

ESSAYS ON PRICING UNDER UNCERTAINTY

A Dissertation

by

DIEGO ALFONSO ESCOBARI URDAY

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

May 2008

Major Subject: Economics

ESSAYS ON PRICING UNDER UNCERTAINTY

A Dissertation

by

DIEGO ALFONSO ESCOBARI URDAY

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Approved by:

Chair of Committee,	Li Gan
Committee Members,	Hae-Shin Hwang
	Steven Puller
	Ximing Wu
Head of Department,	Lawrence J. Oliver

May 2008

Major Subject: Economics

ABSTRACT

Essays on Pricing Under Uncertainty. (May 2008)

Diego Alfonso Escobari Urday, B.A., Universidad Católica Boliviana;

M.S., Texas A&M University

Chair of Advisory Committee: Dr. Li Gan

This dissertation analyzes pricing under uncertainty focusing on the U.S. airline industry. It sets to test theories of price dispersion driven by uncertainty in the demand by taking advantage of very detailed information about the dynamics of airline prices and inventory levels as the flight date approaches. Such detailed information about inventories at a ticket level to analyze airline pricing has been used previously by the author to show the importance of capacity constraints in airline pricing. This dissertation proposes and implements many new ideas to analyze airline pricing. Among the most important are: (1) It uses information about inventories at a ticket level. (2) It is the first to note that fare changes can be explained by adding dummy variables representing ticket characteristics. Therefore, the load factor at a ticket level will lose its explanatory power on fares if all ticket characteristics are included in a pricing equation. (3) It is the first to propose and implement a measure of *Expected Load Factor* as a tool to identify which flights are peak and which ones are not. (4) It introduces a novel idea of comparing actual sales with average sales at various points prior departure. Using these deviations of actual sales from sales under average conditions, it presents is the first study to show empirical evidence of peak load pricing in airlines. (5) It controls for potential endogeneity of sales using dynamic panels.

The first essay tests the empirical importance of theories that explain price dispersion under costly capacity and demand uncertainty. The essay calculates a mea-

sure of an *Expected Load Factor*, that is used to calibrate the distribution of demand uncertainty and to identify which flights are peak and which ones are off-peak. It shows that different prices can be explained by the different selling probabilities. The second essay is the first study to provide formal evidence of *stochastic* peak-load pricing in airlines. It shows that airlines learn about the demand and respond to early sales setting higher prices when expected demand is high and more likely to exceed capacity.

To my little angel Valentina and my wife Carolina.

ACKNOWLEDGMENTS

I am deeply grateful for guidance and support from my committee chair, Dr. Li Gan, and the help of my committee members. Special thanks to Dr. Hae-Shin Hwang who gave me invaluable guidance during my graduate studies at Texas A&M and to Dr. Dennis W. Jansen for his trust and support.

I thank Dr. James Dana for his careful revision of the current version of Chapter II. I also want to show my appreciation to Dr. Volodymyr Bilotkach, Dr. Hae-Shin Hwang, Dr. Paan Jindapon, Dr. Eugenio Miravete, Dr. Carlos Oyarzún, Dr. Joon Park, Dr. Tomas R. Saving, and Dr. Ximing Wu for the comments received along the elaboration of this document. This dissertation greatly benefited from their feedback.

I thankfully acknowledge financial support from the Department of Economics, Dr. Joan Shin, The Lynde and Harry Bradley Foundation, STATA Corporation, contributors to the Maurice Fellowship, and the Private Enterprise Research Center. I also thank Tyffanne Rowan and Pat Nelson for administrative assistance.

Thanks to my friends in College Station; especially Paan Jindapon, Carlos Oyarzún, and Diego Vacaflores. I would like to thank my mother and my father for giving me education opportunities and for teaching me the most important things in life. Thank you for the support, love, and patience of my wife, Carolina. Thanks to God, for blessing my family with little angel Valentina.

Finally, and most of all, I acknowledge a deep debt of gratitude to Dr. Gan. Thank you for opening me your door and sharing your knowledge and enthusiasm towards economics. The completion of this dissertation would not have been possible without him.

TABLE OF CONTENTS

CHAPTER		Page
I	INTRODUCTION	1
II	PRICE DISPERSION UNDER COSTLY CAPACITY AND DEMAND UNCERTAINTY	4
	A. Introduction	4
	B. The data and its main characteristics	9
	1. Fares, inventories and ticket characteristics	11
	2. Representative fare	16
	C. The empirical model	19
	1. Oligopoly model of costly capacity and demand un- certainty	19
	2. Modeling demand uncertainty	26
	3. Calibrating the probability density of demand uncertainty	28
	4. Estimated equation and interpretation	29
	D. Results of the empirical analysis	34
	E. Conclusions	44
III	STOCHASTIC PEAK-LOAD PRICING WITH REAL-TIME DEMAND LEARNING IN THE U.S. AIRLINE INDUSTRY . .	47
	A. Introduction	47
	B. Airline pricing under demand uncertainty	52
	1. Pricing without demand learning	53
	2. Pricing with demand learning	55
	3. Peak-load pricing in airlines	57
	a. Systematic peak-load pricing	57
	b. Stochastic peak-load pricing	58
	C. Data	60
	D. Empirical model	65
	1. Ex-ante distribution of demand uncertainty	66
	a. Sold-out probabilities	67
	b. Expected number of seats sold	68
	c. Calibration	68
	2. Learning the stochastic demand	69

CHAPTER	Page
<ul style="list-style-type: none"> a. Nonparametric estimation of expected sales . . . b. Endogenous threshold estimation c. Dynamic panel with exogenous threshold 	<ul style="list-style-type: none"> 70 71 83
E. Conclusions	87
IV SUMMARY	90
REFERENCES	93
APPENDIX A	99
APPENDIX B	103
APPENDIX C	104
VITA	107

LIST OF TABLES

TABLE		Page
1	Summary statistics	33
2	Estimation results for the censored normal	35
3	Summary of robustness checks	36
4	Nonlinearities in time	39
5	Sensitivity of sold-out probabilities	41
6	Elasticities	43
7	Summary statistics	72
8	Tests for threshold effects: Ratio	78
9	Threshold estimates: Nonparametric	79
10	Regression estimates for the single threshold model: Nonparametric .	81
11	Regression estimates: GMM dynamic panel	87

LIST OF FIGURES

FIGURE		Page
1	Fares and load factors at different days from departure	12
2	Average sales at different days from departure	13
3	Nonparametric regression of daily sales on days prior to departure . .	13
4	Probability that seat h is sold (MCO-LGA)	39
5	Fares and effective cost of capacity at different days prior to departure (AA 323 ATL-DFW)	41
6	Average and standard deviation of fares	62
7	Fares and load factors (Delta 1588 ATL-SJC)	62
8	Estimated nonparametric expected sales	71
9	Confidence interval construction in single threshold model	80
10	Fares and load factors (United 167 BOS-LAX)	99
11	Fares and load factors (Continental 2408 CLE-ORD)	99
12	Fares and load factors (United 7156 IAD-CLE)	100
13	Fares and load factors (Continental 194 LAX-IAH)	100
14	Fares and load factors (American 596 MIA-BOS)	101
15	Fares and load factors (American 1341 MIA-MSY)	101
16	Fares and load factors (US 185 PHL-ORD)	102
17	Fares and load factors (United 579 PIT-ORD)	102

CHAPTER I

INTRODUCTION

Dynamic pricing or most commonly known *yield management*, is used to describe pricing and inventory control decisions. It is important in industries that deal with perishable products such as airlines, where unsold seats perish when the flight leaves the gate. Dana [21] explains that *yield management* in airlines is used to (1) deal with costly capacity and demand uncertainty, (2) implement price discrimination, and (3) implement peak-load pricing. Because of the lack of appropriate data, there exists few empirical understanding on how airlines are actually setting fares and its dynamics as the flight date approaches. This dissertation sets to provide empirical evidence supporting these three roles of dynamic pricing. This analysis takes advantage of a unique U.S. airline's panel disaggregated at the ticket level that contains the evolution of offered fares and seat inventories over a period of 103 days for 228 domestic flights that departed on June 22nd, 2006.

Chapter III tests the empirical importance of the price dispersion predictions of the Prescott [48], Eden [25], and Dana [21] models. Building on Dana [21], it constructs a theoretical model with two empirical predictions with equilibrium price dispersion derived in a setting with costly capacity and demand uncertainty. The theoretical section shows that different fares observed for the same flight can be explained by the different selling probabilities attached to each of these fares. Using information from the *T-100* of the *Bureau of Transportation and Statistics*, this chapter calculates a measure of an *Expected Load Factor*. This measure is then used to

This dissertation follows the style of Journal of Economic Theory.

calibrate the distribution of demand uncertainty under the assumption of normally distributed demand states. The theoretical model, as well as Dana [21], has two empirical predictions: (1) higher prices can be explained by lower selling probabilities, (2) this effect is larger in more competitive routes. Despite the wide applications of this type of models of costly capacity and demand uncertainty to several important market phenomena, there exists little empirical evidence supporting these models. Using the panel of U.S. airline fares and seat inventories, we find evidence that strongly supports both predictions of the models. Higher fares observed close to departure can be explained by aircrafts having less available seats. Moreover, the cost of capacity for those seats is larger than for the seats sold early. Using the *Herfindahl-Hirshman Index* to capture the market structure, it is also shown that the effect of costly capacity on fares is greater in more competitive markets. After controlling for the effect of aggregate demand uncertainty on fares and under the assumption that carriers do not learn about the state of the demand as sales progress, we also obtain evidence of second degree price discrimination in the form of advance-purchase discounts.

Chapter III shows that airlines learn about the demand as sales progress and the departure date nears. Demand learning for airlines is important because in flights when the final demand results to be low, unsold tickets are of little value after departure. Moreover, in flights where the final demand is large, carriers may have to give up important profits when they run out of tickets and some relatively high willingness-to-pay consumers that arrive late are not able to find a seat. Under a price sensitive demand, *stochastic* peak-load pricing suggests that at any point prior departure airlines should set higher fares in expected *peak* flights, where demand is more likely to exceed capacity. Furthermore, in order to promote sales and to avoid having empty seats after departure, lower fares should be set in expected *off-peak* flights. The chapter starts by calibrating the distribution of demand uncertainty un-

der a price commitments assumption. In this scenario, the distribution of demand uncertainty is sufficient to explain why fares increase as departure date nears. Then to see the impact of demand learning on fares, we need a measure of the status of actual sales as compared to sales under average conditions. To do this we use non-parametric techniques to construct a threshold variable to identify different expected demand states at different points prior departure. This threshold variable is utilized to dictate the regime shift in a panel endogenous threshold model. Consistent with the *stochastic* peak-load pricing predictions, the results show that higher fares are set in the *peak* regime when expected demand is large. Lower fares are set in the *off-peak* regime when demand is expected to fall short. To control for potential endogeneity problems regarding part sales levels, the chapter also runs some GMM dynamic panel specifications following Arellano and Bond [1], Arellano and Bover [2], and Blundell and Bond [8].

Finally, chapter IV summarizes the results.

CHAPTER II

PRICE DISPERSION UNDER COSTLY CAPACITY AND DEMAND UNCERTAINTY

A. Introduction

It is widely observed that prices of homogeneous goods within the same market exhibit price dispersion. Some of the most recent evidence includes retail prices for prescription drugs in Sorensen [51], and internet electronic equipment markets in Baye and Morgan [5]. Various models, including search frictions, information asymmetries, and bounded rationality, have been proposed to explain this phenomenon. Here we seek to establish the empirical importance of the price dispersion predictions in the Prescott [48], Eden [25] and Dana [21]’s models.

Prescott [48] considers an example of hotel rooms where sellers set prices before they know the number of buyers, then the equilibrium prices will be dispersed; lower-priced units will sell with higher probability, while higher-priced units will sell with lower probability. Hence, sellers face a tradeoff between price and the probability of making a sell. This same tradeoff is observed in Eden [25], who formalizes Prescott’s model in a setting where consumers arrive sequentially, observe all offers and after buying the cheapest available offer they leave the market. He derives an equilibrium that exhibits price dispersion even when sellers are allowed to change their prices during trade and have no monopoly power. This flexible price version of the Prescott model, developed in Eden [25] and Lucas and Woodford [45], is known as the Uncertain and Sequential Trade (*UST*) model. Dana [21] extends the Prescott model with price commitments for perfect competition, monopoly, and oligopoly and shows that firms offer output at multiple prices. In the oligopoly equilibrium, the market

distribution of prices converges to the Prescott’s distribution as the number of firms approaches to infinity. Moreover, as competition is greater, average price level falls and price dispersion increases. As explained in Eden [28], from the positive economics point of view it does not matter whether prices in the Prescott’s model flexible or rigid. From the point of view of the seller and this paper, both will have the same resulting allocation. In this paper, both the flexible and the rigid version of the model are commonly referred as Prescott-Eden-Dana (*PED* hereafter) models.

Versions of the *PED* model have been applied to solve a variety of economic phenomena, such as wage dispersion and market segmentation (Weitzman [57]), procyclical productivity (Rotemberg and Summers [50]), the role of inventories (Bental and Eden [6]), real effect of monetary shocks (Lucas and Woodford [45]; Eden [26]), destructive competition in retail markets (Deneckere, Marvel, and Peck [23]), advance purchase discounts (Dana [19]), stochastic peak-load pricing (Dana [20]), gains from trade (Eden [28]) and seigniorage payments (Eden [29]). Despite its wide applications, few papers test the empirical predictions of the *PED* models.

This chapter provides a formal test of the *PED* models while helping to explaining price dispersion in the airline industry, which is considered to have one of the most complex pricing systems in the world. We take advantage of a unique U.S. airlines’ panel disaggregated at passenger level that contains the evolution of fares and inventories of seats over a period of 103 days for 228 domestic flights departing on June 22, 2006. The data collection resembles experimental data which controls for most of the product heterogeneities observed in the industry. This represents the perfect control for fences that segment the market allowing our analysis to explain the use of seat-inventory control just under demand uncertainty, costly capacity and price commitments.

Moreover, airlines represent the perfect environment to test the price dispersion

under demand uncertainty and costly capacity. First, air tickets expire at a point in time; once the plane departs carriers can no longer sell tickets. Second, capacity is fixed and can only be augmented at a relatively high marginal cost. Once carriers start selling tickets they are unlikely to change the size of the aircraft.¹ This implies that we can focus on the demand side uncertainty without having to worry about any the uncertainty in the supply given our time frame of study. Moreover, as in the *PED* models, after we control for ticket restrictions that screen costumers, all airplane seats are the same and buyers have unit demands. In order to explain price dispersion we enlarge the definition of airplane seats by an additional 'selling probability' dimension. Once this is achieved, although prices themselves may be dispersed, this dispersion can be explained by appropriately rescaling the price of each unit by its selling probability.

At the risk of over-making this point, let us consider the following example of a perfectly competitive market with zero profits. Each time a carrier sells a seat, the expected marginal revenue is set to be equal to the marginal cost. Because of demand uncertainty, airlines hold inventories of seats that are sold only some of the times. For those seats that are sold only when demand is high, fares must be set higher to compensate for the lower probability of sale. In this chapter we develop a measure of the different selling probabilities. Even though uncertainty is coming from the demand side, we follow the *PED* models and represent this by adjusting the marginal cost of capacity, or ex-ante shadow cost, by these selling probabilities.

By dividing the constant unit cost of capacity by the probability of sale, we obtain the Effective Cost of Capacity (*ECC*), and then we measure the impact of *ECC* on fares. As predicted by Prescott [48] and Eden [25], *ECC* should have a

¹None of the 228 flights in the sample changed the aircraft size.

positive effect on fares. Moreover, as predicted in Dana [21], this effect should be greater in more competitive markets. In this chapter we provide evidence supporting both predictions. On average, a 1 percent decrease in the probability of sale would lead to a 0.219 percent increase in prices. Moreover, this effect was found to be larger in more competitive markets. The reason is straight forward, in a perfectly competitive market where firms have no markups; every dollar increase in the *ECC* will be transferred to prices. On the other hand, in less competitive markets, part of the increase in the *ECC* will be absorbed by the markup.

The findings in this chapter can be additionally motivated as an example of a spot market subject to demand uncertainty and opened to advance purchases. The standard formulation of a spot market subject to uncertain excess demand, assumes either implicitly or explicitly, a tatonnement process that restricts trade until the market-clearing price is found. As pointed out in Dana [21], a spot market subject to price commitments should be opened to advance purchases. As we approach the departure date, the dynamics of fares and inventories in a flight is an example of how the market clearing price is achieved without having to restrict trade in the resolution of uncertainty in the demand. Along the chapter we discuss how the analysis carried out resembles a spot market with price commitments.

By helping to explain one of the sources of price dispersion, this chapter has an important implication for the airline industry as well. Borenstein and Rose [10] calculated that the expected absolute difference in fares between two passengers on a route is 36 percent of the airline's average ticket price. One important source of this price dispersion is the existence of intrafirm price dispersion due to advance-purchase discounts (*APD*). Substantial discounts are generally available to travelers who are willing to purchase tickets in advance. This kind of pricing practices can promote efficiency by expansions in output when demand is elastic or may be the only way

for a firm to cover large fixed costs. Gale and Holmes [35] justify the existence of *APD* in a monopoly model with capacity constraints and perfectly predictable demand. They show that firms using *APD* can divert demand from peak period to off-peak period and achieve a profit-maximizing method of selling tickets. In a similar setting, but with demand uncertainty, Gale and Holmes [34] show that *APD* can promote efficiency by spreading consumers evenly across flights before timing of the peak period is known. In competitive markets, Dana [19] finds that firms may offer *APD* when individual and aggregate consumer demand is uncertain and firms set prices before demand is known. The *PED* models that we test, explain why carriers offer lower priced seats to ‘earlier’ purchasers.² Our results show that one source of the price variation found by Borenstein and Rose [10] comes from the fact that carriers face capacity constraints and have to deal with uncertainty in the demand. Moreover, we find that this source of price dispersion is greater in more competitive markets, result consistent with Borenstein and Rose [10], who also found greater price dispersion in more competitive markets. Our findings represent a refinement of Borenstein and Rose [10]. They attribute this result to price discrimination using a model of monopolistic-competition with certain demand. We argue that if demand uncertainty is considered, part of this price dispersion can be explained by carriers dealing with capacity costs and uncertain demand. The present chapter is the first empirical approach, to my knowledge, that includes uncertainty in the determination of prices in the airline industry.

Despite a number of applications of the *PED* models, few papers test the empirical predictions of the model. Eden [27] provides a test and finds a negative re-

²Note that the term ‘earlier’ used refers to the case when passengers who buy before other passengers, rather than a temporal dimension. Travelers purchasing seats even long before departure may not benefit from *APD* if most of the seats in the airplane have already been sold.

lationship between inventories and output. However, as pointed in the same article, this negative relationship is not necessarily an outcome of the *PED* models. In fact, other models, such as the model of inventory control, would generate the same prediction. Wan [56] tests part the models using data from online book industry. She tests the effect of stock-out probability and search cost on price dispersion and finds evidence that higher stock-out probabilities are associated with higher prices. The *PED* models requires capacity (how many books to store or how many seats on an airplane) to be fixed in the short run. This is less likely to be true for the online book industry than for the airline industry. In addition, Wan [56] does not test the effect of competition on the prices.³

The organization of this chapter is as follows. Section B describes the data and its characteristics. The theoretical motivation and the empirical specification are presented in Section C; first explaining the theoretical motivation, then showing how we model demand uncertainty with an application. Section D explains the empirical results. Finally, Section E concludes the chapter.

B. The data and its main characteristics

The main data source in this chapter comes from data collected on the online travel agency *Expedia.com* for flights that departed on June 22, 2006. It is a panel with 228 cross section observations during 35 periods making a total of 7980 observations. Each cross section observation corresponds to a specific carrier's non-stop flight between a pair of departing and destination cities. The data across time has one observation every three days. The first was gathered 103 days prior to departure, the second 100 days and so on until 7, 4, and 1 day(s) prior to departure, making the 35 observations

³Bilotkach [7] mentions the potential role of the *PED* models in explaining price airline dispersions, but his dataset does not allow him to formally test the model.

in time per flight. As in Stavins [52], the date of the flight is a Thursday to avoid the effect that weekend travel could have. The carriers considered are American, Alaska, Continental, Delta, United and US Airways. The number of flights per carrier was chosen to make sure the share of each of these carriers on the dataset is close to its share on the US airlines' market. For each flight at each time period, this dataset gives us the cheapest available economy class fare and the number of seats that have been sold up to that period.

To calculate the sold out probabilities, the analysis uses a second dataset collected also from *Expedia.com*. Most airlines and online travel agencies do not display sold-out flights on their websites. The reason, according to Roman Blahoski, spokesman of Northwestern, is that they do not want to disappoint travelers. Keeping the online display simple may also be a motive, and according to Dan Toporek, spokesman of *Travelocity.com*, “showing sold-out flights alongside available flights could be confusing.”⁴ Regardless of the reason, this fact allows us to get the information about the sold out probability in each of the routes. We initially make a census of all the available nonstop flights in each of the 81 routes used in the first dataset for seven days from February 2 to February 8 in 2007. The total number of flights is 5,881. The collection is done couple of weeks before the beginning of February when we expect that no flights have yet been sold out, hence *Expedia.com* should show them all. Then, for each of these seven days of the week we check *Expedia.com* once again late at night the day before departure to see whether each of the flights has still tickets available. If the flight is no longer there, we assume that it has already sold all its tickets. This procedure permits us to calculate the sold out probabilities for each of the routes. We interpret this sold out probability as a lower bound because

⁴Both quotes are from David Grossman, “Gone today, here tomorrow,” *USA Today*, August 2006.

i) February is not necessarily a high demand period, and ii) because there may still be some tickets sold the day of the flight that did not enter the computation.

A second important source of data is the $T - 100$ data from the *Bureau of Transportation Statistics*. From the $T - 100$ we obtain a panel containing the yearly average load factors at departure for the same routes as in the main dataset over the period 1990 to 2005. This helped us to calculate the expected number of tickets sold in each route. Moreover, this $T - 100$ gave us the number of enplanements at each endpoint airport to construct some of the instruments.

1. Fares, inventories and ticket characteristics

A typical flight in the sample looks like the American Airlines Flight 323 from Atlanta, GA (ATL) to Dallas-Forth Worth, TX (DFW) depicted in figure 1.⁵ The best way to look at the evolution of seat inventories, in a way that is comparable between flights, is to look at the load factor, defined as the ratio of seats sold at each point in time prior to departure to total seats in the aircraft.⁶ Load factor will go from zero when the plane is empty to one when it is full. In figure 1, the load factor for this flight increases from 0.2, 103 days prior to departure to 0.88 with one day left to depart. The increase is not necessarily monotonic as can be observed when moving from 34 to 31 days prior to departure. This is because some tickets may have been reserved and never bought or maybe bought and cancelled later. In this flight fares initially look fairly stable between \$114 and \$144, but they have a sharp increase during the

⁵By request of Dr. Steven Puller, appendix A shows some additional flights.

⁶Airline's literature defines load factor only once the plane has departed and as the percentage of seats filled with paying passengers. It is calculated by dividing revenue-passenger miles by available seat miles. Here the load factor is defined at each point in time as the flight date approaches. Escobari [30] also uses the ratio of seats sold to total seats at the ticket level to obtain some evidence of peak-load pricing.

last two weeks before departure and peak its maximum at \$279 the last day.




Fig. 1. Fares and load factors at different days from departure

Figure 2 depicts the average fares for the 228 flights in the sample for each of the days prior to departure. The most important characteristic is how fares trend upwards from an average of \$258, 103 days prior to departure to an average of \$473, the last day prior to departure. This means that average fares almost doubled during the period of study.

Figure 3 shows the nonparametric regression of daily sales (as percentage of total capacity) on days prior to departure using 7752 observation over the 228 flights. The bandwidth of 1.14 days is obtained by least squares cross-validation. The figure suggests that as the flight date approaches, more seats get sold. The majority of the seats are being sold during the last month and there seems to be a drop in sales during the last few days close to departure.

It is important to know that inventories evolve not just as a result of sales at the one-way, non-stop flight we are considering. Seats for each city pairs in the sample can be sold as part of a larger trip or as part of a round trip with an extremely large




Fig. 2. Average sales at different days from departure




Fig. 3. Nonparametric regression of daily sales on days prior to departure

amount of possible options. Along this chapter we will be looking at the carriers' optimal pricing decision for the one-way, non-stop flight of June 22 and this will have its own dynamics. This detail is implicit in these types of datasets that look at non transaction data like Stavins [52], McAfee and Velde [46], Chen [17].

The fares used in this chapter are the cheapest fare available at each point in time for a seat in economy class. The cheapest economy class fare at each point in time prior to departure is just the search results found by *Expedia.com* for any other online travel agency or carrier's website when searching for the fare of a given flight.⁷ It is worth pointing out that every time a carrier changes its prices, it also changes some characteristics associated with this fare. To show how fares can be explained with irrelevant ticket characteristics, let's look again at the fares of American Airlines Flight 323 depicted in figure 1. In this example, when the price decreased from \$134 to \$114 between 103 (March 11th) and 100 (March 14th) days prior to departure, the ticket characteristics changed from a 10- to a 14-days-in-advance-purchase-requirement, it changed the first-day-of-travel-requirement from February 11th to March 14th, and some blackout dates were included along with changes in day-and-time-of-the-flight restrictions. None of these restrictions have a real impact on the purchase decision or the effective quality of the ticket unless the traveler knows these characteristics and carries out a detailed analysis evaluating the possibility of canceling the flight later on. If the ticket is bought either 103 or 100 days prior the flight day, having a 10- or a 14-days-in-advance-purchase-requirement is irrelevant. If the passenger has already

⁷Different types of fares sometimes available are the ones travel agencies directly negotiate with airline partners. One example is *ClearanceFares* and *FlexSaver* offered by *Hotwire.com*. These fares come with substantial discounts but impose additional restrictions and involve higher uncertainty. They do not allow changes or refunds and do not allow the traveler to pick the flight times or airline at the moment of booking. Additionally, the traveler cannot earn frequent flyer miles and the fare paid does not guarantee a specific arrival time. Delays can be greater than a day.

decided to fly on June 22 and is buying the ticket either on March 11 or March 14, the first-day of-travel-requirement of February 11 or March 14 are irrelevant as well. Blackouts and day-and-time-of-the-flight restrictions are only important if the traveler decides to change the day of the flight and the new date happens to be exactly in one of the blackout dates. Changing dates will be anyway subject to further restrictions on the tickets available in the new date, and a penalty of 50 plus the differences in fares. The fact is that really few passengers actually know these restrictions even exist since you cannot modify them online and are not printed out in the ticket or the e-ticket. This example also shows that even if the ticket is bought with more than 21 days in advance, it does not necessarily mean it gets the discount of a 21-days-in-advance-purchase-requirement. The same goes along with other restrictions; even if the traveler is willing to accept any blackout or purchase a non-refundable ticket, if only refundable tickets are available, she may well end up buying it, sometimes without knowing the extra benefits. Stavins [52], McAfee and te Velde [46], and Chen [17] also look at these type of fare changes, but do not mention this point. The key point here is that these ticket characteristics that change along with fares are irrelevant for the travelers, and if buying online, it is sometimes impossible for the buyer to change these characteristics. Carriers change these irrelevant tickets characteristics to justify the changes in fares. They do not want to charge two different fares for exactly the same product just because the transactions occurred at different points in time, even if these differences in the product do not have any impact on the purchase decision. In the empirical test we control for the ticket restrictions that do have an impact on the quality of the ticket. Again, a similar assumption has been implicitly made in McAfee and te Velde [46] and Chen [17] and just look at the variations in fares without keeping track of the corresponding variation in irrelevant ticket characteristics. Stavins [52] omits most of these irrelevant ticket characteristics but includes dummy

variables for some advance purchase restrictions. These dummy variables may explain changes in fare, but they do not reflect the underlying force behind why carriers offer advance purchase discounts in the first place. As we argue in this chapter, once the relevant ticket characteristics are controlled for, the key underlying force is seats inventories.

2. Representative fare

A typical concern among people who search to buy tickets online is to know whether or not the fare paid in one place is effectively “the cheapest.” The concern for us is to know if the fares found in *Expedia.com* represent the actual fares offered by the carrier. We want to make sure that the fact that we collected the fare online does not restrict the analysis to just online fares.

The fares reported on different sites are sometimes different. One source of discrepancy comes from the fact that different online travel agencies have different algorithms to report the fares found in the Computer Reservation Systems (CRS). This plays a role when searching complex itineraries that may involve international flights. In our dataset this discrepancy does not arise since we are already restricting the search for a specific flight number on a specific departure date. A second important source of differences comes from variation across purchasing time and seat availability at purchase, the subject matter of this chapter. The third important source of variation arises because different fees and commissions differ across travel agencies. *Expedia.com* charges a lump sum booking fee of \$5 for every one-way ticket, *Travelocity.com* charges \$5 as well, while *Hotwire.com* charges \$6. Other websites like *Priceline.com*, *CheapTickets.com* or *Orbitz.com* allow fees to be a function of the base airfare, the carrier or the destination. For example, fees at *Orbitz.com* range from \$4.99 to \$11.99. “Brick-and-mortar” travel agencies charge even higher fees that can go

up to \$50. Buying on the phone also imposes additional different fees i.e. *CheapTickets.com* charges \$25 while *Travelocity.com* charges \$15.95 for over the phone bookings. Requesting a printed ticket will also impose additional variation. Even the carriers themselves charge different prices for exactly the same ticket. For example US Airways charges no fees if purchased through its website, but charges a \$5 fee for tickets purchased through the airline’s reservation centers and \$10 for tickets issued at the airport or at the city ticket offices. Moreover, the baseline fare may still be different depending on which Computer Reservation System (CRS) the travel agency uses to book its tickets.⁸

Currently, there are four Computer Reservation Systems which store and retrieve travel information used by all travel agents. These are Amadeus, Galileo, Sabre and Worldspan. Airlines pay an average booking fee per segment of \$4.25 when using a CRS, while travel agencies usually obtain CRS at no cost or receive certain payments in exchange for agreeing to use the system. According to the 2005 Report from American Society of Travel Agents (ASTA) [3], the “brick-and-mortar” travel agencies have responded by booking part of their sales using the carriers’ websites and not the CRS. The main source of information of *Expedia.com* is the Worldspan, but as well as *Orbitz.com*, they have established direct connection with airlines’ internal reservation systems to bypass Worldspan and avoid the CRS fees.

While it is difficult to evaluate price differences for exactly the same ticket offered offline, for online markets the information is readily comparable. Chen [17] using a dataset gathered online in 2002 obtained that for quotes found in multiple online sites the differences in prices are on the order of 0.3 to 2.2 percent. Even though not mentioned in her paper, these price differences can be tracked down just by comparing

⁸Additional fees common to all include taxes, special surcharges, segment fees and September 11 security fees.

the different fees charged at each site. Currently, carriers like American, Alaska and United offer a promise that travelers will always find the cheapest fare in its own websites. If the traveler finds a cheaper fare (with more than a \$5 difference), they offer paying back the difference plus additional bonus frequent flyer miles. This shows the carriers' interest on selling through its own websites. In response, *Orbitz.com* and *Expedia.com* adopted similar policies.

Based on all the multiple ways in which fares can potentially differ for exactly the same ticket, we have to come up with a clean measure of a "ticket's fare". The best candidate is each carrier website fare which is directly under the carrier's control and is free of any additional fees imposed by CRS, travel agencies or the same carrier if sold offline. For all the carriers in our sample, the fare found in *Expedia.com* is \$5 more than each carrier's website fare, thus obtaining the carriers' website fare is straight forward. Moreover, it is interesting to know ASTA reported that in 2002 the biggest on-line travel agency was *Expedia.com*, with a market share of 28.7 percent, followed by *Travelocity.com* (28.5 percent) and *Orbitz.com* (21.3 percent).

Regarding online sales, we know that they have been growing significantly during the last couple of years. The ASTA's report in 2005 citing PhoCusWright Inc. as the source, state that for leisure and unmanaged air sales, the overall online sales as a percentage of total sales went up from 30.8 percent in 2001 to 56.2 percent in 2004. Of these sales, 38.3 percent correspond to online travel agencies and 61.7 percent to sales through the airlines web sites.

C. The empirical model

1. Oligopoly model of costly capacity and demand uncertainty

In this section we derive a simple oligopoly model under capacity constraints and demand uncertainty. The predictions of this basic model were already obtained in a more formal environment in Dana [21]. The current derivation extends naturally to our formulation of demand uncertainty and testing procedure in the empirical section.

Let the total number of demand states be $H + 1$. The uncertainty in the demand comes from the fact that each carrier does not know *ex-ante* which demand state may occur. Let N_h be the number of consumers who will arrive at the demand state h , where $h = 0, \dots, H$ and $N_h \leq N_{h+1}$. This ordering implies that all the travelers who arrive at demand state h will also arrive at a higher-numbered demand state $h + 1$. Now, define a batch as the additional number of travelers that arrive at each demand state when compared to the immediate lower demand state, so batch h will be given by $N_h - N_{h-1}$ and the first batch is just N_0 .

Consider the case where consumers' reservation values for homogeneous airplane seats are uniformly distributed $[0, \theta]$, then the demand at state h is given by:

$$D_h(p) = \left(1 - \frac{p}{\theta}\right) N_h \quad (2.1)$$

Each demand state h occurs with probability ρ_h . Given that all demand states have at least N_0 potential travelers, the probability of having N_0 potential travelers arriving is $Pr_0 = \sum_{\kappa=0}^H \rho_\kappa = 1$. In general, the probability that at least N_h potential travelers arrive is the summation of the probabilities of demand states that have at least N_h customers, $Pr_h = \sum_{\kappa=h}^H \rho_\kappa$. This implies that the probability that N_h potential consumers arrive is always as high as the one that N_{h-1} potential consumers arrive, $Pr_h \geq Pr_{h+1}$. Following Prescott [48], the only cost for the carriers is a strictly

positive cost λ incurred on *all* units, regardless whether these units are sold or not. This cost can be interpreted as the unit cost of capacity (or shadow cost), or the cost of adding an additional seat in the aircraft. Unlike Dana [21], we assume that the unit marginal cost of production incurred only on the units that are sold is zero.⁹ Define the *effective cost of capacity* (*ECC*) as $ECC_h = \lambda/Pr_h$. This *ECC* adjusts the unit cost of capacity by the probability that this unit is sold. Since some of the seats will be sold only at higher-numbered demand states, if these units are sold, the effective cost of capacity reflects the costs that should be covered whether or not they are sold. If the unit cost of capacity is \$100, but this unit is sold only half of the times, if it gets sold, the cost that should be covered is \$200.

The number of identical carriers in the market is M . When the demand state is $h = 0$ with the corresponding firm's effective cost of capacity ECC_0 , the standard symmetric Nash equilibrium solution of a Cournot oligopoly competition is:

$$\begin{aligned} p_0 &= \frac{\theta + M \cdot ECC_0}{M + 1} \\ \delta_0 &= D_0(p_0) = \frac{N_0(\theta - ECC_0)M}{\theta(M + 1)} \end{aligned} \quad (2.2)$$

where p_0 is the equilibrium price, and δ_0 is the total amount of seats sold. Note each firm would allocate δ_0/M number of seats at price p_0 . From the second part of Equation 2.2 we obtain that the potential number of passengers that arrive at

⁹In our setting this basically means that the only relevant cost for the carriers is the one incurred when deciding whether or not to hold inventories for an additional seat. The cost that is assumed to be zero is peanuts (or pretzels and soft drinks plus any other marginal cost, i.e. baggage transportation). In the hotel example these marginal costs may include cleaning the room, changing towels, sheets and in many cases the breakfast.

demand state $h = 0$ is:

$$N_0 = \frac{\theta \cdot (1 + M)}{M} \cdot \delta_0 \cdot [\theta - ECC_0]^{-1} \quad (2.3)$$

When the demand state is $h = 1$, according to equation 2.1, the total demand at price p_0 is given by:

$$D_1(p_0) = \left(1 - \frac{p_0}{\theta}\right) \cdot N_1 \quad (2.4)$$

Note that $D_1(p_0) = D_0(p_0)$ since $N_1 = N_0$, i.e., the total amount of seats demanded at price p_0 when $h = 1$ is at least as large as the pre-allocated number of seats d_0 . Dana [21] uses proportioning rationing to assign seats at p_0 . This means that everybody has a equal chance $d_0/D_1(p_0) = N_0/N_1$ to get a seat at p_0 . The residual demand, therefore, is:

$$\begin{aligned} R_1(p|p_0) &= D_1(p) \left(1 - \frac{\delta_0}{D_1(p_0)}\right) \\ &= \left(1 - \frac{p}{\theta}\right) (N_1 - N_0) \end{aligned} \quad (2.5)$$

Again, the symmetric Nash equilibrium solutions if the demand function is $R_1(p|p_0)$ in 2.5 will be:

$$\begin{aligned} p_1 &= \frac{\theta + M \cdot ECC_1}{M + 1} \\ \delta_1 &= M \cdot (N_1 - N_0) \cdot \frac{(\theta - ECC_1)}{\theta \cdot (M + 1)} \end{aligned} \quad (2.6)$$

Compare 2.2 and 2.6, we can see that $p_1 \geq p_0$ given that $Pr_1 \leq Pr_0$. In this

case, from the second part of 2.6 we obtain that the potential number of passengers that arrive at demand state $h = 1$ is given by:

$$N_1 = \frac{\theta \cdot (1 + M)}{M} \cdot \delta_1 \cdot [\theta - ECC_1]^{-1} + N_0 \quad (2.7)$$

If the demand state is $h = 2$, we are interested in the residual demand after those travelers who have bought tickets at price p_0 and p_1 , denoted as $R_2(p|p_0, p_1)$. To find out $R_2(p|p_0, p_1)$, we start with the residual demand after those who bought tickets at p_0 , denoted as $R_2(p|p_0)$, which can be obtained from 2.6:

$$R_2(p|p_0) = \left(1 - \frac{p}{\theta}\right)(N_2 - N_0) \quad (2.8)$$

Travelers who are still in the market after the tickets at p_0 have been sold out will now have the chance to purchase tickets at p_1 . The number of potential consumers who will demand tickets at p_1 is $R_2(p_1|p_0)$, given by 2.8, and the number of tickets available at price p_1 is $R_1(p_1|p_0)$, given by 2.5, $R_2(p_1|p_0) \geq R_1(p_1|p_0)$. We apply the proportional rationing again to get the residual demand $R_2(p|p_0, p_1)$:

$$\begin{aligned} R_2(p|p_0, p_1) &= R_2(p|p_0) \left(1 - \frac{R_1(p_1|p_0)}{R_2(p_1|p_0)}\right) \\ &= \left(1 - \frac{p}{\theta}\right)(N_2 - N_0) \left(1 - \frac{\left(1 - \frac{p_1}{\theta}\right)(N_1 - N_0)}{\left(1 - \frac{p_1}{\theta}\right)(N_2 - N_0)}\right) \\ &= \left(1 - \frac{p}{\theta}\right)(N_2 - N_1) \end{aligned} \quad (2.9)$$

The symmetric Nash equilibrium solution for the residual demand function $R_2(p|p_0, p_1)$ in 2.9 is given by:

$$p_2 = \frac{\theta + M \cdot ECC_2}{M + 1}, \quad \delta_2 = M \cdot (N_2 - N_1) \cdot \frac{(\theta - ECC_2)}{\theta(M + 1)} \quad (2.10)$$

It is important to mention that here carriers are assumed to not observe the seat availability of their competitors. Once carriers sell their portion d_0/M for the first batch N_0 of potential travelers they take the next step which is pricing the second batch $N_1 - N_0$ of consumers. This assumption guarantees that any given carrier does not try to allocate its entire capacity to the first batch at the expense of their competitors. At the end of the derivation once we generalize the findings for a continuum of demand states, this assumption will be no longer needed.

This Cournot pricing strategy at each of the batches may allow the possibility that competitors behave strategically as in a repeated Cournot game where in each subsequent stage of the game firms face each time higher costs given by ECC . Since this is a finitely repeated game, we just obtain the subgame perfect Nash equilibrium by backward induction. Firms will not be able to collude since each subgame is played as a static Cournot game.¹⁰

Proposition 1 generalizes previous discussions to any number of demand states.

Proposition 1 *Let aggregate demand function be given in 2.1. $R_k(p|p_{k-1}, \dots, p_1, p_0)$ is the residual demand when demand state is k and travelers who have bought tickets at lower prices p_0, \dots, p_{k-1} have left the market (as in Eden [25]). We have:*

$$R_k(p|p_{k-1}, \dots, p_1, p_0) = \left(1 - \frac{p}{\theta}\right)(N_k - N_{k-1}) \quad (2.11)$$

Proof When the demand state $k = 1$, according to 2.5, the proposition holds.¹¹ We will prove: if the proposition holds at demand state k , then it must hold at demand

¹⁰The continuum of demand states is like an infinitely repeated game. If collusion is achieved in this scenario, we just require collusion payoffs in each stage game to be a function only of the same stage payoffs for the results in this section to hold. Again, for a stricter derivation of the same results see Dana [21].

¹¹According to 2.9, the proposition also holds for $k = 2$.

state $k + 1$. Suppose the proposition at demand state k holds. When demand state is $k + 1$, according to 2.9, the residual demand after travelers who have bought tickets at lower prices of p_0, \dots, p_{k-1} have left the market is given by:

$$R_{k+1}(p|p_{k-1}, \dots, p_1, p_0) = \left(1 - \frac{p}{\theta} \cdot (N_{k+1} - N_{k-1})\right) \quad (2.12)$$

Therefore, the residual demand after travelers who have bought tickets at lower prices of p_0, \dots, p_{k-1}, p_k have left the market is given by:

$$\begin{aligned} R_{k+1}(p|p_k, p_{k-1}, \dots, p_0) &= R_{k+1}(p|p_{k-1}, \dots, p_1, p_0) \left(1 - \frac{R_k(p_k|p_{k-1}, \dots, p_0)}{R_{k+1}(p_k|p_{k-1}, \dots, p_0)}\right) \\ &= \left(1 - \frac{p}{\theta}\right)(N_{k+1} - N_{k-1}) \left(1 - \frac{\left(1 - \frac{p_k}{\theta}\right)(N_k - N_{k-1})}{\left(1 - \frac{p_k}{\theta}\right)(N_{k+1} - N_{k-1})}\right) \\ &= \left(1 - \frac{p}{\theta}\right)(N_{k+1} - N_k) \end{aligned} \quad (2.13)$$

Note $R_k(p_k|p_{k-1}, \dots, p_0)$ in 2.13 is from 2.11 and $R_{k+1}(p_k|p_{k-1}, \dots, p_0)$ is from 2.13. Equation 2.13 proves Proposition 1. \square

From the residual demand equation of 2.12, it is easy to get that:

$$p_k = \frac{\theta + M \cdot ECC_k}{M + 1}, \quad \delta_k = M \cdot (N_k - N_{k-1}) \cdot \frac{(\theta - ECC_k)}{\theta(M + 1)} \quad (2.14)$$

For the general case, using the second part of 2.14 we obtain that the potential number of passengers that arrive at demand state $h = k$ is given by:

$$N_k = \frac{\theta(1 + M)}{M} \cdot \delta_k \cdot [\theta - ECC_k]^{-1} + N_{k-1} \quad (2.15)$$

By recursive substitution, considering the construction of the ECC for each batch of travelers, and for a continuum and infinite number demand states we can obtain that the number of potential travelers that arrive at demand state h is given by:

$$N_h = \frac{\theta(1+M)}{M} \int_0^h \delta_\omega \left[\theta - \lambda \cdot \left(\int_\omega^\infty \rho_\kappa d\kappa \right)^{-1} \right]^{-1} d\omega \quad (2.16)$$

From these N_h consumers that arrive at demand state h , only $\int_0^h \delta_\kappa d\kappa$ are able to buy a seat. Moreover, notice that the price paid by each group ω is different and given by:

$$P_\omega = \frac{1}{1+M} \left[\theta + M \cdot \lambda \left(\int_\omega^\infty \rho_\kappa d\kappa \right)^{-1} \right] \quad \forall \omega \in [0, h] \quad (2.17)$$

This is just the continuum version of the first part of equation 2.14.¹² We now just use this last equation to derive two testable implications:

$$\frac{\partial p_\omega}{\partial ECC_\omega} = \frac{M}{1+M} > 0, \quad and \quad \frac{\partial \left(\frac{\partial p_\omega}{\partial ECC_\omega} \right)}{\partial M} = \frac{1}{(1+M)^2} > 0 \quad (2.18)$$

The first part of equation 2.18 tells us that when the ECC increases, price also increases. The second part implies that as the market becomes more competitive (larger M), the marginal effect of ECC on fares is greater. Therefore, for a given distribution of demand uncertainty more competitive markets will show greater price dispersion. The expressions in equations 2.18 reduce to a monopoly when $M = 1$ and to a perfectly competitive market when $M \rightarrow \infty$. Note that in a perfectly competitive market, 2.18 predicts that every dollar increase in the ECC is transferred to prices as no markups exist to absorb part of this increase.

¹²Equation 2.17 is analogous to the first equation in p. 1233 in Prescott [48], equation (10) in Eden [25], equation (11) in Dana [19] and more closely related to equation (15) in Dana [21] for an oligopoly case. The benefit from our equation (17) over Dana [21]'s is that by assuming a specific functional form in the demand, price can be isolated on the left hand side of the equation. Dana [21] provides a more general derivation of this result.

2. Modeling demand uncertainty

Let's initially assume that carriers commit to an optimal distribution of prices for each flight before demand is known.¹³ By price commitment we mean that when demand is low, a traveler who arrives early or arrives late will face the same price as long as the carrier has not sold tickets in the meantime. Prices increase only if carriers have been selling tickets. Therefore, the information in the price schedule can be implicitly included in the functional form specified for the selling probability. This basically means that the probabilities are predetermined for each price schedule and the specification of demand uncertainty. The price schedule will be optimal and firms will not want to depart from it as long as they do not start learning about the state of the demand. As mentioned by Dana, useful information about the demand may only be available close to departure or once it is too late for carriers to change fares. Furthermore, as long as carriers do not learn any useful information about the state of the demand during the trading process, we can relax the price rigidity assumption (Eden [25]).

Starting with the simplest scenario where each demand state is equally likely with probability given by $\rho_h = \alpha/m$. This just means that demand states are uniformly distributed $[0, m/\alpha]$ with m being the total number of seats in the aircraft and $\alpha \geq 1$. The last inequality assures that there is a positive probability that the last seat gets sold. Following the intuition from section 1, having m/α demand states is the same as having $m/\alpha = H + 1$ batches $(N_k - N_{k-1})$ of travelers with the first batch N_0 showing up with the highest probability and the subsequent ones showing up each time with a lower probability than the previous one. Assume that the lowest demand

¹³Later in the empirical section we will allow for some deviations from price commitment. In particular, we allow the possibility of current shocks affecting future prices by estimating a dynamic model of Arellano and Bond [1].

state has one consumer buying a ticket ($\delta_0 = 1$) and for subsequent demand states we have one additional buyer each time we move to the next higher demand state ($\delta_k = 1 \forall k$). Because in every demand state there is at least one consumer buying a ticket, the probability of selling the first seat is equal to one. In all but the lowest demand state there are at least two travelers, so the probability of selling the second ticket is given by one minus the probability of the having the lowest demand state, that is $1 - \alpha/m$. In general, the probability that seat h gets sold is given by:

$$Pr_h = 1 - h \frac{\alpha}{m}, \quad h \in \{1, 2, \dots, m\} \quad (2.19)$$

which is just one minus the probability of having any demand state with lower demand than state h given the carrier's price distribution $q(p)$. In this equally likely demand states case, α is a constant that determines the rate at which the probability that the next seat gets sold diminishes.

Assuming that each demand state is equally likely seems too restrictive. Given our construction of demand uncertainty, this would imply that having only one passenger flying is as likely as having the plane at half capacity and that the probability of selling one additional seat decreases linearly. To allow for more flexibility in the characterization of demand uncertainty we consider the case where $\rho_h = \phi_h$, with ϕ being the *pdf* of a normal density that has mean μ and standard deviation σ . From the discussion so far we know that the probability of selling seat h is the summation of the probabilities of all demand states that have at least h travelers. For a continuum of demand states, this is given by $\int_h^\infty \rho_\kappa d\kappa$. Therefore, the probability of selling seat h for the normal density will be:

$$Pr_h = \int_h^\infty \phi_\kappa d\kappa = 1 - \Phi_h, \quad (2.20)$$

with Φ being the *cdf* of a normal distribution.

3. Calibrating the probability density of demand uncertainty

To obtain Pr_h used in calculating the *ECC*, it is necessary to get the values for the parameters α in the uniform distribution and the mean, μ , and standard deviation, σ , in the normal distribution. In this subsection we calibrate the values of these parameters to mimic the demand uncertainty conditions in each of the routes.

A key source of information for the calibration comes from the $T - 100$ data from the *Bureau of Transport Statistics*. We use this dataset to obtain yearly occupancy rates, or load factors at time of departure. This is done in three steps. First, for each of the routes in the sample, we calculate its load factor for the 81 routes in the sample for the period 1990 to 2005, based on the $T - 100$ data. Second, each of these 81 series is used to estimate an ARMA model. Finally, the estimated ARMA model is applied to obtain the 2006 value using a one-step ahead forecast.¹⁴ For routes where the ARMA model predicts a high load factor, meaning that most of the seats are expected to be sold, the calibration procedure will assign higher probabilities to higher demand states. In this case the *ECC* is going to be relatively low for a large majority of the tickets. When the forecasted load factor is low, the probability of selling the last couple of seats is going to fall fast, meaning that the cost of stocking inventories is higher.

The problem with the information obtained from the $T - 100$, however, is that we have a measure of the forecasted value of the average number of tickets sold rather than of the forecasted value of the average number of tickets demanded. This arises because the demand state is censored when transformed to the number of tickets sold.

¹⁴The details of the estimation are available upon request.

Once the aircraft is sold out the $T - 100$ no longer records higher demand states. To overcome this limitation let the underlying demand state h^* , be distributed $N(\mu, \sigma^2)$ with the observed number of seats sold $h = h^*$ if $h < m$ or else $h = m$. Recall here that m is the maximum number of seats available in the airplane. Then the expected number of tickets sold is given by the first moment of the censored normal:

$$E(h) = Prob(h = m) \cdot E(h|h = m) + Prob(h < m) \cdot E(h|h < m) \\ \left(1 - \Phi\left(\frac{m - \mu}{\sigma}\right)\right) \cdot m + \Phi\left(\frac{m - \mu}{\sigma}\right) \cdot \left[\mu - \sigma \frac{\phi((m - \mu)/\sigma)}{\Phi((m - \mu)/\sigma)}\right] \quad (2.21)$$

The expression for $E(h|h < m)$ is obtained from the mean of a truncated normal density. The *pdf* and the *cdf* of the normal density are evaluated at the moment the flight sells out. Hence, the value $\Phi((m - \mu)/\sigma)$ is interpreted as the sold out probability. Using information on the probability that a flight sells out, based on the second dataset obtained from *Expedia.com*, and the expected number of tickets sold, obtained from the ARMA models, we can use 2.21 to obtain values for μ and σ .

Calibrating the value of α in the uniform distribution is simpler. We obtain the analog of equation 2.21, $E(h) = 1 - \alpha/2$, by using the truncated uniform distribution. This equation can be used directly to get α . In this case since we only have to calculate one parameter, the sold-out probabilities are no longer needed. The cost of requiring less information is to have less flexible characterization in which one single parameter affects both the mean and the variance of the distribution of demand states.

4. Estimated equation and interpretation

Following a similar approach as Stavins [52], we estimate a reduced-form model of ln airfare on *ECC*, market concentration, carrier's market share and route-specific factors. The key new variable in our analysis is the *ECC* that measures the effect of costly capacity and demand uncertainty by adjusting the unit cost of capacity by the

probability that the ticket gets sold. The construction of the dataset also allows us to control for all other relevant ticket-specific characteristics as explained in section B. The equation to be estimated is given by:

$$\begin{aligned}
\ln FARE_{ijt} = & \beta_0 + (\gamma_0 + \gamma_1 HHI_j) \times ECC_{ijt} + \beta_1 DAYADV_{ijt} \\
& + \beta_2 DIST_j + \beta_3 DISTSQ_j + \beta_4 ROUSHARE_{ij} \\
& + \beta_5 HHI_j + \vartheta_1 HUB_{ij} + \vartheta_2 SLOT_j + \beta_6 DIFTEMP_j \\
& + \beta_7 DIFRAIN_j + \beta_8 DIFSUN_j + \beta_9 AVEHHINC_j \\
& + \beta_{10} AMEANPOP_j + \mu_i + \nu_{ijt}
\end{aligned} \tag{2.22}$$

where the subscript i refers to the flight, j to the route, and t is time. Dummy variables have estimated coefficients denoted by ϑ , otherwise β . μ_i denotes the unobservable flight specific effect and ν_{ijt} denotes the remainder disturbance. Different error structures will be assumed along the empirical section. Each observation in the sample represents a unique ticket for a carrier on a route. By route we mean a combination of departure and arrival airports on a one-directional trip. $FARE_{ijt}$ is price paid in US dollars. From table 1, the sample mean fare is \$291, with a minimum of \$54 for an American Airlines flight from Dallas Fort Worth, TX to Houston International, TX when at least 80 percent of the plane was empty. The maximum is \$1,224 in a United Airlines flight from Philadelphia International, PA to San Francisco International, CA when there are less than 9 percent of the seats available. The variable of interest in the analysis is ECC which is obtained from $ECC = \lambda / Pr_h$. In particular, when the distribution is uniform as defined in 2.19, we should have:

$$ECC_{ijt} = \frac{\lambda}{Pr_{h_{ijt}}} = \frac{\lambda}{1 - h_{ijt} \frac{\alpha_j}{m_{ij}}} \tag{2.23}$$

where m_{ij} is the total number of seats in the aircraft and $h_{ijt} - 1$ is the number of seats that have already been sold at time t . α_j is the mean of the uniform distribution. ECC is measured in the same units as $FARE$, nevertheless to be able to interpret the magnitude of the coefficient; we initially normalize λ to be equal to one. For the normal density case as presented in 2.20, ECC is given by:

$$ECC_{ijt} = \frac{\lambda}{Pr_{h_{ijt}}} = \lambda \times \left[\int_{h_{ijt}/m_{ij}}^{\infty} \sqrt{2\pi\sigma_j^2} \cdot \exp\left(-(\kappa - \mu_j)^2/2\sigma^2\right) d\kappa \right]^{-1} \quad (2.24)$$

The values for μ_j and σ_j are allowed to change across routes, so they are indexed by route j . h_{ijt} and m_{ij} are directly observable from our dataset.

Now we take a look at three different cases where the ECC should play no role in the pricing decisions and analyze how our construction of this measure respond in each of these cases. In other words, these are the cases where the model of section 1 should predict no price dispersion due to costly capacity and demand uncertainty.

(i) For routes where we expect higher load factors, costly capacity will play a less important role. On the limit, when we expect to sell all the seats in the aircraft in every occasion $E(h) = 1$. In the case for uniform density $\alpha_j = 0$, and from 2.19 we get that the probability of selling the next seat does not decrease with the cumulative number of seats sold, $Pr_h = 1$. For the normal density case $\mu_j \rightarrow \infty$. In both situations, there will be no rising ECC as more seats are sold. Holding inventories of additional seats will have no cost since we know for sure that they will be sold. In summary, $\lim_{E(h) \rightarrow 1} ECC = \lambda$.

(ii) A similar phenomenon would happen if aircrafts had infinite capacity, i.e. no capacity constraints. This can be interpreted as carriers being able to adjust the size of the aircraft anytime before departure at no additional cost. An alternative interpretation could be that the good is not perishable; if the good is not sold today,

it can be sold anytime in the future. Characteristic that does not hold for airline travel since once the plane departs; carriers can no longer sell tickets. Again, we have $\lim_{m \rightarrow \infty} ECC = \lambda$ for both the uniform and the normal.

(iii) Finally, in the case of no demand uncertainty, carriers would just set their capacity levels to match to the certain number of travelers, hence the ECC would play no role, i.e., $\lim_{\sigma \rightarrow 0} ECC = \lambda$ for the normal, but no demand uncertainty holds also for the uniform.

In all three scenarios the price that an airline charges would be same for every seat, and there will be no price dispersion. That is why models omitting demand uncertainty in their interpretations like Borenstein and Rose [10] or Stavins [52] would lead to interpret this variation in prices as price discrimination rather than the effect of the combination between costly capacity and demand uncertainty. Failing to adjust the unit cost of capacity by the probability that the seat gets sold would lead to predict that the shadow cost remains constant, when it doesn't. In addition to ECC , the specification in 2.22 includes the *Herfindahl-Hirshman Index* (HHI) that measures the concentration on the route. HHI is calculated using $ROUSHARE$, which is the carrier's share of total number of seats in all the direct flights on that route, not just the ones from the carriers from which we have fares. Even though similar estimation specifications like in Stavins [52] assumes that HHI is exogenous to airfare estimation, here we provide instruments for both $ROUSHARE$ and HHI . We use $GEOSHARE$ for $ROUSHARE$ and $XFLTHERF$ for HHI , as constructed in Borenstein [9] and Borenstein and Rose [10]. A short explanation of these instruments is given in appendix B and the summary statistics of these two instrument variables are shown in table 1.

The rest of the regressors in the equation are control variables when the estimation is carried out using carrier fixed effects. $DAYADV$ is the number of days prior to

Table 1. Summary statistics

	Mean	Std. Dev.	Min.	Max.	Obs.
<i>FARE(US\$)</i>	291.087	171.879	54.000	1224.000	7933
<i>DAYADV</i>	52.289	30.154	1.000	103.000	7933
<i>DIST</i>	1104.380	620.720	91.000	2604.000	7933
<i>ROUSHASEA</i>	.665	.314	.119	1.000	7933
<i>HHI</i>	.684	.287	.259	1.000	7933
<i>HUB</i>	.737	.440	.000	1.000	7933
<i>SLOT</i>	.298	.458	.000	1.000	7933
<i>DIFTEMP</i>	6.210	4.137	.000	19.000	7933
<i>DIFRAIN</i>	2.010	1.484	.000	4.900	7933
<i>DIFSUN</i>	7.911	8.461	.000	45.000	7933
<i>AVEHHIC(US\$)</i>	35580	4620	25198	53430	7933
<i>AVEPOP</i>	1044072	631862	187704	2897818	7933
<i>GEOSHARE</i>	.674	.324	.025	1.000	7933
<i>XFLTHERF</i>	.708	.285	.252	1.000	7933
<i>ECC-Censored Normal</i>	1.557	.940	1.000	11.668	7933
<i>ECC-Censored Uniform</i>	1.453	1.086	1.005	55.887	7931
<i>LOAD(t = 1)</i>	.881	.153	.227	1.000	228
<i>Sold Out Prob.</i>	.227	.104	.037	.571	81
<i>Forecasted LF</i>	.738	.083	.469	.890	81

departure, while *DIST* and *DISTSQ* are the distance and distance square between the two endpoint airports on a route. *DIFTEMP*, *DIFRAIN*, and *DIFSUN*, are the differences in the average end of October temperature, rain, and sunshine between the two endpoints. They are measured in Fahrenheit degrees, precipitation in inches, and in percentages respectively. Their role is to control for some of the travelers' heterogeneity (i.e. mix of business and tourists). *AVEHHINC* and *AVEPOP* are average median household income in US dollars and average population of the two cities respectively.¹⁵ *HUB* is equal to one if the carrier has a hub in the origin or destination airport, zero otherwise. *SLOT* is a dummy variable equal to one when the number of landings and takeoffs is regulated in either origin or destination airport.¹⁶

¹⁵For cities with more than one airport, the population is apportioned to each airport according to each airport's share of total enplanements. Source: Table 3, *Bureau of Transportation Statistics*, Airport Activity Statistics of Certified Air Carriers: Summary Tables 2000.

¹⁶In some airports like Kennedy (JFK), La Guardia (LGA), and Reagan National (DCA), the U.S. government has imposed limits on the number of takeoffs and landings that may take place each hour. To take into account the scarcity value of acquiring a slot, the variable *SLOT* equals to one if either endpoint of route *j* is one of these airports and zero otherwise.

The summary statistics of all these variables are presented in table 1.

To get an estimate of the unit cost of capacity $\hat{\lambda}$, let $\hat{\gamma}_\iota$ for $\iota = 0, 1$, denote the estimates of γ_ι when the estimation of 2.22 is carried out assuming λ being one. As we have previously seen, one important implication from the perfectly competitive market is that every dollar increase in ECC is passed to prices (see equation 2.18, but assuming $M \rightarrow \infty$). This means that $\partial FARE / \partial ECC = (\hat{\gamma}_0 + \hat{\gamma}_1 HHI) FARE = 1$ when $HHI = 0$. This condition leads to the estimate $\hat{\lambda} = \hat{\gamma} \cdot F\bar{ARE}$, evaluated at the sample mean of $FARE$ and with $\hat{\gamma}_0$ being interpreted as the share of fares that corresponds to ECC . Since there is no reason to believe that λ changes across market structures, we fix it at this value, $\lambda = \hat{\lambda}$. Then, the marginal effect of ECC on fares for any market structure will be obtained from $\partial FARE / \partial ECC = 1 + (\hat{\gamma}_1 / \hat{\gamma}_0) HHI$.

Because of potential changes in costs, Stokey [54] mentioned that the mere presence of price variation over time is not an adequate measure of intertemporal price discrimination. Here we are appropriately controlling for raising marginal costs due to aircraft's capacity constraints under demand uncertainty. Given the construction of the model and under price rigidities, $DAYADV$ is expected to capture the effect of a type of second degree price discrimination named advance purchase discounts.

D. Results of the empirical analysis

The estimates for equation 2.22 using the censored normal construction of the ECC and carrier fixed effects are presented in table 2.¹⁷ The numbers in parentheses are t-statistics calculated using robust standard errors. The first column shows the results when assuming that the effect of ECC on fares does not vary with market concentra-

¹⁷The estimation was carried out with an unbalanced panel of 7933 observations because some fares were no longer available for flights that sold out a couple of days before departure. The missing observations account for less than 0.6 percent of the sample and we do not expect this to bias the estimates.

tion. Consistent with the theoretical predictions, its effect is positive and significant, implying that higher unit costs of capacity increase fares. When this effect is allowed to vary with market concentration in column (2), we find that greater market concentration, as measured by higher values of the HHI , decreases the positive marginal effect. The intuition, again, is that in competitive markets every dollar increase in unit cost of capacity is fully transferred to prices since there are zero markups. In non competitive markets when markups are positive, part of the increase in unit costs of capacity are absorbed by markups and the final effect on prices is lower. All the regression results reported are obtained using the instrument variable $GEOSHARE$ for $ROUSHARE$ and $XFLTHERF$ for HHI , as suggested in Borenstein [9] and Borenstein and Rose [10].

Table 2. Estimation results for the censored normal

Variable	(1)		(2)	
	Coefficient	t-Statistic	Coefficient	t-Statistic
ECC	0.092	(13.470)	0.163	(8.868)
$ECC \times HHI$			-0.091	(-4.388)
$DAYADV$	-0.003	(-12.935)	-0.003	(-12.198)
$DIST$	0.002	(37.285)	0.002	(37.180)
$DISTSQ$	$-3.4e^{-7}$	(-25.577)	$-3.4e^{-7}$	(-25.435)
$ROUSHARE$	0.252	(5.818)	0.254	(5.868)
HHI	-0.079	(-1.660)	0.066	(1.119)
HUB	-0.024	(-1.759)	-0.026	(-1.868)
$SLOT$	-0.246	(-14.445)	-0.253	(-14.755)
$DIFTEMP$	0.003	(2.322)	0.003	(2.341)
$DIFRAIN$	-0.171	(-33.264)	-0.174	(-33.305)
$DIFSUN$	0.004	(5.149)	0.004	(4.987)
$AVEHHINC$	$1.7e^{-5}$	(12.562)	$1.7e^{-5}$	(12.515)
$AVEPOP$	$-1.2e^{-7}$	(-11.844)	$-1.2e^{-7}$	(-11.554)
Carrier FE	Yes		Yes	
Flight FE	No		No	
Period FE	No		No	
$R - squared$	0.482		0.484	

Notes: The results reported here are obtained using $GEOSHARE$ as the excluded instrument variable for $ROUSHARE$ and $XFLTHERF$ as the excluded instrument variable for HHI . The independent variable is $\ln(FARE)$. t -statistics (in parentheses) are based on White robust standard errors. Carrier fixed effects not reported.

Most of the estimates are directly comparable to the ones obtained in Stavins [52] who uses a similar dataset collected in 1995.¹⁸ Even though it is useful to know our

¹⁸The main difference is that Stavins did not have information about seat avail-

estimates are comparable to effects already documented in the literature, in this chapter we are not directly interested in the coefficients of time invariant parameters. Taking advantage of the panel structure of the data, a more suitable specification that will be able to control for unobserved time invariant parameters, but will wipe out these estimates is a model with flight fixed effects. These estimates are presented in table 3. Moving from carrier to flight fixed effects greatly improves the goodness-of-fit as measured by R^2 . In all specifications that include flight fixed effect, R^2 are greater than 0.86.

Table 3. Summary of robustness checks

Variable	(1) Censored Normal		(2) Censored Uniform		(3) Censored Normal Dynamic Panel A&B	
	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	z-Statistic
$LNFARE(-1)$					0.589	(43.103)
ECC	0.175	(11.883)	0.520	(11.512)	0.185	(12.131)
$ECC \times HHI$	-0.134	(-8.058)	-0.519	(-11.503)	-0.122	(-6.403)
$DAYADV$	-0.003	(-24.023)	-0.003	(-25.687)	$-4.3e^{-4}$	(-4.055)
Carrier FE	No		No		No	
Flight FE	Yes		Yes		Yes	
Period FE	No		No		No	
$R - squared$	0.865		0.876		n.a.	

Notes: The independent variable is $\log FARE$, $N=7933$ for columns (1) and (2), and 7472 for column (3) with 228 cross sectional observations in all cases. t -statistics based on White robust standard errors. The construction of the ECC based on the censored normal on columns (1) and (3) and on the censored uniform on column (2).

Table 3 also runs some robustness checks on the construction of the ECC . Column (1) still uses the censored normal, while column (2) constructs the ECC under the censored uniform assumption on the distribution of demand states. Both specifications predict that greater market concentration decreases the positive effect of

ability, thus was unable to control for probability of selling each ticket. Moreover, her dataset had less ticket observations over only twelve routes, while here we have eighty-one routes. Consequently we expect our HHI to be a very good approximation of the market structure. The signs for the estimated coefficients were found to be the same for number of days in advance purchase ($DAYADV$), distance and distance square, market share ($ROUSHARE$), hub, slot, difference in temperature and average household income. The only comparable coefficient sign that does not match is average population. We believe our estimate is a better approximation since she did not adjust average population by the number of airport enplanements as we did. More populated cities get lower airfares.

ECC on fares. However, the magnitude of the effect is very sensitive to the choice of the demand state distribution. The reason why the censored uniform predicts greater marginal effects is simple: it puts excessive weight on lower demand states. The censored uniform predicts that low demand states are as likely as any other demand state. This causes that the *ECC* rises too fast when the first couple of seats are sold, over dimensioning the costs of capacity constraints and demand uncertainty. However, what it's important is to realize that the basic conclusion holds with different specifications of the uncertain demand. Our measure of the selling probability which is used to construct the *ECC* is a function of the number of seats that have already been sold. However, the number of seats that were sold depends on past level of fares. This questions the strict exogeneity assumption about the *ECC*. To account for this potential endogeneity problem, in column (3) we consider a dynamic panel data model where we only have to assume that the explanatory variables are weakly exogenous, plus still instrumenting for the *HHI*. The idea is to difference the regression equation 2.22 to remove any omitted variable created by unobserved flight-specific effects, and then instrument the right and side variables using lag values of the original regression to eliminate potential parameter inconsistency arising from simultaneity bias. The estimates represent GMM in first differences as developed in Arellano and Bond [1]. Here the error term in the model (ν_{ijt} in equation 2.22) may affect future dependent and independent variables. For example, suppose the airline experiences a positive shock at time t that drives up the number of tickets sold. The Arellano and Bond [1] estimate allows fares and number of tickets sold at $t + 1$ to change in response to such a shock, hence the specification is robust to the fact that the amount of seats sold up to this period is a function of prices in the previous periods. The result measure how the exogenous component of *ECC* impacts fares. This specification is robust against deviations from the price commitment as suggested in

Eden [25]. Estimates in column (3) are close to the ones in column (1), supporting the two basic predictions of the theory.

Regarding the exogeneity of ECC , it is important to realize that the argument in this chapter is to analyze whether one way fares respond to a transformation of seat availability on that particular flight. However, one way fares are usually a small portion of the tickets sold. Most of the travelers flying on each of the flights in our dataset bought this leg as part of a round trip ticket, a connecting flight or both. The potential combinations are extremely large and the load factor at each point in time for any of our flights is the result of tickets sold along different combination of legs, maybe even passengers getting a seat with frequent flyer miles. This is an important argument in favor of the exogeneity of ECC and would likely explain why the Arellano and Bond estimates that control for potential endogeneity of ECC do not differ much from the other set of estimates.

Another important result is the coefficient estimate for $DAYADV$, the number of days prior to departure. As discussed in section A, advanced-purchase discounts (APD) have been argued in the literature as a way to divert demand from peak periods to off peak periods (Gale and Holmes [34], [35]; Dana [20]). In column (2), we include $DAYADV$ as a control variable. The coefficient estimate is negative and significant, providing evidence that supports APD . Buying the ticket one day earlier reduces the fare by 87 cents. Having been controlled for the ECC and under the assumptions that carriers cannot learn about the state of the demand, this 87 cents is an appropriate measure of second degree price discrimination in the form of advance purchase discounts. The conditions for this to be considered intertemporal price discrimination are the same as the ones in Dana [19].


To ease the concern that $DAYADV$ may enter into the model nonlinearly, in table 4 we show the results for three additional specifications. The first one, presented

Table 4. Nonlinearities in time

Variable	(1)		(2)		(3)	
	Censored Coefficient	Normal t-Statistic	Censored Coefficient	Normal t-Statistic	Censored Coefficient	Normal z-Statistic
ECC	0.121	(8.155)	0.097	(6.634)	0.096	(6.561)
$ECC \times HHI$	-0.113	(-6.690)	-0.105	(-6.592)	-0.122	(-6.699)
$DAYADV$	-0.010	(-18.502)	-0.024	(-18.802)		
$DAYADVSQ$	$6.4e^{-5}$	(14.920)	$3.8e^{-4}$	(15.800)		
$DAYADVCU$			$-1.9e^{-6}$	(-14.318)		
Carrier FE	No		No		No	
Flight FE	Yes		Yes		Yes	
Period FE	No		No		Yes	
$R - squared$	0.870		0.875		0.880	

Notes: The independent variable is $\ln FARE$, $N=7933$ with 228 cross sectional observations.
 t -statistics based on White robust standard errors.

in column (1), includes a square term for days in advance ($DAYADVSQ$), while the second one, in column (2), includes a cubic term ($DAYADVCU$). A completely flexible model where each time period is allowed to be different with no further restrictions is flight fixed-effects, reported in column (3). Comparing the coefficients reported in table 4 with the ones previously obtained, we conclude that the positive coefficient for ECC (γ_0 in equation 2.22) the negative coefficient for $ECC \cdot HHI$ (γ_1 in equation 2.22) hold. However, magnitude of the estimates of the estimates is somewhat smaller.

Fig. 4. Probability that seat h is sold (MCO-LGA)

To see how the different specifications assign different weights to different demand states, figure 4 shows the probability of selling seat h for the uniform and the normal specifications. The schedules shown are calibrated to match the values for the route Orlando International in Orlando, FL (MCO) to La Guardia in New York, NY (LGA). The 2006 forecasted load factor for this route is 0.82, also higher than the average across routes of 0.74, while the sold out probability was 0.254, higher than the sample average of 0.225. The forecasted value for this route is shown in the figure as the expected number of seats sold $E(h) = 0.822$. Because of the nature of the censored normal, this value is lower than the average of demand states $\mu_j = 0.855$. σ_j and α_j are 0.048 and 0.356 respectively. Note that figure 4 has two different probabilities. The probability that seat h gets sold, ρ_h , measured on the vertical axis and the probability of demand state h , Pr_h , measured as the absolute value of the slope. In an $m = 100$ seat airplane, the censored normal predicts that the 40th passenger will come with a probability $\rho_{0.4} = 0.98$ which obviously does not prevent the next passengers from arriving, whereas the probability that the plane actually departs with exactly 40 passengers is $Pr_{0.4} = 0.21$ percent. Moreover, the area below each of the curves is equal to the expected load factor $E(h)$.

From the estimates under various specifications in tables 2, 3 and 4 it is clear that the main conclusion is robust to various specifications: the effect of ECC is greater in more competitive markets. Now we can extend the analysis to study the magnitude of the effect. Under the assumption of zero markups in perfectly competitive markets, i.e., $HHI = 0$, we have a direct interpretation of the coefficient on ECC . In column (1) of table 3, the coefficient for ECC is 0.175, which means that the unit cost of capacity represents 17.5 percent of the average fare. Given the average fare of \$291, we can calculate the shadow cost of a unit capacity, $\hat{\lambda} = \$50.85$. The marginal effect of ECC on fares is given by $\partial FARE / \partial ECC = 1 + (-0.134 / 0.175) HHI$. When it is

evaluated at the sample mean of HHI (0.684), the marginal effect of ECC on fares is 0.476. This implies that for the average market structure one dollar increase in ECC leads to an increase in 48 cents in fares. When evaluating the effect of ECC on fares at values of HHI of 0.25, 0.50, and 0.75, we get this one is 0.809, 0.618 and 0.427 respectively. For a monopoly carrier from each dollar increase in ECC , 24 cents go to increase prices while 76 cents are absorbed by the markup.




Fig. 5. Fares and effective cost of capacity at different days prior to departure (AA 323 ATL-DFW)

Table 5. Sensitivity of sold-out probabilities

Variable	(1)		(2)		(3)	
	Censored Normal Sold-out Prob.+10%		Censored Normal Sold-out Prob.+20%		Censored Normal Sold-out Prob.+30%	
	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	z-Statistic
ECC	0.290	(11.784)	0.430	(11.387)	0.611	(10.553)
$ECC \times HHI$	-0.203	(-7.245)	-0.301	(-7.037)	-0.451	(-6.923)
$DAYADV$	-0.003	(-20.133)	-0.003	(-18.398)	-0.003	(-17.989)
Carrier FE	No		No		No	
Flight FE	Yes		Yes		Yes	
Period FE	No		No		No	
$R - squared$	0.864		0.863		0.862	

Notes: The independent variable is $\log FARE$, $N=7933$ with 228 cross sectional observations. t -statistics based on White robust standard errors. ECC calculated using different sold-out probabilities. Columns (1), (2), and (3) increase the sold-out probability in each route by a lump sum 10, 20, and 30 percent respectively.

As noted in the construction of the sold out probability, this may be interpreted

as a lower bound rather than an unbiased calculation of it. To see the response of the estimated coefficients to higher sold out probabilities, table 5 provides the estimates when the sold out probability for each of the flights is increased by a lump sum 10, 20 and 30 percent in columns (1), (2) and (3) respectively. Again, the main conclusion of the analysis still holds: greater effect of *ECC* on fares in more competitive markets. However, the magnitude of $\hat{\lambda} = \hat{\gamma}_0 \cdot \bar{FARE}$ changes; as the sold out probability increases, the share of the unit cost of capacity on fares increases as well. This proportion, calculated in table 3 as 17.5 percent, it is now 29.0, 43.0 and 61.1 percent for average sold out probabilities of 32.5 (22.5+10), 42.5 and 52.5 percent respectively. It would be reasonable to believe that this proportion is greater than our original estimate of 17.5 percent in column (1) of table 3. To get an idea of the magnitude, figure 5 presents the same AA flight 323 from ATL to DFW shown in figure 1. The *ECC* was calibrated with the censored normal with $\hat{\lambda} = .611 \cdot 148.14$. It would be difficult to argue about the exact size of the markup, but the ranges we are talking about here look quite reasonable. Moreover, the schedule of *ECC* on figure 5 seems to explain quite well the path followed by fares with the sharp increase for the last couple of seats.

The estimates in table 5 prove robustness in one additional dimension. As the marginal effect of *ECC* on fares is measured by $\partial FARE / \partial ECC = 1 + (\hat{\gamma}_1 / \hat{\gamma}_0) HHI$, we are interested in whether the ratio $\hat{\gamma}_1 / \hat{\gamma}_0$ changes with the sold out probability. In our estimates of column (1) in table 3, this one is -0.76 (-18.80) with the t-statistic in parentheses. For columns (1), (2), and (3) in table 4 this one is -0.70 (-14.63), -0.70 (-13.81), and -0.74 (-13.71) respectively. This provides some evidence that our estimate of the marginal effect of *ECC* on fares is stable, and its magnitude can be obtained with just a lower bound estimate of the sold out probability.

When dropping the assumption of no markups under perfect competition and

without any normalization or knowing the value of λ , we can come with an interpretation of the magnitude of the effect of costly capacity on fares. However, this one is not robust to the magnitude of the sold out probabilities.¹⁹ For our estimates in column (1) in table 3, a one standard deviation increase in the ECC, evaluated at sample means of HHI and fares, increases prices by \$23.77, which corresponds to an increase of 0.14 standard deviations.

Table 6. Elasticities

Variable	(1)		(2)		(3)	
	Censored Normal		Censored Normal		Censored Normal	
	Coefficient	t-Statistic	Coefficient	t-Statistic	Dynamic Panel A&B Coefficient	z-Statistic
$LNFARE(-1)$					0.592	(43.385)
ECC	0.219	(15.644)	0.398	(12.177)	0.397	(10.331)
$ECC \times HHI$			-0.252	(-6.722)	-0.201	(-3.947)
$DAYADV$	-0.002	(-17.985)	-0.002	(-17.691)	$-1.1e^{-4}$	(-0.863)
Carrier FE	No		No		No	
Flight FE	Yes		Yes		Yes	
Period FE	No		No		No	
$R - squared$	0.850		0.855		n.a.	

Notes: The independent variable is $\log FARE$, $N=7933$ with 228 cross sectional observations.
 t -statistics based on White robust standard errors.

Finally, table 6 presents the last set of estimates. These estimates take advantage of the fact that if we take logarithm of ECC , we break its components in two parts. The log of λ will become part of the constant in the regression, while the negative value of the logarithm of the probability that batch h arrives (Pr_h) will keep the same elasticity coefficient as the ECC . In these results the negative value of the logarithm of the probability takes the place of ECC to make the signs comparable to the previous results. Column (1) tells us that a one percent increase in the ECC (or same as one percent decrease in the selling probability), increases fares by 0.219

¹⁹The results follow from the fact that the marginal effect of ECC on $FARE$ is homogeneous of degree zero in λ . The marginal effect holds for any positive value of α :

$$\frac{\partial FARE}{\partial ECC} = \left(\frac{\delta_0}{\alpha \hat{\lambda}} + \frac{\delta_1}{\alpha \hat{\lambda}} HHI \right) FARE \times StdDev(\alpha \hat{\lambda})$$

percent. Once more, as illustrated in columns (2) and (3), the response to *ECC* is greater in more competitive markets.

E. Conclusions

This chapter sets to test the empirical importance of the price dispersion predictions presented in Prescott [48], formalized in Eden [25] and extended in Dana [21]. The basic idea in these theoretical models is that the equilibrium price dispersion can be explained by the different selling probabilities associated with each of the units sold. These selling probabilities play an important role in industries that face capacity constraints and uncertainty about the number of arriving consumers. Although the ideas in Prescott [48] have been extended to multiple areas in the economic literature, few papers attempt to directly test the basic predictions due to the difficulty of coming up with an appropriate measure of the selling probabilities.

In particular, the chapter seeks to find evidence for the two main predictions. i) Lower selling probabilities characterized by higher effective costs of capacity will lead to higher prices. ii) This effect will be larger in more competitive markets. We start building a simple theoretical framework based on Prescott [48], Eden [25] and Dana [21] that contains these two main predictions. The richness of this simple model comes from the fact that it naturally extends to accommodate the calibration of the demand uncertainty and the empirical procedure developed later. The airline industry landscapes the ideal scenario to test this theory. First, because capacity is set and can only be changed at a relatively large marginal cost. Second, the product expires at a point in time, and third, there is uncertainty about the demand. The empirical section takes advantage of a unique dataset that observes the evolution of prices and inventories of seats of 228 flights for over a period of 103 days prior to departure. We

control for ticket restrictions that screen travelers and isolate the effect of the selling probability on prices.

Using the information on seat inventories, plus calculations of the sold out probabilities (based on a second dataset), and the forecasted values of utilization rates (based on a third dataset), we are able to construct the distribution of demand uncertainty for each of the 81 routes in the sample. With this distribution we generate a measure of the selling probability and the effective cost of capacity (*ECC*) for each of the seats in an aircraft. This allows us to test the model by finding out if *ECC* has any effect on the prices, and if so, how this effect varies with market concentration.

Under various specifications, our empirical tests strongly support both predictions of the theory. We show that for the average market structure, when *ECC* increases by one dollar, fares increase by 48 cents, whereas the remaining 52 cents is absorbed by the markup. The elasticity specification tells us that one percent increase in the *ECC* (or same as one percent decrease in the selling probability), increases fares by 0.219 percent. Moreover, price dispersion due to costly capacity under demand uncertainty was found to be greater in more competitive markets. The idea is that more competitive markets have smaller markups, so an increase in marginal costs goes directly to prices. In more concentrated markets where markups are greater, higher costs are partially absorbed by the markup and the effect on fares is smaller. In addition, under the assumption that carriers do not learn about the state of the demand, our results support a second degree price discrimination effect that indicates that buying the ticket one day earlier reduces fares by 87 cents. During the estimation the chapter takes care of various sources of potential endogeneity, building a set of instruments for the market structure and benefiting from the panel structure of the data by running a dynamic model.

Although the dataset collected enjoys some very nice features, it has some draw-

backs that limit extending the results to the airline industry as a whole. The one-way non-stop ticket is only a portion of the tickets sold in each flight, and often it is a small portion. The price schedule posted by carriers as the flight date approaches and tickets are sold is a great example of the *PED* models, but in order for the results in this chapter to hold for the entire industry, we require that the prices of other tickets vary accordingly with the one-way ticket fares. One of the authors believed that this is true, but the other was skeptical. Showing this formally is beyond the scope of this chapter and would require working with a dataset that encompasses more complex itineraries.

CHAPTER III

STOCHASTIC PEAK-LOAD PRICING WITH REAL-TIME DEMAND LEARNING IN THE U.S. AIRLINE INDUSTRY

A. Introduction

The term dynamic pricing, most commonly known as ‘yield management,’ is used to describe pricing and inventory control decisions. It is important in industries that deal with perishable products such as airlines, where unsold seats perish when the flight leaves the gate. Similar examples involve hotel rooms, fashion apparel, cabins on cruise liners, car rentals, entertainment and sporting events, and restaurants. In all these cases the seller can improve its revenues by dynamically adjusting the price of the product rather than committing to a price schedule or a unique price throughout the selling period. Demand uncertainty plays an important role in airline’s dynamic pricing because tickets are sold in advance with prices being set when carriers have limited information about the total number of potential consumers. Moreover, capacity is set in advance and can only be modified at a relatively high marginal cost. When demand is relatively low, unsold tickets are of little value for the carrier. Likewise, when demand is relatively high the airline may give up important profits if some consumers with a relatively high willingness-to-pay have to be rationed. This may be the case of a business traveler who has a relatively large willingness-to-pay for a ticket but arrives when there are no tickets left.

Learning about a price sensitive demand as sales progress and the flight date nears is crucial for airlines to price accordingly. The shadow cost of capacity for the seats on a flight will be different at different points in time prior to departure depending on the expected demand. When the probability that demand will exceed capacity is

large, the shadow cost of capacity is large. Peak-load or congestion pricing, defined as the practice of charging higher prices during peak periods when capacity constraints cause marginal costs to be high, is the pricing strategy that takes into account this shadow cost. Borenstein and Rose [10] provide a clear distinction between two types of peak-load pricing in airlines. *Systematic* peak-load pricing reflects variations in the expected shadow cost of capacity at the time the flight is scheduled and before any ticket is sold, while *stochastic* peak-load pricing that reflects uncertainty about individual flights that is resolved as the flight date nears and tickets are sold. In this paper we control for *systematic* peak-load pricing and analyze the impact of demand learning and *stochastic* peak-load pricing on fares. If at the moment the ticket is sold carriers expect to have a *peak* flight, they will charge higher fares. Moreover, expected *off-peak* flights will be associated with lower fares.

Despite the large theoretical literature on airlines' pricing in economics, marketing and operational research, there exists little empirical understanding on how carriers are actually setting their fares and the dynamics that govern their evolution as the flight date nears. This is the first empirical paper that evaluates the very intuitive predictions of *stochastic* peak-load pricing in airlines and to test whether airlines can reduce the cost of demand uncertainty by responding to early sales. One of the reasons why this hasn't been done before is the lack of adequate data.¹ While most of the empirical research in airlines uses the *Bureau of Transportation Statistics'* *DB1B*, which is a 10% random sample of tickets, recent research has begun analyzing more detailed data that allows tracking day by day decisions by airlines. Stavins [52], and more recently Chen [17], and Bachis and Piga [4] among others look at offered fares by airlines. However, none of these papers has information on inventories of

¹Another could be the challenge of coming up with an adequate empirical test.

seats at each offered fare. This paper takes advantage of a unique U.S. airline’s panel disaggregated at the ticket level that contains the evolution of offered fares and seat inventories over a period of 103 days for 228 domestic flights that departed on June 22, 2006. The data collection resembles experimental data which controls for product heterogeneities and ‘fences’ that segment consumers. This is key, since many price discrimination tools that define ticket characteristics (e.g. saturday-night-stayover) are also used for *stochastic* peak-load pricing to reduce demand uncertainty.²

To test whether airlines learn about the demand and implement a *stochastic* peak-load pricing strategy, the empirical section initially obtains the optimal price schedule under *no* demand learning by calibrating the *ex-ante* distribution of demand states. This is done using information on sold-out probabilities and forecasted values of occupancy rates. The sold-out probabilities are calculated using a second dataset from *Expedia.com*, and the forecasted occupancy rates are obtained using time-series data on occupancy rates from the *T-100* of the *Bureau of Transportation Statistics*. Under the Prescott [48]’s type of models, this distribution should give us the optimal price schedule, which holds through the entire selling horizon as long as airlines do not learn about the demand or if price commitments hold. The basic idea in the testing is to analyze whether airlines deviate from this *ex-ante* optimal price schedule as information about the demand is revealed. To capture the information about the demand that is revealed as sales progress, the paper uses the techniques described in Racine and Li [49] and estimates a nonparametric model using both categorical and continuous data based on flight- and route-level information to explain the evolution of sales. The nonparametric results are used to construct a threshold variable that

²A round-trip ticket that involves *peak* flights may not benefit from a saturday-night-stayover discount even if the stay has a Saturday night. Moreover, if it involves *off-peak* flights, then the discount will have an *stochastic* peak-load pricing component and a price discrimination component.

can distinguish between different expected demand states at different points prior to departure for every flight in the sample. This threshold variable is later utilized to control the regime shift in the estimation of a panel endogenous threshold model as described in Hansen [38] and Hansen [39].

The results are consistent with the *stochastic* peak-load pricing predictions. The panel endogenous threshold estimates found the existence of two pricing regimes. In expected *peak* flights, where demand is more likely to exceed capacity and the shadow cost of a seat is high, airlines set higher fares. For expected *off-peak* flights where demand is expected to be low, the shadow cost of capacity will also be low. Hence, airlines will set lower fares to promote sales. To control for the potential interaction between previous price levels and cumulative sales, the paper also estimates a dynamic panel, where the assumption of strict exogeneity of the regressors is relaxed. The GMM dynamic panel results from the *difference* estimator, as explained in Holtz-Eakin *et. at.* [41] and Arellano and Bond [1], and the *system* estimator, as described Arellano and Bover [2] and Blundell and Bond [8], were found to be consistent with the two pricing regimes and the *stochastic* peak-load pricing predictions found before. Based on the *system* GMM estimates, evaluated at the sample average fare of 291.09 dollars and for a 100 seats airplane, selling one more seat increases fares by 1.21 dollars in an expected *off-peak* flight while increases fares by 1.72 dollars in a expected *peak* flight.

By testing for *stochastic* peak-load pricing and demand learning, this paper explains an important source of price dispersion as well. Borenstein and Rose [10] calculate that the expected absolute difference in fares between two passengers on a route is 36% of the airline's average ticket price. One cost-based source of this price dispersion is *stochastic* peak-load pricing. Even though the figures found in this paper are not directly comparable to this 36%, we find that an increase of one standard

deviation in capacity utilization increases *peak* fares by 15.8% *within* flight standard deviations more than *off-peak* fares. This estimate is after controlling for *systematic* peak load pricing, unobserved flight and route characteristics, and ‘fences’ that restrict consumers that are commonly used as price discrimination tools.

By focusing on the role of the evolution of inventories on dynamic pricing, this paper is able to identify three components in the evolution of fares as the flight date nears. First, the *stochastic* peak load pricing component as the difference in fares between *peak* and *off-peak* flights. Second, the effect of demand uncertainty and costly capacity on fares as explained in the theoretical works by Prescott [48], Eden [25], Dana [21], and more recently by Deneckere and Peck [24]. The importance of this effect to explain price dispersion has been previously documented empirically in Escobari and Gan [32]. Finally, the third component, advance purchase discounts. This last component is consistent with the price discrimination argument in Dana [19], where the existence of second degree price discrimination takes the form of advance purchase discounts. Moreover, it is also consistent with the existence of advance purchase discounts under an uncertain peak demand period in Gale and Holmes [34] and advance purchase discounts with perfectly predictable peak demand times in Gale and Holmes [35]. It is important to mention that both of the works by Gale and Holmes do not consider the shadow cost of capacity and there is no cost-based price variation. The predicted price dispersion suggests price discrimination.

The organization of the paper is as follows. Section B presents a model of pricing under demand uncertainty that extends to the empirical testing. The data is explained in section C. The empirical model is presented in section D. Finally, section E concludes the paper.

B. Airline pricing under demand uncertainty

Airline pricing has three basic characteristics that make its study fascinating. First, capacity is fixed and can only be augmented at a relatively high marginal cost. It is unlikely for carriers to change the size of the aircraft once they have already started selling tickets. Doing so would involve a large rescheduling of the fleet and airport slots. Second, air tickets expire at a point in time; once the plane departs carriers can no longer sell tickets. Tickets that haven't been sold by then have little value to the carrier. On the other hand, the carrier may still want to reserve a certain number of seats if it expects to have last-minute travelers who are willing to pay substantially higher fares (see Lin [44]). Finally, the third characteristic, there is uncertainty in the demand. This becomes crucial because airlines sell in advance and fares have to be set when carriers have limited information about the total number of potential passenger that will show up to get a ticket. Under this basic scenario, it is key for carriers to learn about the final state of demand as tickets are sold and the departure date nears. If information about a price sensitive demand becomes available, airlines will want to adjust their prices accordingly to maximize profits. These characteristics, common to various industries originated a large amount of theoretical literature on optimal pricing of a perishable non-renewable asset with stochastic demand. However, there is still few empirical understanding about how actual prices are set. This is the first paper that provides evidence of *stochastic* peak-load pricing in airlines. Moreover, it is also the first to presents empirical evidence that shows that airlines learn about the demand as the departure date nears.

Airline pricing, nevertheless, is much more sophisticated than dealing with an inventory of seats that expire at a point in time. Usually air tickets involve complex itineraries and carriers exploit 'fences' such as saturday-night-stayover requirement,

minimum- and maximum-stay, nonrefundable purchases, frequent flyer miles, black-outs, days in advance requirements, or volume discounts to segment consumers. The nature of the dataset used in this paper resembles experimental data, and jointly with the econometric techniques employed, it controls for all these fencing devices and complex itineraries to allow us focusing on the pricing of an inventory of seats that expire at departure. Therefore, this overview of airline pricing under demand uncertainty discusses just this case. We begin with the case where there is *no* demand learning and capacity is costly. We then explain the implications for pricing decisions when airlines learn about the demand through the information contained in early sales. At the end of the section we discuss about the implications of two cost-based sources of price dispersion for airlines: *systematic* and *stochastic* peak-load pricing.

1. Pricing without demand learning

The simplest model that explains dispersed prices for a homogeneous good under costly capacity and demand uncertainty is Prescott [48]. He considers a model of hotel rooms where prices are set *ex-ante* —before the total number of buyers is known—. Motivated with the airlines’ problem, the Prescott model assumes that there is a stochastic demand n for homogeneous airline seats with a probability distribution function $F(n)$. Consumers are identical and purchase only one seat if the price is lower than a reservation value \bar{p} . The equilibrium prices will be dispersed with $H(p)$ being the equilibrium number of seats priced at p or below. Travelers observe all prices and buy the less expensive unit available. In equilibrium, a seat priced p will be vacant with probability $F[H(p)]$. Let λ be the unit cost of capacity incurred on *all* units, whether these units are sold or not. In a perfectly competitive market, the zero expected profit condition implies that expected revenue should be equal to the unit cost of capacity, $[1 - F[H(p)]] \cdot p = \lambda$. This last equation can also be written as

$$p = \frac{\lambda}{1 - F[H(p)]} \equiv ECC \quad (3.1)$$

for all $p \in [\lambda, \bar{p}]$. Any price offered in equilibrium must be equal to the unit cost of capacity divided by the probability that a unit offered at that price will be sold. Dana [21] interprets this last term as the *effective cost of capacity (ECC)*, which is the revenue the carrier must earn if the seat is sold in order to cover the unit capacity cost it incurs whether or not the seat is sold. The intuition from this result is simple. Consider the case where there are two equally likely demand states and the cost of holding a seat in the aircraft is \$100. If a given seat in the aircraft is only sold during the high demand state, the carrier has to charge a fare equal to $\$100/0.5 = \200 , to compensate the times the seat is not sold during the low demand state.

The key implication from the Prescott model is that lower-priced units will be sold with higher probability and higher-priced unit with lower probability. Therefore, sellers face the trade-off between price and the probability of making a sell. Even though Equation 3.1 is constructed for a perfectly competitive market, Dana [21] derives an analog of Equation 3.1 for perfect competition, monopoly, and oligopoly. In all cases the key implication is the same. However, in noncompetitive markets, the effect of *ECC* on fares has to be adjusted by the size of the markup.

Prescott's model was later formalized by Eden [25] in a setting where consumers arrive sequentially, observe all offers and after buying the cheapest available offer they leave the market. Eden derives an equilibrium that exhibits price dispersion even when sellers are allowed to change their prices during trade and have no monopoly power. It is important to realize that the absence of price commitments alone is not enough to generate *stochastic* peak-load pricing. Information about the final state of the demand has to be revealed to price accordingly. Prescott's "hotels"

model, as pointed out by Eden [25] and Lucas and Woodford [45], has an interesting time-consistency property. Deneckere and Peck [24] explain that even if firms could sequentially compete by choosing a price after each market transaction, fares will still follow Equation 3.1. The requirement for fares to depart from the original price schedule predicted by Equation 3.1 is that some additional information about the final state of the demand is revealed as the flight date approaches.

2. Pricing with demand learning

If fares can adjust as information about the demand is revealed over time, the predictions from the dynamic pricing literature are very intuitive. As explained in Lin [44], when an airline sells seats for the same class, the fares offered will be different depending on the time to departure and the current seat inventory. The airline has incentives to promote sale when departure time is approaching and inventories are high. On the other hand, the airline may still want to reserve some seats if it expects to have some last-minute travelers willing to pay substantially higher prices. Lin [44] presents a theoretical dynamic pricing model where customers arrive in accordance with a conditional Poisson process. Sellers learn about the final state of the demand as sales move forward and the optimal price adjusts dynamically to maximize expected total revenue. The results indicate that higher prices should be set when demand is expected to be larger. A similar result is found in Gallego and van Ryzin [36] and Kincaid and Darling [42], where at a given point in time prior to departure optimal prices will be higher if the inventory is lower, signaling a higher demand state.

There is a key difference between the models presented in Kincaid and Darling [42], Gallego and van Ryzin [36], and Lin [44], and the Prescott-type of models. In the first ones all costs related to the production of the seat are sunk, so the value for the seller for an unsold item is zero. If demand is expected to be low, fares are

allowed to drop, result that explains the ‘last minute deals’ or cheap fares that airlines offer in some flights in order to promote sales. On the other hand, Prescott-type of models assume that capacity is costly; airlines have to be able to cover the unit cost of capacity adjusted by the probability of sale for each of the seats. Moreover, since there is no demand learning, fares will always increase as sales progress, prediction that can be easily seen from Equation 3.1. The simplification of *no* costly capacity in the first type of models is not realistic, at least for airlines pricing, while the existence of *no* demand learning in the second type seems also restrictive. However, as information about the demand becomes available, pricing accordingly gains significance while dealing with costly capacity loses attractiveness. This difference in predicted outcomes as information about the demand becomes available will let us identify the existence of demand learning.

The basic information carriers use to learn about the final state of the demand is *realized* demand up to a given point prior to departure. This is how many seats have been sold up to a given point in time, which contains some information about the speed of selling tickets and can be used to predict whether final demand may exceed capacity. Models that take into account *realized* demand in its pricing policies include Gallego and van Ryzin [36], Chatwin [16], and Lin [44]. The exact nature of how information about current sales is taken into account to forecast the final state of the demand is not necessarily important in this paper. Airlines may have very different ways to use information about early sales to adjust price later on, however, all these models have the same testable prediction. At a given point in time, lower inventory levels, signaling an expected higher demand, results in higher prices. Likewise, if time passes by and no seats have been sold, this is evidence of moving into a low demand state and lower fares should be set. As will be pointed out below, this prediction from operational research literature is equivalent to the *stochastic* peak-load pricing

prediction found in economics journals.

3. Peak-load pricing in airlines

Classic peak-load pricing models under certainty (*systematic* peak-load pricing, e.g. Boiteaux [11], Steiner [53], Hirshleifer [40], Williamson [58]) and peak-load pricing models under uncertainty (*stochastic* peak-load pricing, e.g. Brown and Johnson [12], Visscher [55], Carlton [14]) they all suggest charging higher prices during peak times and lower prices during off-peak times. However, as explained in McAfee and te Velde [46] these models poorly suit airline pricing. One important reason is that they assume the existence of a spot market, where all consumers in the *peak* demand pay a higher price and all consumers in the *off-peak* demand pay a lower price. However, advance purchases and different expectations about the demand prior to departure are an important ingredient in the pricing problem.

For the airlines, fluctuations in the demand can be broken down into two parts. The *deterministic* component that refers to fluctuations in the demand which are known to carriers before selling starts, and the *stochastic* component of the demand for a flight, that is orthogonal to all information carriers have at the time of scheduling. As explain in Borenstein and Rose [10], these two components of demand give rise to two different types of peak-load pricing in airlines.

a. Systematic peak-load pricing

As explained in Boreintein and Rose [10], fluctuations in capacity utilization across flights imply different opportunity costs of providing airline service. As a result, prices will depend on when a particular customer travels. For flights that depart during peak times of the day, peak days of the week, or special holidays, most of the aircrafts will be in use, so the expected shadow cost of aircraft capacity will be

high. On the other hand, during off-peak times the shadow cost may be close to zero. At the airport level, peak times congestion can also be associated to higher costs.³ *Systematic* peak-load pricing will reflect variations in these costs, with prices being higher in peak periods and lower in off-peak periods.

In order for carriers to follow a *systematic* peak-load pricing strategy, they require prior knowledge of peak flights (or peak periods) since *systematic* refers to variations in the demand known to carriers when they create their flight schedule. In their empirical study, Borenstein and Rose [10] control for *systematic* peak-load pricing under the assumption that this one is correlated with the variability in airlines' fleet utilization rates and airports' operation rates. In the present paper we control for variation in prices due to *systematic* peak-load prices as well, but neither this paper nor Borenstein and Rose [10] are able to measure the effect of this type of peak-load pricing on fares. Escobari [31] provides empirical evidence supporting *systematic* peak-load pricing by looking at variations in *ex-ante* known demand intensities. Moreover, he estimates a congestion premia and provides support for the main empirical prediction in Gale and Holmes [35], less discount seats on peak periods.

b. Stochastic peak-load pricing

The existence of *stochastic* peak-load pricing depends on two conditions, the degree of price flexibility that carriers have over the selling horizon, and the information about the demand that is revealed as sales progress. Price flexibility alone is not enough because even if carriers can adjust prices, if no new information about the demand

³Using simulations for the Minneapolis-St. Paul airport, Daniel [22] finds that airport congestion pricing would reduce net social costs by about 24% by smoothing out demand of landings and takeoffs. In a more recent paper, Brueckner [13] finds that airport congestion is fully internalized under a monopoly carrier, but not in under a oligopoly.

is learned, they will set the same prices they set before sales begun. If these two conditions are satisfied, Crew and Kleindorfer [18] explain that the optimal *stochastic* peak-load pricing will depend on the probability at the time the ticket is sold that demand will exceed capacity and the expected shadow cost if this happens. That is, the shadow cost of capacity will depend on the expectations about the demand. Under a price sensitive demand, in an expected *peak* flight, when demand is expected to be high and more likely to exceed capacity, the shadow cost will be large and carriers will set higher fares. Likewise, in an expected *off-peak* –low expected demand– flight where the capacity constraint is unlikely to be binding, carriers will set lower fares in order to promote sales. An interpretation of the shadow cost of capacity is related to the opportunity cost of selling the next seat of the aircraft. In an expected *peak* flight this one is high because if the seat is not sold now, it can potentially be sold later on at a relatively high price.

The price sensitive demand faced at a flight level captures an oligopolistic market with some consumers deciding not to travel at higher prices and others and other consumers deciding to shift to different flights. This demand shifting is central in peak-load pricing models, because given capacity constraints it is a way to increase output. In a model where the peak time is unknown, Dana [20] shows that the equilibrium price dispersion that arises from demand uncertainty, costly capacity and price rigidities can achieve the same efficient demand shifting than peak-load pricing. In this case, under *stochastic* peak-load pricing, under no price rigidities and once carriers learn whether the flight is peak, the shifting should be at least as large as in Dana [20].

As explained in Dana [21] ‘yield management,’ is utilized (1) to implement peak-load pricing, (2) to implement third degree price discrimination using ‘fences’ to separate consumers, and (3) to deal with the pricing under uncertain demand for

a perishable asset. Some theoretical yield management model such as Kincaid and Darling [42], Gallego and van Ryzin [36], Chatwin [16], or Lin [44] that explain the dynamic pricing of inventories that expire at a point in time are consistent with the *stochastic* peak-load pricing for airlines presented in Borenstein and Rose [10]. That is, a lower inventory level at a point in time —signaling an expected *peak* demand— rises the optimal price. Moreover, time spent without selling tickets —signaling a shift to a lower expected demand state— lowers the price. Even though these yield management models do not talk about peak-load pricing, because as explained in McAfee and te Velde [47], they focus at the operational and revenue management decision level, they do not represent a competing interpretation of the observationally equivalent empirical implication. They can be interpreted as *stochastic* peak-load pricing models once motivated with the existence of a shadow cost of capacity. Furthermore, the price sensitive demand they model can capture the demand shifting argument in peak-load pricing.

C. Data

The paper has two main sources of data, the Online Travel Agency (OTA) *Expedia.com*, used to build two datasets, and the *T-100* from the *Bureau of Transportation Statistics* used to construct a third dataset. The first dataset from *Expedia.com* is a panel with 228 cross-sectional observations and 35 observations in time, making a total of 7980 observations. Each cross sectional observation is a specific carrier’s non-stop one-way flight in one of the 81 routes considered, where a route is a pair of departing and destination cities. The observations in time start 103 days prior to departure and were gathered every three days up until one day prior to departure, making the 35 observations in time. All flights depart the same date, Thursday June

22, 2006. The carriers considered are American, Alaska, Continental, Delta, United, and US Airways. The number of flights per carrier was chosen to make sure that the share of each of these carriers is close to its share in the US airlines' market. This dataset has similar characteristics the one used in Stavins [52], but with two important differences. First, the data here is a panel and second, it has information about seat availability at each fare, where fare is the cheapest available economy class fare. The only two previous papers that work with such a detail information on prices and inventories are Escobari [31] and Escobari [32]. The panel structure allows to control for unobservable cost differences. The second dataset from *Expedia.com* was collected to obtain an estimate of the sold out probabilities for each of the 81 routes.

The third dataset comes from the *T-100* obtained from the *Bureau of Transportation Statistics*. This is a panel containing average load factors at departure for the same 81 routes over the period 1990 to 2005. This dataset will be useful to estimate the expected number of tickets sold in each route, used to derive the *ex-ante* demand uncertainty.

Figure 6 illustrates the evolution of average fares and the standard deviation of fares for the 228 flights at different number of days prior to departure. There is an important increase in posted fares during the last ten days. Furthermore, there is also an increase in the standard deviation of fares, indicating that fares will be more dispersed close to departure. This is some indication that fares will be increasing more for some flights than for others, including the possibility that fares in some flights do not increase at all during the last days. This would be the case of the flight depicted in figure 7.

This flight presented in figure 7 also illustrates the intuition behind the *stochastic* peak-load pricing we are testing in this paper. This is flight Delta 1588, covering the 2111 miles between Atlanta, GA (ATL) and San Jose, CA (SJC) with a Boeing 737-




Fig. 6. Average and standard deviation of fares




Fig. 7. Fares and load factors (Delta 1588 ATL-SJC)

800 that has a total capacity of 199 economy class seats, departs at 7:54 p.m. and arrives at 10:00 p.m. Figure 7 shows the evolution of fares, inventories of seats and the expected evolution of seat inventories for a period of 103 days prior to departure. As required by Prescott [48], Eden [25], and Dana [21], fares represent the cheapest available fare for a given flight at each point in time prior to departure. A detailed explanation of why the fares from *expedia.com* used in this paper are representative for the industry is presented in Escobari and Gan [32] section 2.2. The evolution of inventories is best viewed as the ratio of available seats to total seats in the aircraft. We refer to this ratio as the load factor, which is a ticket level load factor. The airline literature defines load factor only once the plane departed as the percentage of seats filled with paying passengers. Our load factor varies for each posted fare and will go from zero when the plane is empty to one when it is full. In this Delta flight 1588, the load factor went from 0.235, 103 days prior to departure to 0.995 one day prior to departure. An interesting feature on this particular flight is that load factor is increasing monotonically. The decrease between 67 and 61 days prior to departure may be because some tickets have been reserved and never bought or maybe bought but cancelled later. The exact calculation of the expected load factor will be explained below. For now, it is just important to know that it is a measure of how the carrier, Delta, expects sales to evolve over time for this particular flight under normal conditions and price commitments. If at a given point in time the actual load factor is significantly above expected load factor, it is reasonable for the carrier to believe that demand will exceed capacity and a *stochastic* peak-load pricing strategy would suggest charging higher fares. This is exactly what happened during the period between 94 and 67 days prior to departure. Load factor was relatively high as compared to the expected evolution of load factor under “average” conditions. At this consumers’ arrival rate, demand would exceed capacity. Even if the arrival rate of

future consumers is independent of this high arrival rate and just follows an “average” arrival rate, most likely demand would still exceed capacity if Delta keeps the same price schedule. The optimal *stochastic* peak-load pricing would be to set higher fares, which is exactly what they do. When load factor decreases between 67 and 61 days prior to departure, fares also decrease. Furthermore, notice that during the last month of sales load factor increased even to higher levels, but this increase can be explained by the expected load factor, so there is no reason to charge higher fares at this point.

The evolution of sales in each flight in the sample is the result of tickets being bought across a huge number of potential itineraries, where the observed leg may be just part of a larger trip. What is important to realize is that the fare charged by the carrier is the carrier’s response to the level of inventories, and this one has its own dynamics. Here we are just making explicit what previous studies the work with non-transactions data implicitly assumed, e.g. Stavins [52], Chen [17], and McAfee, R.P., te Velde [46]. It is reasonable to believe that fares for more complicated itineraries vary accordingly with the one way fare. Bachis and Piga [4] explain how some European carriers price all its legs independently, so there is no extra charge for one-way tickets. Actually, observing higher fares on one-way tickets is perfectly consistent with the predictions in Prescott [48], Eden [25], and Dana [21], where earlier purchasers benefit from lower fares. The idea is simple, a round-trip fare is the combination of two parts. Because the return date is further away from the purchase date, the second part is being bought with more days in advance than the first, with the presumably lower load factor. Hence, a round-trip ticket, measured as the summation of these two parts will be less expensive than just multiplying the first part of the ticket by two.

D. Empirical model

The empirical model developed in this section is based on two testing procedures for *stochastic* peak-load pricing. The first one is just based on analyzing how the pricing of the next available ticket differs depending on the expectations of demand. The testing will study whether an expected peak or an expected off-peak flight follow different pricing strategies. The second testing procedure is built on the models developed by Prescott [48], Eden [25], and Dana [21], where price dispersion exists in a setting with capacity constraints and demand uncertainty. The key feature in these models is that there is *no* demand learning as sales progress. Therefore, we start building the price schedule based on these models and then we test how pricing differs as carries learn about the demand. For the second testing procedure we derive the *ex-ante* distribution of demand uncertainty, which is before any information about actual sales is revealed. This will let us calculate the *effective cost of capacity*, which should have a positive impact on fares.

Common to both testing procedures, using nonparametric techniques we then develop a measure of the evolution of the expected number of seats sold. This measure is exogenous to the actual evolution of sales, so by comparing actual sales with expected sales at any point prior to departure we can obtain information on the likelihood that demand will exceed capacity. An endogenous panel threshold model is then estimated to separate between expected peak and expected off-peak flight, allowing for different pricing strategies in different regimes. To control for potential enogeneity, the empirical section closes with the estimation of dynamic panels with a exogenous selection of the threshold.

1. Ex-ante distribution of demand uncertainty

The *ex-ante* distribution of demand uncertainty refers to the distribution of arriving consumers known to the carrier before any ticket is sold. Based on this distribution, Prescott [48] showed that equilibrium prices will be dispersed. In this subsection we calibrate the *ex-ante* distribution of demand uncertainty. Under price commitments or if no information about the final state of the demand is revealed as tickets get sold, this *ex-ante* distribution of demand uncertainty should explain the observed price dispersion.

There exists uncertainty in the demand because carriers do not know *ex-ante* the total number of passengers that will buy tickets. Consider the case of having an infinite number of demand states. Let N_h be the number of consumers who arrive at demand state h , where $h = 0, \dots, \infty$ and $N_h \leq N_{h+1}$. This last inequality imply that consumers who arrive at demand state h will also arrive at a higher-numbered demand state $h+1$. Define a batch as the additional number of travelers who arrive at demand state h when compared to the immediate lower demand state $h-1$, therefore batch h is given by $N_h - N_{h-1}$ with the first batch given by N_0 .

Each demand state h occurs with probability ρ_h . Because all demand states have at least N_0 travelers, the probability that N_0 travelers arrive is $Pr_0 = \int_0^\infty \rho_\kappa d\kappa = 1$. In general, the probability that N_h travelers arrive is given by $Pr_h = \int_h^\infty \rho_\kappa d\kappa$, the summation of all demand states that have at least N_h consumers. Assume that each batch has one consumer buying a ticket, hence the probability of selling seat h is the summation all demand states that have at least h travelers buying a seat. Additionally, when demand states are normally distributed $\rho_h = \phi_h$, with ϕ being the *pdf* of a normal distribution, the probability of selling seat h is given by equation 2.20, with Φ the *cdf* of a normal distribution. This is given the equilibrium distribution

of prices, so this schedule of selling probabilities will hold as long as carriers do not depart from price commitments.

This Pr_h corresponds to the $F[h(p)]$ in Equation 3.1. To derive a measure of the *effective cost of capacity* and its impact on fares, we will calibrate the distribution of demand uncertainty at a route level. To do this we follow Escobari and Gan [32] and assume normally distributed demand states. The key feature that allows the calibration process is that demand states are censored when transformed to tickets sold. Once the aircraft is sold out, higher demand states are no longer observed. To get the values of the mean μ and the standard deviation σ , at the route level, for the normally distributed demand states we first need two pieces of information, the sold-out probabilities and the expected number of tickets sold for each of the routes.

a. Sold-out probabilities

The sold out probabilities for each of the 81 routes are obtained using the second dataset from *Expedia.com*. The fact that allows calculating these sold-out probabilities is that airlines and online travel agencies do not display their sold-out flights on their websites.⁴ First, a couple of weeks in advance when no flight was expected to be sold-out yet, we made a census of all the available non-stop flights in each of the 81 routes during seven days between February 22 and February 8, 2007. The total number of flights was 5,881. Then, late the night before each of those seven days, we counted the number of flights still available at each route. If a flight was no longer there, it was assumed to be sold-out. The calculated sold out probability is just the

⁴The reason, according to Roman Blahoski, spokesman of Northwest Airlines, is that they do not want to disappoint the travelers. Keeping the online display simple may also be a motive, and according to Dan Toporek, spokesman of *Travelocity.com*, “showing sold-out flights alongside available flights could be confusing.” Both of these quotes are from David Grossman, “Gone today, here tomorrow,” *USA Today*, August 2006.

ratio of sold out flights to total number of flights for each route.

b. Expected number of seats sold

The expected number of seat sold are calculated using the *T-100* from the *Bureau of Transportation Statistics*. From the *T-100*, we obtain the average load factors at departure time for the 81 routes over the period 1990 to 2005. Each of these 81 series is used to estimate an ARMA model. Then, using a one-step forecast we obtain the expected number of seats sold for 2006. For routes where the expected number of seats sold is high, meaning that most of the seats are expected to be sold, the calibration procedure will assign higher probabilities to higher demand states. The details of the estimation are available upon request.

c. Calibration

Let the underlying demand state h^* be distributed $N(\mu, \sigma^2)$ and let m be the total number of seats in the aircraft. The number of seats sold h is equal to demand state h^* before the plane sells out, $h = h^*$ if $h < m$, and equal to total number of seats in the aircraft, $h = m$, otherwise. The expected number of tickets sold is given by the first moment of the censored normal given in equation 2.21. $E(h|h < m)$ comes from the mean of a truncated density and the *pdf* and *cdf* are evaluated at the moment the flight sells out. Therefore, $\Phi((m - \mu)/\sigma)$ is interpreted as the sold out probability. With information on the sold-out probabilities obtained in subsection a and the information on the expected number of tickets sold obtained in subsection c, we use Equation 2.21 to obtain the values of $\hat{\mu}_j$ and $\hat{\sigma}_j$ at the route level.

2. Learning the stochastic demand

As carriers learn about the state of the demand they may want to depart from any price commitments to increase their profits. The way carriers use actual bookings to infer about the state of the demand can be complex and may differ across carriers, but once some information is revealed, the outcome predicted by the *stochastic* peak-load pricing is simple. *Stochastic* peak-load pricing suggests charging higher fares in expected peak flights, while charging lower fares in expected off-peak flights. To test if this is true, the first step is to separate between expected peak and expected off-peak flights.

Under ‘normal’ conditions, let’s say, when a flight is not expected to be peak nor off-peak, sales should have a natural evolution over time as the flight date approaches. The rate at which tickets are sold need not be constant in time and may differ from route to route or across carries. If tickets are sold faster than the ‘normal’ rate and at a given point prior to departure there are less seats left unsold than under ‘normal’ conditions, it would be reasonable for the carrier to believe that this is a peak flight. Clearly, this expected peak flight was not known to the carrier *ex-ante*, before the flight was opened for booking.

To test for the existence of demand learning with the corresponding *stochastic* peak-load pricing as the response to information about the final state of the demand we take the following steps. First, using nonparametric techniques we come up with a measure of the evolution of sales under ‘normal’ or average conditions. Then we estimate a panel endogenous threshold model to see whether there are different pricing regimes when the expectation of demand differs. The threshold variable that dictates the regime switch is the ratio of actual sales to expected sales at a point in time prior to departure. Higher sales relative to normal sales would be evidence of a

peak-demand flight. Finally, to control for potential endogeneity in the regressors, we estimate a dynamic panel with an exogenous distinction of expected peak and expected off-peak flights.

a. Nonparametric estimation of expected sales

In this section we come up with a measure of the evolution of sales under average or normal conditions. That is, we estimate an exogenous measure of expected sales as the departure date nears for each of the flights in the sample. This measure of the evolution of sales for each flight is expected to be captured by the flight, carrier and the route's characteristics. Consider the following nonparametric model of cumulative sales on flight, carrier, and route characteristics.

$$LOAD_{ijt} = g(DAYADV_{ijt}, \mathbf{X}) + \eta_{ijt} \quad (3.2)$$

The subscript i refers to flight, j to route, and t is time. Equation 3.2 is a panel estimated using kernel methods for mixed datatypes as explained in Racine and Li [49], and Li and Racine [43]. The dependent variable is $LOAD$, defined as the total number of seats sold divided by the total number of seats in the aircraft. The explanatory variables include the number of days in advance $DAYADV$ and \mathbf{X} . This \mathbf{X} has two flight level characteristics: departure time $DEPTIME$, and route concentration as measured by the *Herfindahl-Hirshman Index (HHI)*. Table 7 provides the summary statistics of these variables and detailed description is included in Appendix C.

$E(LOAD_{ijt}|DAYADV, \mathbf{X})$, which is the evolution of the expected cumulative sales for flight i , is obtained first by estimating Equation 3.2 using the observations from all other routes except the route from flight i as train data. Then flight i 's

characteristics are used as evaluation data. This means that Equation 3.2 is estimated 81 times, once for each route, and evaluated 228 times at the corresponding flight's characteristics. To illustrate part of the results, the estimated nonparametric expected sales at different points prior to departure and for different trip distances is shown in figure 8. This was done using all datapoints as train data and with the remaining variables held constant at their median values for the evaluation points. On average longer legs have larger load factors and travelers also decide to book earlier.

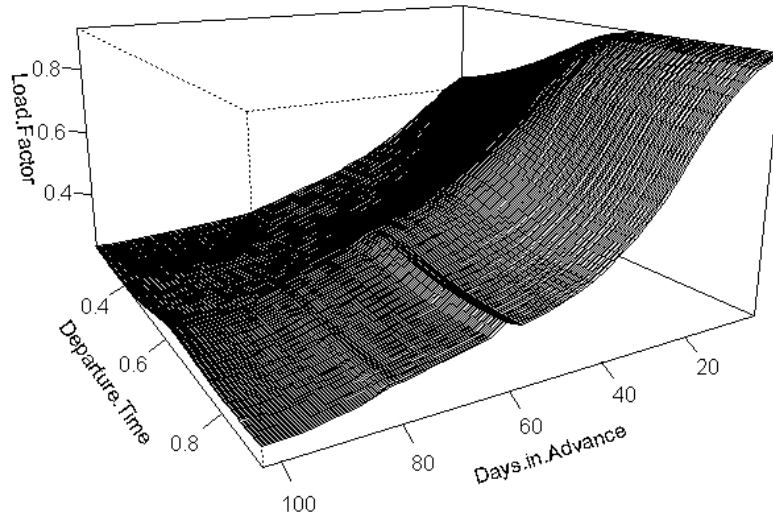


Fig. 8. Estimated nonparametric expected sales

b. Endogenous threshold estimation

Under the existence of menu costs, the benefits from switching pricing strategies may still be lower than the costs. Therefore even if carriers get to learn about the state of the demand, some degree of price flexibility is necessary in order to have *stochastic* peak-load pricing. Moreover, demand is never fully learned; as sales take place and some information is revealed about the demand, there will always be some uncertainty

Table 7. Summary statistics

Variables	Mean	Std. Dev.	Min.	Max.	Obs.
<i>For the nonparametric estimation</i>					
<i>LOAD</i>	.509	.252	.012	1.000	7933
<i>DAY ADV</i>	52.289	30.154	1.000	103.000	7933
<i>DIST</i>	1104.380	620.720	91.000	2604.000	7933
<i>HHI</i>	.684	.287	.259	1.000	7933
<i>For the calibration of demand uncertainty</i>					
Forecasted LF	0.739	0.083	0.469	0.890	81
Sold-out probability	0.227	0.104	0.037	0.571	81
<i>For the endogenous panel threshold estimation</i>					
<i>FARE</i> ^a	291.087	171.879	54.000	1224.000	7933
<i>LOAD</i>	0.509	0.252	0.012	1.000	7933
<i>ECC</i>	1.557	0.940	1.000	11.668	7933
$E(LOAD DAY ADV, \mathbf{X})$	0.506	0.219	0.089	0.991	7933
$E(LOAD)$	0.509	0.181	0.299	0.882	7933
$\hat{S}_{nonparam.}$ (% dev.)	-0.014	0.184	-0.931	1.040	7933
$\hat{S}_{average}$ (% dev.)	0.001	0.389	-0.961	1.943	7933

Notes: ^a The standard deviation for *FARE* between flights is 152.933, and within is 78.751.

remaining about its final state. Under this scenario, carriers will want to wait until they have enough evidence toward having an expected peak or an expected off-peak flight before deciding to switch its pricing strategies. This suggest the existence of different pricing regimes for different expected demand states rather than a continuum of fully adjustable fares sensitive to every new peace of information about the expected final state of the demand. In this subsection we estimate and endogenous panel threshold model to test for the existence of different pricing regimes. The different regimes are given by the different expectations of the final state of the demand.

The $E(LOAD_{ijt}|DAY ADV_{ijt}, \mathbf{X}_{ijt})$ estimated in the previous section is a measure of the expected evolution of sales for flight i under average conditions and it is independent to the actual evolution of sales given by $LOAD_{ijt}$. This is because the observations of the load factor of flight i are not included in the nonparametric estimation of $E(LOAD_{ijt}|DAY ADV_{ijt}, \mathbf{X}_{ijt})$ for the same flight i . By independent we mean that if flight i is expected to be a peak-flight, is independent from the average

of the other $-i$ flights from being peak-flights. Therefore, the ratio

$$S_{ijt} = \frac{LOAD_{ijt}}{E(LOAD_{ijt}|DAYADV_{ijt}, \mathbf{X}_{ijt})} \quad (3.3)$$

contains the necessary information to know whether at time t prior to departure actual sales are high, low or about the same as compared to sales under average conditions. The basic information contained in equation 3.3 is how far actual sales are from sales under average conditions. Therefore, deviations from these mean or average conditions are important. One other potential specification would be to calculate

$$S_{ijt} = \frac{LOAD_{ijt} - E(LOAD_{ijt})}{E(LOAD_{ijt})}, \quad (3.4)$$

but this last specification of S would give exactly the same results as when using 3.3.

In a simplified version of this test, $E(LOAD_{ijt}|DAYADV_{ijt}, \mathbf{X}_{ijt})$ can be replaced by the average load factor across flights at each point in time prior to departure.⁵ If at a given point prior to departure the ratio in equation 3.3 is relatively large, it would be reasonable for carriers to think they are in a peak period and that expected demand will be greater than the allocated capacity. On the other hand, low values indicate that sales are low relative to average or normal sales and it would be reasonable for airlines to think they are in an off-peak period and some seats may be left unsold.

The threshold variable S_{ijt} has two interesting properties. Recall that the dataset was constructed in a way that all flights share the same departure date, hence they also share the same dates prior to departure. If, for example, sales are higher/lower during weekends, this should affect all flights and will change both, $LOAD_{ijt}$ and $E(LOAD_{ijt}|DAYADV_{ijt}, \mathbf{X}_{ijt})$, but the threshold variable S_{ijt} should remain unchanged. Here carriers are assumed to know whether specific dates affect sales (e.g.

⁵Both of these robustness checks were pointed out by Professor Hwang during the preliminary examination.

weekends) and take this into account in their calculations of expected demand. Higher or lower sales on a given point in time common to all flight and known to the carriers will have no impact on the definition of expected peak and expected off-peak flight. What is even more important, the construction of this ratio allows us to control for *systematic* peak-load pricing. During *ex-ante* known congested periods, *stochastic* peak-load pricing suggests that carriers will charge higher fares. As explained in Borenstein and Rose [10], this type of peak-load pricing arises at an airport or fleet level. Here, the most likely capacity constraint is given by the total number of aircrafts. As a result *systematic* peak-load pricing should affect *all* flights while the ratio S_{ijt} remains unchanged. The drawback in this approach is that we will not be able to measure the effect of *systematic* peak-load pricing on fares. For an estimation of the congestion premia on fares due to *systematic* peak-load pricing, see [31].

This section estimates a threshold model to test whether carriers have different pricing strategies for different expected states of the demand. The threshold variable that will control the shift between expected peak and expected off-peak flights is S_{ijt} . To avoid an arbitrary selection of the number of pricing regimes and selection of the threshold(s), we estimate the model using the panel threshold regression methods with individual-specific fixed effects of Hansen [38]. The equation to be estimated has the form

$$\begin{aligned} \ln(FARE)_{ijt} = & \delta_0 DAYADV_{ijt} + \delta_1 \ln(ECC)_{ijt} \cdot I(\hat{S}_{ij,t-1} \leq \gamma) \\ & + \delta_2 \ln(ECC)_{ijt} \cdot I(\gamma < \hat{S}_{ij,t-1}) + \nu_{ij} + \varepsilon_{ijt} \end{aligned} \quad (3.5)$$

where $I(\cdot)$ is the indicator function, S_{ijt} is the threshold variable and γ is the threshold. Moreover, ν_{ij} is the unobserved carrier- and flight-specific effect, ε_{ijt} is error term, and as before the subscripts i denotes flight, j is route and t is time. Another

way of writing Equation 3.5 is

$$\ln(FARE)_{ijt} = \delta_0 DAY ADV_{ijt} + \begin{cases} \delta_1 \ln(ECC)_{ijt} + \nu_{ij} + \varepsilon_{ijt} & \text{if } \hat{S}_{ij,t-1} \leq \gamma \text{ (off-peak)} \\ \delta_2 \ln(ECC)_{ijt} + \nu_{ij} + \varepsilon_{ijt} & \text{if } \gamma < \hat{S}_{ij,t-1} \text{ (peak)}. \end{cases}$$

For the case of Equation 3.5, the observations are divided into two pricing regimes depending on whether the threshold variable $S_{ij,t-1}$ is smaller or larger than the threshold. The regime-independent variables $DAY ADV$, is included to control for a time trend. Even though Equation 3.5 is illustrated for only one threshold, the actual estimation process test for the existence of up to three thresholds, allowing for up to four different pricing regimes. In the absence of regime changes, Equation 3.5 follows the form suggested by the theory under *no* demand learning in Equation 3.1.

Given the construction of the dataset we perfectly control for important sources of price dispersion observed in the industry (e.g. saturday-night stayover, minimum and maximum stay, different connections/legs, fare class, refundability). Furthermore, estimating the model using flight fixed effects allows controlling for unobservable time invariant characteristic, which includes all the time invariant control variables included in Stavins [52] (e.g. flight, carrier, and route characteristics). Flight fixed effects should also control for most of the *systematic* peak-load pricing as well, since by definition, this type of peak-load pricing arises at the aircraft level and should affect equally every seat on that flight. If *systematic* peak-load pricing affects prices across the same aircraft differently (e.g. less discount seats) then the departure time variable, $DEPATIME_{ijt}$, used in the estimation of Equation 3.2 should take care of it. The main coefficient of interest is the Effective Cost of Capacity ECC . Prescott [48]’s type of models predict a positive effect of ECC on fares. However, this is true

under no demand learning or under price commitments. With the specification of Equation 3.5, the coefficient on ECC is allowed to be different across flights and at different points prior to departure, depending on the expectations of the demand. As predicted by *stochastic* peak-load pricing, higher expected demand states will be associated with a greater impact of ECC on fares, while lower expected demand states will be associated with lower or even a negative coefficient on ECC .

The empirical specification is estimated as a constant elasticity model in $\ln - \ln$ form. This is because both variables $FARE$ and ECC are measured in dollars. Moreover, recall that $ECC = \lambda/Pr$, then estimating the equation using the logarithm of ECC allows separating its components in two. $\ln(\lambda)$ goes as part of the regression intercept while the coefficient on $\ln(Pr)^{-1}$ remains the same as the coefficient on $\ln(ECC)$. We can then interpret this coefficient as the impact of a percentage increase in ECC or a percentage decrease in the selling probability, Pr , on fares. The interpretation of this elasticity measure does not require knowing the value of λ . An alternative specification replaces $\ln ECC$ with $LOAD$ in Equation 3.5. The *stochastic* peak-load pricing analysis follows the same logic as with $\ln ECC$, however, the interpretation is somehow different. Here a change in $LOAD$ represents an increase capacity utilization.

In order to estimate the nonlinear specification in Equation 3.5 we follow the procedure proposed in [38]. First, to eliminate the unobserved carrier- and flight-specific effects, for a given γ and for each flight we obtain the deviations from the time averages. Stacking the data over all flights we obtain $Y = \mathbf{V}(\gamma)\delta + \varepsilon$, where Y and $\mathbf{V}(\gamma)$ are just the stacked fixed-effects transformation just explained on $\ln(FARE)$ and the set of explanatory variables respectively. Notice the values of the explanatory variables are a function of the value of the threshold. For any given γ , the vector of slope coefficients δ can be estimated by ordinary least squares to obtain $\hat{\delta}(\gamma)$.

Chan [15] and Hansen [39] recommend the estimation of γ by least squares, hence its estimator is

$$\hat{\gamma} = \arg \min_{\gamma} Y'(I - \mathbf{V}(\gamma)'(\mathbf{V}(\gamma)\mathbf{V}(\gamma)')^{-1}\mathbf{V}(\gamma)')Y. \quad (3.6)$$

After $\hat{\gamma}$ is found, the estimate for the slope coefficients is $\hat{\delta}(\hat{\gamma})$. Then the next step is to find out if the threshold is statistically significant. The null hypothesis of no threshold in Equation 3.5 can be characterized by $H_0: \delta_1 = \delta_2$. As explained in Hansen [38], classical tests have non-standard distributions because under the null γ is not identified. Therefore, we follow Hansen [37] and simulate the asymptotic distribution of the likelihood test by bootstrapping. The likelihood ratio to test H_0 is based on $F_1 = (SSE_0 - SSE_1(\hat{\gamma}))/\hat{\sigma}^2$, where SSE_0 is the sum of squared errors under the null after the fixed-effects transformation is made. Similarly, SSE_1 is the sum of squared errors of the fixed-effects transformation made on Equation 3.5. For a larger number of thresholds the idea is similar, with the important characteristic that sequential estimation is consistent. Therefore, in order to test for the number of thresholds, we allow for sequentially zero, one, two, and three thresholds. As in Hansen [38], the observations are first sorted on the threshold variable and the search of the threshold is restricted to specific quantiles. The more quantiles the finer the grid to which the search is limited. Bootstrapping simulates the asymptotic distribution of the likelihood ratio test. This likelihood ratio is used to test whether the threshold is statistically significant under the null of no threshold. When rejecting the null, one more threshold is included.⁶

Table 8 provides the results that test for the number of thresholds: the test statistics F_1 and F_2 , along with the bootstrap p-values and critical values. From the

⁶The estimation used 400 quantiles and 300 bootstrap replications for each of the bootstrap tests.

bootstrap p-values, the null of no threshold for the one threshold model is rejected at a 1% level in all specifications. However, no evidence of further thresholds is found. The results for the three thresholds model are not reported since non of the second thresholds was found to be significant.

Table 8. Tests for threshold effects: Ratio

	Average				Nonparametric			
	<i>LOAD</i>		$\ln(ECC)$		<i>LOAD</i>		$\ln(ECC)$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Test for a single threshold</i>								
F_1	41.184	30.376	20.683	35.499	37.892	41.525	47.190	47.906
p-value	0.940	0.585	1.000	0.685	0.050	0.055	0.055	0.055
Bootstrap critical values								
10%	133.531	70.878	163.987	83.595	32.998	34.825	40.935	39.822
5%	149.573	81.996	180.336	95.220	37.612	41.539	49.515	49.762
1%	174.258	110.502	212.933	110.032	49.592	56.726	59.246	63.687
<i>Test for a double threshold</i>								
F_2	44.862	20.949	11.086	2.014	36.038	30.297	22.609	21.401
p-value	0.833	0.787	0.906	1.000	0.087	0.733	0.320	0.320
Bootstrap critical values								
10%	109.002	53.401	56.947	41.378	32.741	27.063	33.939	36.473
5%	113.930	61.787	64.138	44.135	39.018	34.328	43.196	46.279
1%	150.697	71.721	107.076	72.550	59.374	65.577	55.646	55.956

Notes: Because non of the second thresholds was found significant, the tests for triple thresholds are not reported. Odd numbered columns have 6732 observations across 198 flights and even numbered columns have 7128 across 216 flights.

Because the original dataset is unbalanced and the testing procedure implemented in this section only allows for balanced panels, we work with two subsets of the data. The first one has 198 flights over 35 time periods (covering a period of 100 days prior to departure) is reported in the even numbered columns of table 8. The second has 216 flights over 34 time periods (103 days prior to departure) and is reported in the even numbered columns of table 8. Moreover, the first four columns were calculated using the simple average for $E(LOAD)$, while the last four columns use the nonparametric specification for the calculation of the expected load factor.

The point estimates for the thresholds in the specifications where the threshold is significant are presented in table 9, along with the asymptotic 95% confidence intervals. When using the simple average for the calculation of $E(LOAD)$, table 8 indicates that none of the specifications returned a significant threshold estimate.

The specifications from columns (5) to (8) in table 8 which use the nonparametric specification for $E(LOAD)$ were found to have a significant threshold estimate for the one-threshold model. As expected, all point estimates lie around one and the confidence intervals are very tight, indicating little uncertainty about the nature of this division. The results indicate the existence of two pricing regimes. The first pricing regime occurs when $\gamma < \hat{S}_{ij,t-1}$. Notice that in this regime actual sales are relatively larger than sales under average conditions, hence we call this the peak period pricing regime. The second regime is characterized by $\gamma \geq \hat{S}_{ij,t-1}$. This will be referred as the off-peak period pricing regime since actual sales are relatively lower than sales under average conditions.

The confidence interval construction shown in figure 9, tabulated for specification in column (3) of table 9, provides further insights for the threshold results. The point estimate is the value of γ at which the likelihood ratio is equal to zero. The confidence interval $[\bar{\gamma}, \underline{\gamma}]$, are the values for γ for which the likelihood ratio lies beneath the straight line. Moreover, there are no other major dips in the likelihood ratio, which would be evidence of a third pricing regime.

Table 9. Threshold estimates: Nonparametric

	<i>LOAD</i>		$\ln(ECC)$	
	(1)	(2)	(3)	(4)
γ	0.999	1.011	0.978	0.977
Asymptotic 95% confidence interval				
$\bar{\gamma}$	0.978	0.984	0.976	0.976
$\underline{\gamma}$	1.187	1.012	0.991	0.995

Notes: % percentage deviations. The test for a triple threshold not reported, since non of the second thresholds was found significant. Columns (1) and (3) have 6732 observations across 198 flights and columns (2) and (4) have 7128 across 216 flights.

The regression estimates for the single threshold model are presented in table 10. The first noticeable result is that columns (1) and (2) are very similar, while (3) and (4) also look alike. Thus, the two balanced subsamples yield very similar results. The figures in parentheses are White-robust t-statistics. The regime-independent coeffi-




Fig. 9. Confidence interval construction in single threshold model

cient *DAYADV*, included as a control for a time trend is highly significant in all four specifications. The coefficient on *DAYADV* in column (4) means that after controlling for capacity constraints and demand uncertainty, route, carrier and flight characteristics, ticket characteristics that segment consumers and *systematic* and *stochastic* peak-load pricing, buying a ticket one day in advance reduces the ticket price by 57.7 cents.⁷ This is a measure of second degree price discrimination in the form of advance-purchase requirements. As pointed out in Dana [19], for advance-purchase discounts to be classified as discriminatory, it is necessary to define an appropriate measure of costs. Prices are considered discriminatory when the price markups over costs are different for different consumers. In this analysis, the costs for different seats in the same aircraft may be different due to the existence of uncertain demand and

⁷This one is calculated using the average fare for the subsample used in the estimation of column (4). This is $285.85 \times -2.018/10^3 = -0.577$ dollars.

costly capacity. These different costs are captured by ECC . Finally, recall that we are fully controlling for other sources of second degree price discrimination such as Saturday-night-stayover.

Table 10. Regression estimates for the single threshold model: Nonparametric

Regressor	$LOAD$		$\ln(ECC)$	
	(1)	(2)	(3)	(4)
<i>Regime-independent coefficients</i>				
$DAYADV/10^3$	-1.562 (-8.660)	-1.357 (-8.040)	-2.308 (-18.240)	-2.018 (-16.183)
<i>Regime-dependent coefficients</i>				
$LOAD_{ijt} \cdot I(\gamma \geq S_{ij,t-1})$	0.249 (6.400)	0.222 (6.282)		
$LOAD_{ijt} \cdot I(\gamma < S_{ij,t-1})$	0.343 (10.424)	0.307 (9.911)		
$\ln(ECC)_{ijt} \cdot I(\gamma \geq S_{ij,t-1})$			-0.051 (-1.663)	-0.062 (-2.122)
$\ln(ECC)_{ijt} \cdot I(\gamma < S_{ij,t-1})$			0.152 (11.085)	0.136 (9.182)
SEE	289.167	285.581	291.428	287.979

Notes: The independent variable is $\ln(FARE)$. t-statistics in parentheses based on White-robust standard errors. All regressions are estimated with flight fixed effects, not reported. Even numbered columns have 6732 observations across 198 flights and odd numbered columns have 7128 across 216 flights. $I(\gamma < S_{ijt})$ is referred as the peak period, while $I(\gamma \geq S_{ijt})$ is the off-peak period.

The variables we are mostly interested in are the regime-dependent. From table 10 we observe that $LOAD$ in columns (1) and (2) and $\ln(ECC)$ in columns (3) and (4) are all highly significant and have a positive effect on fares in the peak regime. For the off-peak regime only the specification for the load factor has significant coefficients. In all four specifications the off-peak period regime, $\gamma \geq \hat{S}_{ij,t-1}$, has a lower coefficient than the peak period regime, $\gamma < \hat{S}_{ij,t-1}$. We know from the results in table 8 that the coefficients in both regimes are significantly different. The results from column (2), evaluated at the subsample average fare of 285.85 dollars indicate that in a 100 seat aircraft, having one seat less available increases fares by 63.5 cents in an expected off-peak flight while increases fares by 87.7 cents in an expected peak flight. Columns (3) and (4) require some additional care. The effect of ECC on fares as predicted by Prescott [48]’s type of models is positive. However, as sales progress and carriers learn about the state of the demand, the coefficient on ECC will be

the outcome of two different type of models. The Prescott [48]’s type and *stochastic* peak-load pricing. The later only predicts that fares will be larger during expected peak flights. Thus the only requirement on our regime-dependent coefficients is that the expected peak regime should have a larger coefficient than the expected off-peak regime. When capacity is not costly and expected demand is smaller than allocated capacity, carriers will be willing to sell the last seats in the aircraft for any price above the operating marginal cost (e.g. baggage transportation, soft drink and pretzels). Consequently the last seats could be priced very low and the coefficient on ECC could be negative indicating lower fares for later seats. However, the results are consistent with having costly capacity and provide important evidence supporting Prescott [48]’s type of models, already documented in Escobari and Gan [32]. Columns (3) and (4) show that fares respond positively to ECC in both peak and off-peak regimes. Furthermore, there is also an important evidence supporting the existence of *stochastic* peak-load pricing with the peak regime coefficient being greater than in the off-peak regime.⁸

Fares will be increasing at a higher rate during expected peak regimes. Carriers forecasting that demand will be greater than allocated capacity will set higher fares to increase their profits and sell the remaining available capacity to travelers with higher valuations. If price commitments were to prevail or if carriers do not learn about the state of the demand, the flight will still sell-out in a high demand period. However in the absence of *stochastic* peak-load pricing existing capacity will be allocated to travelers that arrive first and not necessarily to travelers with higher valuations sorted by higher prices. On the other hand, when a low demand flight is expected, fares will

⁸In this case a direct interpretation of the coefficient on ECC , as we did with $LOAD$, would not be entirely correct since ECC is constructed based on an *ex-ante* distribution of demand uncertainty.

increase at a lower rate. This is consistent with cheap fares offered close to departure and ‘last minute deals’. Airlines offer this kind of tickets when demand falls short and allocated capacity is likely to remain underutilized.

c. Dynamic panel with exogenous threshold

In the previous endogenous threshold estimation we used the methods described in Hansen [38] and Hansen [39] to identify two pricing regimes. This procedure developed for non-dynamic balanced panels required us to assume strict exogeneity of the regressors and to work with two balanced panels, subsets of the original unbalanced dataset. In this subsection we take care of these two issues. We will reestimate the model as suggested in Equation 3.1 to test for the existence of demand learning, but this time using the dynamic panel techniques as developed in Holtz-Eakin *et. at.* [41], Arellano and Bond [1], Arellano and Bover [2], Blundell and Bond and [8]. This will let us work with the entire unbalanced panel and relax the assumption of strict exogenous regressors. We will assume that the switch between the two pricing regimes is the same as the one estimated in the previous part. Specifically, the equation to be estimated is

$$\begin{aligned} \ln(FARE)_{ijt} = & \alpha \ln(FARE)_{ij,t-1} + \beta_1 DAYADV_{ijt} + \beta_2 PEAK_{ij,t-1} \\ & + (\delta_0 + \delta_1 PEAK_{ij,t-1}) \cdot \ln(ECC)_{ijt} + \nu_{ij} + \varepsilon_{ijt} \end{aligned} \quad (3.7)$$

The idea is the same as in the estimation of Equation 3.5. This means analyzing the effect of the effective cost of capacity on fares under no demand learning as suggested by Equation 3.1, while allowing for the existence of different pricing regimes when the expectation of future demand differs. As found in section b, here we allow for the coefficient of ECC on fares to have two possible values that represent the expected peak and expected off-peak regimes. The division between these two regimes

is assumed to be the same as before. Then the variable that dictates the shift is $PEAK_{ijt} = I(S_{ijt} > 1)$ in the ratio specification of S (same as $PEAK_{ijt} = I(S_{ijt} > 0)$ in the percentage deviation construction of S), with the 1 selected intuitively when actual sales are larger than sales under average conditions. Hence, $PEAK$ takes the value of one when the flight is expected to be a peak flight and is zero if it is expected to be an off-peak. These two pricing regimes will be significantly different if the interaction coefficient δ_1 is statistically significant. Then, during an expected off-peak flight the effect of ECC on fares will be δ_0 , while in an expected peak flight it will be $\delta_0 + \delta_1$. As before $DAYADV$ controls for any time trend. The coefficient on the lagged dependent variable, $\ln(FARE_{ij,t-1})$, is not of direct interest, but allowing for dynamics in the underlying process may be crucial for recovering consistent estimates of the other parameters. As in the previous section, for a second specification $\ln(ECC)$ will be replaced with $LOAD$.

The reason why a dynamic estimation is important is because both the *effective cost of capacity*, ECC , and the load factor, $LOAD$, are functions of cumulative sales. But the number of tickets that have already been sold —cumulative sales— depend on previous price levels. So there is reason to believe that the assumption of strict exogeneity of the regressors may be violated. The way the panel estimator presented in this section controls for endogeneity is by using ‘internal instruments’. We assume that the explanatory variables are only ‘weakly exogenous’, which means that the cumulative sales can be affected by current and past realization of fares, but must be uncorrelated with future realizations of the error term. Weak exogeneity does not mean that consumers do not take into account expected future changes in fares in their decisions to buy or not a ticket; it just means that future (unanticipated) shocks in fares do not influence current cumulative sales or the decision to buy a ticket. We will assess the validity of this weak exogeneity assumption below.

To estimate Equation 3.7, we first take first-differences to eliminate carrier- and flight-specific effects. Then the resulting equation requires instruments to deal with the potential endogeneity of the explanatory variables and with the problem that the construction of the new error term, $\varepsilon_{ijt} - \varepsilon_{ij,t-1}$, is correlated with the lagged dependent variable, $\ln(FARE_{ijt-1}) - \ln(FARE_{ijt-2})$. The GMM *difference* panel estimator that we will report constructs its moment conditions under the assumptions that the error term, ε , is not serially correlated, and that the explanatory variables are weakly exogenous. Then the moment conditions used for the *difference* estimator are:

$$E[y_{ij,t-s}(\varepsilon_{ijt} - \varepsilon_{ij,t-1})] = 0 \quad \text{for } s \geq 2; t = 3, \dots, T, \quad (3.8)$$

$$E[\mathbf{W}_{ij,t-s}(\varepsilon_{ijt} - \varepsilon_{ij,t-1})] = 0 \quad \text{for } s \geq 2; t = 3, \dots, T. \quad (3.9)$$

where y_{ijt} is the natural logarithm of fare and \mathbf{W}_{ijt} is the set of explanatory variables other than the lagged logarithm of fare.

Blundell and Bond [8] point out an statistical shortcoming with this *difference* estimator. When the explanatory variables are persistent over time, lagged levels of these variables are weak instruments for the regression equation in differences. To reduce the potential biases and imprecision associated with the usual difference estimator we employ the *system* estimator suggested in Blundell and Bond [8]. This *system* estimator combines the regression in differences with the regression in levels. The instruments for the regression in differences are the same as above. The instruments for the regression in levels are the lagged differences of the corresponding variables. The validity of these instruments relies on the following additional assumption: There is no correlation between the differences of the right-hand side variables in Equation 3.7 and the flight-specific effects, but there may be correlation between the levels of the right-hand side variables and the flight-specific effects. Then, for the

regression in levels included as a second part of the system the additional moment conditions are:

$$E[(y_{ij,t-s} - y_{ij})(\nu_{ij} + \varepsilon_{ijt})] = 0 \quad \text{for } s = 1, \quad (3.10)$$

$$E[(\mathbf{W}_{ij,t-s} - \mathbf{W}_{ij,t-s-1})(\nu_{ij} + \varepsilon_{ijt})] = 0 \quad \text{for } s = 1. \quad (3.11)$$

To address the validity of the instruments we consider two specification tests suggested in Arellano and Bond [1], Arellano and Bover [2], and Blundell and Bond [8]. To test the overall validity of the instruments we provide a Sargan test of over-identifying restrictions, which analyzes the sample analogs of the moment conditions used in the GMM estimation. To test the hypothesis that the error term, ε_{ijt} , is not serially correlated, we test whether the differenced error term is second-order serially correlated.

The dynamic panel results in table 11 show two sets of estimates. The first three columns were calculated when the expected load factor is simply the average, while columns (4) to (6) use the nonparametric estimate of the expected load factor. The first set of estimates shows that there is not significant difference between peak and off-peak periods. Presumably because the simple average is not a good approximation of average conditions. The set of estimates does show that there is a positive and significant difference between the two periods. This can be appreciated by the significant coefficient in the interaction, $LOAS_{ijt} \cdot PEAK_{ij,t-1}$, variable. Moreover, besides the *differences* and the *system* estimators described above, for comparative purposes, table 11 also reports the panel estimates when the estimation is done in *levels* using flight fixed effects. Additionally, table 11 gives the p-values the differenced equation exhibit no second-order serial correlation (valid specification).

Based on the *system* GMM estimates of column (6), evaluated at the sample average fare of 291.09 dollars and for a 100 seats airplane, imply that having one

Table 11. Regression estimates: GMM dynamic panel

	Levels (1)	Average First diff. (2)	System (3)	Levels (4)	Nonparametric First diff. (5)	System (6)
<i>Load Factor</i>						
$\ln(FARE)_{ij,t-1}$	0.948 (255.948)	0.613 (44.734)	0.597 (17.591)	0.948 (254.811)	0.603 (43.820)	0.586 (17.197)
$DAYADV_{ijt}/10^3$	-0.137 (-1.188)	1.407 (6.897)	1.093 (3.457)	-0.344 (-3.603)	1.198 (5.778)	0.844 (2.571)
$PEAK_{ij,t-1}$	-0.044 (-3.364)	-0.048 (-1.537)	-0.065 (-1.451)	-0.035 (-3.035)	-0.068 (-3.156)	-0.083 (-3.307)
$LOAD_{ijt}$	0.137 (4.924)	0.587 (13.421)	0.548 (7.551)	0.079 (3.942)	0.479 (11.014)	0.418 (6.067)
$LOAD_{ijt} \cdot PEAK_{ij,t-1}$	0.039 (1.426)	0.047 (0.947)	0.052 (0.627)	0.060 (2.612)	0.150 (4.114)	0.173 (3.549)
Serial correlation test ^a (p-value)		0.743	0.731		0.752	0.739

Notes: The dependent variable is $\ln(FARE)$. t-statistics in parentheses based on White robust standard errors. $PEAK = I(S > 1)$. ^a The null hypothesis is that the errors in the first-difference regression exhibit no second-order serial correlation (valid specification).

less seat available increases fares by 1.22 dollars in an expected *off-peak* flight while increases fares by 1.72 dollars in a expected *peak* flight. Moreover, as information about the final state of the demand becomes available, the results are consistent with *stochastic* peak-load pricing with higher fares being set in expected peak demand periods and lower fares set in expected off-peak periods.

The regressions satisfy the specification tests. There is no evidence of second order serial correlation. Regarding the sign on $DAYADV$, this one is no longer comparable with the one reported in table 10 because of the existence of the lagged dependent variable in the dynamic panel regressions. This also explain why the sign on $DAYADV$ is so volatile.

E. Conclusions

One important source of uncertainty for airlines is that they have limited information about the demand at the moment of scheduling a flight. Because tickets are sold in advance, prices should be set in an environment of uncertainty about the total number of arriving consumers. Having a good approximation of the expected demand as sales

progress is key because (1) seats left unsold have little value after departure, and (2) carriers may forego important profits if the flight sells out and some consumers that would have paid even higher prices have to be rationed.

To respond to demand uncertainty carriers charge different prices as the flight date approaches. This dynamic pricing strategy utilized by airlines can be separated into three different components. Peak-load pricing, price discrimination and dealing with costly capacity and demand uncertainty. Correctly identifying the existence of peak-load pricing is central. While price discrimination is just a transfer from buyers to sellers with no gain to society, peak-load pricing is beneficial because it assures that only high valuation consumers get the good and can expand output through demand shifting. This paper sets to find evidence of *stochastic* peak-load pricing in airlines, where variation in prices can be explained by variations in the shadow cost attached to each seat.

Following Escobari and Gan [32], we initially calibrate the *ex-ante* —before any ticket is sold— distribution of demand uncertainty using information on sold-out probabilities and forecasted values of occupancy rates. After controlling for restrictions that segment consumers (e.g. saturday-night-stayover, minimum and maximum stay, different connections/legs, fare class, refundability), the calibrated *ex-ante* demand distribution is enough to construct the optimal distribution of prices under the Prescott-type of models where there is price rigidities or *no* demand learning. Using nonparametric techniques we then construct a threshold variable that is used as a proxy to identify different expected final demand states at different points prior to departure. Using this threshold variable we estimate a panel endogenous threshold model to test whether carriers depart from price commitments as information about the demand arrives and sales progress. The result identified two different pricing regimes. Consistent with the predictions of *stochastic* peak-load pricing, in the ex-

pected peak flight regime when sales are larger than usual, fares will be higher. In the expected off-peak flights where sales are lower than usual, fares will be lower. To control for potential endogeneity of the regressors and the interaction between cumulative sales and previous level of prices, we also estimate a dynamic panel model. The results also supported the existence of demand learning and *stochastic* peak-load pricing in airlines.

CHAPTER IV

SUMMARY

In Chapter II, we test the empirical importance of the price dispersion predictions presented in Prescott [48], formalized in Eden [25] and extended in Dana [21]. The basic idea in these theoretical models is that the equilibrium price dispersion can be explained by the different selling probabilities associated with each of the units sold. These selling probabilities play an important role in industries that face capacity constraints and uncertainty about the number of arriving consumers. Although the ideas in Prescott [48] have been extended to multiple areas in the economic literature, few papers attempt to directly test the basic predictions due to the difficulty of coming up with an appropriate measure of the selling probabilities.

In particular, the chapter seeks to find evidence for the two main predictions. i) Lower selling probabilities characterized by higher effective costs of capacity will lead to higher prices. ii) This effect will be larger in more competitive markets. Using the information on seat inventories, plus calculations of the sold out probabilities, and the forecasted values of utilization rates, we are able to construct the distribution of demand uncertainty for each of the 81 routes in the sample. With this distribution we generate a measure of the selling probability and the effective cost of capacity (*ECC*) for each of the seats in an aircraft. This allows us to test the model by finding out if *ECC* has any effect on the prices, and if so, how this effect varies with market concentration.

Under various specifications, our empirical tests strongly support both predictions of the theory. We show that for the average market structure, when *ECC* increases by one dollar, fares increase by 48 cents, whereas the remaining 52 cents is absorbed by the markup. The elasticity specification tells us that one percent

increase in the *ECC* (or same as one percent decrease in the selling probability), increases fares by 0.219 percent. Moreover, price dispersion due to costly capacity under demand uncertainty was found to be greater in more competitive markets.

Chapter III tests whether carriers learn about the demand and price accordingly as the departure date nears and sales progress. Demand learning is important because tickets are sold in advance and prices should be set in an environment of uncertainty about the total number of arriving consumers. To reduce the cost of demand uncertainty carriers charge different prices as the flight date approaches and dynamically adjust prices depending on the expectation of demand.

Following the procedure developed in Chapter II, Chapter III initially calibrates the *ex-ante* –before any ticket is sold– distribution of demand uncertainty using information on sold-out probabilities and forecasted values of occupancy rates. After controlling for restrictions that segment consumers (e.g. saturday-night-stayover, minimum and maximum stay, different connections/legs, fare class, refundability), the calibrated *ex-ante* demand distribution is enough to construct the optimal distribution of prices under the Prescott-type of models where there is price rigidities or *no* demand learning. Using nonparametric techniques we then construct a threshold variable that is used as a proxy to identify different expected final demand states at different points prior departure. The result identified two different pricing regimes consistent with the predictions of *stochastic* peak-load pricing. In the expected peak flight regime when sales are larger than usual, fares will be higher. In the expected off-peak flights where sales are lower than usual, fares will be lower.

The results in this dissertation, even though motivated initially in the airline industry, can be easily extended to industries that deal with demand uncertainty and costly capacity. Some examples involve hotel rooms, fashion apparel, cabins on cruise liners, car rentals, entertainment and sporting events, and restaurants. Because of

the richness of the data, there are multiple ways that it can be used to continue studying airline pricing. Two other papers that take advantage of this kind of data are Escobari [31] who finds evidence of *systematic* peak-load pricing, and Escobari and Jindapon [33] who look at the dynamics of price discrimination when carriers offer refundable and non-refundable fares.

REFERENCES

- [1] M. Arellano, S. Bond, Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations, *Rev. Econ. Stud.* 58 (1991), 277-297.
- [2] M. Arellano, O. Bover, Another look at the instrumental variable estimation of error-components models, *J. Econometrics* 68 (1995), 29-51.
- [3] ASTA Government Affairs, National Commission to Ensure Consumer Information and Choice in the Airline Industry, Upheaval in travel distribution, impact on consumers and travel agents, Report to Congress and the President, 2005
- [4] E. Bachis, C. Piga, On-line price discrimination with and without arbitrage conditions, Loughborough University, Department of Economics Working Paper. Presented at the HIOG, Savannah, GA, April, 2007.
- [5] M. Baye, J. Morgan, Price dispersion on the lab and on the internet: Theory and evidence, *RAND J. Econ.* 35 (2004), 449-466.
- [6] B. Bental, B. Eden, Inventories in a competitive environment, *J. Polit. Economy* 101 (1993), 863-886.
- [7] V. Bilotkach, Understanding price dispersion in the airline industry: Capacity constraints and consumer heterogeneity, In *Advances in airline economics*, vol. 1, pp 329-345, edited by Darin Lee, Elsevier Sci, New York, 2006.
- [8] R. W. Blundell, S. Bond, Initial conditions and moment restrictions in dynamic panel data models, *J. Econometrics* 87 (1998), 115-143.

- [9] S. Borenstein, Hubs and high fares: Dominance and market power in the U.S. airline industry, *RAND J. Econ.* 20 (1989), 344-365.
- [10] S. Borenstein, N.L. Rose, Competition and price dispersion in the U.S. airline industry, *J. Polit. Economy* 102 (1994), 653-683.
- [11] M. Boiteaux, La tarification des demandes en point: application de la theorie de la vente au cout marginal, *Revue Generale de l'electricite* 58 (1949), 321-340
- [12] G. Brown, B. Johnson, Public utility pricing and output under risk, *Amer. Econ. Rev.* 59 (1969), 119-128.
- [13] J.K. Brueckner, Airport congestion when carriers have market power, *Amer. Econ. Rev.* 92 (2002), 1357-1375.
- [14] D.W. Carlton, Peak load pricing with stochastic demand, *Amer. Econ. Rev.* 67, (1977) 1006-1010.
- [15] K.S. Chan, Consistency and limiting distributions of the least squares estimator of a threshold autoregressive model, *Ann. Statist.* 21 (1993), 520-533.
- [16] R.E. Chatwin, R.E., Optimal dynamic pricing of perishable products with stochastic demand and a finite set of prices, *Europ. J. Operational Res.* 125 (2000), 149-174.
- [17] J. Chen, Differences in average prices on the internet: Evidence from the online market for air travel, *Econ. Inquiry* 44 (2006), 656-670.
- [18] M.A. Crew, P.R. Kleindorfer, *The economics of public utility regulation*, MIT Press, Cambridge, MA, 1986.

- [19] J. Dana Jr., Advance-purchase discounts and price discrimination in competitive markets, *J. Polit. Economy* 106, (1998) 395-422.
- [20] J. Dana Jr., Using yield management to shift demand when the peak time is unknown, *RAND J. Econ.* 30 (1999a), 456-474.
- [21] J. Dana Jr., Equilibrium price dispersion under demand uncertainty: The roles of costly capacity and market structure, *RAND J. Econ.* 30 (1999b), 632-660.
- [22] J.I. Daniel, Congestion pricing and capacity of large hub airports: A bottleneck model with stochastic queues, *Econometrica* 63 (1995), 327-370.
- [23] R. Deneckere, H. Marvel, J. Peck, Demand uncertainty and price maintenance: Markdown as destructive competition, *Amer. Econ. Rev.* 87 (1997), 619-641.
- [24] R. Deneckere, J. Peck, J., Dynamic competition with random demand and costless search: A theory of price posting, University of Wisconsin, Department of Economics Working Paper, 2005.
- [25] B. Eden, Marginal cost pricing when spot markets are complete, *J. Polit. Econ.* 98 (1990), 1293-1306.
- [26] B. Eden, The adjustment of prices to monetary shocks when trade is uncertain and sequential, *J. Polit. Economy* 102 (1994), 493-509.
- [27] B. Eden, Inflation and price adjustment: An analysis of microdata, *Rev. Econ. Dynam.* 4 (2001), 607-636.
- [28] B. Eden, Uncertain market conditions and gains from trade: A sequential trade model with heterogeneous buyers, Vanderbilt University, Department of Economics Working Paper, 2005.

- [29] B. Eden, International seigniorage payments, Vanderbilt University, Department of Economics Working Paper, 2007.
- [30] D. Escobari, Are airlines price discriminating? Tourist versus business travelers, Texas A&M University, Department of Economics Working Paper, 2005. Presented at the IIOC, Boston, MA, April, 2006.
- [31] D. Escobari, Systematic peak-load pricing, congestion premia and demand diverting: Empirical evidence, Texas A&M University, Department of Economics Working Paper, 2006.
- [32] D. Escobari, L. Gan, Price dispersion under costly capacity and demand uncertainty, NBER working paper 13075, Cambridge, MA, 2007.
- [33] D. Escobari, P. Jindapon, Optimal refund contracts and advance purchases in airlines, University of Alabama, Department of Economics, Finance and Legal Studies Working Paper, 2008.
- [34] I. Gale, T. Holmes, The efficiency of advance-purchase discounts in the presence of aggregate demand uncertainty, *Int. J. Ind. Organ.* 10 (1992), 413-437.
- [35] I. Gale, T. Holmes, Advance-purchase discounts and monopoly allocation of capacity, *Amer. Econ. Rev.* 83 (1993), 135-146.
- [36] G. Gallego, G.J. van Ryzin, Optimal dynamic pricing of inventories with stochastic demand over finite horizons, *Manage. Sci.* 40 (1994), 999-1020.
- [37] B.E. Hansen, Inference when the nuisance parameter is not identified under the null hypothesis, *Econometrica* 64 (1996), 413-430.
- [38] B.E. Hansen, Threshold effects in non-dynamic panels: Estimation, testing, and inference, *J. Econometrics* 93 (1999), 345-368.

- [39] B.E. Hansen, Sample splitting and threshold estimation, *Econometrica* 68 (2000), 575-603.
- [40] J. Hirshleifer, Peak loads and efficient pricing: Comment, *Quart. J. Econ.* 72 (1958), 451-462.
- [41] D. Holtz-Eakin, W. Newey, H. Rosen, Estimating vector autoregressions with panel data, *Econometrica* 56 (1988), 1371-1395.
- [42] W.M. Kincaid, D.S. Darling, An inventory pricing problem, *J. Math. Ana. Applications* 7 (1963) 183-208.
- [43] Q. Li, S. Racine, *Nonparametric econometrics*, Princeton University Press, Princeton, NJ, 2007.
- [44] K.Y. Lin, Dynamic pricing with real-time demand learning, *Europ. J. Operational Res.* 174 (2006), 522-538.
- [45] R.E. Lucas, M. Woodford, Real effects of monetary shocks in an economy with sequential purchases, NBER working paper 4250, Cambridge, MA, 1993.
- [46] R.P. McAfee, V. te Velde, Dynamic pricing in the airline industry, in *Handbook of economics and information systems*, vol. 1, pp 527-570, edited by T.J. Hendershott, Elsevier Sci, New York, 2006.
- [47] J.I. McGill, G.J. van Ryzin, Revenue management: Research overview and prospects, *Transp. Sci.* 33 (1999), 233-256.
- [48] E. Prescott, Efficiency of the natural rate, *J. Polit. Economy* 83 (1975), 1229-1236.

- [49] J.S. Racine, Q., Li, Nonparametric estimation of regression functions with both categorical and continuous data, *J. Econometrics* 119 (2004), 99-130.
- [50] J.J. Rotemberg, L.H. Summers, Inflexible prices and procyclical productivity, *Quart. J. Econ.* 105, (1990) 851-874.
- [51] A. Soresen, Empirical price dispersion in retail markets for prescription drugs, *J. Polit. Economy* 108 (2000), 833-850.
- [52] J. Stavins, Price discrimination in the airline market: The effect of market concentration, *Rev. Econ. Statis.* 83 (2001), 200-202.
- [53] P.O. Steiner, Peak loads and efficient pricing, *Quart. J. Econ.* 71 (1957), 585-610.
- [54] N. Stokey, Intertemporal price discrimination, *Quart. J. Econ.* 93 (1979), 355-371.
- [55] M.L. Visscher, Welfare-maximizing price and output with stochastic demand: Comment, *Amer. Econ. Rev.* 63 (1973), 224-229.
- [56] M. Wan, Demand uncertainty, search costs, and price dispersion: Evidence from the online book industry, University of Texas at Dallas, The School of Management Working Paper, 2007.
- [57] M. Weitzman, A theory of wage price dispersion and job market segmentation, *Quart. J. Econ.* 104 (1989), 121-137.
- [58] O.E. Williamson, Peak load pricing and optimal capacity under constraints, *Amer. Econ. Rev.* 56 (1966), 810-827.

APPENDIX A

SOME FLIGHTS

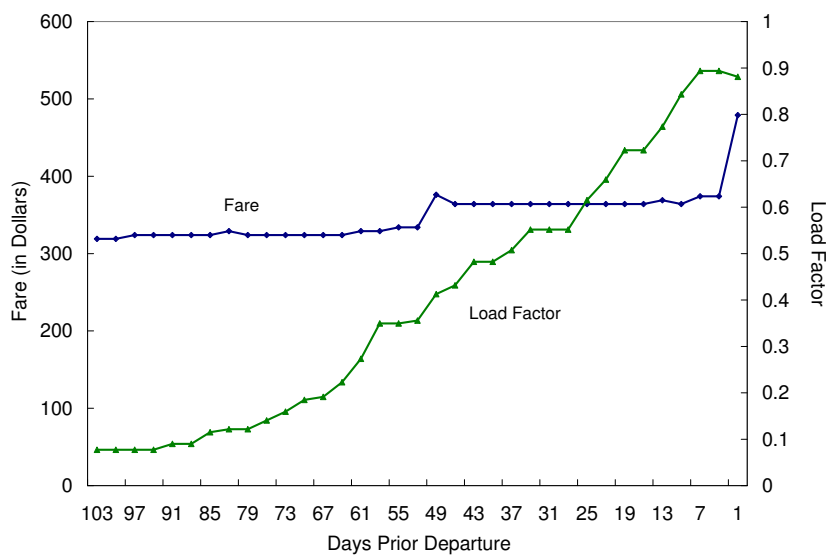


Fig. 10. Fares and load factors (United 167 BOS-LAX)




Fig. 11. Fares and load factors (Continental 2408 CLE-ORD)




Fig. 12. Fares and load factors (United 7156 IAD-CLE)




Fig. 13. Fares and load factors (Continental 194 LAX-IAH)




Fig. 14. Fares and load factors (American 596 MIA-BOS)




Fig. 15. Fares and load factors (American 1341 MIA-MSY)

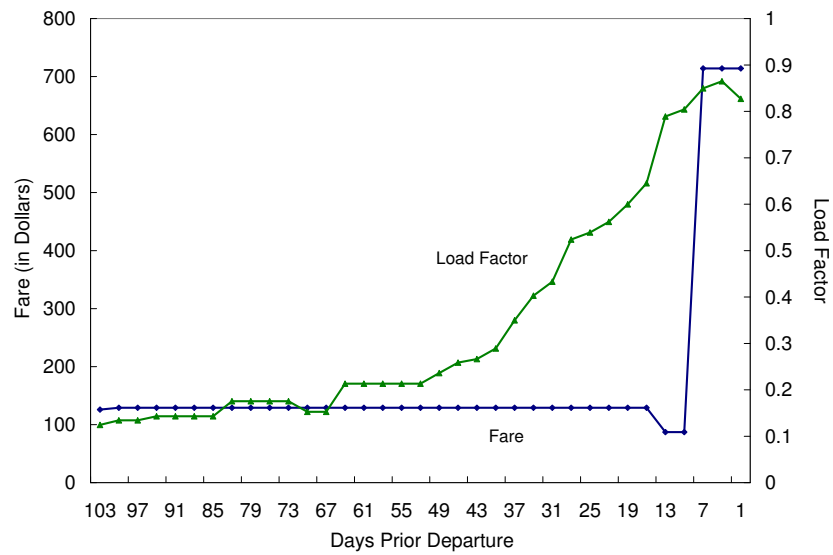


Fig. 16. Fares and load factors (US 185 PHL-ORD)




Fig. 17. Fares and load factors (United 579 PIT-ORD)

APPENDIX B

INSTRUMENTS

The construction of the instruments follow Borenstein [9] and Borenstein and Rose [10]. In particular, the instrument for *ROUSHARE* is called *GEOSHARE*, defined as:

$$GEOSHARE = \frac{\sqrt{ENP_{x1} \cdot ENP_{x2}}}{\sum_y \sqrt{ENP_{y1} \cdot ENP_{y2}}},$$

with y indexes all airlines and x indexes the observed airline. ENP_{y1} and ENP_{y2} are airline y 's average daily enplanements at the two endpoint airports during the second quarter of 2006. The instrument for *HHI* is called *XFLTHERF*:

$$XFLTHERF = ROUSHARE^2 + \frac{HHI - ROUSHARE^2}{(1 - ROUSHARE)^2} \cdot (1 - ROUSHARE)^2.$$

This instrument assumes that the concentration of the flights on a route that is not performed by the observed airline is exogenous with respect to the price of the observed carrier. More on these instruments can be found in Borenstein [9] and Borenstein and Rose [10].

APPENDIX C

VARIABLE DESCRIPTION

$FARE_{ijt}$: Price in US\$ paid for the one-way airfare.

$LOAD_{ijt}$: Load factor, defined as total number of seats sold at time t divided by total number of seats in the aircraft.

ECC_{ijt} : *Effective cost of capacity*, calculated by dividing costly capacity, λ (initially normalized to one), by the probability that this seat will be sold. For the *censored normal* case this one is given by

$$ECC_{ijt} = \frac{\lambda}{Pr_{h_{ijt}}} = \lambda \cdot \left[\int_{h_{ijt}/m_{ij}}^{\infty} \sqrt{2\pi\sigma_j^2} \cdot \exp\left(-(\kappa - \mu_j)^2/2\sigma^2\right) d\kappa \right]^{-1}$$

m_{ij} is the total number of seats in the aircraft and h_{ijt} is the number of seats that have already been sold. The values for μ_j and σ_j are obtained from the calibration procedure in section c.

$DAYADV_{ijt}$: Number of days in advance the ticket was purchased.

S_{ijt} : Threshold variable, defined as the ratio of actual seats sold to expected number of seats sold. $S_{ijt} = LOAD_{ijt}/E(LOAD_{ijt}|DAYADV_{ijt}, \mathbf{X}_{ijt})$.

$PEAK_{ijt}$: Variable equal to one if flight i is expected to be a peak flight, $PEAK_{ijt} = I(\gamma < S_{ijt})$.

$DEPTIME_j$: Time of the day the flight departed.

$DIST_j$: Nonstop mileage between the two endpoint airports on a route.

ROUSHARE_{ij}: Carrier's share on the route based on total number of seats in direct flights for the day of the flight.

HHI_j: *Herfindahl-Hirshman Index* of concentration on the observed route, with *ROUSHARE* used as the measure of market share of each carrier.

$$HHI_j = \sum_{i=1}^N ROUSHARE_{ij}^2$$

HUB_{ij}: Variable equal to one if the carrier has a hub in the origin or destination airports.

SLOT_j: In some airports like Chicago O'Hare (ORD), Kennedy (JFK), La Guardia (LGA), and Reagan National (DCA), the U.S. government has imposed limits on the number of takeoffs and landings that may take place each hour. To take into account the scarcity value of acquiring a slot, the variable *SLOT* equals to one if either endpoint of route *j* is one of these airports and zero otherwise.

DIFTEMP_j: Absolute difference in average end of October temperatures, measured in Fahrenheit degrees, between the departure and destination cities.

DIFRAIN_j: Absolute difference in average end of October precipitation, measured in inches, between the departure and destination cities.

DIFSUN_j: Absolute difference in average end of October sunshine, measured in percentage, between the departure and destination cities.

AVEHHINC_j: Average of the median household income in the two cities.

AVEPOP_j: Average population in the two cities. For cities with more than one airport, the population is apportioned to each airport according to each airport's share of total enplanements. Source: Table 3, Bureau of Transportation

Statistics, Airport Activity Statistics of Certified Air Carriers: Summary Tables 2000.

$AA_j, AL_j, CO_j, DE_j, UN_j, US_j$: Variables equal to one if the carrier on route j is American, Alaska, Continental, Delta, United, or US Airways respectively.

VITA

Diego Alfonso Escobari Urday was born in La Paz, Bolivia. He received his Bachelor of Arts degree in Economics from Universidad Católica Boliviana in 2000, and a Master of Science degree also in Economics from Texas A&M University in 2004. He received his Doctor of Philosophy degree in May 2008 under the supervision of Dr. Li Gan. His research interests include industrial organization, applied econometrics, and public economics.

Dr. Escobari Urday may be reached at the Department of Economics, San Francisco State University, San Francisco, CA 94132. escobaridiego@hotmail.com