142.4	138.2	153.1
22	160.4	148.1	165.2
23	161.5	141.6	159.4

We next examine the relationship between airline seat prices and time and how the nature of this relationship depends on available seat capacity. For each time period ($Time_{it}$, a week in this case) and for each flight (twenty-three flights in all), we estimate the following regression model for the price variable: $Price_{it} = \alpha_i + \beta_i Time_{it} + \varepsilon_{it}$.

To determine how the preceding relationship varies with seat capacity, we let

$$\beta_i = \nu + \eta \cdot C_i + \delta \text{ Summer_Dummy} + \lambda \text{ Sunday_Dummy} + \mu_i$$

where C_i denotes the capacity of flight i and where Summer_Dummy = {1 if the flight is in the summer and 0 if it is in the winter} and Sunday_Dummy = {1 if the flight is on a Sunday and 0 if it is on Monday} are two control variables that we introduce because they could affect the rate at which price changes. The results of the regression are given in Table 2.

Table 2: Regression results: Price vs. capacity

Regression Analysis			Price Slope – Capacity	
Price – Time			Regression	
Flight	Intercept	Price	Variable	Estimate
Number	(α_i)	Slope (β_i)		(standard error)
1	12.1	0.856	Intercept	6.597
2	14.2	0.485		(2.569)
3	17.4	1.799		
4	12.8	1.382	Capacity	–4.058
5	20.5	0.203		(1.967)
6	12.0	0.426		
7	7.0	0.821	Summer_Dummy	–1.148
8	7.6	0.492		(0.623)
9	18.1	0.780		
10	30.2	–0.302	Sunday_Dummy	0.051
11	51.5	0.542		(0.617)
12	48.8	0.485		
13	12.5	2.497	R-Square	25.8%
14	21.6	0.267		
15	60.6	0.645		
16	44.7	1.293		
17	13.7	5.430		
18	13.5	5.618		
19	142.8	–1.031		
20	128.1	–0.564		
21	42.3	0.706		
22	48.1	0.889		
23	53.3	0.679		
Mean		1.061		
Slope				

We see from Table 2 that, in the price versus time regression, twenty of the twenty-three slope coefficients (β_i) are positive and the overall mean is positive (1.061). Thus, price increases over time. We also note from the price-slope-versus-capacity regression in Table 2 that the coefficient of the capacity variable is significantly different from zero and negative, implying that the rate at which price changes decreases as capacity increases.

Another way to look at the relationship between price and capacity is to regress price against remaining capacity. For each time period ($Time_{jt}$, a week in this case) and for each pair of cities (six pairs), we estimate the following six regression models for each city pair (four flights per city for five city pairs and three flights per city pair for a single flight) for the remaining capacity variable ($RemCap_{jt}$, $j = 1, 2, \dots, J$).

$$Price_{jt} = \alpha_i + \gamma_j RemCap_{jt} + \delta Summer_Dummy + \lambda Sunday_Dummy + \xi_{jt}.$$

Table 3: Regression results: Price vs. remaining capacity

Route	No. of Flights	γ - Effect of Remaining Capacity on Price (standard error)	R-square
London (Stansted) – Edinburgh	4	-0.316 (0.026)	75.4%
East Midlands – Edinburgh	4	-0.155 (0.011)	74.6%
London (Stansted) – Rome (Ciampino)	4	-0.128 (0.023)	85.8%
East Midlands – Barcelona	4	-0.25 (0.029)	87.7%
Liverpool – Alicante	4	-0.137 (0.098)	65.2%
London (Luton) – Malaga	3	-0.145 (0.041)	49.6%

Based on the previous result, our hypothesis is that the parameters γ_j are negative. The results of the regression are given in Table 3. We see from Table 3 that, for price versus remaining capacity, all of the slope coefficients (γ_j) are negative with an overall mean of

-0.189. In addition, the fit of the models as measured by the R-Squared values suggests that this relationship between capacity and price is well captured.

To summarize, our main descriptive empirical findings are that ticket prices increase over time and that the rate of increase varies negatively with remaining seat capacity. Also, we find clear evidence for at least two segments of consumers that vary according to their arrival times. In addition, we observe that easyJet does not resort to any last-minute deals to clear capacity. The interesting question we address next is whether and under what conditions these pricing policies are optimal.

4. Model Development

We consider a one-way airline route between two cities with a monopoly service provider. We assume two segments of customers: Higher valuation and lower valuation consumers. We define higher valuation consumers as business travelers (denoted by B) and lower valuation consumers as tourist travelers (denoted by T) although as Talluri and van Ryzin (2004, p. 517) note, in practice the distinction between business and leisure travels is not so clear cut. The consumers can arrive in two, three, or four time periods depending on whether the firm offers last-minute deals, on the duration of the sales period, and on consumer characteristics (as specified later). The tourists' utility from the travel is uniformly distributed over an interval of $(0, \alpha)$. The business travelers' utility from the travel is distributed uniformly over an interval of $(\alpha, \bar{\alpha})$.

Both tourist and business consumers arrive in the first two periods. Tourist consumers who did not purchase tickets in period 1 or 2 also arrive during the third period. Since our main interest is in the proportion of business and tourist travelers, we assume that all tourists

and a fraction γ of the business segment arrive during the first period. The remainder $(1 - \gamma)$ of the business segment appears in the second period (alternatively, we could have made the dual assumption that business consumers arrive in both periods and only a fraction of the tourists arrive in the first period). The business consumers who arrive during the second period have uncertainty with respect to their utility. With probability θ , the business segment learns during period 2 that business meetings that require air travel will be held at the destination city and its utility is distributed over the interval of $(\alpha, \bar{\alpha})$. With probability $(1 - \theta)$, these business meetings are not held and the utility from the air travel thus equals zero. To create a clear-cut segmentation between business travelers and tourists, we assume that the upper bound of the valuation of the business traveler is much higher than that of the tourist; in other words, $\alpha < \bar{\alpha}/2$.

There are two events in period 1: first the airline announces the price and then consumers decide whether to buy. There are three events in period 2: the airline announces the price, uncertainty about the state of the business passengers is resolved, and, finally, consumers decide whether to buy tickets. In period 3 (if it exists), the firm may announce a price and tourist consumers who have a higher valuation than this price purchase tickets. Let $f_i(x)$ ($i = B, T$) be the density of consumer distribution. Because the tourist customers are distributed uniformly in the interval $(0, \alpha)$, if the price in period 1, 2, or 3 is p , then the tourists whose utility is greater than the price will buy tickets. Thus, the proportion of tourists who buy seats

at price p is represented by $\int_p^{\alpha} f_T(x)dx = (\alpha - p)/\alpha$. Similarly, the proportion of business

passengers buying tickets at price p is represented by $\int_p^{\bar{\alpha}} f_B(x)dx = (\bar{\alpha} - p)/(\bar{\alpha} - \alpha)$.

When we define the number of tourist passengers as M_T and the number of business passengers as M_B , their respective demands at price p are given by $M_T (\alpha - p) / \alpha$ and $M_B (\bar{\alpha} - p) / (\bar{\alpha} - \alpha)$. To simplify notations, we normalize the market sizes as follows: $N_T = M_T / \alpha$ and $N_B = M_B / (\bar{\alpha} - \alpha)$. An important part of our model is that consumers are forward looking, and therefore will decide to purchase a ticket in period i ($i = 1, 2, 3$) if the utility in the period will be positive and higher than the utility of purchasing in other periods. Therefore, the following equations represent demand in the three periods:

$$\begin{aligned}
 \text{period 1} \quad D_1 &= \begin{cases} N_T L_{T1} (\alpha - p_1) + \gamma M_B & \text{if } p_1 < \alpha \\ \gamma N_B L_{B1} (\bar{\alpha} - p_1) & \text{if } \bar{\alpha} > p_1 > \alpha \end{cases} \\
 \text{period 2} \quad D_2 &= \begin{cases} N_T L_{T2} (\alpha - p_2) + (1 - \gamma) M_B & \text{if } p_2 < \alpha \\ (1 - \gamma) N_B L_{B2} \theta (\bar{\alpha} - p_2) & \text{if } \bar{\alpha} > p_2 > \alpha \end{cases} \\
 \text{period 3} \quad D_3 &= \begin{cases} N_T L_{T3} (\alpha - p_3) & \text{if } p_3 < \alpha \\ 0 & \text{if } p_3 > \alpha \end{cases}
 \end{aligned}$$

L_{ji} is an identity function that equals one if and only if segment j ($j = T, B$) purchases the product in period i ; it is zero otherwise. Without loss of generality, let the marginal cost of supplying the seat be zero. Let the capacity (the number of airline seats) be fixed at C . Restrictions on capacity play a dominant role in our analysis, as will be explained later.

5. Analysis

We start this section with analysis of the simpler case in which tourist consumers do not arrive during the third period. Analysis of this case will provide the necessary intuition for analysis of the more complicated cases. Later, in section 5.2, we analyze the extended model

in which there are three periods and forward-looking tourist consumers, who, when making their first-period purchasing decisions, take into account the option of waiting until the third period and purchasing tickets at a reduced price if such tickets are available.

5.1. Two-period game

In this case, we consider two scenarios that relate to the proportion of business and tourist consumers: if the business segment that arrives during the first period is large enough

($\gamma N_B > N_T$), then, regardless of capacity, the airline charges a high price and sells only to

business consumers in both periods and completely ignores the tourist segment. Optimal prices

are given by $p_1 = p_2 = \bar{\alpha} - \frac{C}{N_B(\theta + \gamma(1-\theta))}$ for $C < C_1$ and $p_1 = p_2 = \bar{\alpha}/2$ for $C_1 \leq C$ where

$C_1 = N_B \bar{\alpha} (\theta + \gamma(1-\theta))/2$ is the optimal quantity to sell to the business segment. The proof of this case is the same as the proof of Proposition 1, *mutatis mutandis*. All proofs are given in

the appendix, which is available to download at easyjetpricing.homestead.com.

If the tourist segment is large enough ($N_T > \gamma N_B$), however, the airline should consider three cases. In the first case, capacity is so low that the airline sells only to business

consumers. In the second case, capacity is binding but high enough that the airline can

effectively discriminate so it sells to both markets. In the third case, capacity is high enough

and is not effectively binding. To set the boundaries of these cases, we define $S(C)$ as

$S(C) = \alpha N_T / 2 + N_B [\gamma(\bar{\alpha} - \alpha) + (1 - \gamma)\theta\bar{\alpha}] / 2 - C$. We later show that $S(C)$ is the capacity

shortage or the difference between demand (at unconstrained prices) and capacity.

Proposition 1: *With low capacity, $C < C_2$, the firm sells only to the business segment. With intermediate capacity, $C_2 < C < C_3$, the optimal pricing scheme is to increase the price over time so as to discriminate between the two segments while restricting demand for **both** segments. With excess capacity, $C > C_3$, the*

airline price discriminates between the two segments while adjusting the first-period price to take into account early-arriving business consumers.

In the following, we expand on the intuition behind this result.

Prices and profits with low capacity: When capacity is low, the airline is better off selling exclusively to business consumers, who have higher valuations and thus will pay more. This case can be divided into two subsections: low capacity and very low capacity. Consider the extreme case in which the airline has only one seat available. Obviously, the airline would rather sell it to a business consumer who has the higher valuation. As we increase capacity, this policy remains valid until capacity exceeds the optimal quantity to sell to the business segment C_1 . To find the value of C_2 , note that, when capacity equals C_1 , the airline charges the optimal monopoly price $p = \bar{\alpha}/2$ and sells only to the business segment. When capacity increases further, the airline does not immediately decrease the price in period 1 to capture more demand from the tourist segment and compensate by increasing price in period 2 to the business segment. The airline employs this policy only when *the additional revenue from the tourist segment compensates for the loss in revenue from the business segment*. Up to this capacity (defined as C_2), the airline keeps the price constant at $p = \bar{\alpha}/2$ in both periods and serves only the business segment. We derive the value of C_2 in the appendix.

Prices and profits with intermediate capacity: Recall that, without capacity constraints, if an airline could perfectly discriminate between the two segments, it would charge monopoly prices $\alpha/2$ to the tourist segment and $\bar{\alpha}/2$ to the business segment. One might have expected that when capacity increases above C_2 the airline would continue to sell to business consumers at a price of $p = \bar{\alpha}/2$ and start to sell to tourist consumers at a reduced price as fillers to increase utilization of the airplane. However, this approach is problematic; if

the airline reduces the price to attract tourists, seats are “unfortunately” sold at the reduced price to business consumers who arrive during the first period. Therefore, the optimal behavior is to decrease the price in period 1 below $\alpha < \bar{\alpha}/2$ and increase the price in period 2 above $\bar{\alpha}/2$. In this section, we show that the optimal pricing scheme is to increase the price over time so as to discriminate between the two segments when the capacity is intermediate in value, $C_2 < C < C_3$, where C_3 is given by $S(C_3) = 0$ (i.e., for any capacity for which $C > C_3$, there is no shortage).

Prices and profits with high capacity: Finally, when the capacity constraint is not binding, the airline discriminates between the two segments by charging monopoly prices. With these prices, it is straightforward to compute the overall demand D :

$D = \alpha N_T / 2 + N_B [\gamma(\bar{\alpha} - \alpha) + (1 - \gamma)\theta\bar{\alpha}] / 2$. It is now clear that $S(C)$ is indeed the difference between demand and capacity ($D - C$), or, the capacity shortage.

In the high-capacity case, the airline practices price discrimination between the two segments. In the intermediate case, the airline adjusts the level of prices for both segments. That is, it does not serve the business consumer first and use the tourists as a buffer in case it has some excess capacity. Rather, it restricts the demand for both segments (by raising appropriate prices) so as to equate capacity to expected demand. Only in the low-capacity case does the airline forgo the tourist segment and serve the business segment exclusively. We also note that, when capacity is high, prices in the two periods are independent of capacity and thus the *difference* between the prices in the two periods is independent of capacity as well. When capacity is low or medium, prices do depend on capacity.

The airline does not sell all seats in two different scenarios. First, as one expects, the airline does not sell all seats when capacity is very high ($C > C_3$). More interestingly, when

capacity is relatively low ($C_1 < C < C_2$), the airline's optimal policy is to keep some seats unsold. The reason is that, when capacity equals C_1 , the airline charges the optimal monopoly price and sells only to the business segment. If capacity increases, the airline does not instantly decrease the price in period 1 to capture more demand from the tourist segment and compensate by increasing the price to the business segment. Rather, it employs this policy only when the capacity is large enough that the additional revenue from the tourist segment compensates for the loss in revenue from the business segment. Thus we have the following result:

Result 1a: *For a given market size, there are two capacity ranges, $C_1 < C < C_2$ and $C > C_3$, that the airline should never choose for its operating capacity. In these capacity zones, as capacity increases, neither the price nor the demand change and thus the airline pays more for its capacity while its revenue remains unchanged.*

5.2. Analysis of the three-period game: Last-minute deals

We observe that, in practice, last-minute deals are occasionally offered, often at very low prices. If the airline decides to engage in such offers, either directly or via a reseller, it can set a new price that will attract the lower end of the tourist segment that did not purchase tickets in period 1. Last-minute deals are often made very close to the actual flight time. For example, in some European airports, one can buy tickets at greatly reduced prices for same-day flights. Thus, in actual practice as well as in our models, last-minute deals are rendered irrelevant for the business segment. If the price in period 3 (the last-minute period) is low, then the airline has to worry about consumers from the tourist segment waiting to buy tickets in period 3 instead of buying them in period 1. Indeed, some will; the question is how to fence the higher-utility consumers out of this segment. High-utility tourist consumers do not wait for last-minute deals because of uncertainty with respect to the existence of such deals. There are two

sources for this uncertainty. First, consumers are uncertain with respect to the airline's policy as the airline might randomize with respect to offering last-minute deals. The second source is uncertainty with respect to the actual arrival of business consumers in period 2. We begin by analyzing the case of consumers' uncertainty with respect to the airline's policy and continue with the case where the source of uncertainty is arrival of the business segment.

5.2.1. Modeling uncertainty with respect to the airline's strategy

We model the uncertainty of consumers regarding airline strategy with the help of an additional parameter, β . Consumers' expectations are such that with probability β the airline will offer last-minute deals and with probability $(1 - \beta)$ the airline will not. We start with a single-shot, three-period game and continue with a repeated game in which each round is composed of a three-period game. To separate the two uncertainties (firm strategy and consumer arrival), in this section we treat θ as the proportion of business travelers who arrive during the second period. In the next section (5.2.2), we treat θ as consumers' arrival uncertainty.

When $C'_3 < C$, the airline has to consider two cases. First, the airline should consider the intermediate case $C'_3 < C < C_4$ in which the airline is constrained during period 3 and thus fills all seats with last-minute consumers. The second is the large-capacity case in which capacity exceeds C_4 and the airline is not constrained, in which case it sells to only some of the remaining tourist consumers. The following proposition summarizes our findings with respect to last-minute deals. The solution of the last-minute model and proofs of proposition 2 is given in the appendix:

Proposition 2: *As long as consumers are not certain that the airline will offer last-minute deals ($\beta < 1$), the airline should always offer a last-minute deal. As β increases, the airline has less incentive to offer last-minute deals. If $\beta = 1$, the airline does not offer a last-minute deal.*

With asymmetric information, the optimal policy is to offer last-minute deals. This policy is optimal because the airline succeeds in discriminating between two classes within the tourist segment based on valuation. This discrimination is possible, however, **only** with the existence of consumer uncertainty. Consumers with higher valuation will not care to wait for the last-minute deal as this causes them uncertainty with respect to flight availability. The tourist consumers with lower valuation will indeed wait, hoping to buy at the last minute if tickets are still available and knowing that the flight could sell out.

If the probability β that the airline will offer a last-minute deal is very low, then consumers behave myopically in that they hardly take into account the possibility of a price reduction in the third period. That obviously increases the attractiveness of such a deal for the airline. However, as the probability β increases and consumers expect that the airline may offer a last-minute deal, the price reduction in the third period diminishes and therefore the additional profit from employing a deal decreases and is actually equal to zero for $\beta = 1$.

Interestingly, first-period prices with the last-minute deal can be greater or less than the constrained first-period price in the two-period model. In the large-capacity case, our intuition about the direction of prices holds and the airline charges a first-period price that is higher than the corresponding price from the two-period model. Under intermediate capacity, price depends on consumers' uncertainty, β . The following example provides additional intuition into this result. Consider the following values: $M_B = 50$, $M_T = 150$, $\gamma = 0.4$, $\alpha = 1$, $\bar{\alpha} = 2.2$,

$C = 98$, and $\theta = 0.3$. We then analyze both the two-period setting and the last-minute model for two different values of consumer uncertainty, $\beta = 0.55$ and $\beta = 0.9$.

As shown in Tables 4a and 4b, the airline serves twenty-eight business consumers; twenty arrive in the first period and eight arrive in the second. In Table 4a, the difference between the two-period model and the last-minute model is the treatment of the tourist segment. In the latter game, the airline raises the first-period price by 17% and lowers the third-period price by 5%. This 22% difference yields more tourists being served (seventy compared with sixty-five) at a higher average price (0.58) and greater profits relative to the case in which no last-minute deal is offered. As shown in Table 4b, the airline succeeds in discriminating between the tourist segments but does so at prices that are lower than the optimal price of the tourist segment in the two-period model.

Table 4a: Low Consumer Uncertainty, $\beta = 0.55$

	P_1	D_1	P_2	D_2	P_3	D_3	Profits
Two Periods	0.57	20B, 65T	1.1	8B	–	–	57
Last-Minute Deal	0.67	20B, 25T	1.1	8B	0.54	45T	63

Table 4b: High Consumer Uncertainty, $\beta = 0.9$

	P_1	D_1	P_2	D_2	P_3	D_3	Profits
Two Periods	0.567	20B, 65T	1.1	8B	–	–	57
Last-Minute Deal	0.565	20B, 25T	1.1	8B	0.53	45T	58

Note that in all cases prices to tourists are higher than they would be in the absence of the business segment. In this example, the latter price would have been 0.5 while the *lowest* price charged in the last-minute deal is 0.52. This leads us to the next proposition:

Result 2a: *In both the two-period and the last-minute game, tourist consumers subsidize the business segment. In the absence of business consumers, the tourist pays a lower price and in the absence of tourists the business consumer pays a higher price.*

If the airline could plan for the arrival of each segment, it would be better off if tourist and business consumers arrived separately. In such a case, the airline would be able to perfectly discriminate between these two segments. Unfortunately for the airline and the tourists, this is not the case. Business and tourist consumers arrive together and therefore the airline raises the price above what it would have charged in the case of separate arrival. Note that, for a fixed capacity, Result 2a confirms Dana's (1999) result that tourist travelers subsidize business consumers. However, this is not the case when capacity changes; when $C_2 < C < C_3$, an increase in capacity causes the number of tourists to increase and the price to business consumers to decrease. On the other hand, when capacity is at point C_2 , the airline moves from selling only to business consumers at a relatively low price of $p = \bar{\alpha}/2$ to practicing price discrimination between the two segments by selling to business consumers at a much higher rate and to tourists at a much lower rate. Thus, at this capacity, business consumers actually suffer from having the tourists join the flight because business consumers pay a higher price. Note also that in general tourists subsidize the business segment in terms of flight frequencies and possible destinations. Absent tourists, the number of flight would be reduced.

One can raise a question about the way the consumer learns about the probability that the airline will offer a last-minute deal and whether this uncertainty will be resolved over time. Even though our game involves three periods, it is played only once. For learning to occur, the game should be repeated. In such a case consumers would have an opportunity to learn about the firm's policy and thus it might not be optimal for the firm to employ last-minute deals. In that case, we can model a repeated game in which consumers in a steady state know the probability of the firm employing a last-minute deal (β) and the firm optimally chooses β .

Note that, even though consumers can learn the probability of a last-minute deal being offered, they still cannot know whether the firm is going to offer a last-minute deal for a specific flight. The equilibrium in such a game depends on capacity. Under low capacity, the firm is indeed better off not offering last-minute deals. Under large capacity, however, the firm optimally will include last-minute deals as part of the equilibrium.

5.2.2. Modeling uncertainty with respect to the business segment's arrival

During the first two periods, consumers maximize their utility and the airline maximizes profits based on its expectations regarding the arrival of business consumers. However, when the airline and consumers make third-period decisions, they already know how many business consumers have arrived and therefore how many seats are left for sale during this period. Note that we make two simplifying assumptions to keep the analysis tractable. First, we assume that either all business consumers arrive with probability θ or none of them arrive with probability $(1 - \theta)$. Second, we assume that no business consumers arrive during the first period; in other words, $\gamma = 0^1$. The following proposition summarizes our results (for the proof, see Appendix):

Proposition 3: *For intermediate capacity range ($C_{12} < C < C_{23}$) the equilibrium is such that the airline sells to tourist consumers during the first period at a relatively high price and decreases the price in the third period. In this case, if the business demand materializes in the second period, there are no seats left for the third period; if the business demand does not materialize, the airline is capacity-constrained in the third period and thus partially price-discriminates within the tourist segment.*

For low or high capacity range ($C < C_{12}$ or $C > C_{23}$) the airline cannot price discriminate within the tourist segment by offering last minute deal.

¹ We also modeled the case in which each business consumer faces this uncertainty individually but only a numerical solution can be achieved. Assuming that there is no heterogeneity among business consumers regarding the uncertainty allows us to capture the main trade-off and keep the solution tractable.

Consider the case in which capacity is very high; in that instance, strategic consumers who know the capacity of the flight realize that the airline will always have seats for sale in the third period and therefore will choose to wait and purchase tickets during the third period. This consumer behavior forces the airline to charge unconstrained profit-maximizing prices during periods 2 and 3 and to sell none during period 1; in other words, the airline would not benefit from price discrimination within the tourist segment. On the other hand, consider the case where the initial capacity is very low; in this case, the airline again is better off not selling tickets during the first period at all and then offering a second-period price that is set to sell all available seats to some of the business consumers if they arrive. If they do not arrive, the airline then has an opportunity to sell to tourist consumers during the third period. Just like in the previous case, the airline cannot price-discriminate within the tourist segment in this case.

In contrast to the preceding cases, the airline employs price discrimination between the two segments, as well as within the tourist segment, in the case of intermediate capacity. In this case, the optimal policy is to sell to some tourists during period 1 and then to sell to business consumers during period 2 if they arrive. The third-period policy in such cases depends on the arrival of business consumers.

The intriguing conclusion from this proposition is that there are only limited cases in which the airline will price-discriminate within the tourist segment: *Ex ante*, the firm tries to price-discriminate only if capacity is bounded. *Ex post*, the firm will succeed in price-discriminating within the tourist segment only if business consumers do not arrive.

So, for example, we might have expected that the airline would reserve some seats for the third period under larger capacity so it could discriminate within the tourist segment. In the optimal strategy, however, this is not the case. As soon as capacity increases, the airline should sell entirely to the business segment in the second period and to the tourist segment in the third

without price-discriminating within the tourist segment. Note the centrality of the assumption of strategically playing consumers. It is easy to ascertain that the firm can always employ full price discrimination if consumers are not looking forward and playing strategically.

5.3. The effects of duration: When to release ticket for sale

Looking at easyJet's flights, it is obvious that they release tickets for sale at different periods before actual flight times. In this section, we show the conditions under which the duration decision is a strategic variable and analyze the effect of duration on the ability of the firm to price-differentiate among its customers.

Consider the case where consumers arrive earlier. Should the firm introduce tickets earlier? The most likely segment to purchase the ticket earlier—if the firm does release tickets earlier—is a price-sensitive segment. However, if their valuations come from the same distribution as the rest, then there is no advantage in selling tickets earlier. Obviously if capacity is tight and these consumers' valuations are below other consumers' valuations, the firm should not start selling tickets earlier. If, on the other hand, capacity is very large, the firm again should not start selling tickets earlier as the firm's optimal strategy is to offer monopoly prices. Thus, the only case when it matters is when capacity is at an intermediate level. Recall that in our model time is a discrete rather than a continuous parameter and thus the duration decision is simply adding a period before what we earlier called period 1. We define this period as period 0. Next, consider the case where a fraction $(1-\delta)$ of the M_T consumers arrives at period 0 and has a valuation for the flight that is drawn from a uniform distribution $(0, \underline{\alpha})$ where $\underline{\alpha} < \alpha/2$. We call this segment Low-Tourist. The rest of the tourist segment (δM_T) arrives in period 1 with a valuation for the flight that is drawn from a uniform

distribution $(\underline{\alpha}, \alpha)$. We call this segment High-Tourist. The following proposition summarizes what happens in this case.

Proposition 4: *For an intermediate capacity range, $C_L < C < C_U$, the equilibrium is such that the airline sells to Low-Tourist consumers during the period 0 and increases prices and sell to High-Tourist consumers and business consumers in periods 1 and 2 respectively. If the business demand in the second period materializes, there are no seats left for the third period; if the business demand does not materialize, the airline is capacity-constrained in the third period and thus partially price-discriminates within the Low-Tourist segment.*

This proposition and the discussion preceding it lead to the following result:

Result 4a: *For an intermediate capacity level, the larger the number of segments (that differ in their price sensitivity), the longer the duration of the period in which tickets are offered for sale.*

The airline would choose to sell earlier only in cases in which selling earlier enables the firm to better discriminate price inter-temporally. However, note that period 0 does not replace last-minute deals. It only allows the airline to refine its price discrimination within the lowest-valuation segment.

6. Conclusions and extensions

To summarize, our paper has shed some insight into the pricing strategy of low-cost short-haul airlines. In our model, the airline faces two or three segments of consumers that gain different utilities from the flight. We find empirically that the assumption of the existence of two or three segments is consistent with the data, as is the point at which price in each period is a function of the remaining capacity at that point in time. We examined the phenomenon of last-minute discounting by introducing into the analysis a third period in which the airline has unsold seats. For an intermediate capacity level, uncertainty with respect to the arrival of the business segment will cause the firm to offer last-minute deals and thus

partially price-discriminate within the tourist segment. The same is true with uncertainty with respect to the actual behavior of the firm: if consumers are uncertain whether the firm will offer last-minute deals, then, in equilibrium, both in a one-shot game and in a repeated game, the firm will, with some probability, offer such deals. In addition, for an intermediate capacity level, we found that the larger the number of segments (that differ in price sensitivity), the longer the duration of the period in which tickets are offered for sale.

We have modeled consumers as risk-neutral to capture the main point we wanted to make. In deliberating the effect of adding risk aversion, consider for example Proposition 3 in the intermediate-capacity case. In solving this case (in the appendix), we define x to be the tourist with the highest utility who will purchase a ticket in period 3 at price p_3 . Thus this consumer is indifferent between purchasing a ticket in the first period at price p_1 and purchasing a ticket at the last minute at price p_3 . The firm then sets the price in the third period as $x/2$ (if it is not constrained, it sets the monopoly price for this residual segment). With risk aversion, the consumer who is indifferent between purchasing a ticket in the first period at price p_1 and purchasing a ticket at the last minute at price p_3 will have a lower utility x' (equivalently, the same marginal consumer x will demand a lower third-period price to remain indifferent). Since $p_3 = x/2$, this will lower the price in the third period. Thus risk aversion will require deeper cuts for the last-minute deals.

The effects of competitive entry into a city-pair market could affect our results. The way to have a feel for the effect is to note that direct competition by another low-cost carrier in the same city-pair route would reduce demand for the incumbent airline's seats for these routes, or, alternatively, would create excess capacity. Our scenarios for large capacity would not change as, obviously, the airline had excess capacity before the competitor's entry. If,

however, capacity is constrained, then consider Proposition 3. It can be easily verified (see the appendix for details) that the first-period price will increase but second- and third-period prices will decrease with the implied change in capacity. Thus the response of the airline to entry is to lower the slope of the price curve over time between the first and second period; that is, the price discount for early buying is reduced, but the last-minute deal is more pronounced. The reason is that the airline responds to entry by trying to attract the more profitable business segment by lowering its price and limiting the demand of the tourist segment by increasing the price in the first period. If business demand does not materialize, the airline cuts its price deeply to attract tourists with a last-minute deal. This adds to our claim that the tourist segment subsidizes the business segment as it can also be verified that the average price to the tourist segment (in periods 1 and 3) is higher in the post-entry market structure. Thus, with added competitive pressure, the airline responds by lowering the price to the business segment and raising the average price paid by tourists.

References

- Biyalogorsky, E., Gerstner, E., Weiss, D., and Xie, J. (2005). The economics of service upgrades. *Journal of Service Research*, 7(3), pp. 234–244
- Biyalogorsky, E., Carmon, Z., Fruchter, G., and Gerstner, E. (1999). Overselling and opportunistic cancellations. *Marketing Science*, (18), pp. 605–610.
- Borenstein, S., and Rose, L.R. (1994). Competition and price dispersion in the U.S. airline industry. *Journal of Political Economy*, (102), pp. 653–683.
- Carpenter, G.S., and Hanssens, D.M. (1994). Market expansion, cannibalization, and international airline pricing strategy. *International Journal of Forecasting*, (10), pp. 313–326.
- Chintagunta, P.K., Jain, D.C., and Vilcassim, N.J. (1991). Investigating heterogeneity in brand preferences in logit models for panel data. *Journal of Marketing Research*, (28), pp. 417–428.
- Dana, J.D. (1999). Using yield management to shift demand when the peak time is unknown. *Rand Journal of Economics*, pp. 456–474.
- Dana, J.D. (1998). Advance-purchase discounts and price discrimination in competitive markets. *Journal of Political Economy*, (106), pp. 395–422.
- Desai, P., Koenigsberg, O., and Purohit, D. (2004). Strategic decentralization and channel coordination. *Quantitative Marketing and Economics*, 2(1), 5–22.
- Done, K. (2004). Ryanair’s dream run comes to an end. *Financial Times*, January 29.
- Elmaghraby, W., and Keskinock, P. (2003). Dynamic pricing in the presence of inventory considerations; Research overview, current practices, and future directions. *Management Science*, 49 (10), pp. 1287-1309.
- Feng, Y., and Gallego, G. (1995). Optimal starting times for end-of-season sales and optimal stopping times for promotional fares. *Management Science*, (41), pp. 1371–1391.
- Gallego, G., and van Ryzin, G.J. (1997). A multi-product dynamic pricing problem and its applications to network yield management. *Operations Research*, (45), pp. 24–41.
- Gallego, G., and van Ryzin, G.J. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, (40), pp. 999–1020.
- Jerath, K., S. Netessine, and S. Veeraraghavan. (2007). Revenue management with strategic customers: Last minute selling and opaque selling. Working paper, Wharton School of Business.
- Kretsch, S.S. (1995). Airline fare management and policy. in Jenkins, D. (ed.), *The Handbook of Airline Economics*. Columbus, OH: McGraw-Hill, pp. 477–482.
- Liu Q. and G. Van Ryzin. (2006). Strategic capacity to induce early purchases. Working paper, Columbia Business School.

- McGill, J.I., and van Ryzin, G.J. (1999). Revenue management: Research overview and prospects. *Transportation Science*, (33), pp. 233–256.
- Morrison, S.A., and Winston, C. (1990). The dynamics of airline pricing and competition. *American Economic Review*, (80), pp. 389–393.
- Ovchinnikov A. and J. M. Milner (2007). Revenue management with end-of-period discounts in the presence of customer learning. Working paper, Rotman School of Business.
- Shulman J. and A Coughlan. (2007). Used goods, not used bads: Profitable secondary market sales for a durable goods channel. *Quantitative Marketing and Economics*, 5(2), 191-210.
- Talluri, K.T., and van Ryzin, G.J. (2004). *The Theory and Practice of Revenue Management*. Norwell, MA: Kluwer Academic Publishers.
- Watherford, L.R., and Pfeifer, P.E. (1994). The economic value of using advance booking of orders. *Omega*, (22), pp. 105–111.
- Wedel, M., Desabro, W.S., Bult, J.R., and Ramaswamy, V. (1993). A latent class Poisson regression model for heterogeneous count data. *Journal of Applied Econometrics*, (8), 397–411.
- You, P.S. (1999). Dynamic pricing in airline seat management for flights with multiple legs. *Transportation Science*, (33), pp. 192–206.