

University of Texas Rio Grande Valley

ScholarWorks @ UTRGV

Economics and Finance Faculty Publications
and Presentations

Robert C. Vackar College of Business &
Entrepreneurship

9-2017

Airport, airline and departure time choice and substitution patterns: An empirical analysis

Diego Escobari

The University of Texas Rio Grande Valley, diego.escobari@utrgv.edu

Follow this and additional works at: https://scholarworks.utrgv.edu/ef_fac



Part of the [Finance Commons](#)

Recommended Citation

Escobari, Diego. "Airport, airline and departure time choice and substitution patterns: An empirical analysis." *Transportation Research Part A: Policy and Practice* 103 (2017): 198-210. <https://doi.org/10.1016/j.tra.2017.05.034>

This Article is brought to you for free and open access by the Robert C. Vackar College of Business & Entrepreneurship at ScholarWorks @ UTRGV. It has been accepted for inclusion in Economics and Finance Faculty Publications and Presentations by an authorized administrator of ScholarWorks @ UTRGV. For more information, please contact justin.white@utrgv.edu, william.flores01@utrgv.edu.



Munich Personal RePEc Archive

Airport, airline and departure time choice and substitution patterns: An empirical analysis

Escobari, Diego

The University of Texas Rio Grande Valley

12 June 2017

Online at <https://mpra.ub.uni-muenchen.de/79857/>

MPRA Paper No. 79857, posted 24 Jun 2017 05:29 UTC

Airport, Airline and Departure Time Choice and Substitution Patterns: An Empirical Analysis*

(*Transportation Research Part A*, Forthcoming)

Diego Escobari[†]

June 12, 2017

Abstract

This paper uses the random-coefficients logit methodology that controls for potential endogeneity of prices and allows for general substitution patterns to estimate various demand systems. The estimation takes advantage of an original ticket-level revealed preference data set on travels from the New York City area to Toronto that contains prices and characteristics of not only flight choices but also of all non-booked alternative flights. Consistent with having higher valuations, our results show that travelers buying closer to departure have a higher utility of flying. Moreover, consumers' heterogeneity decreases as the flight date nears. At the carrier level, we identify which carriers have more price-sensitive consumers and which carriers face greater competition. In addition, the results suggest that our multi-airport metropolitan area can be considered as a single market and that JFK and Newark are relatively closer substitutes. Overall, consumers are more willing to switch to alternative carriers than between airports or departure times.

Keywords: Airline choice; Airport choice; Departure time choice; Substitution patterns; Airline demand; Elasticities

*The author appreciates useful comments by Chad Kwon and Alejandro Rojas. Four anonymous referees provided insightful comments that helped improve the paper. Stephanie C. Reynolds provided outstanding assistance with the data.

[†]Department of Economics & Finance, The University of Texas Rio Grande Valley, Edinburg, TX 78539, Phone: (956) 665-3366, Fax: (956) 665-5020, Email: diego.escobari@utrgv.edu, URL: <http://faculty.utrgv.edu/diego.escobari>

1 Introduction

Estimating a demand system is at the heart of the analysis of flights as differentiated products. In airline markets, passengers need to make a number of decisions, including the choice of airport, carrier, and departure time. It is not only important to estimate the demand at each of these levels, but also to obtain precise estimates of the substitution patterns. At the airport level, studying choice and substitution patterns is key to address the definition of a market in areas served by multiple airports. This is linked to the assessment of airport congestion, regulation, pricing of gates and airways, airport expansions, and improving access to airports. At the carrier level, the importance arises when assessing the level of competition between carriers and when implementing pricing strategies. At the departure time level, the question is important when focusing on congestion (e.g., stochastic and systematic peak-load pricing) and demand-shifting across alternative departure times.

In this paper we use an original ticket-level revealed preference data set on direct travels from the New York City area to Toronto with information on prices and characteristics of not only the flights selected by passengers, but also of all non-booked alternative flights. The availability of information on prices and on unchosen alternatives solves two problems typically faced by most studies of air travel demand.¹ This allows us to estimate passengers' preferences by having the same information on the products available to them when they booked a flight. In addition, we control for ticket characteristics that serve to implement systematic and stochastic peak-load pricing (see, e.g., Borenstein and Rose, 1994; and Escobari, 2009, 2012) and also serve to segment consumers and price discriminate (see, e.g., Escobari and Jindapon, 2014).

We model flights as differentiated products and estimate various demand systems using the random-coefficients logit methodology. Our estimation approach helps overcome various challenges. First, we want the estimation of the substitution patterns across flights

¹For example, Harvey (1987), Pels et al. (2001), Hess and Polak (2006), and Koster et al. (2011) have no information on prices, while Pels et al. (2003), Hess and Polak (2005b), and Pels et al. (2009) have only average prices. Moreover, Hess et al. (2007) explain that survey data—which is used in most revealed preference data studies—has major issues that arise because of the often low quality of the information in the unchosen alternatives.

to be consistent with economic theory.² Second, passengers face a large number of flights, which implies a large number of parameters to be estimated. The logit model proposed by McFadden (1973) solves this dimensionality problem by projecting products onto a space of characteristics. One concern in McFadden’s approach is the Independent of Irrelevant Alternatives (IIA) property that means that substitution patterns (cross-price elasticities) are driven entirely by market shares and not by product characteristics. Alternative methods such as the nested logit were aimed at relaxing the IIA assumption, but they still have the constraint that products need to be classified a priori as in the Deaton and Muellbauer (1980)’s almost ideal demand system. The random-coefficients discrete-choice models of demand that we employ initially follow Berry (1994) and Berry, Levinsohn and Pakes (1995, henceforth, BLP), but extend the selection of instruments to include instruments from the dynamic panels literature and Chamberlain’s (1987) optimal set of instruments as suggested in Reynaert and Verboven (2014). The BLP random-coefficients demand methodology maintains the advantage of the logit by allowing for a large number of products. In addition, it allows for general substitution patterns that take into account heterogeneity of consumers’ tastes which produce more realistic own- and cross-price elasticities (Nevo, 2000b). The estimation approach also allows us to control for potentially endogenous prices while retaining the benefits of alternative discrete-choice models. Controlling for potentially endogenous prices is key given that airlines jointly compete in prices, departure times, and even in the selection of the airport.

The results show evidence that the utility of flying is greater for travelers who buy closer to departure, which is consistent with having higher valuations. Moreover, we also find that as the flight date nears, consumers’ heterogeneity increases. Our elasticity estimates show that at the airline level, United has the most price-sensitive travelers followed by American and Delta. In addition, we find evidence that Air Canada and Delta appear to have differentiated themselves more from the rest of the carriers while American, United and Continental appear to be relatively closer competitors. At the airport level, the cross-price elasticity estimates suggest that the multi-airport metropolitan area that comprises

²An example of demand systems consistent with economic theory includes the almost ideal demand system (AIDS) of Deaton and Muellbauer (1980). Njegovan (2006) uses an AIDS to estimate elasticities for air-travel demand.

Newark, JFK and La Guardia can be considered as a single market. The results also show that travelers from Newark are the least price sensitive and that JFK and Newark appear to be relatively closer substitutes. Overall, consumers are more willing to switch to an alternative carrier than to switch between airports or to alternative departure times.

The study of airport choice has long been of interest to researchers and it is generally agreed that access time and flight frequency are dominant factors explaining airport demand. Estimating the choice and the degree of substitution between competing airports is important as airport planners need to know if, for example, particular investments will increase market share. Regarding the methodology, most previous studies used multinomial logit (e.g., Skinner, 1976; Harvey, 1987) or nested logit (e.g., Pels et al., 2000, 2001, 2003), while there are some that employed probabilistic choice set multinomial logit (e.g., Başar and Bhat, 2004) and weighted conditional logit (e.g., Ishii et al., 2009). In terms of data, previous work used either revealed preference data (e.g., Pels et al., 2001, 2003; Başar and Bhat, 2004; Zhang and Xie, 2005; Hess and Polak, 2005a,b, 2006; Ishii et al., 2009; Pels et al., 2009), stated preference data (e.g., Skinner, 1976; Harvey, 1987; Prousaloglou and Koppelman, 1999; Hensher et al., 2001; Zhang and Xie, 2005; Hess et al., 2007; Loo, 2008; Ishii et al., 2009; Hess, 2010; Koster et al., 2011; Marcucci and Gatta, 2011; de Luca, 2012; and Drabas and Wu, 2013), or mixed data (Ortúzar, J. de D. and Simonetti, 2008).³ In terms of methodology, the study that is closest to ours is the demand side of Berry and Jia (2010) who estimate a structural model with aggregate data.⁴

Studying airline choice and substitution patterns is important because airlines need to know the degree in which increasing their prices shifts passengers to alternative carriers. Zhang et al. (2010) show that the degree of substitution between carriers affects the concession revenue sharing between an airport and its airlines. The degree of substitution between airlines also affects the vertical relationship between airports and airlines (Fu et al., 2011), the potential vertical collusion between airports and airlines (Barbot, 2009; Barbot et al., 2013), the internalization of congestion costs and congestion pricing (Brueckner, 2002; Mayer and Sinai, 2003; Morrison and Winston, 2007; and Rupp, 2009),

³Bilotkach et al. (2012) study airport choice and its effect on airfares using a natural experiment.

⁴Escobari and Mellado (2014) is also closely related, but they do not estimate substitution patterns and they use a simple logit model that does not control for endogeneity of prices.

and competition and the role of de-hubbing on prices (Tan and Samuel, 2016).

The organization of the paper is as follows. Section 2 describes and discusses the benefits and assumptions behind the data. The empirical model for airlines is proposed in Section 3. Section 4 starts by describing the instruments and then discusses the results. Section 5 concludes.

2 Description and Discussion of the Data

We have an original ticket-level revealed preference data set with information on prices, sales and product characteristics of not only the flights in which sales occurred, but also on prices and product characteristics of all contemporaneous non-booked alternatives.⁵ This means that our data set contains travelers' choices and the same information available to them on all available flight options when they booked the trip. Having this information is consistent with discrete-choice random-utility models in which a potential traveler arrives to the market, observes all available products and characteristics (including prices), and then decides to buy the ticket that offers the highest utility or decides not to buy any ticket if the highest utility is from the outside good. When buying airline tickets, travelers can easily observe prices and characteristics of all available options via online travel agencies.

Following a similar strategy as Escobari (2012) and Escobari and Jindapon (2014), the data set was collected from the online travel agency Expedia.com by keeping track of posted prices and seat inventories at different times prior to departure. Our collection strategy has important advantages over similar data sets obtained from computer reservation systems (e.g., Stavins, 2001). First, we have a panel which allows us to control for observed and unobserved flight-specific characteristics. Second, we recorded inventories which are key to identify when sales occur, and third, we also recorded prices when sales do not occur. Our detailed information on prices at the ticket level represents an important improvement over previous work on air travel demand that uses revealed preference data. As explained in

⁵Previous studies that use revealed preference data are almost entirely based on surveys (see, e.g., Proussalogloua and Koppelman, 1999; Pels et al., 2000, 2003; Zhang and Xie, 2005; Hess and Polak, 2006; Hess et al., 2007; Loo, 2008; Koster et al., 2011). Hess et al. (2007) explain that survey data is generally collected from departing passengers, which has major issues that arise because of the low quality of the information in the unchosen alternatives.

Pels et al. (2009), a common problem to most studies of airline demand is the poor quality of air fares data. Hence, most of previous studies either ignore prices or use average prices.⁶

To make the problem tractable we focus on flights departing from the New York City area and arriving in Toronto. This single city pair during our period of study generated over half-a-million observations. Previous studies that focused on a particular city pair include Bilotkach (2007) and Ishii et al. (2009), while studies that focused on a geographical area include Pels et al. (2001, 2003), Başar and Bhat (2004), and Bilotkach et al. (2012). One benefit from the New York City area is that it has three international airports that offer direct flights to Toronto. This will allow us to capture the degree of substitution between airports. The John F. Kennedy International Airport (JFK) and La Guardia Airport (LGA) are both located in Queens, New York, while the Newark Liberty International Airport (EWR) is located Newark, New Jersey. All arrivals are at the Toronto Pearson International Airport (YYZ), which is the only big airport that serves Toronto.

The sample includes directional non-stop one-way economy-class tickets. We assume trips with one or more stops are of a significantly different quality as well as tickets that belong to a different class (e.g., first class).⁷ Our approach to study a single leg follows the theoretical literature on pricing and airline demand where most of the work is built on single-leg models (see, e.g., Gallego and van Ryzin, 1994; and Dana, 1998).⁸ One limitation from this approach will arise when trying to generalize the results to round-trip itineraries. For example, round-trip prices depend on the dates of both trips (see, e.g., Escobari et al., 2017). We do not attempt to generalize the results to round-trip tickets; however, we believe our setting is not significantly different, for example, from the empirical literature in airlines that assumes that one-way and round-trip tickets follow the same structure and that the round-trip fare is just the one-way fare multiplied by two (see, e.g., Borenstein

⁶For example, Harvey (1987), Pels et al. (2001), Hess and Polak (2006), and Koster et al. (2011) do not have information on fares. Moreover, Ishii et al. (2009) use an approximation of fares, while Pels et al. (2003), Hess and Polak (2005b), Pels et al. (2009) only have average fares.

⁷Direct flights between the New York City area and Toronto take about 98 minutes, so it is reasonable to argue that tickets that involve one or more stops are not a desired alternative for travelers.

⁸Escobari et al. (2017, section 2.2) explain that collecting round-trip data from online travel agencies suffers from a “curse of dimensionality,” as the alternative departing and returning combinations grows exponentially with the number of available flights.

and Rose, 1994, p. 677; and Gerardi and Shapiro, 2009, p. 5).

The carriers in the sample are American Airlines, Air Canada, Continental, Delta, LAN Airlines, and United. In addition, the sample includes all of the 317 direct service flights that departed between December 19 and December 24, 2008. For each flight in the sample we fixed the departure date and recorded fares and inventories every three days starting at 40 days in advance up until 1 day prior to departure. Following Escobari (2012, p. 710), sales are recorded as the difference between beginning-of-period and end-of-period seat inventories. Hence, for example, if the number of available seats decreased from 17 to 16 on a particular flight, we record a sale and assume that the sale occurred at the beginning-of-period one-way posted price. Note that this one-way price is not necessarily the price at which the sale took place as the seat might have been bought, for example, with frequent flyer miles, as part of a round-trip ticket, or as part of a longer itinerary (e.g., a passenger going from Miami to Toronto connecting at a New York City area airport). We argue that the observed one-way fares are relevant to our analysis because these one-way prices are the base for the prices of other tickets that offer the same available seat. The simplest example to illustrate this point comes from Borenstein and Rose (1994) and Gerardi and Shapiro (2009) who assume that round-trip fares are just one-way fares multiplied by two.⁹ A second example would involve carriers pricing each leg independently, such that the final price of the ticket is just the summation of the prices of each of the legs—Bachis and Piga (2011) explain that European Low Costs Carriers follow this practice. In both of these examples, there is a perfect correlation between one-way prices and the prices of the other two types of tickets. While we do not draw conclusions beyond one-way fares, these examples illustrate how fairly common assumptions in the literature might enable us to extend some of our results to other types of tickets.¹⁰

⁹None of the papers that work with the popular Bureau of Transportation Statistics' DB1B data set know which portion of the ticket's price corresponds to each of the legs (see, e.g., Gerardi and Shapiro, 2009).

¹⁰This point can also be illustrated with the following example. Let the observed one-way price be p_{ow} and sales (demand) on the same flight be q_{ow}^d . In a simplified scenario, we are interested in estimating $\frac{\partial q_{ow}^d}{\partial p_{ow}}$. However, this marginal effect might be channeled through prices of other tickets for the same seat, for example, a round-trip ticket price p_{rt} such that $\frac{\partial q_{ow}^d}{\partial p_{rt}} \frac{\partial p_{rt}}{\partial p_{ow}}$. That is, p_{rt} is affected by p_{ow} , and sales q_{ow}^d are affected by p_{rt} . If we assume, as in Gerardi and Shapiro (2009), that $p_{rt} = 2 \cdot p_{ow}$, then $\frac{\partial p_{rt}}{\partial p_{ow}} = 2$ and

When recording a sale we follow Escobari (2012, 2014) and assume that the reduction in inventories comes from a passenger buying a ticket on the observed flight from the New York City area to Toronto. If 100% of the passengers buy direct trips, there is no need to make this assumption. However, the reduction in inventories might be, for example, from an American Airlines (AA) passenger flying from Miami (MIA) to Toronto with a stopover at JFK (MIA→JFK→YYZ). This is a concern if the fraction of passengers buying a longer itinerary that involves the observed JFK→YYZ flight is relatively large. A simple way to motivate our approach is to view an AA passenger flying MIA→JFK→YYZ as simply buying two legs, the AA ticket MIA→JFK and the AA ticket JFK→YYZ. This is still a passenger who has a demand for the AA ticket JFK→YYZ; we observe this purchase and we observe the fraction of the price that corresponds to JFK→YYZ. Notice that not all flights have the same potential for capturing connecting passengers and this can bias our estimates. Following the same example, we define a flight’s “relevant network” as the number of flights that feed passengers to our observed JFK→YYZ flight (i.e., flights arriving from other destinations to JFK in the hours leading to the JFK→YYZ flight departure) and flights departing from YYZ that can potentially connect the JFK→YYZ passengers to additional destinations. Note that this “relevant network” is flight specific and constant over time. One key advantage from our data set is that we have a panel and observe multiple purchases for the same flight. Hence, including flight (or product) dummy variables in the estimation will absorb any affect that the particular flight’s “relevant network” might have on the demand.

An interesting element when focusing on one-way tickets is that network yield management can help us generate exogenous price changes to identify the price coefficients. Note that the observed one-way ticket price can be affected by pricing at the network level (network yield management). For example, the price on a JFK→YYZ flight could be held high to protect the seats for connecting passengers that generate higher network value. The carrier will still offer one-way tickets and we observe this higher price that reflects the higher network value. This is also in line with Belobaba’s (1989) Expected Marginal Seat Revenue (EMSR). The price variation in one-way tickets is the result of opening and

our estimates can also help us learn something about how demand responds to prices of round-trip tickets,
 $\frac{\partial q_{ow}^d}{\partial p_{rt}}$.

closing of booking classes, and the observed prices are equal to the EMSR of selling the same seat on a more sophisticated itinerary that includes the same observed leg. The idea is simple, the observed variation in one-way prices reflects the relative scarcity of seats that can be sold as one-way tickets or as part of longer itineraries.

[Table 1, here.]

The summary statistics are presented in Table 1 with Panel A reporting on the main variables in the analysis. The price is in US\$ and corresponds to the fare between the corresponding New York City area airport and Toronto. The variable Days in Advance is the number of days prior to departure in which fares were recorded, while Sale is a dummy variable equal to one if a sale occurs, zero otherwise. A salient feature in this panel is the substantial observed price dispersion. The maximum price is about 16 times larger than the minimum price, which is consistent with the relatively large standard deviation of fares. The sample size of 560,244 observations times the average of the variable Sale is equal to 10,708. This figure corresponds to the total number of tickets sold. On average, every time a ticket is sold we also recorded about 52 additional fares from competing alternative flights.

Panel B disaggregates prices by airline. The largest carrier in the route is Air Canada with 30.6% (97/317) of the flights, followed by United (26.5%) and American (21.1%). LAN Airlines offered only 4 (1.3%) flights. In addition, this panel illustrates the substantial differences in average prices across carriers. For example, the average price for Delta is 38.0% greater than the average price in the route, and the difference between the average price for Delta (\$234.36) and LAN Airlines (\$130.45) is over \$100. It is valuable to understand that for some carriers (e.g., Air Canada) this is a core route, while for some other carriers (e.g., LAN Airlines) this route is a continuation of its service from Santiago, Chile to New York. LAN Airlines probably offers service in this route as an alternative to just have its aircraft wait in the JFK airport for the return trip to Santiago. Hence, LAN Airlines might be more willing to fly with lower load-factors and at lower prices. An interesting element in panel B is that there is substantial price dispersion within the same carrier, with the ratio of the highest price to the lowest price being the largest for United (15.5) and American (14.1), followed by Continental (13.0) and Air Canada (9.1).

In panels C and D the disaggregation of prices is by airport and by departure time, respectively. The airport with the most flights is La Guardia with more than half of the flights in the sample. JFK offers the least number of flights and has the highest average fare—about \$44.6 higher than the average fare in the route. Differences in fares across airports might be the results of the internalization of airport congestion costs (see, e.g., Brueckner, 2002; and Mayer and Sinai, 2003) or additional differences in costs or different consumer types across airports. From panel D, we observe that average fares are higher in the afternoon and in the evening, perhaps as a result of systematic peak-load pricing at the flight level (see, e.g., Escobari, 2009).

3 Empirical Model for Airlines

In our empirical model for airlines, we observe $t = 1, 2, \dots, T$ markets with each having $i = 1, 2, \dots, I_t$ consumers. Moreover, there are $j = 1, 2, \dots, J$ products in each market for which we observe aggregate quantities, average prices, and product characteristics. The definition of the J products in a market depends on the level of aggregation that we use to obtain aggregate quantities and average prices. In particular, we define the j options to be airports, carriers, and departure times. The idea is that different product definitions will allow us to assess substitution patterns at different levels (i.e., airport, carrier, and departure time). The importance of aggregation is also motivated by the existence of a large number of flights (317). If each flight were to be considered as a separate product and consumers are flexible in their choice of departure dates, then the matrix of own- and cross-price elasticities (of dimensions 317×317) would be very difficult to estimate and interpret. Hence, for example, if we aim at capturing substitution patterns across airports, we will aggregate flights at the airport level. For expositional purposes, in the description of the model we refer to j to denote a flight.

We define a market as the air transportation between the New York City area and Toronto. This is consistent with Brueckner et al. (2014) who provide strong evidence that city-pairs, rather than airport-pairs, are the appropriate market definition for the analysis of passenger air transportation in many (but not all) large metropolitan areas. Moreover, we model different times to departure as different markets in which consumers arrive,

observe all available options and then decide to purchase the ticket that gives them the highest utility or decide not to buy any ticket and leave the market. This is consistent with Williams (2014), who shows that increasing prices over time provides little incentives for consumers to wait to purchase later. In addition, he shows that only a small transaction cost is needed to convince consumers not to wait.

Following Nevo (2000b, 2001), we model individual i 's indirect utility from traveling in flight j and in market t with the following quasilinear form:

$$u_{ijt} = \alpha_i(y_i - p_{jt}) + x_{jt}\beta_i + \xi_{jt} + \varepsilon_{ijt}, \quad (1)$$

where y_i is the income of consumer i , p_{jt} is the price in flight j and in market t , x_{jt} is the vector of K observable non-price characteristics of flight j in market t , ξ_{jt} captures unobserved (by the econometrician) flight characteristics, and ε_{ijt} is the remainder stochastic term with zero mean. Moreover, α_i is the marginal utility of income, and β_i is the vector of individual-specific taste coefficients.

Note that quasilinearity in equation (1) implies that the indirect utility function is free of any wealth effects. For airline tickets, this is a reasonable assumption as tickets usually represent a small proportion of consumer i 's income. Petrin (2002), for example, estimates a model that includes wealth effects in the estimation of demand for cars where the assumption of no wealth effects might no longer work. In addition, equation (1) models the existence of unobserved product characteristics ξ_{jt} that are identical across consumers, but can change across markets and can capture elements of differentiation across flights.

To model the distribution of consumer taste parameters, we have that individual characteristics consist of two components, observed consumers' heterogeneity H_i and unobserved additional characteristics v_i . Because the estimation uses aggregate data, we do not need to know individual characteristics. We only need to assume that we know something about the distributions of consumers' heterogeneity H_i while the remainder characteristics v_i do not contain such information. Characteristics v_i can include whether the traveler is a tourist or a business traveler, which is something that we do not explicitly model.¹¹ Individual

¹¹Berry and Jia (2010) model a discrete number of types rather than the continuous heterogeneity modeled here. The advantage in a discrete number of types is the reduced computational burden which is important when the number of products becomes large. In Berry and Jia (2010), their two types can be described as business and tourist travelers.

characteristics are formally modeled as:

$$\begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \Pi H_i + \Sigma v_i, \quad v_i \sim P_v^*(v), \quad H_i \sim \hat{P}_H^*(H), \quad (2)$$

where H_i is a $d \times 1$ vector of variables that captures observed travelers' heterogeneity and v_i captures additional unobserved characteristics. We denote the parametric distribution of v_i by $P_v^*(v)$, and because we will be estimating the nonparametric distribution of consumers' heterogeneity H_i , we will use the notation $\hat{P}_H^*(H)$ for its distribution. Π is a $(K + 1) \times d$ matrix of coefficients that capture how taste characteristics change with heterogeneity H_i , and Σ is a $(K + 1) \times (K + 1)$ matrix of parameters.

To complete the demand system if consumers decide not to fly and leave the market, the specification of the outside good has the following indirect utility:

$$u_{i0t} = \alpha_i y_i + \xi_{0t} + \pi_0 H_i + \sigma_0 v_{i0} + \varepsilon_{i0t}. \quad (3)$$

Without an outside good, the quantities of tickets purchased would be unchanged if all carriers were to simultaneously increase prices for all available flights. In addition, without additional assumptions, we cannot identify the mean utility of the outside good, ξ_{0t} , and the coefficients π_0 and σ_0 are not separately identifiable from coefficients on an individual-specific constant term in equation (1). We follow an approach that is equivalent to normalizing the utility of the outside good to be equal to zero. That is, we set ξ_{0t} , π_0 , and σ_0 to be all equal to zero, which works because the term $\alpha_i y_i$ is common to all flights and therefore vanishes.

If we combine equations (1) and (2) we have:

$$\begin{aligned} u_{ijt} &= \alpha_i y_i + \delta_{jt}(x_{jt}, p_{jt}, \xi_{jt}; \theta_1) + \mu_{ijt}(x_{jt}, p_{jt}, v_i, H_i; \theta_2) + \varepsilon_{ijt}, \\ \delta_{jt} &= x_{jt} \beta - \alpha p_{jt} + \xi_{jt}, \\ \mu_{ijt} &= [-p_{jt}, x_{jt}](\Pi H_i + \Sigma v_i), \end{aligned} \quad (4)$$

where the vector $\theta_1 = (\alpha, \beta)$ contains the linear parameters in the model, while the vector $\theta_2 = (\Pi, \Sigma)$ contains the nonlinear parameters. We define $\theta = (\theta_1, \theta_2)$ as the vector containing both sets of parameters. In equations (4), $[-p_{jt}, x_{jt}]$ is a row vector and the indirect utility is expressed as the summation of $\alpha_i y_i$, the mean utility δ_{jt} , and the mean-zero heteroscedastic deviations from mean utility $\mu_{ijt} + \varepsilon_{ijt}$.

Following various models that explain pricing and sales in airline markets (see, e.g., Deneckere and Peck, 2012), consumers are assumed to have unit demands—they buy a single ticket on the flight that gives them the highest utility. Because an individual i is defined as a vector of consumers' heterogeneity and product-specific shocks, this implicitly defines the set of individual attributes that leads to the choice of a ticket in flight j . Let this set be:

$$A_{jt}(x_{\cdot t}, p_{\cdot t}, \delta_{\cdot t}; \theta_2) = \{(H_i, v_i, \varepsilon_{i0t}, \dots, \varepsilon_{iJt}) | u_{ijt} \geq u_{illt}\} \quad \forall l = 0, 1, \dots, J, \quad (5)$$

where $x_{\cdot t} = (x_{1t}, \dots, x_{Jt})'$ are the observed characteristics, $p_{\cdot t} = (p_{1t}, \dots, p_{Jt})'$ are the prices, and $\delta_{\cdot t} = (\delta_{1t}, \dots, \delta_{Jt})'$ are the mean utilities associated with all of the available flights. A_{jt} in equation (5) defines individuals who chose flight j in market t . If no two flights give exactly the same level of utility, the market shares for the j th flight can be obtained as the integral over the mass of consumers in A_{jt} :

$$s_{jt}(x_{\cdot t}, p_{\cdot t}, \delta_{\cdot t}; \theta_2) = \int_{A_{jt}} dP^*(H, v, \varepsilon) = \int_{A_{jt}} dP_\varepsilon^*(\varepsilon) dP_v^*(v) d\hat{P}_H^*(H), \quad (6)$$

where $P^*(\cdot)$ denotes the population distribution functions. The second equality follows from the assumption of independence of H , v , and ε . Nevo (2000a) and Knittel and Metaxoglou (2014) explain how to obtain the integral in equation (6) numerically given the assumptions on the distribution of the individual attributes.

One common approach to evaluate the integral in equation (6) is to assume that consumers' heterogeneity is fully captured by the random shock ε_{ijt} .¹² Then equation (1) reduces to:

$$u_{ijt} = \alpha(y_i - p_{jt}) + x_{ij}\beta + \xi_{jt} + \varepsilon_{ijt}. \quad (7)$$

If we additionally assume that ε_{ijt} in equation (7) is i.i.d. type I extreme-value, the market share of flight j in market t is:

$$s_{jt} = \frac{\exp(x_{jt}\beta - \alpha p_{jt} + \xi_{jt})}{1 + \sum_{k=1}^J \exp(x_{kt}\beta - \alpha p_{kt} + \xi_{kt})}. \quad (8)$$

Because income is assumed to be common to all options, it drops out of the equation. In this case the substitution patterns are greatly restricted as the following price elasticities

¹²That is, $\theta_2 = 0$, $\alpha_i = \alpha$, and $\beta_i = \beta$.

of the market shares hold:¹³

$$\eta_{jkt} = \frac{\partial s_{jt}}{\partial p_{kt}} \frac{p_{kt}}{s_{jt}} = \begin{cases} -\alpha p_{jt}(1 - s_{jt}) & \text{if } j = k \\ \alpha p_{kt} s_{kt} & \text{otherwise.} \end{cases} \quad (9)$$

There are two concerns with the elasticities presented in equations (9). Because market shares s_{jt} are usually small, then $\alpha(1 - s_{jt})$ is close to constant. This means that the own-price elasticities are proportional to own price. This predicts higher markups for lower priced flights and the functional form of the indirect utility directly determines the patterns of the own-price elasticity.

The second concern occurs in the cross-price elasticities and arises from the i.i.d. structure of the random shock ε_{ijt} . This is known as the Independence of Irrelevant Alternatives (IIA) property and implies that substitution patterns are proportional to market shares. In the airline context, we would have that if United in the Newark Liberty Airport and Delta in the La Guardia Airport have the same market shares, an increase in prices from Air Canada in the Newark Liberty airport would have the same substitution towards United (in Newark Liberty) and Delta (in La Guardia). However, intuitively, we would expect more passengers to shift to flights from United in the same airport (Newark Liberty) that saw Air Canada prices rise.

The nested logit allows more flexible substitution patterns by dividing products into groups, but the groups need to be defined a priori. Our approach to allow more realistic substitution patterns without classifying products a priori involves keeping the i.i.d. extreme-value distribution assumption in ε_{ijt} and using μ_{ijt} in equation (4) to introduce correlation between flight choices. This correlation is a function of flight and consumer characteristics, hence similar substitution patterns will follow from similar flights or similar consumer characteristics.

Under the assumption that ε_{ijt} in equation (4) is i.i.d. type I extreme-value, we can write $s_{ijt} = \exp(\delta_{jt} + \mu_{ijt}) / [1 + \sum_{k=1}^J \exp(\delta_{kt} + \mu_{ikt})]$ as the probability of individual i buying a ticket in flight j in market t . Then the price elasticities of the market shares of

¹³Appendix A presents the derivation of these elasticities.

equation (6) are:

$$\eta_{jkt} = \frac{\partial s_{jt}}{\partial p_{kt}} \frac{p_{kt}}{s_{jt}} = \begin{cases} -\frac{p_{jt}}{s_{jt}} \int \alpha_i s_{ijt} (1 - s_{ijt}) d\hat{P}_H^*(H) dP_v^*(v) & \text{if } j = k \\ \frac{p_{kt}}{s_{jt}} \int \alpha_i s_{ijt} s_{ikt} d\hat{P}_H^*(H) dP_v^*(v) & \text{otherwise.} \end{cases} \quad (10)$$

Note that the own-price elasticities ($j = k$) and the cross-price elasticities ($j \neq k$) are not driven by functional form. Moreover, they change across individuals and flights. The flexible substitution patterns in equations (10) involve using simulation methods to compute the integral in equation (6). We use the estimation methods proposed by Berry (1994) and Berry et al. (1995) which additionally control for endogeneity of prices. Appendix B presents a summary of the estimation methods we employ.

4 Estimation

4.1 Instruments

The GMM estimator that we employ needs a set of exogenous instruments Z . The exogenous variables in the demand specification (e.g., the number of days in advance) will serve as their own instruments. However, price can not serve as an instrument because it is potentially endogenous as it might be correlated with the error term. The most common cause of this correlation is that carriers set prices knowing more about the error term than the econometrician. This occurs, for example, if there are unobserved flight characteristics that affect the demand for airline tickets. Note that our data consist of posted prices, which means that airlines set prices first and then consumers make their purchase decisions based on these fixed prices. This might make the potential endogeneity of prices less of a concern if there are unobserved ticket characteristics that affect prices because at least they are determined prior to the revelation of the traveler's decision to buy an airline ticket. In a panel data setting, this argument would help support the idea that our posted prices might be predetermined (or weakly exogenous) rather than endogenous. This contrasts with transaction data in which the identification of the demand is more difficult because demand and supply variables are jointly determined from the point of view of the econometrician.

For the selection of Z , we employ Chamberlain's (1987) optimal set of instruments which

consists of the expected value of the derivatives of the structural error term with respect to the parameter vector, evaluated at an initial estimate of the parameters. Reynaert and Verboven (2014) find that the use of Chamberlain’s (1987) optimal set of instruments reduces bias and drastically improves the efficiency and stability of the parameter estimates. In addition, we also use a second set of instruments. Taking advantage of the panel structure of the data that gives us repeated observations of the same markets, we follow Anderson and Hsiao (1981) and Arellano and Bond (1991) and use lagged values of fares as instruments. An advantage is that these are instruments that vary by flight and over time. Escobari (2012) uses the same set of instruments in his estimation of demand for air travel.

4.2 Results

Table 2 displays the results obtained from the model based on equation (4) using equation (6) to compute the predicted market shares. The different specifications across columns are aimed at capturing estimates for different product definitions. With the goal of assessing the degree of substitution across different airlines, column 1 aggregates flights at the airline level. Likewise, columns 2 and 3 aggregate flights at the airport level and at the departure time level respectively. The aggregation of prices and sales in column 1 implies, for example, that an American Airlines flight that departs in the morning is the same product as an American Airlines flight that departs in the afternoon. For the specification in the third column we divide departure times in Morning (flights that depart before noon), Afternoon (flights that depart between noon and 5:00 p.m.), and Evening (flights that depart after 5:00 p.m.).

The revealed preference data set that we employ does not have information on the purpose of the trip (e.g., leisure versus business), so we include the number of days in advance in the vector of observable exogenous characteristics x_{jt} that affect utility. This can help control for the purpose of the trip as it is reasonable to argue that leisure travelers are more likely to buy in advance. Moreover, the data set does not contain membership on frequent-flier programs, which is known to affect airline choice. This can bias the constant estimates because, for example, an AAdvantage member is more likely to have a higher American constant than non-members. To allow for different constants across products (but still the same constant across heterogeneous buyers who buy the same product), all specifications

include product dummy variables that control for time-invariant characteristics.

BLP-based models typically use demographics to capture observed heterogeneity, however in our data, flights in all markets share the same demographics. Hence, in our specifications, observed heterogeneity H_i is drawn from the distribution of actual sales. Note that because marginal utilities β_i change by individual i , panel A in Table 2 reports the means of the distributions of marginal utilities, $\bar{\beta}$. Furthermore, panels B and C present the effects of heterogeneity around the means of the β s. In particular, panel B reports the standard deviations of the distribution of the marginal utilities, σ_β , which can be interpreted as the effects of unobserved heterogeneity. On the other hand, panel C presents the effects of observed heterogeneity, H_i , on the slope parameters.

[Table 2, here.]

From panel A, we observe that the mean price coefficients have the expected negative sign across all three specifications and that the marginal effects are statistically significant at the airline (column 1) and at the departure time (column 3) levels. The mean coefficient on days in advance is also negative and statistically significant at the airport and at the departure time levels. This suggests that the utility of flying is greater for those travelers who buy closer to departure. Travelers who typically buy closer to departure are more likely to have higher valuations (see, e.g., Dana, 1998), which is consistent with obtaining a higher utility of traveling.

The statistically significant coefficients in panels B and C are evidence of consumers' heterogeneity and that travelers have differentiated demands. Along with the widely documented price dispersion in airline markets, this is consistent with price discrimination practices in which airlines segment heterogeneous travelers to extract more consumer surplus (see, e.g., Escobari and Jindapon, 2014). In panel B, the estimates of the standard deviations are not statistically significant for price, but they are for days in advance in two of the product specifications (columns 1 and 3). These positive and statistically significant coefficients indicate that consumers' unobserved heterogeneity is lower closer to departure. Panel C reports the estimates of the interaction term of days in advance with observed consumers' heterogeneity H_i . The estimates show a positive and statistically significant effect only at the airport level (column 2). Consistent with the estimates in panel B, this

positive effect is evidence that observed consumers' heterogeneity, as driven by variation in sales, decreases as the departure date nears.

[Table 3, here.]

Table 3 presents the estimates of the own- and cross-price elasticities that correspond to the product definition at the airline level. These estimates are from the same specification of column 1 on Table 2 and were obtained using equations (10) with the integrals being approximated by simulations as presented in Vincent (2015). Following the same notation as in equations (10), the entry (j,k) in the table corresponds to row j and column k . This captures the percentage change in quantity demanded of carrier j when there is a one percent increase in the price of carrier k . The diagonal elements ($j = k$) are the own-price elasticities, while the off-diagonal elements ($j \neq k$) are the cross-price elasticities. As suggested in equations (10), the model does not imply constant elasticities so we use average values to evaluate them.

The figures along the main diagonal suggest that United travelers are the most price sensitive, followed by American and Delta. The least price sensitive are from Air Canada. Note that the aggregation of flights at the carrier level entails a fairly broad product definition which can help explain the relatively large own-price elasticities when compared to the cross-price elasticities. Carriers with similar characteristics are viewed as relatively closer competitors and are expected to have larger substitution patterns. Hence, when comparing across columns, it is reasonable to argue that for this particular route between New York City and Toronto, United, American, and Continental are relatively closer competitors, while Air Canada and Delta appear to have differentiated themselves more from the other carriers. This can explain the low cross-price elasticities for Delta. When comparing equations (9) and (10), we can observe that the structure imposed by the logit model forces all of the off-diagonal elements within the same column to be the same. Hence, we can assess the flexibility of our approach by examining the variation of the cross-price elasticities in each column. This flexibility in equations (10) also allows to explain the observed asymmetry in the cross-price elasticities.

Note that because we have data on one-way fares, the estimates should be interpreted with care if trying to generalize to round-trip itineraries. This is because most of the

“legacy carriers” (e.g., American, Delta, LAN Airlines) have different fare structures when considering one-way and round-trip tickets. For example, the one-way fare structure can be more expensive because corporate travelers who buy one-way tickets are likely to have higher valuations than tourists who are more likely to buy round-trip tickets. Hence, focusing on one-way tickets means that our demand estimates are more likely to capture the behavior of corporate travelers.

[Table 4, here.]

Tables 4 and 5 present the own- and cross-price elasticities for alternative product definitions. Table 4 follows the specification in column 2 of Table 2 and aggregates at the airport level. All three airports are operated by The Port Authority of New York & New Jersey. In terms of location, JFK and La Guardia are both located in Queens, while Newark is located west in New Jersey. All three are fairly close as there are less than 10 miles between La Guardia and JFK, 26 miles between La Guardia and Newark, and about 33 miles between JFK and Newark. In terms of total passenger traffic, the largest is JFK (53 million passengers in 2014), followed by Newark (35.6 million) and La Guardia (27 million).

The own-price elasticities presented on the main diagonal of Table 4 suggest that travelers from La Guardia are the most price sensitive, followed by JFK and Newark. Estimating cross-price elasticities at the airport level help decide which airports warrant being grouped as a single market in a multi-airport metropolitan area. The cross-price elasticities in Table 4 show evidence of demand-shifting across airports, which suggest that these three New York City area airports can be considered as a single market.¹⁴ In addition, the cross-price elasticities in columns 2 and 3 are relatively larger for Newark and JFK, implying that those two airports appear to be relatively closer substitutes. Capturing cross-price elasticities at the airport level also has policy implications, for example, when deciding on improvements in access time to an airport, improvements in flight frequency or when assessing flight delays (see, e.g., Bishop et al., 2011). Our estimates on the degree of demand-shifting across airports can also be valuable to policy makers when deciding to implement particular airport

¹⁴Brueckner et al. (2014) present a method to group airports into a single market based on the comparability of incremental competition effects from nearby airports.

congestion pricing strategies.

[Table 5, here.]

Table 5 presents the elasticity estimates at the departure time level following the specification of column 3 on Table 2. The relatively high own-price sensitivity for evening flights (when compared to morning and afternoon flights) can be explained by fewer flights being offered during the evening (17.3% of the flights in the sample). From the relatively low cross-price elasticities, we can infer that consumers have a relatively low willingness to switch to alternative departure times, which is consistent with a broad product definition. For example, the small cross-price elasticities in the first column imply that a price increase in flights departing during the evening has a negligible shift of consumers to flights that depart in the afternoon or in the morning. These relatively low cross-price elasticities also imply that pricing strategies aimed at solving congestion problems at particular departure times are unlikely to succeed.

Our estimates of own-price elasticity estimates are close to estimates found by previous studies.¹⁵ Overall, when comparing the magnitudes of the estimates across Tables 3 through 5, we observe that consumers are more price sensitive at the carrier level, followed by airport and departure time. We interpret this as evidence that travelers are more willing to switch to an alternative carrier than to switch between airports or to alternative departure times.

5 Conclusion

In this paper we use the random-coefficients logit methodology to estimate various demand systems and examine the airport, airline, and departure time choice and substitution patterns. Our estimation approach controls for potential endogeneity of prices and allows for general substitution patterns that take into account the heterogeneity in consumers' tastes. Following Escobari (2012) and Escobari and Jindapon (2014), we take advantage of an original ticket-level revealed preference data set gathered from an online travel agency.

¹⁵Using 204 own-price elasticity observations from 37 previous studies, Brons et al. (2002) explain that own-price elasticities ranged between -3.20 and 0.21 with a mean elasticity of -1.146 .

One key advantage from these data is that it contains the same information on prices and characteristics of all flights available to consumers at the moment of booking.

The results using our data from the New York City area are consistent with the widely observed price dispersion and the price discrimination practices in the industry. We find that the utility of flying is greater for those travelers who buy closer to departure. Moreover, our estimates show that travelers' heterogeneity decreases as the flight date nears.

When looking at the own- and cross-price elasticity estimates, we find significant differences across airlines, airports and departure times. At the carrier level, the results show that United, American and Continental are relatively close competitors. On the other hand, Air Canada and Delta appeared to have differentiated themselves more from the other carriers. At the airport level, we find that travelers from La Guardia are the most price sensitive. In addition, our estimates of demand-shifting across airports provide evidence that the New York City area airports can be considered as a single market and that Newark and JFK are relatively closer substitutes when compared to La Guardia. The relatively low cross-price elasticities at the departure time level suggest that higher prices during peak times aimed at solving congestion problems are more likely to reduce overall demand for travel rather than shift passengers to less congested periods. Overall, our results show that consumers are more willing to switch to an alternative carrier than between airports or departure times.

The combination of our original data set and the random-coefficients logit estimation provide valuable information on the degree of competition between airports, between carriers, and between flights offered at different departure times. This characterization of the behavior of passengers can be valuable to local authorities, regulators, airport planners and airlines. One potential limitation of our study is that we do not model the choice of travel destination. While we can argue that this plays no role for business travelers, it might be important for tourist travelers whose destination might not be set a priori. This is still an open question for future research.

Appendices

A Elasticities

To obtain the elasticities of demand $\frac{\partial s_{jt}}{\partial p_{kt}} \frac{p_{kt}}{s_{jt}}$ in equations (9) we need to calculate $\frac{\partial s_{jt}}{\partial p_{kt}}$ first. Let $W_{jt} = \exp(x_{jt}\beta - \alpha p_{jt} + \xi_{jt})$, so we can write equation (8) as $s_{jt} = W_{jt}/[1 + \sum_{k=1}^J W_{kt}]$. Then,

$$\frac{\partial s_{jt}}{\partial p_{kt}} = \frac{\partial W_{jt}/\partial p_{kt}}{1 + \sum_{k=1}^J W_{kt}} + \left(\frac{-W_{jt}}{(1 + \sum_{k=1}^J W_{kt})^2} \right) \frac{\partial W_{kt}}{\partial p_{kt}}.$$

For the own-price elasticities, $j = k$, we have:

$$\begin{aligned} \frac{\partial s_{jt}}{\partial p_{jt}} &= \frac{-\alpha W_{jt}}{1 + \sum_{k=1}^J W_{kt}} + \left(\frac{-W_{jt}}{(1 + \sum_{k=1}^J W_{kt})^2} \right) (-\alpha W_{jt}) \\ &= -\alpha s_{jt} + \alpha s_{jt}^2 \\ &= -\alpha s_{jt}(1 - s_{jt}) \end{aligned} \quad (\text{A.1})$$

Hence, $\eta_{jkt} = \frac{\partial s_{it}}{\partial p_{jt}} \frac{p_{jt}}{s_{jt}} = -\alpha s_{jt}(1 - s_{jt}) \frac{p_{jt}}{s_{jt}} = -\alpha p_{jt}(1 - s_{jt})$. For the cross-price elasticities, $j \neq k$, we have:

$$\begin{aligned} \frac{\partial s_{jt}}{\partial p_{kt}} &= \frac{0}{1 + \sum_{k=1}^J W_{kt}} + \left(\frac{-W_{jt}}{(1 + \sum_{k=1}^J W_{kt})^2} \right) (-\alpha W_{kt}) \\ &= \alpha \left(\frac{-W_{jt}}{1 + \sum_{k=1}^J W_{kt}} \right) \left(\frac{-W_{kt}}{1 + \sum_{k=1}^J W_{kt}} \right) \\ &= \alpha s_{jt} s_{kt}. \end{aligned} \quad (\text{A.2})$$

Then $\eta_{jkt} = \frac{\partial s_{it}}{\partial p_{kt}} \frac{p_{kt}}{s_{jt}} = \alpha s_{jt} s_{kt} \frac{p_{kt}}{s_{jt}} = \alpha p_{kt} s_{kt}$. For the elasticities in equations (10), the steps to obtain $\frac{\partial s_{ijt}}{\partial p_{ikt}}$ are the same. Then we need to integrate across individuals weighting by its probability in the population.

B Estimation

The estimation follows Berry (1994), Berry et al. (1995), and Nevo (2000b) to obtain consistent estimates of the parameters in the model presented in section 3. As in Berry (1994), we use instruments and compound the error term to form the GMM objective function to be minimized:

$$\min_{\theta} \|s(x, p, \delta(x, p, \xi; \theta_1); \theta_2) - S\| \quad (\text{B.1})$$

The market shares $s(\cdot)$ are the ones defined in equation (6), and S are the observed market shares. Direct minimization of the objective function in equation (B.1) is difficult because most parameters (θ_1, θ_2) enter nonlinearly, making the minimization process very costly.

There are many flight-level characteristics that are unobserved ξ_j along with various individual traveler characteristics that are also unobserved, $(H_i, v_i, \varepsilon_i)$. This comes on top of having potential correlation between airline prices and the unobserved econometric error term. To take this into account we use Chamberlain's (1987) optimal set of instruments and lags of prices for the instruments Z . With $\omega(\cdot)$ being a function of the model parameters, the population moment conditions are then:

$$E[Z_m \omega(\theta^*)] = 0, \quad m = 1, 2, \dots, M, \quad (\text{B.2})$$

The population parameters are denoted by θ^* and its two-step GMM estimates are:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \omega(\theta)' Z \Phi^{-1} Z' \omega(\theta). \quad (\text{B.3})$$

In the first step, we obtain Φ as the (consistent) estimate of $E[Z' \omega \omega' Z]$. As in Berry (1994), when defining the error term in equation (B.1), we use the structural error ξ_{jt} . Moreover, to implement equation (B.3), the error terms need to be written as a function of the parameters and the data. This involves solving for each market the implicit system of equations $s(\delta_t; \theta_2) = S_t$. This is done by computing $s(\delta_t; \theta_2)$ using equation (6). The ε s are integrated analytically under the assumption that $P_\varepsilon^*(\varepsilon)$ follows an extreme-value distribution. However, because the other two integrals in equation (6) cannot be evaluated analytically, we use Monte Carlo integration to approximate them with,

$$s_{jt}(x_t, p_t, \delta_t; \theta_2) = \frac{1}{ns} \sum_{i=1}^{ns} s_{ijt} = \frac{1}{ns} \sum_{i=1}^{ns} \frac{\exp(\delta_{jt} + \mu_{ijt})}{\sum_{k=1}^J \exp(\delta_{kt} + \mu_{ikt})}. \quad (\text{B.4})$$

The simulations assume a number ns of individuals (draws) and use the variables that have random coefficients as well as draws from $\hat{P}_v^*(v)$ and $P_H^*(H)$. The system of nonlinear equations $s(\delta_t; \theta_2) = S_t$ is then solved numerically using the contraction mapping proposed in Berry et al. (1995).¹⁶

Berry et al. (2004) provide the asymptotic distribution theory for this estimator. They allow for three sources of error: the sampling error in estimating market shares, the simulation error in approximating the shares predicted by the model, and the underlying model

¹⁶Please see Berry (1994) and Berry et al. (1995, 2006) for more details.

error. They show that the limiting distribution of the parameter estimator is normal provided that the size of the consumer sample, n , and the number of simulation draws, ns , grow at a large enough rate relative to the number of products J . In particular, the estimator will be consistent if $J \log J/n$ and $J \log J/ns$ converge to zero as J increases. For asymptotic normality at rate \sqrt{J} , we require J^2/n and J^2/ns to be bounded.

References

- Anderson, T. and Hsiao, C. (1981). Estimation of dynamic models with error components. *Journal of the American Statistical Association*, 76(375):598–606.
- Arellano, M. and Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies*, 58(2):277–297.
- Bachis, E. and Piga, C. A. (2011). Low-cost airlines and online price dispersion. *International Journal of Industrial Organization*, 29(6):655–667.
- Başar, G. and Bhat, C. (2004). A parameterized consideration set model for airport choice: an application to the San Francisco Bay area. *Transportation Research Part B: Methodological*, 38(10):889–904.
- Barbot, C. (2009). Airport and airlines competition: Incentives for vertical collusion. *Transportation Research Part B: Methodological*, 43(10):952–965.
- Barbot, C., D’Alfonso, T., Malighetti, P., and Redondi, R. (2013). Vertical collusion between airports and airlines: An empirical test for the European case. *Transportation Research Part E: Logistics and Transportation Review*, 57(1):3–15.
- Belobaba, P. P. (1989). Application of a probabilistic decision model to airline seat inventory control. *Operations Research*, 37(2):183–197.
- Berry, S. (1994). Estimating discrete-choice models of product differentiation. *Rand Journal of Economics*, 25(2):242–262.
- Berry, S., Carnall, M., and Spiller, P. (2006). *Airline hubs: Costs, markups and the implications of customer heterogeneity*, ed. Darin Lee, volume 1 of *Advances in Airline Economics: Competition Policy and Antitrust*, pages 183–214. Amsterdam: Elsevier.
- Berry, S. and Jia, P. (2010). Tracing the woes: An empirical analysis of the airline industry. *American Economic Journal: Microeconomics*, 2(3):1–43.
- Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

- Berry, S., Linton, O. B., and Pakes, A. (2004). Limit theorems for estimating the parameters of differentiated product demand systems. *Review of Economic Studies*, 71(3):613–654.
- Bilotkach, V. (2007). Asymmetric regulation and airport dominance in international aviation: Evidence from the London-New York market. *Southern Economic Journal*, 74(2):505–523.
- Bilotkach, V., Gorodnichenko, Y., and Talavera, O. (2012). Sensitivity of prices to demand shocks: A natural experiment in the San Francisco Bay area. *Transportation Research Part A: Policy and Practice*, 46(7):1137–1151.
- Bishop, J. A., Rupp, N. G., and Zheng, B. (2011). Flight delays and passenger preferences: An axiomatic approach. *Southern Economic Journal*, 77(3):543–556.
- Borenstein, S. and Rose, N. L. (1994). Competition and price dispersion in the us airline industry. *Journal of Political Economy*, 102(4):653–683.
- Brons, M., Pels, E., Nijkamp, P., and Rietveld, P. (2002). Price elasticities of demand for passenger air travel: A meta-analysis. *Journal of Air Transport Management*, 8(3):165–175.
- Brueckner, J. K. (2002). Airport congestion when carriers have market power. *American Economic Review*, 92(5):1357–1375.
- Brueckner, J. K., Lee, D., and Singer, E. (2014). City-pairs versus airport-pairs: A market-definition methodology for the airline industry. *Review of Industrial Organization*, 44(1):1–25.
- Chamberlain, G. (1987). Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of Econometrics*, 34(3):305–334.
- Dana, Jr., J. D. (1998). Advance-purchase discounts and price discrimination in competitive markets. *Journal of Political Economy*, 106(2):395–422.
- de Luca, S. (2012). Modelling airport choice behaviour for direct flights, connecting flights and different travel plans. *Journal of Transport Geography*, 22:148–163.

- Deaton, A. and Muellbauer, J. (1980). An almost ideal demand system. *American Economic Review*, 70(3):312–326.
- Deneckere, R. and Peck, J. (2012). Dynamic competition with random demand and costless search: A theory of price posting. *Econometrica*, 80(3):1185–1247.
- Drabas, T. and Wu, C.-L. (2013). Modelling air carrier choices with a segment specific cross nested logit model. *Journal of Air Transport Management*, 33(1):8–16.
- Escobari, D. (2009). Systematic peak-load pricing, congestion premia and demand diverting: Empirical evidence. *Economics Letters*, 103(1):59–61.
- Escobari, D. (2012). Dynamic pricing, advance sales, and aggregate demand learning in airlines. *Journal of Industrial Economics*, 60(4):697–724.
- Escobari, D. (2014). Estimating dynamic demand for airlines. *Economics Letters*, 124(1):26–29.
- Escobari, D. and Jindapon, P. (2014). Price discrimination through refund contracts in airlines. *International Journal of Industrial Organization*, 34(3):1–8.
- Escobari, D. and Mellado, C. (2014). *The choice of airport, airline, and departure date and time: Estimating the demand for flights*, volume 4 of *Advances in Airline Economics: The Economics of International Airline Transport*, pages 177–198. Emerald Group Publishing.
- Escobari, D., Rupp, N., and Meskey, J. (2017). Price discrimination and focal points for tacit collusion: Evidence from the airline industry. Working Paper. East Carolina University. Available at SSRN: <https://ssrn.com/abstract=2815279>.
- Fu, X., Homsombat, W., and Oum, T. H. (2011). Airport-airline vertical relationships, their effects and regulatory policy implications. *Journal of Air Transport Management*, 17(6):347–353.
- Gallego, G. and van Ryzin, G. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8):999–1020.

- Gerardi, K. and Shapiro, A. (2009). Does competition reduce price dispersion? new evidence from the airline industry. *Journal of Political Economy*, 117(1):1–37.
- Harvey, G. (1987). Airport choice in a multiple airport region. *Transportation Research Part A: General*, 21(6):439–449.
- Hensher, D. A., Stopher, P. R., and Louviere, J. J. (2001). An exploratory analysis of the effect of numbers of choice sets in designed choice experiments: An airline choice application. *Journal of Air Transport Management*, 7(6):373–379.
- Hess, S. (2010). Evidence of passenger preferences for specific types of airports. *Journal of Air Transport Management*, 16(4):191–195.
- Hess, S., Adler, T., and Polak, J. W. (2007). Modelling airport and airline choice behaviour with the use of stated preference survey data. *Transportation Research Part E: Logistics and Transportation Review*, 43(3):221–233.
- Hess, S. and Polak, J. W. (2005a). Accounting for random taste heterogeneity in airport choice modelling. *Transportation Research Record*, 1915:36–43.
- Hess, S. and Polak, J. W. (2005b). Mixed logit modelling of airport choice in multi-airport regions. *Journal of Air Transport Management*, 11(2):59–68.
- Hess, S. and Polak, J. W. (2006). Exploring the potential for cross-nesting structures in airport-choice analysis: A case-study of the Greater London area. *Transportation Research Part E: Logistics and Transportation Review*, 42(2):63–81.
- Ishii, J., Jun, S., and Van Dender, K. (2009). Air travel choices in multi-airport markets. *Journal of Urban Economics*, 65(2):216–227.
- Knittel, C. R. and Metaxoglou, K. (2014). Estimation of random-coefficient demand models: Two empiricists’ perspective. *Review of Economics and Statistics*, 96(1):34–59.
- Koster, P., Kroes, E., and Verhoef, E. (2011). Travel time variability and airport accessibility. *Transportation Research Part B: Methodological*, 45(10):1545–1559.

- Loo, B. P. (2008). Passengers airport choice within multi-airport regions (MARs): Some insights from a stated preference survey at Hong Kong International Airport. *Journal of Transport Geography*, 16(2):117–125.
- Marcucci, E. and Gatta, V. (2011). Regional airport choice: Consumer behaviour and policy implications. *Journal of Transport Geography*, 19(1):70–84.
- Mayer, C. and Sinai, T. (2003). Network effects, congestion externalities, and air traffic delays: Or why not all delays are evil. *American Economic Review*, 93(4):1194–1215.
- McFadden, D. (1973). *Conditional logit analysis of qualitative choice behavior*, ed. P. Zarembka, pages 105–142. *Frontiers of Econometrics*. New York: Academic Press.
- Morrison, S. A. and Winston, C. (2007). Another look at airport congestion pricing. *American Economic Review*, 97(5):1970–1977.
- Nevo, A. (2000a). Mergers with differentiated products: The case of the ready-to-eat cereal industry. *Rand Journal of Economics*, 31(3):395–421.
- Nevo, A. (2000b). A practitioner’s guide to estimation of random-coefficients logit models of demand. *Journal of Economics & Management Strategy*, 9(4):513–548.
- Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.
- Njegovan, N. (2006). Elasticities of demand for leisure air travel: A system modelling approach. *Journal of Air Transport Management*, 12(2):33–39.
- Ortúzar, J. de D. and Simonetti, C. (2008). Modelling the demand for medium distance air travel with the mixed data estimation method. *Journal of Air Transport Management*, 14(6):297–303.
- Pels, E., Nijkamp, P., and Rietveld, P. (2000). Airport and airline competition for passengers departing from a large metropolitan area. *Journal of Urban Economics*, 48(1):29–45.
- Pels, E., Nijkamp, P., and Rietveld, P. (2001). Airport and airline choice in a multiple airport region: An empirical analysis for the San Francisco Bay area. *Regional Studies*, 35(1):1–9.

- Pels, E., Nijkamp, P., and Rietveld, P. (2003). Access to and competition between airports: a case study for the San Francisco Bay area. *Transportation Research Part A: Policy and Practice*, 37(1):71–83.
- Pels, E., Njegovan, N., and Behrens, C. (2009). Low-cost airlines and airport competition. *Transportation Research Part E: Logistics and Transportation Review*, 45(2):335–344.
- Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110(4):705–729.
- Prousaloglou, K. and Koppelman, F. S. (1999). The choice of air carrier, flight, and fare class. *Journal of Air Transport Management*, 5(4):193–201.
- Reynaert, M. and Verboven, F. (2014). Improving the performance of random coefficients demand models: The role of optimal instruments. *Journal of Econometrics*, 179(1):83–98.
- Rupp, N. G. (2009). Do carriers internalize congestion costs? Empirical evidence on the internalization question. *Journal of Urban Economics*, 65(1):24–37.
- Skinner, R. E. (1976). Airport choice: An empirical study. *Journal of Transportation Engineering*, 102(4):871–883.
- Stavins, J. (2001). Price discrimination in the airline market: The effect of market concentration. *Review of Economics and Statistics*, 83(1):200–202.
- Tan, K. M. and Samuel, A. (2016). The effect of de-hubbing on airfares. *Journal of Air Transport Management*, 50:45–52.
- Vincent, D. W. (2015). The Berry-Levinsohn-Pakes estimator of the random-coefficient logit demand model. *Stata Journal*, 15(3):854–880.
- Williams, K. R. (2014). Dynamic airline pricing and seat availability. Working Paper. Yale University.
- Zhang, A., Fu, X., and Yang, H. G. (2010). Revenue sharing with multiple airlines and airports. *Transportation Research Part B: Methodological*, 14(8-9):944–959.

Zhang, Y. and Xie, Y. (2005). Small community airport choice behavior analysis: A case study of GTR. *Journal of Air Transport Management*, 11(6):442–447.

Table 1: Summary Statistics

VARIABLES	(1) #Flights	(2) Mean	(3) SD	(4) Min	(5) Max
<i>Panel A. Whole Sample:</i>					
Price	317	169.8	133.3	65	1,075
Days in Advance	317	19.00	12.43	1	40
Sale	317	0.0191	0.137	0	1
<i>Panel B. Price by Airline:</i>					
American	67	192.38	194.72	76	1,075
Air Canada	97	144.57	69.15	81	736
Continental	42	155.21	114.51	77	1,001
Delta	23	234.36	203.77	105	953
LAN Airlines	4	130.45	25.48	123	220
United	84	173.42	105.25	65	1,008
<i>Panel C. Price by Airport:</i>					
Newark Liberty	113	150.73	101.27	65	1,001
La Guardia	163	171.87	133.96	76	1,075
John F. Kennedy	41	214.46	186.27	87	1,075
<i>Panel D. Price by Departure Time:</i>					
Morning	142	142.06	95.94	65	1,075
Afternoon	120	192.68	162.87	77	1,075
Evening	55	191.55	130.96	81	953

Notes: The sample size is 560,244. Fares measured in US\$.

Table 2: Demand System Estimates

Model:	(1)	(2)	(3)
Substitution Patters:	Airline	Airport	Departure Time
<i>Panel A. Means of the distributions of marginal utilities ($\bar{\beta}$):</i>			
Price/100	-1.101 [‡] (0.640)	-0.903 (0.662)	-0.559* (0.165)
Constant	-1.759 [‡] (1.001)	-2.431 [‡] (1.474)	-2.945* (0.710)
Days in Advance	-0.288 [†] (0.136)	-0.202 (0.207)	-0.239 [‡] (0.126)
<i>Panel B. Standard deviations of the distributions of marginal utilities (σ_{β}):</i>			
Price/100	0.352 (0.365)	0.368 (0.327)	$8.85e^{-8}$ (2.037)
Days in Advance	0.223 [†] (0.098)	0.142 (0.119)	0.172 [†] (0.073)
<i>Panel C. Interactions with observed Heterogeneity:</i>			
Days in Advance	0.008 (0.006)	0.0112 [‡] (0.0067)	0.000542 (0.00671)
Number of Observations	420	252	252
Number of Markets	84	84	84
Halton draws	500	500	500

Notes: The table reports GMM parameter estimates. All regressions control for product fixed effects. ‡ significant at 10%; † significant at 5%; * significant at 1%. The numbers in parentheses are asymptotically robust standard errors.

Table 3: Median Own- and Cross-Price Elasticities at the Airline Level

	(1)	(2)	(3)	(4)	(5)
	United	Delta	Continental	Air Canada	American
United	-2.7503	$1.15e^{-4}$	0.0354	0.00971	0.0234
Delta	0.0304	-2.1485	0.0338	0.0100	0.0205
Continental	0.0270	$9.83e^{-5}$	-1.8121	0.0100	0.0193
Air Canada	0.0234	$9.20e^{-5}$	0.0316	-1.4269	0.0179
American	0.0315	$1.05e^{-4}$	0.0341	0.0100	-2.2141

Notes: The table reports GMM parameter estimates. The sample includes 420 observations in 84 markets.

Table 4: Median Own- and Cross-Price Elasticities at the Airport Level

	(1)	(2)	(3)
	La Guardia	John F. Kennedy	Newark Liberty
La Guardia	-1.6873	0.0154	0.0132
John F. Kennedy	0.0117	-1.4435	0.0143
Newark Liberty	0.0112	0.0160	-1.2311

Notes: The table reports GMM parameter estimates. The sample includes 420 observations in 84 markets.

Table 5: Median Own- and Cross-Price Elasticities
at the Departure Time Level

	(1) Evening	(2) Afternoon	(3) Morning
Evening	-2.7100	0.00589	0.0131
Afternoon	$9.29e^{-5}$	-1.4545	0.0131
Morning	$9.29e^{-5}$	0.00589	-0.8833

Notes: The table reports GMM parameter estimates. The sample includes 420 observations in 84 markets.