

Targeting leisure and business passengers with unsegmented pricing*

Marco Alderighi [†]

Università della Valle d'Aosta

Università Bocconi

Marcella Nicolini [‡]

Università di Pavia

Claudio A. Piga [§]

Keele University

Rimini Centre Economic Analysis

January 25, 2016

Abstract

We analyze the fare setting strategy of a leading European low-cost carrier, Ryanair, which, until recently, adopted an unsegmented pricing policy (all tickets belong to a single fare class). We show that, to account for different demand characteristics, the company adjusts the two main components governing the dynamics of posted fares, namely *time* (the number of days before departure) and *capacity* (the current number of available seats). We find that: 1) in routes with a strong presence of leisure (business) traffic, fares are set to be less (more) responsive to the time component; 2) in schedules more suitable for leisure (business) travellers, fares are set to be less (more) responsive to the capacity component.

JEL Classification: D22, L11, L93.

Keywords: Dynamic pricing, Unsegmented pricing, Market segmentation, Ryanair.

*We wish to thank Paul Dobson, Carmine Garzia, Gerardo Patriotta, and Federico Visconti for their valuable comments. Piga gratefully acknowledges financial support from the Leverhulme Trust Research Grant RPG-051.

[†]e-mail: m.alderighi@univda.it

[‡]Corresponding author. Università di Pavia, Via San Felice 5, 27100 Pavia (PV), Italy; e-mail: marcella.nicolini@unipv.it

[§]e-mail: c.piga@keele.ac.uk

1 Introduction

Since their appearance in North America and European countries, low-cost carriers have challenged traditional airline companies, by proposing a business model that is based on a cost leadership strategy. Such an approach has become a marketing lever that credibly promotes a brand image based on the strong adherence to the low-cost concept, which implies that the firm strives to be never knowingly undersold by its competitors.

As a cost advantage normally derives from the exploitation of scale economies and learning effects, a cost leader company usually sells a “standardized, no-frills product” (Porter, 1985: 13). The most obvious limitation of standardization is that the firm may not be able to attract those demand segments with a high willingness to pay for a differentiated product or a personalized service. Nonetheless, there is evidence that an increasing number of business travellers choose low-cost carriers, an aspect that raises the question as to whether the intrinsic simplicity of the pricing strategy associated to the standard low-cost business model is suited to the task of taking full advantage of the opportunities offered by the business demand segment. While it is true that in recent years some of these companies have innovated their original pricing approach based on an unsegmented pricing policy (i.e., charging a single-fare class where all seats carry the same characteristics) by embracing a segmented pricing strategy (i.e., charging multiple-fare classes), in this paper we show that even before they did so, low-cost carriers were able to target business travellers by exploiting some crucial features of their Revenue Management (RM) system. Indeed, RM constitutes a central organizational function in companies that have to set the prices of highly perishable services under uncertainty (McGill and Van Ryzin, 1999). It encompasses a set of processes and techniques that are ultimately responsible for the price offered to the final consumers, and in this sense it can be used to target different segments of demand.

In this study we propose an in-depth analysis of the Revenue Management (RM) system implemented by the leading European low-cost airline, Ryanair. Since its inception in the 1980s’, this carrier has been recognized to be the most adherent to a pure cost leadership strategy which has been tightly associated to an exclusive application of an unsegmented pricing approach (Klophaus et al., 2012); however, in mid-2014, Ryanair started offering a business-class fare, which, relative to the basic-class, includes a bundle of additional services each carrying a fixed price to be added to the basic-class fare. Because RM is responsible for the definition of the latter fare, even if our analysis is based on data collected before the introduction of the business-class fare, it can provide insights into how the company fine-tunes its fares depending on specific routes and flights characteristics.

Our analysis hinges on the collection of primary price data taken from the Ryanair website over an 18 months' period, containing information on posted fares and available seats retrieved on a selection of days before departure, and covering scheduled flights for 42 international routes with one endpoint in the UK. This is complemented by secondary data based on official statistics on market characteristics coming from the International Passenger Survey (IPS), which provides route information on the composition of passengers in terms of their travel motivation. For each flight in the sample, we also identify the time and the day of departure, and, from the IPS, the (quarterly average) share of business travellers carried by all companies on the city-pair. We use the latter variables to stratify data according to two dimensions. Flights operated on routes where the share of business passengers is below or above the median of the sample, are respectively labelled as "leisure routes" and "business routes". Flights leaving in early morning or late afternoon during weekdays are assigned to the "business hours" category, the remaining flights to the "leisure hours" category.

We report evidence indicating that Ryanair adapts its RM techniques to the market's characteristics, thus effectively carrying out a segmentation strategy. In practice, the two main components governing the dynamic of posted fares, namely *time* (the number of days before departure) and *capacity* (the current number of available seats), are adjusted to account for different compositions in demand characteristics. We find that: 1) in leisure (business) routes, fares are set to be less (more) responsive to the time component; 2) in leisure (business) hours, fares are set to be less (more) responsive to the capacity component.

Our paper is mainly related to that of Salanti et al. (2012), who study the pricing behaviour of EasyJet. Indeed, our findings confirm their result that the time component is less prominent on leisure routes than on business routes. However, there is no existing evidence exploring the simultaneous use of capacity and time components to manage different consumer segments. Other papers have highlighted, through econometric techniques, that LCCs adjust their pricing policy to market conditions. Bilotkach et al. (2015) find that yield management interventions are less effective in routes with predominantly leisure traffic, while the degree of competition does not produce any substantial difference. Alderighi et al. (2015), applying a similar methodology to ours, show that route length, competitive pressure, market volatility, seasonality and time schedule affect both time and capacity components. Malighetti et al. (2009, 2010) mainly focus on the time component of the pricing policy of EasyJet and find that route length, route frequency and airport dominance play a significant role, while competitive conditions are less relevant. Finally, some works base their analysis on a case study investigation of LCC behaviour. Forsyth (2003) recognized that LCCs may target business

travellers through the use of multiple fares. Other studies in this stream of research show that LCCs have followed different ways in order to attract business travellers including some forms of hybridization (Lawton and Solomko, 2005) and the use of multiple fares (Fageda et al., 2015).

The rest of the paper is organized as follows. In Sections 2 and 3 we briefly review the airline literature on revenue management and we provide the main hypotheses to be evaluated. Section 4 offers a brief description of Ryanair’s business model and strategic positioning. Section 5 describes the data and Section 6 presents the methodology. Section 7 provides some descriptive results and the econometric evidence. Section 8 concludes the paper.

2 Revenue management in theory and practice

McGill and Van Ryzin (1999) identify four key areas for RM in the airline sector: forecasting, overbooking, seat inventory control and pricing. Forecasting plays a critical role in predicting the probability of different states of demand but it is an antecedent process that is largely reflected in the implementation of the other three components; overbooking is the practice of selling a number of tickets larger than the available seats accounting for the fact that some passenger show up for the flight. It is more often employed by Full-Service Carriers (henceforth, FCCs) than LCCs. The two remaining components are paramount in this work, since they are closely examined in the empirical section.

Seat inventory control defines the availability of seats for different booking classes. Even if the airline adopts unsegmented pricing and sells only one fare class with tickets of the same homogenous characteristics, it is still possible for the airline to pre-assign seats with different fare levels to distinct ‘buckets’, each denoting a group of consecutive seats that the airline wants to sell at a given price. In line with the work by Talluri and van Ryzin (2004), the airline decides the buckets’ size (i.e., the number of seats in a bucket) as well as their fare level ex-ante when a flight is first put on sale, based on its demand forecast. The outcome is a full pricing plan for all seats on sale, detailing how the fare will change as the plane fills up. Based on the theoretical model in Dana (1999a), the optimal pricing plan is monotonically increasing because the cost of a seat varies inversely with the probability of selling it.

To test whether the airlines in practice follow this recommendation, it is necessary to know the load factor at the time a fare is posted. Such an information is generally proprietary and very secretive, but Alderighi et al. (2015), Clark and Vincent (2012), and Escobari (2012) used online data to overcome this difficulty. Their empirical findings provide strong support

to the notion that fares increase as a plane fills up. Dana (1999a) also predicts that the price dispersion generated by the optimal pricing plan is larger in more competitive routes. Alderighi et al. (2015) test this prediction by looking at the slope of Ryanair's pricing plan, and find it is on average steeper (i.e., more dispersed) in less competitive routes.

Dana (1999b) offers a complementary explanation as to why the slope of the pricing plan may vary to generate an efficient equilibrium response to uncertainty about the distribution of consumers' departure time preferences. For instance, it is generally assumed that it is more costly for business travellers to choose a flight departing at an inconvenient time. If there is uncertainty on which flight will have peak demand, the airline can profitably alter the mix of low and high price seats and exploit the fact that lower-priced units stock out at the peak time before they stock out at the off-peak time. Conversely, if it is highly likely that a flight will be the off-peak one, then it is profitable to increase the number of low-priced seats to induce some consumers, who would prefer to flight at an alternative time, to switch flight.

Pricing, the fourth component of RM, is strongly intertwined with seat inventory. In particular, to test whether a carrier engages in inter-temporal price discrimination, it is necessary to control that any change in fares over time is not simply induced by the evolution of the flight's load factor. The theoretical literature on Advance Purchase Discounts (APD) has identified customers' heterogeneity in terms of demand uncertainty as the main reason why the airlines may want to change their fares over time. Individuals with different travel motivations learn about their need to travel at various intervals before the actual travel date; notably, business travellers are very likely to book a flight very close of its actual departure, while leisure travellers tend to plan well in advance.

Because travel motivation is not observable by the carrier, APD provide a simple way to screen consumers by their demand uncertainty. Gale and Holmes (1993) demonstrate that in a monopoly with capacity constraints and perfectly predictable demand, APD form an efficient mechanism to carry out second degree price discrimination because they are used to separate customers with different evaluation of the flight: only the high-type will buy the non-discounted fares, while the low-types, anticipating the price hike, will make arrangements to book as soon as their need to travel is revealed. By doing so, the airlines price discriminate across customers on the basis of their price elasticity and time valuation. Möller and Watanabe (2010) compare APD (where prices are increasing over time) with clearance sales (where prices decrease over time), and illustrate how the former are more appropriate when a consumer faces no or little risk of being rationed.

One often observed way to manage the complexity of dealing simultaneously with all of the

above aspects is to offer a portfolio of segmented fare categories, from economy to business, each with a different set of tickets and in-flight characteristics (free meal, refundability, cancellation, change of dates, etc). This is typically what traditional carriers do. Our empirical strategy is aimed at shedding light on how the same RM problems were tackled by a prominent low-cost carrier committed to an unsegmented pricing policy.

3 Hypotheses

The previous discussion has highlighted two important drivers of airline fares: inter-temporal price discrimination and the aircraft seat availability. In the remainder of the paper, the former will be referred to as the ‘time component’, and the latter as the ‘capacity component’. Both components are under the control of a revenue manager; both are assumed to be adjusted to account for the mix of customer types in a market. The joint empirical investigation of the hypotheses developed in this section should therefore provide some indication of how a carrier’s RM system can be effectively used to appropriate the value created by a higher-level low-cost strategy, even in circumstances where varying market conditions may be managed via the application of a single pricing policy.

Leisure and business travellers differ along two dimensions. First, business travellers usually assign a greater value to a flight and therefore their demand is more price inelastic than that of leisure travellers. Second, business travellers tend to plan less in advance than leisure ones (Talluri and van Ryzin, 2004). Therefore, heterogeneity between the two traveller segments allows to inter-temporally segment the market, applying higher fares to clients arriving late, mostly business travellers; and cheaper fares to those arriving early, mostly leisure travellers (Desiraju and Shugan, 1999; Courty, 2003; Netessine, 2006). When the share of business travellers is high (i.e., on a business route), the carrier will therefore tend to choose a pricing profile with large price hikes in the proximity of the departure date.

Hypothesis 1 (*leisure/business routes*): *on leisure (business) routes, fares tend to be less (more) strongly affected by the time component of RM.*

Flights early in the morning (6am-10am) or late in the evening (6pm-10pm) in the week days (i.e. in a business time) are appealing for business travellers (Borenstein and Netz, 1999). Indeed, they allow to easily visit a destination within a working day, i.e., departure in the morning with the return scheduled in the evening. Leisure travellers usually prefer a more comfortable schedule and generally have more flexible preferences in terms of departure

date and time. As a result, the demand of the leisure segment is more price elastic due to the larger substitutability between flights across routes, departure dates and departure and arrival times.

Because the wide majority of LCCs' customers travel for leisure purposes, and network considerations make it necessary to operate flights at 'inconvenient times' for them, the carrier faces the problem of inducing some of the leisure travellers to choose flights operated during business hours. Based on Dana (1999b), one way to achieve this is by modifying the slope of the pricing plan; in business hours, the first buckets of seats should be assigned a lower than average fare level, to attract leisure passengers, while the last buckets should be priced significantly higher since they will be most likely sold to business customers whose need to purchase is revealed only once the flight is sufficiently full. In leisure time flights, because the expected demand is high, the carrier can achieve a higher yield by setting a higher fare for the first buckets, without increasing its fares too sharply for the last buckets. The upsides of such a strategy are that the carrier guarantees itself a smoother occupancy rate across the two types of flights (business vs. leisure hour) with higher occupancy rates also in the off-peak period (business time), as well as higher fares when unexpected high demand materialises, e.g., if and when business travellers choose their most convenient flight.

Hypothesis 2 (*leisure/business hours*): *In leisure (business) hour flights, the capacity component of RM, where fares depend on the remaining seats in a flight, tends to be less (more) prominent.*

Finally, as the route and hour dimensions concern different unrelated aspects of flight supply, we expect the carrier to independently use these two pricing policies.

Hypothesis 3 (*leisure/business routes and hours*): *For flights in both leisure (business) routes and hours, the capacity and time components of RM tend to be less (more) jointly prominent; in mixed cases (e.g., business routes and leisure hours), the dominant component is consistent with Hypotheses 1 and 2 (see: Figure 1).*

4 A business case: Ryanair

4.1 Ryanair's business model

Since the deregulation of the US and European airline sector, LCCs have gained considerable traffic volumes at the expenses of FSCs and regional carriers (Barrett, 2011). According to

the International Civil Aviation Organization (ICAO), in 2012 LCCs supplied about one third of the overall seat capacity on continental routes both in U.S. (31%) and in Europe (37%).

By adopting the business model pioneered by Southwest Airlines, Ryanair rapidly became Europe's largest carrier and it established its clear low-cost 'firm identity' based on "continued improvements and expanded offerings of its low-fares service... while maintaining a continuous focus on cost-containment and operating efficiencies" (www.ryanair.com/doc/investor/Strategy.pdf). Indeed, all the characteristics which allow a carrier to get and maintain a cost leadership strategy in the airline sector are satisfied (Alamari and Fagan, 2005; Bilotkach et al., 2010): one type of aircraft, secondary airports, quick boarding time, high number of rotations, web-based selling strategy, no loyalty scheme, automated check-in, no food, no premium cabin, no connection guaranties, one class of tickets without price discrimination based on multiple service and cabin classes, no specific restrictions like minimum stay requirements and Saturday night stay-over, no overbooking. These features define Ryanair's business model, a model which heavily contrasts with that of FSCs, based on a number of customized services, such as business client support, business and economy cabins, free on board services, selective VIP lounge access, travel restrictions, multiple reservation classes, multiple fares simultaneously available, etc.

These striking differences are also reflected in the way the two types of carriers design their RM systems (Malighetti et al., 2009). Ryanair's fare structure is simpler than that adopted by FSCs, an aspect that by itself supports the view of a necessary coordination between a firm's higher-level strategy and its pricing system: up until mid-2014, it adopted only one reservation class, and consequently, one single fare which is intended for all customers (unsegmented pricing). It also rules out travel agents' commissions, connecting flights, code-sharing agreements and frequent flyers' programmes. FSCs' fare structure is much more complex: it usually includes 11-13 different reservation classes which account for different ground and in-flight services and flexibility requirements (segmented pricing). This differentiation strategy is an instrument to more easily segment the market and target business travellers (Cento, 2009). Nonetheless, it might not be sufficient, given the increased heterogeneity in customers' preferences (Teichert et al., 2008).

4.2 Ryanair's strategic positioning

Ryanair's choice of keeping its RM system as simple as possible has promoted the fidelization of leisure passengers (Hsu, 2006); however, Ryanair's network covers many non-strictly leisure markets where it faces a strong financial incentive to design a competitive strategy aimed at

attracting business travellers, whose willingness to pay is generally considered to be higher than the market average. Thus, the company faces the dilemma of maintaining its core commitment to low fares without compromising its well-established firm identity.

Ryanair’s strategic positioning as an exemplary LCC is beyond discussion and has been consistently pursued throughout its history: over the years, it has continued to adhere strictly to a cost leadership strategy (Porter, 1980, 1985; Barrett, 2011), offering a “no-frills service” (Porter, 1985: 13) combined with an “aggressive pricing” strategy (Porter, 1980: 36). The main consequence is that leisure travellers have formed its largest market demand segment. As its network expanded, so as to cover many capital cities and important economic areas, the number of non-strictly leisure destinations also grew, and with it the opportunity to tap into the profitable demand segment of business travellers, whose willingness to fly low-cost was, at least partly, motivated by private companies’ travel policies aiming at containing costs but also by the high punctuality that many LCCs, and Ryanair in particular, have guaranteed (Dobson and Piga, 2013).

Given its great prolonged success, it seems unlikely that Ryanair did not attempt to take advantage of such an opportunity, even within the constraints imposed by the low-cost business model, and in particular by the standardization of its fares implied by the unsegmented pricing approach. In line with the approach discussed above, we hypothesize the carrier responded by adjusting its pricing policy (i.e., its RM system) in those markets with a higher potential demand by business travellers. A closer investigation of Ryanair’s standard fares offers therefore an invaluable opportunity to gain insights into the capability role played by a company’s pricing system, its compatibility with the company’s wider strategy and its ability to set prices that can be effectively used to extract value from premium customers, even when the wider strategy does not explicitly considers segmentation as one of its objectives (Porter, 1980, 1985).

5 Data

5.1 Experimental design

The data used in the study derive from two different sources. Primary data on repeated prices for single flights and their characteristics were retrieved directly from Ryanair’s website using a web crawler, i.e., a programme that automatically launches the online queries necessary to book a flight to a given destination. Secondary data detailing the composition of a route’s passengers by reason of travel (leisure vs. business) were obtained from the International

Passenger Survey (henceforth, IPS), a quarterly survey collected by the UK Office of National Statistics.

We consider only international flights operated by Ryanair with one endpoint in the UK and the other in one of the following countries: Austria, Belgium, France, Germany, Ireland, Italy, the Netherlands, Norway, Spain, Sweden over the period January 2004 - June 2005. The collection strategy exploited a feature of Ryanair’s website: during the sample period, it was possible to purchase up to 50 seats using a single query.¹ Therefore, at the time of a query one could learn the exact number of available seats (up to 50) on a flight; this information is central for the identification of the capacity component of RM.

The web crawler, after issuing a query, retrieved all the information shown on the returned web page. To obtain information regarding the seat availability of a flight at a specific point in time, the web crawler followed this algorithm based on three steps.

Step 1. Issue a query for 50 seats for a specific flight. The flight is due to depart D days from the date of the query, where D assumes the following values: 1, 4, 7, 10, 14, 21, 28, 35, 42, 49, 56, 63 and 70. The variable containing the information on the days to departure is labelled *BookingDay* in the subsequent analysis.

Step 2. If the airline’s site returns a valid fare, this can be interpreted as follows: “ D booking days prior to departure, there are at least 50 seats available on the flight”. In this case we can not retrieve any more precise information regarding the observed number of available seats, which is thus censored at the level of 50. We store this information in a variable labelled *AvailableSeats*, which in this case assumes the value of 50. We also retain the corresponding value of the fare posted for the query of 50 seats, which we label *TopFare*. Additionally, we collect the information on the fare for a single seat, which is saved in the *Fare* variable. Finally, we store the value of *BookingDay* and all the other flight’s details (see below).

Step 3. If the site fails to return a valid fare for that flight, the web crawler infers that there are fewer than 50 seats available. It then searches the highest number of seats in a query that returns a valid fare. This value defines the number of seats available D days before a flight’s departure; it is stored in *AvailableSeats*. In this case, *TopFare* corresponds to the unit price at which the airline was willing to sell all the remaining seats in a single transaction. As in the previous case, we also store the fare for a single seat in *Fare*.

By repeating this procedure every day, we track the seats and the associated fares (both *Fare* and *TopFare*) of a flight at specific time intervals defined by *BookingDay*. Given our experimental design, for every daily flight we have up to 13 prices charged on different days

¹This is no longer possible, as the query can be for a maximum of 25 seats as of September 2014. However, the mechanism illustrated here still applies.

before the day of departure, which allow the identification of the evolution of fares over time (time component of RM in Hypothesis 1).

5.2 Data structure

The web crawler also saved the departure date, the scheduled departure and arrival time, the origin and destination airports and the flight identification code. These variables are combined to obtain the panel identifier, which corresponds to a single flight between two destinations at a given time of the day on a specific day during January 2004 - June 2005. The time dimension of the panel is given by the time before departure (i.e., the booking day).

As we need to classify the routes according to the predominant type of customers, we use the data provided by the IPS database: it contains a random sample of around 2 percent of passengers entering/leaving the U.K. by air and provides quarterly information on expenditure levels and passenger characteristics, including the purpose of the journey. For each route, we aggregate the survey information across carriers to measure the percentage of passengers traveling for a specific reason (business or leisure). Such information is used to determine the passenger mix on the route, and then to classify flights into different routes categories.

After merging this information, we end up with data for 42 of the 154 routes that Ryanair operated to these countries over the sample period; in some cases, we consider more than one flight code per route when the airline operated more than one flight per day.

5.3 Dependent variable

Fare, i.e., the posted price of one seat retrieved by the crawler, represents the dependent variable in the empirical analysis. The fares collected are net of add-ons and other fees, such as charges for the use of credit cards as methods of payment. All fares do not include tax and handling fees. Excluding taxes and fees does not affect our results as the fixed per-passenger tax affects all tickets and thus does not modify the pricing plan based on the flight's capacity. Fares are for a one-way flight and are quoted in Sterling. Focusing only on the outward leg from the UK is not problematic, as it is widely acknowledged that European LCCs charge prices for each leg independently (Bachis and Piga, 2011).

Although from mid-2014 Ryanair started offering a business fare, during the period considered in our analysis the carrier was strictly following an unsegmented pricing strategy, still currently in place for non-business fares, where all tickets carry the same penalties for a name, date and/or route variation: these charges are so high relative to the average price of a ticket that it is often cheaper to buy a new ticket than change it. All tickets permit the same free

in-flight hand baggage allowance (max 10 kg) with a fixed fee for each checked baggage (max 15 kg per item). Ryanair started charging a fixed fee for check-in and luggage only in 2006, that is, after our sample period.

5.4 Independent and stratification variables

The independent variables necessary to test our hypotheses relate to the capacity and time dimensions along which the pricing strategy is designed. *AvailableSeats*, which represents the unsold space on a flight on a given booking day, identifies the capacity component of RM. This variable is censored, as its value spans from 1 to 50; additionally, it is endogenous. Its inclusion calls for the implementation of an appropriate econometric specification, described in the next Section.

The time component of RM is modeled by including a set of dummies corresponding to the different days of booking (*BookingDay*). Additionally, as we know the day in which the booking has taken place, we can control for specific features related to the period of the year during which the query was carried out by including a set of dummies for the month (*MonthOfBooking*) and the day of week (*DayOfWeekOfBooking*) in which the data was retrieved. Finally, some ancillary regressions include other dummy variables for the week number (*WeekNumberOfDept*), the hour (*HourOfDept*) and the route (*Route*) of the scheduled departure.

We stratify the sample based on the hour type (business vs. leisure hours) and the route type (business vs. leisure routes). The first layer is obtained using the scheduled departure time of flights. A flight is operated in *BusinessHours* if its scheduled departure is from 6:00 a.m. to 9:59 a.m. and from 6:00 p.m. to 9:59 p.m. during weekdays. The remaining flights are operated in *LeisureHours*. The second layer is obtained by classifying flights in accordance with travellers motivation. More specifically, for each flight in the sample, we have computed the share of business travellers carried by all companies on the city-pair, using quarterly data from the International Passenger Survey (IPS) collected by the UK Office for National Statistics. Flights operated on routes where the share of business passengers is below or above the median of the sample (which is 37.1 per cent), are respectively labelled as *LeisureRoutes* and *BusinessRoutes*. It is worth mentioning that the second stratification variable refers to passengers carried by all companies flying on the selected city-pair, not just Ryanair. Moreover, the choice of valuating departures and arrivals at city level implies that also passengers usually flying from airports not served by Ryanair or by other low-cost carriers are included in the analysis. This choice leads to a significant passenger heterogeneity among route types: in *LeisureRoutes*, the share of business travellers is 21.3 per cent, while

in *BusinessRoutes* it is equal to 49.6 per cent. This stratification choice thus allows us to take into account both effective and potential traffic of business and leisure passengers when testing whether Ryanair implements tailored pricing policies to attract and/or extract value from these market segments.

Table 1 shows the descriptive statistics as well as the correlation matrix between the dependent variable and other variables of interest for the censored sample, i.e. for *AvailableSeats* < 50. The correlations between the variables are mostly significant and have the expected signs.

6 Methods

In our analysis we adopt two different approaches to identify the main traits of the pricing strategy of Ryanair. In the first one, descriptive statistics are applied to analyze how fares change across route and hour types. Thus, we split the sample along these two dimensions and investigate how fares evolve according to the availability of seats and the time at which the fare was posted. A comparison of fare trends in the different cases can provide *prima facie* evidence of the hypothesized differential pricing behaviour. In the second approach we adopt econometric analysis to separate fare changes induced by variations in the flight’s remaining capacity from time effects that are unrelated to the actual observed evolution of sales. We can thus investigate how Ryanair used these two components to target different customer segments. Our basic econometric approach consists in the estimation of the following pricing equation:

$$\ln Fare = \alpha_0 + \alpha_1 AvailableSeats + \alpha_2 BookingDay + \alpha_3 MonthOfBooking + u, \quad (1)$$

where the dependent variable, *Fare*, is considered in logs. In equation (1) we include the main independent variables, *AvailableSeats* and the set of *BookingDay* dummies as well as a set of *MonthOfBooking* dummies. The classic ordinary least squares (OLS) model is inappropriate in this context, as *AvailableSeats* has two features which need special attention. First, *AvailableSeats* spans from 1 to 50 and is thus censored, due to the retrieving procedure implemented by the web crawler. This censoring induces a bias in the estimates, and needs to be corrected. Second, it is endogenous, due to an omitted variable problem: some unobserved determinants of the airline pricing behavior may be correlated with a flight’s time-invariant factors. More importantly, the distribution of bucket sizes, which constitutes the pre-set capacity-based pricing plan, may be altered by idiosyncratic, *discretionary* interventions of the airline’s revenue

manager (Bilotkach et al. 2015).²

To deal with endogeneity, we choose two instruments. Their validity depends on the extent they are correlated with the endogenous independent variable, *AvailableSeats*, and uncorrelated with the error term of the pricing equation, u . The first instrument, *BookOnHolidays*, is a dummy variable equals to one if the fare is posted during a holiday period (i.e., main UK Bank Holidays and the week before and after Christmas and Easter). Its effect on *AvailableSeats* may be driven by the fact that the ticket purchasing activity in such periods is likely to be different from non-holiday periods (e.g., when on holiday a person is less willing to spend time planning future trips). Its effect on *AvailableSeats* may be driven by the fact that the ticket purchasing activity in such periods is likely to be different from non-holiday periods (e.g., when on holiday a person is less willing to spend time planning future trips).

The second instrument, *LagMeanSlope*, is the flight-specific and booking day-specific expected slope of the pricing plan. It is constructed as the difference between *TopFare* and *Fare* divided by *AvailableSeats*, using data referring to the same flight identification code and booking date of flights departing on the same day of the week of the three preceding weeks.³ We choose this lag to capture the fact that pricing plans may change with the day of the week, i.e., the pricing plan for a Monday may be different from the one for a Wednesday. *LagMeanSlope* is expected to increase with occupancy, since the fare plan is convex (see: Table 2). Descriptive statistics for both instrumental variable as well as their correlation with *AvailableSeats* are reported in Table 1.

To handle censoring and endogeneity, we follow the procedure 17.4 in Wooldridge (2002), as applied in Alderighi et al. (2015):

1. We estimate a Tobit specification pooling all observations:

$$AvailableSeats = \beta_0 + \beta_1 BookingDay + \beta_2 LagMeanSlope + \gamma \mathbf{X} + v, \quad (2)$$

where \mathbf{X} includes week, route, day of booking and time of departure dummies.

2. We retrieve the residuals \hat{v} for the selected sub-sample.

²For instance, the manager may *discretionally* decide to expand the size of lower-priced buckets if the flight is not selling as expected; that is, fares may be reduced (increased) when *AvailableSeats* is higher (lower) than expected. The *true* fare setting model should also include the RM analyst’s discretionary intervention, *RMI*. As this is unobserved, its effect is included in the error term u ; endogeneity is thus due to an omitted variable problem resulting from the correlation between *AvailableSeats* and *RMI*. Therefore, estimation using OLS should produce a downward bias in the coefficient for *AvailableSeats*.

³Even if the denominator is censored and equal to 50, this ratio can still be interpreted as the slope of the pricing plan, because *TopFare* represents the value of the 50th seat ahead of the one being sold, whose value is given by *Fare*.

3. On the selected sub-sample, we estimate a modified version of (1), where we include \hat{v} among the regressors to correct for sample selection. As *AvailableSeats* is endogenous, we adopt an Instrumental Variable Two-Stage Fixed Effect (IVFE) estimator, using as instruments *HolidayPeriod* and *LagMeanSlope*.

To estimate (1), given the structure of our data, we focus on a panel where the identifier is the single flight (defined by a combination of departure date and flight code) and the time dimension is given by the time before departure (i.e., the booking day). This panel structure also allows us to control for all unobserved characteristics which are specific to the single flight, such as, for instance, market structure and distance, which do not change during the limited period captured by the booking days. Moreover, the fixed-effects approach allows us to control for possible strategic effects at the route level, where, for example, the airline can opt to implement temporary capacity limits, i.e., reduce the number of daily flights.

The choice of instruments is validated by the tests presented in Tables 4-6: the Hansen’s J statistic for overidentifying restrictions and the Kleibergen-Paap LM statistic, which tests whether the equation is identified.⁴ To anticipate our results, both tests, as well as the weak instruments tests not reported, strongly support our choice of instruments.

7 Results

7.1 Descriptive analysis

Tables 2 and 3 show the average value of *Fare*, broken down by route type and, respectively, each of the two RM effects, i.e., available seats and booking days. Because these statistics do not control for important flight characteristics, differences in the mean fare levels across different route and hour types are not particularly informative. However, both Tables provide important insights into the effects of the two RM components (i.e., capacity and time) on fare changes. Table 2 reports the mean *Fare* value across different *AvailableSeats* categories. It describes how the capacity component operates: average prices in each sub-sample tend to increase as the plane fills up. This aspect has been largely neglected in the literature due to the difficulty to match information on a flight’s seat inventory with offered fares.

Table 3 reports the mean *Fare* value across different booking days. Two different samples are considered: the whole sample (upper part of the table), and a sub-sample (lower part of

⁴The joint null hypothesis of Hansen’s J statistic is that the instruments are valid. If the test fails to reject the null hypothesis, then all instruments used are considered exogenous. As for the Kleibergen-Paap LM statistic, a rejection of the null indicates that the matrix of reduced-form coefficients is full column rank and the model is identified.

the table), containing only the fares for flights with less than 50 seats available. This Table also confirms the importance of controlling for capacity. Indeed, in the top part, where we do not control for seat capacity, average prices tend to increase constantly as the departure date nears. However, in the bottom part, each row exhibits a U-shaped temporal profile: fares start relatively high, then decrease and then increase in the last 7 to 10 days. Thus, comparing fares without knowing how many seats are left on a plane is likely to lead to biased estimates of the time effect.

The combined evaluation of both Tables is also revealing; the capacity effect operates in a strictly monotonic manner, while the time effect is U-shaped. However, capacity is sold sequentially over time, and so if the temporal effect was only a reflection of the change in remaining seats, then we should also observe a strictly monotonic temporal effect. The fact that we do not is suggestive of possible temporal effects that are independent of the evolution of a flight’s capacity. Indeed, consistent with Hypothesis 1, in business routes we observe more variability over time, with fares displaying a more pronounced U-shaped behaviour, while in leisure routes the temporal profile seems flatter. Moreover, in accordance to Hypothesis 2, in business hours, fares are more affected by seat availability than in leisure hour.

Because of the simultaneous occurrence of these two effects it is difficult to provide sound conclusions based on descriptive statistics, alone. In the next sub-section we rely on econometric analysis to better identify the two effects.

7.2 Econometric analysis

The econometric procedure from the previous section is applied separately to each sub-sample of flights in, respectively, business vs. leisure routes and business vs. leisure hours, as well as their pairwise combinations. If the estimates reveal that the capacity and the time dimensions play a different role in each sub-sample, we can conclude that the RM at Ryanair was designed to segment its customer basis even if it adopted an unsegmented pricing policy.⁵

The coefficients of *AvailableSeats* in Table 4 (see also: Figure 2) suggest that seat inventory management does not differ significantly between business and leisure routes, although in each sub-sample *AvailableSeats* is an important driver of fares. The negative coefficient implies that, on average, in both sub-samples every time an extra seat is sold the price increases by about 3.0%. Conversely, in line with Hypothesis 1, the evidence points towards sharp differences in

⁵The estimates from the Tobit and the first stage of the IVFE regressions for the first two hypotheses are reported in the Appendix. It is noteworthy, however, that these estimates are found to be very similar, which suggests that the censoring of *AvailableSeats* induced by the data collection strategy is properly accounted for by the inclusion of the Tobit residuals in the main panel estimation.

the temporal profile of fares: in the last two weeks prior to departure, i.e., during the period when most business travellers learn about their need to travel, fares tend to rise more sharply in business routes, as expected. For instance, all things equal, fares taken one and ten days from departure are, respectively, about 27% ($0.777 - 0.507$) and 12% ($0.169 - 0.051$) higher in business routes, relative to the base case of fares posted 21 days prior to departure. Overall, the evidence reported in Table 4 suggests that Hypothesis 1 is supported by the empirical evidence.

In Table 5 (see also: Figure 3), which distinguishes flights according to their departure time and day, the difference between the coefficients of the *BookingDay* dummies is negligible; for instance, after controlling for capacity effects, fares evolve in a similar manner during the two weeks preceding the departure. As Hypothesis 2 predicts, the role of remaining capacity is paramount for flights in business hours; indeed, as revealed by the the coefficients of *AvailableSeats*, such flights tend to have a steeper pricing profile, so that every time a seat is sold, fares on average increase by about 3.6%, as opposed to only 2.9% for flights in leisure hours. In practice, this can be achieved by setting lower fares for the first buckets posted for sale, as also indicated by Table 2: when less than 10 seats remain, the average fare is similar in the two sub-samples of business and leisure hours, but in business hour flights it tends to be lower the more seats remain to sell. Such a profile is consistent with the attempt to smooth demand across flights in leisure and business hours, by inducing leisure travellers to choose flights in business hours.

To test Hypothesis 3, the two previous pairs of sub-samples are further combined to obtain the four sub-samples reported in Table 6. The previous results are robust to this division of the sample: the temporal dimension of RM is important for flights in business routes, while the capacity dimension plays a more crucial role in flights departing during business hours. However, both dimensions are found to be concurrent and significant drivers of fares for flights in business routes and hours. By the same token, a generally flatter temporal and capacity profile characterizes flights in leisure routes and hours. Overall, the results in Tables 4-6 lend support to the view that the adoption of an unsegmented pricing policy does not hinder a firm's ability to organize its RM activities to take advantage of the opportunities that different markets conditions offer.

8 Conclusions

Although simpler than the one adopted by FSCs, this study shows that the RM approach by LCCs hinges on a multidimensional set of factors, the level of which can be adjusted according to the characteristics of the markets the airline serves. The outcome is a flexible arrangement where the carrier can differentiate the way it determines the fares it offers to its customers even if officially it is committed to an unsegmented pricing approach. RM allowed Ryanair to tap into the profitable segment of business travellers despite it posted only a single fare at each point in time. The empirical analysis illustrates a firm's ability to implement an adaptive response of its RM system to market conditions that enhanced its ability to extract value from its customers without contravening to the rules imposed by the business model.

This paper relies on data collected in 2004 and 2005, a period when Ryanair did not offer add-on services. Until mid-2014, Ryanair continued to operate an unsegmented RM approach, although in previous years it had started offering such add-on services as priority boarding, extra luggage, and reserved seats that are generally considered to be valued mostly by business and wealthier leisure travellers. In mid-2014, the carrier introduced a business class, which, in addition to the previous add-on services, also bundles the option of date changes, free airport check-in and fast-tracking through the airport security process. Arguably, although the 2014 changes facilitated and quickened the online purchase process of the bundled services for some time-constrained customers, they do not seem to have fundamentally altered the central role of the RM methods leading to the setting of the basic fare level, to which the business fare is linked.

Finally, in this paper we have limited our analysis to Ryanair. This choice allowed us to derive some important conclusions on the behaviour of the European largest and most successful low-cost carrier, and to describe how it targets different users even under unsegmented pricing. Our empirical findings suggest a very simple behavioural rule: business routes (hours) call upon a higher time (capacity) component. Some anecdotal evidence collected from practitioners in the field suggests that this sophisticated pricing scheme is the result of a trial-and-error strategy more than a deliberative choice of the carrier. If this is the case, we expect that also other (successful) LCCs are implementing such a pricing scheme. The behavioural rule we have unveiled in our research could be useful for young LCCs for improving their pricing policies and for LCCs in general for pricing of new routes. In this respect, it will be the object of future research to investigate whether this pricing rule has been fully accounted by other carriers and whether it has been affected by the current practice of many LCCs, which have embraced a process of hybridization, to charge multiple fares.

References

- [1] Alamari F, Fagan S. 2005. Impact of the adherence to the original low-cost model on the profitability of low-cost airlines. *Transport Reviews* **25**(3): 377-392.
- [2] Alderighi M, Nicolini M, Piga CA. 2015. Combined effects of load factors and booking time on fares: Insights from the yield management of a low-cost airline. *Review of Economics and Statistics* **97**(4): 900-915.
- [3] Bachis E, Piga CA. 2011. Low-cost airlines and online price dispersion. *International Journal of Industrial Organization* **26**(6): 655-667.
- [4] Barrett SD. 2011. Ryanair and the Low-cost Revolution, in J. F. O’Connell, G. Williams, *Air Transport in the 21st Century: Key Strategic Developments*. Chp. 3, 113-128.
- [5] Bilotkach V, Gaggero A, Piga CA. 2015. Airline Pricing under Different Market Conditions: evidence from European Low Cost Carriers. *Tourism Management* **47**: 152-163.
- [6] Bilotkach V, Gorodnichenko Y, Talavera O. 2010. Are airlines price-setting strategies different? *Journal of Air Transport Management* **16**: 1-6.
- [7] Borenstein S, Netz N. 1999. Why do all the flights leave at 8 am?: Competition and departure-time differentiation in airline markets *International Journal of Industrial Organization* **17**: 611-640.
- [8] Cento A. 2009. *The airline industry: challenges in the 21th century*, Springer Science & Business Media. Physica-Berlag Heidelberg: Berlin.
- [9] Clark R, Vincent N. 2012. Capacity-contingent pricing and competition in the airline industry, *Journal of Air Transport Management* **24**: 7-11.
- [10] Courty P. 2003. Ticket pricing under demand uncertainty. *Journal of Law and Economics* **63**(2): 627-652.
- [11] Dana J. 1999a. Equilibrium Price Dispersion Under Demand Uncertainty: the Roles of Costly Capacity and Market Structure. *RAND Journal of Economics* **30**(3): 632-660.
- [12] Dana J. 1999b. Using yield management to shift demand when peak time is unknown. *RAND Journal of Economics* **30**(4): 456-474.

- [13] Desiraju R, Shugan SM. 1999. Strategic service pricing and yield management. *Journal of Marketing* **63**(1): 44-56.
- [14] Dobson PW, Piga CA. 2013. The Impact Of Mergers On Fares Structure: Evidence From European Low-Cost Airlines. *Economic Inquiry* **51**(2): 1196-1217.
- [15] Escobari D. 2012. Dynamic Pricing, Advance Sales, and Aggregate Demand Learning in Airlines. *Journal of Industrial Economics* **60**(4): 697-724.
- [16] Fageda X, Suau-Sanchez P, Mason KJ. 2015. The evolving low-cost business model: Network implications of fare bundling and connecting flights in Europe. *Journal of Air Transport Management* **42**: 289-296.
- [17] Forsyth P. 2003. Low-cost carriers in Australia: experiences and impacts. *Journal of Air Transport Management* **9**: 277-284.
- [18] Gale I, Holmes T. 1993. Advance-Purchase Discounts and Monopoly Allocation of Capacity. *American Economic Review* **83**: 135-46.
- [19] Hsu G. 2006. Jacks of all trades and masters of none: audiences' reactions to spanning genres in feature film production. *Administrative Science Quarterly* **51**: 420-450.
- [20] Klophaus R, Conrady R, Fichert F. 2012. Low cost carriers going hybrid: Evidence from Europe. *Journal of Air Transport Management* **23**: 54-58.
- [21] Lowton T C, Solomko S. 2005. When being the lowest cost is not enough: Building a successful low-fare airline business model in Asia, *Journal of Air Transport Management* **11**: 355-362.
- [22] Malighetti P, Paleari S, Redondi R. 2009. Pricing strategies of low-cost airlines: The Ryanair case study, *Journal of Air Transport Management* **15**(4): 195-203.
- [23] Malighetti P, Paleari S, Redondi R. 2010. Has Ryanair's pricing strategy changed over time? An empirical analysis of its 2006-2007 flights, *Tourism Management*, **31**: 36-44.
- [24] McGill J, van Ryzin G. 1999. Revenue Management: Research Overview and Prospects. *Transportation Science* **33**(2): 233-256.
- [25] Möller M, Watanabe M, 2010. Advance Purchase Discounts and Clearance Sales. *Economic Journal* **120**:1125-1148.

- [26] Netessine, S. 2006. Dynamic pricing of inventory/capacity with infrequent price changes. *European Journal of Operational Research* **174**: 553-580.
- [27] Porter ME. 1980. *Competitive strategy*. Free Press: New York.
- [28] Porter ME. 1985. *Competitive advantage*. Free Press: New York.
- [29] Salanti A, Malighetti P, Redondi R. 2012. Low-cost pricing strategies in leisure markets, *Tourism Management*, **33**: 249-256.
- [30] Talluri KT, van Ryzin G. 2004. *Theory and Practice of Revenue Management*. Springer: New York.
- [31] Teichert T, Shehu E, von Wartburg I. 2008. Customer segmentation revisited: The case of the airline industry. *Transportation Research Part A* **42**: 227-242.
- [32] Wooldridge JM. 2002. *Econometric analysis of cross section and panel data*. MIT Press, Cambridge MA, first edition.

TABLES AND FIGURES

Figure 1: Hypothesis 3: RM Strategy combines time and capacity components

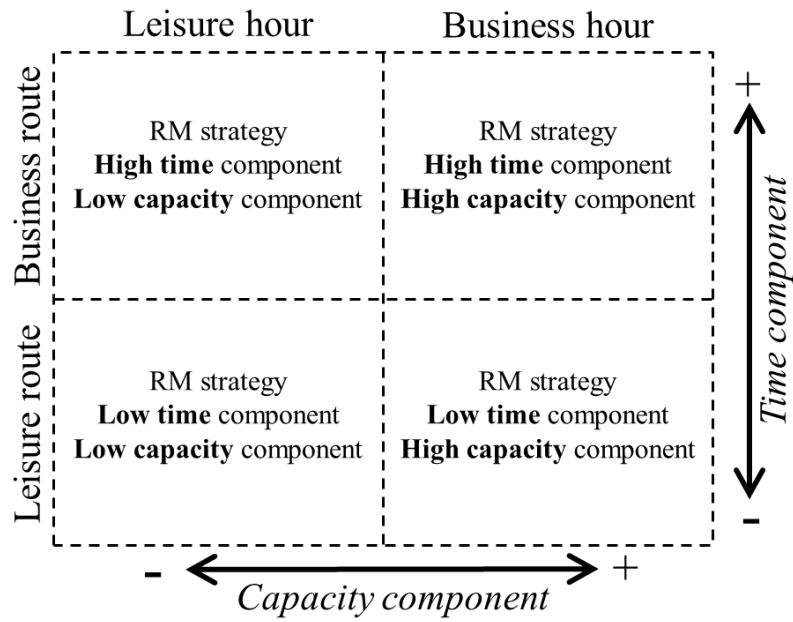


Figure 2: Business and Leisure Routes

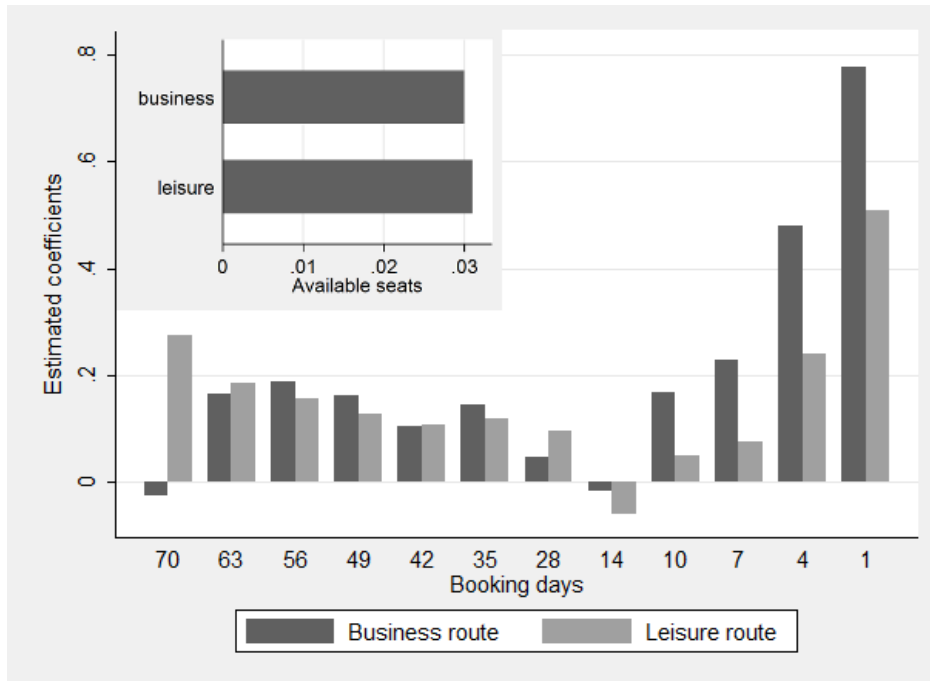


Figure 3: Business and Leisure Hours

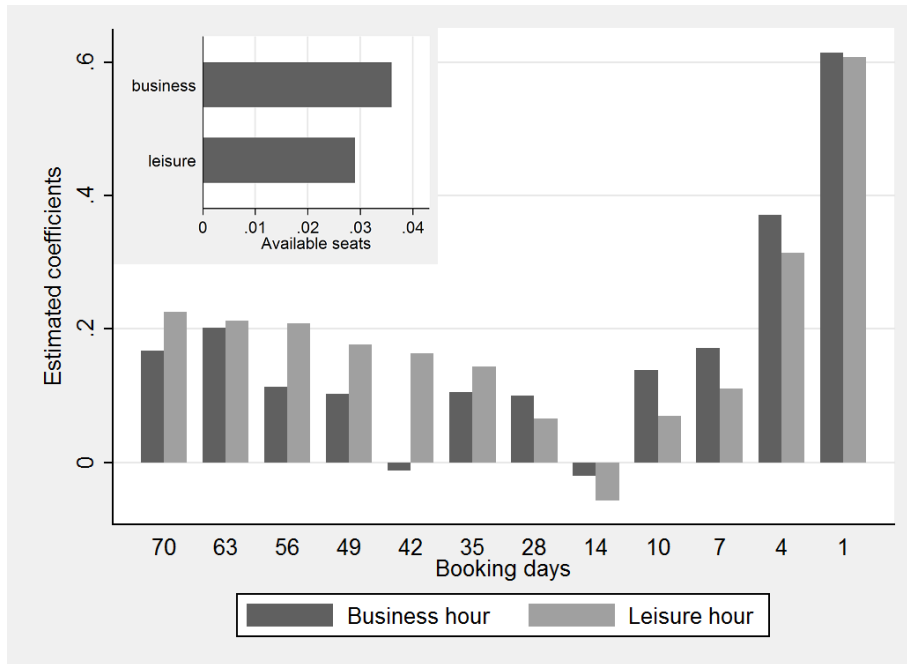


Table 1: Descriptive statistics and correlation matrix

	Mean	S.D.	(1)	(2)	(3)	(4)	(5)	(6)
(1) Fare	65.57	42.91						
(2) AvailableSeats	26.21	13.62	-0.60***					
(3) BookingDay	11.01	11.09	-0.18***	0.33***				
(4) BusinessRoute	0.48	0.50	-0.12***	-0.02***	-0.09***			
(5) BusinessHour	0.35	0.48	-0.06***	0.03***	-0.04***	-0.04***		
(6) BookOnHolidays	0.22	0.42	0.04***	0.02***	-0.01	-0.01**	-0.01**	
(7) LagMeanSlope	3.05	2.63	0.05***	-0.22***	-0.22***	0.02***	-0.01***	0.000

Note: Statistics computed for Seats<50. Pearson correlations are reported. Significant at *10%, ** 5%, and *** 1%.

Table 2: Descriptive statistics: mean *Fare* by *AvailableSeats* category

		AvailableSeats							
		1-9	10-19	20-29	30-39	40-49	≥50	1-49	Total
Route type	Business	116.7	74.0	53.0	40.5	33.8	16.6	60.1	27.0
	Leisure	123.9	86.5	65.3	52.7	44.6	23.8	70.5	35.8
Hour type	Business	119.4	79.4	56.7	44.7	35.7	16.0	62.3	25.7
	Leisure	121.0	80.8	60.9	48.3	41.9	23.5	67.4	35.5
Route/hour type	Biz/Biz	120.2	74.3	50.3	38.2	31.5	13.7	60.1	22.8
	Lei/Biz	118.6	84.4	61.7	49.2	38.7	18.3	64.4	28.6
	Biz/Lei	114.8	73.9	54.2	41.6	35.0	18.8	60.1	29.9
	Lei/Lei	126.3	87.6	67.5	55.0	48.6	28.2	74.3	41.1

Table 3: Descriptive statistics: mean *Fare* by *BookingDay*

		Booking Day										
		1	4	7	10	14	21	28	35	42	49-70	Total
All seats												
Route type	Business	81.8	55.0	36.9	31.8	22.5	19.4	17.1	16.2	15.2	14.5	27.0
	Leisure	92.0	63.2	47.6	42.9	33.6	29.6	27.0	25.5	24.4	23.8	35.8
Hour type	Business	79.8	54.6	36.7	32.4	22.5	18.8	16.8	15.5	14.5	13.9	25.7
	Leisure	91.4	61.9	45.9	40.6	32.0	28.9	26.0	25.0	23.8	23.3	35.5
Route/hour type	Biz/Biz	76.3	52.0	33.1	28.1	18.3	15.2	13.3	12.3	11.7	11.1	22.8
	Lei/Biz	83.3	57.2	40.3	36.7	26.6	22.1	20.0	18.5	17.1	16.5	28.6
	Biz/Lei	85.2	56.9	39.4	34.3	25.4	22.5	19.8	18.9	17.7	16.9	29.9
	Lei/Lei	98.3	67.5	52.9	47.3	38.9	35.2	32.2	30.7	29.6	29.1	41.1
		Booking Day										
Seats<50		1	4	7	10	14	21	28	35	42	49-70	Total
Route type	Business	94.2	69.0	53.4	47.1	39.7	39.5	41.8	50.5	55.5	66.6	60.1
	Leisure	101.8	76.1	64.9	60.4	55.4	57.4	61.1	61.7	66.7	77.1	70.5
Hour type	Business	95.1	71.9	56.7	51.0	44.2	44.1	48.5	49.4	52.1	56.0	62.3
	Leisure	99.5	72.9	60.8	55.7	50.3	52.3	55.0	61.1	67.1	80.6	67.4
Route/hour type	Biz/Biz	94.8	71.6	54.3	46.5	38.2	36.9	41.0	48.6	58.5	60.6	60.1
	Lei/Biz	95.4	72.2	58.7	54.7	49.1	49.4	53.2	49.9	49.1	53.1	64.1
	Biz/Lei	93.9	67.8	53.0	47.3	40.6	40.8	42.1	51.3	54.4	69.8	60.1
	Lei/Lei	105.7	78.5	68.7	64.0	59.2	61.8	65.2	66.9	73.5	84.5	74.3

Table 4: Estimates on Business and Leisure Routes

	Route type			
	Business		Leisure	
AvailableSeats	-0.030***	(0.001)	-0.031***	(0.001)
BookingDay1	0.777***	(0.047)	0.507***	(0.036)
BookingDay4	0.479***	(0.040)	0.240***	(0.033)
BookingDay7	0.230***	(0.034)	0.075***	(0.028)
BookingDay10	0.169***	(0.029)	0.051**	(0.023)
BookingDay14	-0.015	(0.024)	-0.059***	(0.019)
BookingDay28	0.046*	(0.025)	0.095***	(0.021)
BookingDay35	0.144***	(0.036)	0.119***	(0.030)
BookingDay42	0.106**	(0.053)	0.108***	(0.042)
BookingDay49	0.163***	(0.060)	0.129***	(0.047)
BookingDay56	0.189**	(0.087)	0.158***	(0.051)
BookingDay63	0.166**	(0.066)	0.186***	(0.058)
BookingDay70	-0.024	(0.108)	0.274***	(0.063)
Tobit residual	-0.000	(0.001)	0.001	(0.001)
DUMMIES:				
MonthOfBooking	YES		YES	
Number of obs.	27,716		30,870	
R2	0.617		0.542	
Excluded instruments:	2		2	
KP LM stat.	$\chi^2(2) = 258.9^{***}$		$\chi^2(2) = 343.3^{***}$	
Hansen J stat.	$\chi^2(2) = 0.040$		$\chi^2(2) = 0.000$	

Note: The dependent variable, *Fare*, is the natural log of the fare obtained from a query for one seat. Bootstrap Standard Errors (SE) are reported in parenthesis, clustered by route and week. 250 repetitions. Significant at *10%, ** 5%, and *** 1%.

Table 5: Estimates on Business and Leisure Hour

	Hour type			
	Business		Leisure	
AvailableSeats	-0.036***	(0.002)	-0.029***	(0.001)
BookingDay1	0.615***	(0.054)	0.608***	(0.036)
BookingDay4	0.371***	(0.048)	0.314***	(0.031)
BookingDay7	0.171***	(0.041)	0.110***	(0.026)
BookingDay10	0.138***	(0.033)	0.069***	(0.021)
BookingDay14	-0.021	(0.028)	-0.058***	(0.016)
BookingDay28	0.100***	(0.032)	0.065***	(0.018)
BookingDay35	0.105**	(0.043)	0.143***	(0.025)
BookingDay42	-0.013	(0.066)	0.163***	(0.035)
BookingDay49	0.102	(0.074)	0.177***	(0.040)
BookingDay56	0.113	(0.093)	0.209***	(0.053)
BookingDay63	0.202**	(0.089)	0.212***	(0.051)
BookingDay70	0.168	(0.115)	0.226***	(0.060)
Tobit residual	0.001	(0.001)	0.000	(0.001)
DUMMIES:				
MonthOfBooking	YES		YES	
Number of obs.	20,397		38,189	
R2	0.593		0.542	
Excluded inst.:	2		2	
KP LM stat.	$\chi^2(2) = 343.1^{***}$		$\chi^2(2) = 393.5^{***}$	
Hansen J stat.	$\chi^2(2) = 0.098$		$\chi^2(2) = 0.007$	

Note: The dependent variable, *Fare*, is the natural log of the fare obtained from a query for one seat. Bootstrap Standard Errors (SE) are reported in parenthesis, clustered by route and week. 250 repetitions. Significant at *10%, ** 5%, and *** 1%.

Table 6: Combining Route and Hour Dimensions

	Route/Hour type							
	Biz/Biz		Lei/Biz		Biz/Lei		Lei/Lei	
AvailableSeats	-0.035***	(0.002)	-0.035***	(0.002)	-0.027***	(0.002)	-0.029***	(0.001)
BookingDay1	0.795***	(0.081)	0.529***	(0.080)	0.794***	(0.053)	0.484***	(0.047)
BookingDay4	0.538***	(0.071)	0.282***	(0.068)	0.471***	(0.046)	0.210***	(0.041)
BookingDay7	0.288***	(0.057)	0.109*	(0.059)	0.216***	(0.040)	0.053	(0.036)
BookingDay10	0.221***	(0.048)	0.094**	(0.048)	0.152***	(0.031)	0.025	(0.029)
BookingDay14	0.020	(0.042)	-0.039	(0.040)	-0.027	(0.027)	-0.066***	(0.024)
BookingDay28	0.078	(0.053)	0.103**	(0.046)	0.020	(0.029)	0.089***	(0.023)
BookingDay35	0.101	(0.070)	0.090	(0.060)	0.139***	(0.039)	0.123***	(0.034)
BookingDay42	-0.085	(0.120)	-0.003	(0.092)	0.143***	(0.055)	0.142***	(0.045)
BookingDay49	0.138	(0.108)	0.050	(0.111)	0.165**	(0.070)	0.148***	(0.055)
BookingDay56	0.183	(0.132)	0.023	(0.134)	0.168	(0.111)	0.180***	(0.060)
BookingDay63	0.164	(0.120)	0.197	(0.136)	0.160**	(0.075)	0.176***	(0.068)
BookingDay70	-0.076	(0.212)	0.289*	(0.169)	0.004	(0.139)	0.256***	(0.068)
Tobit residual	0.001	(0.001)	0.002	(0.002)	-0.001	(0.001)	0.001	(0.001)
DUMMIES:								
MonthOfBooking	YES		YES		YES		YES	
Number of obs.	9,069		11,328		18,647		19,542	
R2	0.644		0.549		0.607		0.549	
Excluded instr.:	2		2		2		2	
KP LM stat.	$\chi^2(2) = 163.7^{***}$		$\chi^2(2) = 188.8^{***}$		$\chi^2(2) = 165.3^{***}$		$\chi^2(2) = 242.0^{***}$	
Hansen J stat.	$\chi^2(2) = 0.344$		$\chi^2(2) = 0.000$		$\chi^2(2) = 0.000$		$\chi^2(2) = 0.023$	

Note: The dependent variable, *Fare*, is the natural log of the fare obtained from a query for one seat. Bootstrap Standard Errors (SE) are reported in parenthesis, clustered by route and week. 250 repetitions. Significant at * 10%, ** 5%, and *** 1%.

TECHNICAL APPENDIX

Estimates from the Tobit and the first stage of the IVFE regressions for the first two hypotheses are presented in Tables 7 and 8, respectively. These are not meant for publication.

Table 7: Tobit model (dependent variable: *AvailableSeats*)

	Route type				Hour type			
	Business		Leisure		Business		Leisure	
LagMeanSlope	-2.650***	(0.157)	-2.301***	(0.123)	-2.664***	(0.138)	-2.290***	(0.111)
BookingDay1	-31.700***	(0.496)	-29.847***	(0.468)	-34.008***	(0.542)	-30.236***	(0.382)
BookingDay4	-26.009***	(0.494)	-24.966***	(0.463)	-26.993***	(0.526)	-25.998***	(0.389)
BookingDay7	-20.575***	(0.457)	-20.494***	(0.429)	-21.526***	(0.486)	-20.267***	(0.354)
BookingDay10	-15.616***	(0.401)	-15.889***	(0.382)	-17.649***	(0.453)	-16.025***	(0.321)
BookingDay14	-9.676***	(0.341)	-10.238***	(0.319)	-10.361***	(0.358)	-9.851***	(0.269)
BookingDay28	9.299***	(0.381)	9.438***	(0.322)	10.069***	(0.419)	9.059***	(0.280)
BookingDay35	18.532***	(0.583)	16.823***	(0.481)	18.375***	(0.641)	17.236***	(0.416)
BookingDay42	25.211***	(0.747)	23.645***	(0.637)	25.494***	(0.830)	23.839***	(0.537)
BookingDay49	30.104***	(0.985)	28.386***	(0.756)	30.465***	(1.022)	28.602***	(0.685)
BookingDay56	31.194***	(1.236)	30.923***	(0.861)	30.837***	(1.187)	31.261***	(0.786)
BookingDay63	31.514***	(1.362)	31.946***	(0.939)	31.518***	(1.229)	32.041***	(0.899)
BookingDay70	31.915***	(1.542)	34.491***	(1.208)	32.985***	(1.328)	33.699***	(1.059)
Constant	67.691***	(3.624)	78.259***	(3.299)	70.047***	(9.629)	65.036***	(2.896)
DUMMIES:								
DayOfWeekOfBooking	YES		YES		YES		YES	
WeekNumberOfDept	YES		YES		YES		YES	
HourOfDept	YES		YES		YES		YES	
Route	YES		YES		YES		YES	
Number of obs.	146,267		147,854		119,775		174,346	
Pseudo R2	0.178		0.175		0.184		0.170	

Note: Standard errors clustered by route and week. Significant at *10%, ** 5%, and *** 1%.

Table 8: First stage (dependent variable: *AvailableSeats*)

	Route type				Hour type			
	Business		Leisure		Business		Leisure	
HolidayPeriod	0.386***	(0.085)	0.109***	(0.029)	0.287***	(0.069)	0.155***	
LagMeanSlope	-2.320***	(0.016)	-2.212***	(0.007)	-2.364***	(0.019)	-2.164***	(0.001)
BookingDay1	-30.973***	(0.111)	-29.413***	(0.040)	-31.272***	(0.100)	-29.855***	(0.047)
BookingDay4	-24.924***	(0.114)	-24.516***	(0.044)	-25.739***	(0.106)	-24.666***	(0.041)
BookingDay7	-19.836***	(0.096)	-20.158***	(0.034)	-20.605***	(0.088)	-19.895***	(0.036)
BookingDay10	-14.372***	(0.107)	-15.463***	(0.040)	-15.595***	(0.099)	-14.952***	(0.029)
BookingDay14	-9.163***	(0.070)	-10.025***	(0.024)	-9.703***	(0.064)	-9.588***	(0.024)
BookingDay28	8.442***	(0.085)	9.182***	(0.027)	9.156***	(0.080)	8.719***	(0.023)
BookingDay35	16.392***	(0.138)	16.404***	(0.040)	16.465***	(0.116)	16.568***	(0.034)
BookingDay42	22.188***	(0.175)	23.001***	(0.053)	22.854***	(0.159)	22.787***	(0.045)
BookingDay49	26.091***	(0.271)	27.560***	(0.066)	27.260***	(0.220)	27.229***	(0.055)
BookingDay56	27.272***	(0.300)	29.999***	(0.083)	27.394***	(0.302)	29.854***	(0.060)
BookingDay63	27.571***	(0.489)	30.895***	(0.112)	28.009***	(0.426)	30.554***	(0.068)
BookingDay70	27.711***	(0.437)	33.309***	(0.119)	29.183***	(0.409)	32.046***	(0.068)
Tobit residual	0.838***	(0.005)	0.955***	(0.002)	0.860***	(0.004)	0.931***	(0.001)
DUMMIES:								
MonthOfBooking	YES		YES		YES		YES	
Number of obs.	27,716		30,870		20,397		38,189	
R2	0.948		0.984		0.952		0.976	

Note: Standard errors clustered by route and week. Significant at *10%, ** 5%, and *** 1%.