

On-line Booking and Revenue Management: Evidence from a Low-Cost Airline

MARCO ALDERIGHI

Università della Valle d'Aosta, Italy.

CLAUDIO A. PIGA*

Loughborough University UK, RCEA Rimini

Using unique data on a low-cost airline posted prices and seat availability, this study sheds some light on whether the airline's actual practice of yield management techniques conforms with some predictions from economic models of peak-load pricing under demand uncertainty. On the one hand, robust support is found to the notion that prices increase as the seat availability decreases; on the other, theoretical models that do not account for stochastic peak-load pricing fail to capture an important source of dispersion in the data.

Keywords: Inter-temporal pricing, competition, price dispersion.

JEL Classifications: D43, L10, L93

1 Introduction

In Europe, the liberalisation process of the airline industry started in 1987 and developed gradually, granting progressively more rights to European carriers to operate within the European market, until 1997 when permission was granted to European carriers to operate domestic flights in member countries other than their home market. In 2004, a last legislative package was issued by the Commission with the aim to create a Single Pan-European Sky by integrating the air management structures of the member countries.¹

The European liberalisation facilitated the entry of many Low-Cost Carriers (hence, LCC), whose business model was somehow adapted from the one pioneered in the U.S. by Southwest Airline. In particular, two airlines, Ryanair and EasyJet, proved to be so successful that by 2004

*Piga gratefully acknowledges the financial support from the Leverhulme Research Fellowship Grant Award, reference RF/7/RFG/2007/0022. Thanks to Li Gan, Massimo Filippini, Rico Maggi, Roberto Patuelli, Steve Puller, Steve Wiggins, and Manuel Hernandez for useful comments and feedback. The authors bear exclusive responsibility for any remaining error. Corresponding author: c.a.g.piga@lboro.ac.uk

© 2010 Claudio A. Piga. Licenced under the Creative Commons Attribution-Noncommercial 3.0 Licence (<http://creativecommons.org/licenses/by-nc/3.0/>). Available at <http://rofea.org>.

¹Regulations (EC) No 549/2004, 550/2004, 551/2004 and 552/2004.

Ryanair's capitalization on the London Stock Exchange overtook that of a leading Full Service Carrier (FSC), British Airways.

A central element in the business model followed by most Low-Cost Carriers is represented by their almost total reliance on the Internet as a distribution channel. In this article data on fares and available seats were retrieved from the Ryanair' web site in order to shed some light on some largely unexplored aspect of this airline's yield management system, i.e., the strategy the airline follows to set on-line fares under varying conditions of demand uncertainty. A notable innovation in this study is the possibility to combine fares with the number of seats available at the time the fare was retrieved.

The simple business model pursued by the airline enable us to test some of the theoretical predictions derived in Dana (1999). An important characteristic of Dana's model is that the prices are set before the actual realisation of demand is known; in practice, based on a set of probabilities for each possible state of nature, the firm has to decide the level of prices and the associated number of seats it will sell for each possible realization of demand. For instance, for the case of two-states demand (low and high), the firm will set two prices and the corresponding number of tickets available at each price. The analysis consider the case in which the firm operates in a perfectly competitive market, and the case where the firm is a monopoly.

Dana shows that regardless of the market structure, the firm should determine different "batches" of seats, and that fares should increase as fewer batches remain unsold. That is, the profile of fare should be an increasing function of the number of sold seats. We test this hypotheses by estimating a pricing profile linking a flight's seat occupancy with offered fares, and find a positive relationship where, on average, an extra sold seat induces an increase of 3-4% in posted fares.

Another important feature of Dana's model is the commitment of the firm to the schedule of prices it sets before demand is known. That is, once the price and size of each batch of seats is decided, the airline will strictly adhere to it and will not modify it even if it wanted to do so. Such a price rigidity may arise from advertisement or promotions, or because the firm must incur a very high cost in tracking the evolution of demand for all the flights it operates and adjust fares to reflect demand conditions. Price commitments enforces a price system in which customers who either arrive early or arrive late when demand is low can buy at lower prices, while customers who arrive late when demand is high pay a higher price. An individual consumer at the time of purchase may only observe the price of each firm's least expensive remaining unit, but *ex-post* consumers will have paid different prices solely because of the random order in which they were served, which reflect a different level of capacity utilization. That is, the model does not assume that individuals with a higher willingness to pay arrive at a late stage. Therefore, the model assume inter-temporal price discrimination away.

In Dana (1999), "stochastic peak-load" pricing, i.e., the adjustment of fares due to an update on the airline's information on demand conditions, is ruled out by the assumption that it is

too costly for the airline to keep track of the evolution of demand on its flights. That is, the commitment assumed by Dana (1999) is incompatible with possible fare updates the airline may implement if at some point in time, say 2-3 weeks before take-off, it finds out that a flight is selling better than expected. This implies that in practice we should observe no consistent departure from the pricing profile when, at specific points in time, the airline may observe that only a limited number of seats remain to be sold.

Our results reveal that fare updating seems to take place consistently, leading us to conclude that “stochastic peak-load” pricing is an integral part of the pricing strategy pursued by one of the largest European airline, in contrast with the modelling assumption made in Dana (1999). This is an important contribution of the paper, since the evidence on the impact of stochastic peak-load pricing on fares is scant given the difficulty to obtain relevant data on seat occupancy.²

To sum up, the findings in Dana (1999) lead to the formulation of two hypotheses, that we put to a test in this paper. First, the relationship between fares and inventories is on average (non-strictly) monotonically increasing. That is, we should expect that posted fares increase as the number of available seats falls. Second, in Dana (1999) airlines irrevocably commit themselves to distributions of prices and seat “buckets” for each flight before learning demand. Commitment implies that firms will not change or update their pricing decisions on the basis of actual bookings, i.e., on how well the flight is selling. Since firms commit to a monotonically increasing pricing profile, fare reductions between two consecutive booking days when available seats remain stable or fall, can be interpreted as evidence of updating taking place. Updating should be more likely as the time of departure approaches and as the number of unsold seats increases.

The analysis also allows to gain new insights into the inter-temporal profile of posted fares, a topic that has been largely explored in the existing literature (Pitfield, 2005; Mantin and Koo, 2009). While previous studies showed a monotonic increase of fares as the date of departure approaches, our estimates suggest a *U*-shaped relationship which is consistent with a pricing strategy where the airlines try to adjust fares over time to meet the demand of different segment of buyers, that are heterogenous in their travel motivations.

In the remainder of the paper, a brief introduction to the literature on yield management is offered in Section 2, followed by the data presentation, the econometric approach and the results.

2 Literature Review

Setting airfares and allocating aircraft seats is a complex process. Airlines have to deal with demand fluctuations, consumer heterogeneity, and the uncertainty about when and where passengers want to travel. In addition, aircraft capacity is limited and the nature of the product

²See Puller et al. (2009) for a study of a related problem in the case of fares offered by Traditional, Full Service Carriers.

perishable, as unsold seats cannot be offered once the flight has departed (Alderighi et al., 2004).

To deal with these challenges, airlines have developed a set of techniques known as yield or revenue management (Weatherford and Bodily, 1992). Alderighi et al. (2004) distinguish between traditional and simplified yield management. The former is the one developed and implemented by the FSC to cope with the new competitive environment that followed the liberalization process. The latter defines the set of techniques implemented by the LCC. In both cases, a central issue is the need to define and price certain product characteristics in order to accommodate passengers' heterogeneity and different willingness to pay. Traditional companies, aware of travellers' different preferences, have tried to meet such heterogeneity by offering a differentiated product with a large variety of in-flight and ground services. Different airfares based on the different levels of service quality are therefore offered for the same flight (Puller et al., 2009). In addition, to ensure that each segment of travellers acquires its required level of service, companies apply "fences" such as minimum stays at the travel destination, penalties for ticket cancellation or travel date change, or purchase time limits. FSC offer such differentiated products through reservation classes that reflect the market segmentation. To each fare class a certain number of seats must be allocated in order to optimally accommodate the total demand (Puller et al., 2009). This crucial forecasting activity is known as inventory control, and it is applied to all flights operated by each airline in its own network. In particular, purchase time limit is a "fence" that has gained more and more importance within the yield management associated with the pricing by LCC.

The conventional wisdom holds that carriers tend to attach monotonically increasing airfares to sequential booking classes in order to cope with the uncertainty over demand (Dana, 1999). McGill and Van Ryzin (1999) refer to the latter practice as "low-before-high fares" and explain that it is due to the assumption that booking requests arrive in strict fare sequence, from the lowest to the highest as the date of departure nears. Many scholars have devoted their attention to the existence of such airfare dynamics both from a theoretical (Belobaba, 1987; Gale and Holmes, 1993, 1992; Dana, 1998) and an empirical point of view (Borenstein and Rose, 1994; Stavins, 2001; Giaume and Guillou, 2004; Pels and Rietveld, 2004; Piga and Bachis, 2007; Pitfield, 2005). Belobaba (1987), for example, explains that monotonic fares respond to a situation in which transaction costs of adjusting prices to the incoming information about the actual demand are high for FSC, especially in the context of complex hub-and-spoke systems. Gale and Holmes (1993) argue that in a monopoly with capacity constraints and perfectly predictable demand, advance-purchase discounts (ADP, hereafter) are used to divert demand from peak periods to off-peak periods in order to maximize profits. In doing that airlines price discriminate across customers on the basis of their price elasticity and time valuation. Similarly, when the demand is uncertain, APD help to improve profitability by spreading customers evenly across flights before the peak period is known (Gale and Holmes, 1992). The trade-off faced by a trav-

eller with uncertain demand between buying early (risking that she might not need to use the ticket) and buying late (and risk being rationed) is also central in Möller and Watanabe (2009), where APD appears to be a particularly suitable pricing strategy for airlines. Finally, Dana (1998) maintains that in competitive markets where prices are set before the demand is known firms find convenient to implement the "low-before-high-fares" principle in order to cope with uncertain consumer demand.

From an empirical view-point, Stavins (2001) was the first to develop a model in which purchase restrictions and time of booking prior to departure were used as explanatory variables. Although the main objective of her study was to identify the relationship between price dispersion and concentration, her estimates also confirmed the idea that such ticket restrictions as the 14 days requirement, exert a negative and significant effect on fares. Giaume and Guillou (2004) applied the same model to flights leaving from Nice (France) to several European destinations, finding further support for the monotonic property. More recently, Escobari (2006) has complemented Stavin's model with the load factors at the moment of ticket purchases concluding that airfares' monotonic increases over time are due to peak load pricing rather than inter-temporal discrimination. What emerges from the past contributions is the ubiquity of monotonically increasing fares that is assumed to hold even in the simplified yield management developed by the LCC, with fares becoming more and more expensive over time. Such a received wisdom is challenged in Piga and Bachis (2007), who present evidence indicating that for some airlines the early booking fares may be higher than those available from four to two weeks prior to departure. It would therefore seem that the monotonic property does not adequately and fully describe the time profile of many LCCs' pricing schemes when on-line daily fares are used for the analysis. This is probably related to the easiness with which fares can be changed online, due to low menu costs (Smith et al., 1999). Digital markets possess characteristics that do not appear compatible with a monotonic temporal increase of the offered airfares. It has been argued for example that search and menu costs are very low on the Internet. Customers and competitors are thought to be able to easily track down companies' prices and find the cheapest fare available (Baye et al., 2004; Brynjolfsson and Smith, 2000; Smith and Brynjolfsson, 2001). A strictly monotonic increase of fares over time does not seem to be compatible with the airline market where demand uncertainty forces the companies to adjust their fares according to demand and makes tacit collusion difficult to sustain.

3 Data Collection

Our analysis is based on primary data on fares and secondary data on routes traffic, where a route is identified in this study as an airport-pair combination.³ The fares in this study were

³Previous studies on pricing behaviour in the U.S. Airlines industry have used different cohorts of the same dataset, i.e., the Databank of the U.S.A. Department of Transportation's Origin and Destination Survey, which is a 10 percent yearly random sample of all tickets that originate in the United States on

collected using an “electronic spider”, which connected directly to the website of Ryanair. The selection of this site was motivated by the fact that it was possible to obtain information about the number of seats left at the time of the query (see below).

The dataset includes daily flights information from August 2003 up to, and including, June 2005. In addition to UK domestic fares, routes to the following countries were surveyed: Austria, Belgium, France, Germany, Ireland, Italy, Netherlands, Norway, Spain, Sweden, as well as the UK, whose domestic routes were also considered.

In order to account for the heterogeneity of fares offered by airlines at different times prior to departure, every day we instructed the spider to collect the fares for departures due, respectively, 1, 4, 7, 10, 14, 21, 28, 35, 42, 49, 56, 63 and 70 days from the date of the query. Henceforth, these will be referred to as “booking days”.⁴ The return flight for both types of directional journey was scheduled one week after the departure. For those routes where an airline operates more than one flight per day, all fares for every flight were collected. Thus, for every daily flight we managed to obtain up to 13 prices that differ by the time interval from the day of departure. The main reason to do so was to satisfy the need to identify the evolution of fares - from more than two months prior to departure to the day before departure – which has been noted to be very variable for the case of LCC (Pels and Rietveld, 2004; Giaume and Guillou, 2004).

The collection of the airfares has been carried out everyday at the same time: in addition to airfares we collected the time and date of the query, the departure date, the scheduled departure and arrival time, the origin and destination airports and the flight identification code. Fares were collected before tax and handling fees for each one-way trip.

To complement the price data with market structure characteristics, secondary data on the traffic for all the routes and all the airlines flying to the countries indicated above was obtained from the UK Civil Aviation Authority (henceforth, CAA).⁵ For each combination of company, route and departure period (i.e., month/year), the CAA provided the number of monthly seats, the number of monthly passengers and the monthly load factors. These were broken down at the flight identification code level, that is, for each flight operated by all the airlines in a given month and route.

3.1 Retrieving Data On Seats Availability

The collection strategy exploited a feature of the Ryanair’s website: during the data collection period, Ryanair allowed bookings of up to 50 seats using a single query. This made it possible

U.S. carriers (Borenstein, 1989; Borenstein and Rose, 1994; Evans and Kessides, 1993, 1994; Kim and Singal, 1993; Lederman, 2008)

⁴For instance, if we consider London Stansted-Rome Ciampino as the route of interest, and assume the query for the flights operated by a given airline was carried out on March 1st 2004, the spider would retrieve the prices for both the London Stansted-Rome Ciampino and the Rome Ciampino-London Stansted routes for departures on 2/3/04, 5/3/04, 8/3/04, 11/3/04 and so on.

⁵See www.caa.co.uk

to learn if, for a specific flight code on a route, if at the time of the query fewer than 50 seats were available for booking. The web-spider worked as follows: it issued a query for 50 seats for a flight, identified by a specific flight code on a specific route.⁶ The flight was due to depart “X” days from the date of the query. If the airline’s site returned a valid fare, then we interpreted this finding as follows: “X” days prior to departure, there were at least 50 seats available on the flight. We could not however retrieve any precise information regarding the actual number of available seats, which is thus censored at the level of 50. The fare for 50 seats was saved by the spider.

More interestingly, if the site failed to display a valid fare for that flight, the programme inferred that there were fewer than 50 seats available and then started a search to obtain the highest number of seats (N) in a query that returned a valid fare. We interpret this as the maximum number of seats available, and the fare retrieved corresponds to the unit price at which the airline is willing to sell all the N seats in a single transaction. The spider created a record containing the seats and fare info, plus the flight code, the route’s endpoints airports, the time and date of the flight and of the query. “X” took the following values of the booking days indicated above.

By repeating the same operation every day, for each daily flight we could track the seats and the associated fare from 70 up to the day before departure. That is, for each flight code, we have daily data from September 2003 until June 2005. We consider a variety of routes and in some cases more than one flight code per route when the airline operated more than one daily flight. To simplify the data analysis, the procedure considered only flights departing from an airport within the UK, and arriving at either a domestic or an international airport. Thus, all fares were in Sterling.

There is an important distinction that has to be made with regards to the interpretation of the fares retrieved using the procedure previously outlined. If, for example, the spider returned 28 left seats for a given booking day, then the retrieved fare would correspond to the posted fare for a booking of 28 seats, i.e., for the number of seats that would close the flight. While this fare is interesting in itself, we complemented it with the fare obtained by running the same query for a single seat, which resulted from using the spider without the search algorithm for the available seats. This second set of one-seat fares enables us to evaluate the gradient of the pricing profile.

4 Econometric Methodology

The type of intra-firm price dispersion described in Dana (1999) arises because prices are set under uncertainty, so that in equilibrium a firm will set a price that is inversely related to the probability of selling an extra unit of the product. As a flight fills up, the probability of selling an

⁶Such a feature is not available any more, as only 15 seats at most can now be booked simultaneously. It is not possible to know why the airline changed this feature; a likely explanation may be that group bookings rarely exceeds 15 seats.

extra seat becomes smaller; hence, the firm will ex-ante set and commit to a fare that increases with the number of sold seats.

To gain further insights on the hypotheses formulated in the Introduction, we estimate the following equation:

$$P_{ib}^t = \beta Q_{ib}^t + \gamma b + \beta U_i^t + \delta A_{ib}^t + \alpha_i^t + \xi_{ib}^t, \quad (1)$$

where P_{ib}^t denotes the price posted b days before the date of departure t , Q_{ib}^t the number of seats sold b days before the date of departure t , b is a set of booking days' dummies, U_i^t is a set of dummies aimed at capturing a shift in the relationship between P_{ib}^t and Q_{ib}^t , A_{ib}^t denotes observations corresponding to cases in which the airline is posting a promotional (i.e., very low) fare, α_i^t are flights' fixed effects and ξ_{ib}^t is a white-noise error.

The data is structured as a panel with the unit of analysis being a daily flight, identified by a specific code, and the temporal dimension represented by the different booking days. In other words, each panel group tracks the evolution of fares and remaining capacity as the date of departure nears. To account for the possibility that the price-quantity relationship is flight (or route) specific, each model adopts a Fixed-Effects (FE) estimator.⁷

The characteristics of the collected dataset present the investigator with the need to tackle two interesting econometric problems. First, the theoretical model in Dana (1999) jointly determines the size of the "buckets of seats" and the level of price at which they will be put on sale. That is, while the fixed-effect estimator accounts for the possible correlation between Q_{ib}^t and α_i^t , the former may still be correlated with ξ_{ib}^t . We address the possible endogeneity issue by using an Instrumental Variable (IV) technique whose features are related to the other econometric problem.

Second, in addition to being endogenous, Q_{ib}^t is also censored. Following Wooldridge (2002, 643-644), I first obtain the predicted values from the following Tobit model:

$$Q_{ib}^t = \vartheta + \theta W_{ib} + \mu_{ib}, \quad (2)$$

where W_{ib} represents a set of demand shifters listed in the next subsection. The predicted values thus obtained are then used as the instrumental variable in the estimation of (1).

4.1 Variable Definition

We now define the demand shifters used in the estimation of (2), and then proceed to identify the variables capturing the possible updating of the airline's pricing profile.

A set of dummy variables for each booking day, for each day of the week and for the actual time of departure was included among the regressors in W_{ib} ; we therefore aim to capture possible differing demand conditions over these dimensions.

⁷The same concern may apply to the promotional variable, although the same methodological solution apply.

A dummy for promotional pricing was also included, as well as one identifying routes whose endpoints are both defined as “bases” by Ryanair. In such routes, we expect to observe a higher level of traffic, as a base operates in ways that are similar to the hubs in a hub-and-spoke system.

The other regressors used in (2) were: the relative size of the citypair to which the route is part, route length, the number of routes Ryanair operates within the citypair, the total number of flights all companies operate in the route and in the citypair, and the number of flights Ryanair operated in the route and in the citypair. All these variable are likely to shift demand for a flight.

A central element of this research pertains to the extent to which the airline modifies its ex-ante pricing profile to manage unexpected realization of demand. That is, we are interested in evaluating whether the airline engages in “stochastic peak-load pricing” Borenstein and Rose (1994). To this purpose, we have considered whether, at given points in time, a change in price is observed whenever a flight’s number of remaining seats falls below the following predetermined thresholds of 30, 25 and 20 seats. More precisely, for each of these thresholds, we created three dummies if the thresholds are reached 21, 14 and 10 days prior to departure. These variables are all equal to 1 if the threshold is reached at the earliest time (i.e., 21 days), but can have different values if the threshold is reached at a later date. The effect of such dummies is therefore cumulative and so their overall effect is given by the sum of the estimates of each dummy.

We expect that their cumulative effect is larger the earlier a lower threshold is reached. That is, 20 seats remaining three weeks from departure constitutes a clear signal that the demand for the flight is high, and that therefore the airline can adjust its original pricing profile upwards. So the magnitude of the updating in price should be inversely related with the number of left seats. The earlier the threshold is reached, the more likely it is that the airline detects the unexpected increase in demand and that its fares are modified accordingly.

5 Results

All estimates were derived using both a standard panel Fixed Effect estimator and an Instrumental Variables (IV) approach to take into account the endogeneity of “Sold Seats”. The evidence reveals that the latter variable maintains a positive sign in all specifications, thereby providing support to the theoretical results in (Dana, 1999) that fares should increase as capacity fills up. The gradient of the pricing profile is about 4% in Table 1, i.e., each extra seats is sold at price which is about 4% higher than the price of the previous seat. This value remains stable even after controlling for booking days and the presence of promotional fares offered for a specific flight. The inclusion of the latter variable induces an increase of the R^2 from about 0.46 to 0.71. Adding a set of “Booking Days” dummies generally improves the goodness of fit and reveals a quite interesting property. All things equal, when the effect of seat occupancy is taken into account, the time profile of fares appears to be *U*-shaped, with the minimum occurring between 21 and 14 days before a flight’s departure and noticeable price hikes being imposed just a few days from take-off. This finding contradicts other studies that do not control for seat

Table 1: Fixed Effect (FE) and Instrumental Variable (IV) FE Estimates without updating. Dependent variable:ln(fare).

	FE	FE IV	FE	FE IV	FE	FE IV
Sold seats	0.044 ^a	0.043 ^a	0.036 ^a	0.042 ^a	0.030 ^a	0.040 ^a
Booking Days1			0.232 ^a	0	0.360 ^a	0
Booking Days4			-0.026	-0.230 ^a	0.085	-0.233 ^a
Booking Days7			-0.228 ^a	-0.399 ^a	-0.118 ^b	-0.384 ^a
Booking Days10			-0.222 ^a	-0.362 ^a	-0.130 ^b	-0.349 ^a
Booking Days14			-0.369 ^a	-0.474 ^a	-0.214 ^a	-0.380 ^a
Booking Days21			-0.277 ^a	-0.337 ^a	-0.167 ^a	-0.261 ^a
Booking Days28			-0.180 ^a	-0.207 ^a	-0.120 ^b	-0.164 ^a
Booking Days35			-0.104 ^b	-0.110 ^a	-0.076	-0.085 ^a
Booking Days42			-0.064	-0.061	-0.029	-0.025
Booking Days49			-0.025	-0.017	0.01	0.022
Booking Days56			0.031	0.045	0.035	0.056 ^c
Booking Days63			-0.017	-0.019	-0.007	-0.01
Promotional Pricing					-4.625 ^a	-4.563 ^a
Constant	2.850 ^a	2.871 ^a	3.172 ^a	3.164 ^a	3.240 ^a	3.228 ^a
R ²	0.42	0.42	0.47	0.46	0.72	0.71
N	89461	89461	89461	89461	89461	89461

Note: ^a, ^b, ^c denote 1%, 5% and 10% significance levels, respectively.

occupancy, where the temporal fare profile is generally shown to be monotonically increasing (Pitfield, 2005; Mantin and Koo, 2009).

For robustness, Table 2 reports estimates obtained using the full sample of fares, i.e., the sample where all available flights are included even if the number of available seats does not fall below the threshold of 50. Previous results are confirmed, as the pricing profile remains positive and with a gradient varying from about 4.6 to about 8% .

The estimates in Table 3 show that when we introduce the possibility that the airline updates its pricing profile to account for positive demand realisations, the coefficients for “Sold Seats”, “Booking Days” and “Promotional Pricing” remain stable. Most importantly, those flights in which 20 or less seats were available 21, 14 and 10 days from departure on average record fares that are about 22% (10.7%+6.5%+4.9% in the IV regression) higher. The effect decreases as the number of available seats increases. When the number of available seats increases to 25, the overall impact on fares diminishes, as expected, to 15% (0.152%+7.0%-7.2% in the IV

Table 2: Fixed Effect (FE) and Instrumental Variable (IV) FE Estimates without updating. Full sample. Dependent: Ln(fare)

	FE	FE IV	FE	FE IV	FE	FE IV
Sold seats	0.046 ^a	0.084 ^a	0.004 ^a	0.083 ^a	0.012 ^a	0.061 ^a
Booking Days=1			2.158 ^a	-	1.370 ^a	-
Booking Days=4			1.724 ^a	0.245 ^a	1.009 ^a	0.065 ^a
Booking Days=7			1.052 ^a	-0.175 ^a	0.555 ^a	-0.225 ^a
Booking Days=10			0.952 ^a	-0.047 ^c	0.491 ^a	-0.147 ^a
Booking Days=14			0.184 ^a	-0.540 ^a	0.170 ^a	-0.282 ^a
Booking Days=21			0.003	-0.370 ^a	0.101 ^a	-0.129 ^a
Booking Days=28			-0.074 ^b	-0.257 ^a	0.051 ^a	-0.059 ^a
Booking Days=35			-0.070 ^a	-0.159 ^a	0.038 ^b	-0.014
Booking Days=42			-0.048 ^b	-0.099 ^a	0.042 ^a	0.012
Booking Days=49			-0.113 ^a	-0.156 ^a	-0.003	-0.026 ^b
Booking Days=56			-0.130 ^a	-0.161 ^a	-0.033 ^a	-0.050 ^a
Booking Days=63			-0.079 ^a	-0.105 ^a	-0.020 ^a	-0.035 ^a
Promotional Pricing					-3.909 ^a	-4.028 ^a
R ²	0.16	0.05	0.16	0.07	0.72	0.62
N	408771	408335	408771	408335	408771	408335

Note: ^a, ^b, ^c denote 1%, 5% and 10% significance levels, respectively.

regression); it falls to 3.4% when 30 seats or less were available 21 days or less before a flight's departure. It would appear, therefore, that when demand is high (a situation reflected by fewer seats available well before a flight's departure), for fixed booking times, fares appears to be pushed upwards relative to flights which are less likely to be filled up to capacity. This finding suggests a tendency by the airline to update its pricing rule in response to particularly favorable demand conditions.

6 Conclusions

In this study we have illustrated some important features that characterise the yield management system of the largest European Low-Cost Airline. This was possible thanks to an original way used to retrieve information on fares and available seats using the airline's web site.

A main result is the identification of the shape of the pricing profile adopted by the airline, that is, how fares change as the flight fills up. Consistent with the prediction of the economic model of yield management proposed in Dana (1999), we estimate a positive relationship between fares and sold seats, which is stable across regressions. To our knowledge, no other study has estimated such a relationship.

Table 3: Fixed Effect (FE) and Instrumental Variable (IV) FE Estimates with updating. Dependent: Ln(fare)

	FE	FE IV	FE	FE IV	FE	FE IV
Sold seats	0.030 ^a	0.041 ^a	0.029 ^a	0.038 ^a	0.028 ^a	0.037 ^a
Booking Days=1	0.274 ^a	0	0.261 ^a	0	0.264 ^a	0
Booking Days=4	-0.003	-0.231 ^a	-0.022	-0.241 ^a	-0.021	-0.248 ^a
Booking Days=7	-0.208 ^a	-0.379 ^a	-0.232 ^a	-0.402 ^a	-0.235 ^a	-0.416 ^a
Booking Days=10	-0.221 ^a	-0.341 ^a	-0.251 ^a	-0.376 ^a	-0.259 ^a	-0.398 ^a
Booking Days=14	-0.317 ^a	-0.397 ^a	-0.342 ^a	-0.431 ^a	-0.337 ^a	-0.440 ^a
Booking Days=21	-0.277 ^a	-0.316 ^a	-0.278 ^a	-0.330 ^a	-0.257 ^a	-0.319 ^a
Booking Days=28	-0.187 ^a	-0.194 ^a	-0.189 ^a	-0.209 ^a	-0.175 ^a	-0.203 ^a
Booking Days=35	-0.073	-0.088 ^a	-0.079 ^c	-0.090 ^a	-0.084 ^c	-0.093 ^a
Booking Days=42	-0.021	-0.021	-0.026	-0.024	-0.03	-0.027
Booking Days=49	0.018	0.024	0.016	0.023	0.013	0.022
Booking Days=56	0.042	0.061 ^c	0.043	0.059 ^c	0.039	0.056 ^c
Booking Days=63	-0.005	-0.009	-0.005	-0.01	-0.006	-0.009
Promotional Pricing	-4.630 ^a	-4.578 ^a	-4.628 ^a	-4.578 ^a	-4.634 ^a	-4.582 ^a
Days=21; Seats=30	0.171 ^a	0.185 ^a				
Days=14; Seats=30	0.034 ^c	0.005				
Days=10; Seats=30	-0.044 ^b	-0.146 ^a				
Days=21; Seats=25			0.132 ^a	0.152 ^a		
Days=14; Seats=25			0.088 ^a	0.070 ^a		
Days=10; Seats=25			0.014	-0.072 ^a		
Days=21; Seats=20					0.082 ^a	0.107 ^a
Days=14; Seats=20					0.076 ^a	0.065 ^a
Days=10; Seats=10					0.117 ^a	0.049 ^a
Constant	3.314 ^a	3.252 ^a	3.336 ^a	3.289 ^a	3.342 ^a	3.308 ^a
R ²	0.72	0.71	0.72	0.71	0.72	0.71
N	89461	89461	89461	89461	89461	89461

Note: ^a, ^b, ^c denote 1%, 5% and 10% significance levels, respectively.

Failing to control for seat occupancy is also likely to bias the temporal profile of fares, i.e., how fares vary as the date of departure approaches. While the existing literature has generally posited a strictly monotonic relationship, our estimates indicate that the temporal profile appears to be *U*-shaped. This finding is consistent with the following interpretation, based on the existence of buyers with different motivation to travel. Very early bookers are those who

need to reach a specific destination and have very little flexibility on their departure date: their demand is thus slightly inelastic and they are willing to pay a moderately high price to secure a place on a specific plane.⁸ Early-intermediate bookers are normally those who do not fix their departure day ex-ante and shop around across different departure days and destinations: this substitutability makes their demand highly elastic.⁹ Finally, last-minute bookers are generally those whose choice of destinations and travel dates are fixed, so that they are more likely to be willing to pay a high fare.¹⁰

The second main hypothesis tested in this study looks more closely at another characteristics of the model in Dana (1999), according to which airlines commit to a fixed and unchangeable pricing profile. This implies that fares do not reflect any new information the airline may receive on the evolution of a flight's demand. However, our analysis reveals shifts in the pricing profile which appear to be inconsistent with the commitment hypothesis. Such shifts are likely motivated by the airline's desire to adjust its fares upward to reflect positive realization of demand. That is, the possibility that the airline engages in "stochastic peak-load pricing" is not rejected by our econometric analysis.

References

- Alderighi, M., Cento, A., Nijkamp, P. and Rietveld, P. (2004), The Entry of Low-Cost Airlines, *Discussion papers, n.74 (3)*, Tinbergen Institute.
- Baye, M., Morgan, J. and Scholten, P. (2004), Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site, *Journal of Industrial Economics* LII(4), 463–496.
- Belobaba, P. P. (1987), Airline Yield Management. An Overview of Seat Inventory Control, *Transportation Science* 21, 63–73.
- Borenstein, S. (1989), Hubs and High Fares: Airport Dominance and Market Power in the U.S. Airline Industry, *Rand Journal of Economics* 20, 344–365.
- Borenstein, S. and Rose, N. L. (1994), Competition and Price Dispersion in the U.S Airline Industry, *Journal of Political Economy* 102, 653–683.
- Brynjolfsson, E. and Smith, M. D. (2000), Frictionless Commerce? A Comparison of Internet and Conventional Retailers, *Management Science* 46, 563–585.
- Dana, J. D. (1998), Advance-Purchase discounts and price discrimination in competitive markets, *Journal of Political Economy* 106(2), 395–422.
- Dana, J. D. (1999), Equilibrium price dispersion under demand uncertainty: the roles of costly capacity and market structure, *Rand Journal of Economics* 30(4), 632–660.

⁸Think, for instance, of a family planning the Summer or Winter holiday, something which is generally done well in advance of departure.

⁹This could correspond to the market for short-term breaks.

¹⁰This market segment is typically associated with passengers traveling for business purposes.

- Evans, W. N. and Kessides, I. N. (1993), Localized Market Power in the U.S. Airline Industry, *Review of Economics and Statistics* 75, 66–75.
- Evans, W. N. and Kessides, I. N. (1994), Living by the "Golden Rule": Multimarket Contact in the U.S. Airline Industry, *The Quarterly Journal of Economics* 109, 341–366.
- Gale, L. I. and Holmes, T. J. (1992), The Efficiency of advanced-purchase discounts in the presence of aggregate demand uncertainty, *International Journal of Industrial Organization* 10, 413–437.
- Gale, L. I. and Holmes, T. J. (1993), Advance-Purchase Discounts and Monopoly Allocation capacity, *American Economic Review* 83(1), 135–146.
- Giaume, S. and Guillou, S. (2004), Price Discrimination and Concentration in European Airline Market, *Journal of Air Transport Management* 10, 305–310.
- Kim, E. H. and Singal, V. (1993), Mergers and Market Power: Evidence from the Airline Industry, *American Economic Review* 83, 549–569.
- Lederman, M. (2008), Are Frequent-Flyer Programs a Cause of the "Hub-Premium"?, *Journal of Economics and Management Strategy* 17, 35–66.
- Mantin, B. and Koo, B. (2009), Dynamic price dispersion in airline markets, *Transportation Research, Part E* 45, 1020–1029.
- McGill, J. I. and Van Ryzin, G. J. (1999), Revenue Management: Research Overview and Prospects, *Transportation Science* 33, 233–256.
- Möller, M. and Watanabe, M. (2009), Advance Purchase Discounts and Clearance Sales, *Economic Journal*.
- Pels, E. and Rietveld, P. (2004), Airline pricing behaviour in the London-Paris market, *Journal of Air Transport Management* 10, 279–283.
- Piga, C. A. and Bachis, E. (2007), Pricing strategies by European traditional and low cost airlines: or, when is it the best time to book on line?, in D. Lee (ed.), *Advances in Airline Economics. The Economics of Airline Institutions, Operations and Marketing*, Elsevier, Amsterdam, Holland, pp. 319–344.
- Pitfield, D. E. (2005), Some Speculations and Empirical Evidence on the Oligopolistic Behaviour of Competing Low-Cost Airlines, . *Journal of Transport Economics and Policy* 39, 379–390.
- Puller, S. L., Sengupta, A. and Wiggins, S. (2009), Testing theories of Scarcity pricing and Price Dispersion in the Airline Industry, *N. 15555, december*, NBER Working Paper Series.
- Smith, M. D. and Brynjolfsson, E. (2001), Consumer Decision-Making at an Internet Shopbot: Brand still Matters, *Journal of Industrial Economics* 49, 541–558.
- Stavins, J. (2001), Price Discrimination in the Airline Market: The Effect of Market Concentration, *The Review of Economics and Statistics* 83(1), 200–202.
- Weatherford, L. R. and Bodily, S. E. (1992), A Taxonomy and Research Overview of

Review of Economic Analysis 3 (2010) 272–286

Perishable-Asset Revenue Management: Yield Management, Overbooking, and Pricing,
Operations Research 5, 831–844.

Wooldridge, J. (2002), *Econometric Analysis of Cross Section and Panel Data*, 1st edn, MIT Press, Cambridge, USA.