

Air Ticket Sales as Bids from Airline Alliances

Marc Ivaldi*, Milena Petrova[†], Miguel Urdanoz[‡]

January 26, 2015

Abstract

Motivated by the higher price sensitivity and service homogenisation in the airline industry in recent years, we propose a new methodology to deal with transaction prices and to estimate the effect of alliances in the US domestic market. The assumption that airlines compete on price allows us to take advantage of the observational equivalence between Bertrand competition and the reverse English auction. We then apply an MLE method, developed by Paarsch (1997) for estimating auctions, to recover the distributional characteristics of air fares using a sample of airline tickets from the US domestic market. This procedure allows us to benefit from the heterogeneity of individual prices while most studies have used average prices, which would have involved a loss of information and a potential bias. We find that an alliance operating in a market is associated with prices on average 18.9 percent higher. Additionally, we find the standard deviation of ticket prices to be 4.3 percent higher, which is likely related to more efficient revenue management practice by alliance partners operating together in the same market.

JEL classification: R48, L40, L93.

Keywords: airlines, alliances, airfares.

*Toulouse School of Economics; marc.ivaldi@tse-fr.eu

[†]PhD Candidate, Toulouse School of Economics; milenajpetrova@gmail.com

[‡]Toulouse Business School; m.urdanoz@tbs-education.fr

The authors would like to thank David Salant, Darin Lee, Diego Escobari and Paul Scott for their invaluable comments and suggestions. We are also thankful to the participants at the 2014 EARIE Conference in Milan, the 2014 ITEA Conference in Toulouse, and the 2014 TSE Student Workshop in Toulouse. Any remaining errors are our own.

1 Introduction

An airline alliance is an agreement between two or more carriers that allows them to cooperate. More specifically, alliance partners can market their partner's tickets and collaborate in supplying a product (what is known as *codeshare*), they can coordinate their schedules, and they can share frequent-flyer programmes and promotional campaigns. International alliances, such as SkyTeam, Star and OneWorld, are very well-known but they are different in a number of crucial ways from the US domestic market alliances. Access to international markets is regulated, hence partnering with a foreign competitor allows an airline to reach the foreign market and this is complementary to its home country operations. Additionally, these type of agreements generally receive anti-trust immunity.² The US domestic market, on the other hand, is not regulated and any carrier is free to provide service between any two cities.³ The networks of alliance carriers can overlap and market access and complementarity of service are no longer arguments for alliance formation. A number of concerns have been raised as regarding to the potential anti-competitive effects of alliances, namely their effects on prices, traffic and market entry. Alliance proponents have argued that improved operating efficiency, frequency of service and network reach would compensate consumers for any potential anti-competitive effects. The Department of Justice does not grant domestic alliances anti-trust immunity and has followed their operation closely since their formation in the 1980's. In this paper, we would like to focus on one particular competition aspect, the prices of alliance partners in overlapping markets, and to provide a new methodology that we believe is useful for understanding this important matter.

While there exists previous academic work on the effect of airline alliances on prices, we believe that the more dynamic conditions in recent years in the US domestic market oblige us to revisit the issue and previous results. Moreover, we propose a new methodology that is motivated by the increasing competitiveness of this market, an issue that has received much attention from industry experts.⁴ Our model emphasizes the homogenisation of service and the increased price sensitivity in the airline market. We assume the ticket sales process is adequately approximated by Bertrand competition, where the consumer makes his choice based on price and the competing airlines try to offer the best price they can. Our methodology exploits the strategic equivalence between the Bertrand game and a reverse English auction. A *reverse English auction* has the "inverse" set-up of a regular English auction - the auctioneer is the buyer, and

²For example, the Department of Justice gives anti-trust immunity to the following SkyTeam airlines: Delta/Air France-KLM/Alitalia/Czech/Korean.

³One exception is airports where traffic congestion is monitored by the Department of Transportation, thereby the number of available slots and their allocation is regulated. However, only a few airports are affected by this issue.

⁴To be discussed at length in the following pages.

sellers compete in offering prices at which they would be willing to supply the good or service. This interpretation allows us to consider the observed air fares as winning bids and we propose to analyse their distribution by methods pertaining to the econometrics of auctions.

Our methodology is different from what has been done previously, and we believe it contributes to the literature by allowing for a more comprehensive treatment of the price data. While previous empirical studies on the impact of airline alliances investigate average fares, that is to say fares aggregated over passengers, per airline, per market and per period, we work with individual transaction prices. The proposed auction game allows us to model the distribution of prices and this way our analysis benefits from the information we can extract from the full distribution rather than just the mean. We avoid the risk of distorting the results that using an aggregated variable could create when one is not always able to control for the impact of averaging in the estimation method. We implement our method on the DB1B data direct service duopoly markets operated by legacy carriers during Q3 of 2008. We find that the presence of an alliance in a duopoly market (i.e., the two players in that market are in an alliance) is associated with a 18.9 percent higher fares. Additionally, we find a positive effect of alliances on the standard deviation of prices of 4.3 percent, demonstrating an improved ability of the alliance partners to price discriminate.

In the rest of this section, we briefly discuss how the DB1B data has been used in the literature, we review the research on airline alliances, and we describe the recent changes in the competitive environment of the airline industry motivating our model.

The DB1B survey

The US Department of Transportation (DOT) publishes a comprehensive data source, the Airline Origin and Destination Survey (DB1B), since 1993. It is a 10 percent sample of all airline tickets sold in the US domestic market with information on the price, the origin and destination, and the itinerary details of the passengers. The DB1B, thus, is the standard data set used for analysing any issue that pertains to the US airline industry.

Each observation in the DB1B data set corresponds to an actual individual sale. The main feature of the data is that while the observation contains the fare paid for the sale (the transaction price), only the *market*⁵ and *product*⁶ characteristics of the ticket sale

⁵The relevant *market* is defined as the directional city pair: the origin and the destination of the flight.

⁶The *product* characteristics are the operating airline, and whether the flight is direct or connecting.

are recorded. The *flight*⁷ characteristics and *advance purchase*⁸ for that ticket sale are not available in the DB1B. To illustrate this issue, consider the following example. Let's take the sale of a direct return ticket by United Airlines between Chicago and Seattle sold 5 days in advance for a 9 AM outbound flight on Monday and 6 PM inbound flight on Wednesday. From the DB1B, we are not able to tell if the outbound flight was on Monday or on Wednesday, at 9 AM or at 9 PM, and similarly for the inbound flight. We are also not be able to tell of the ticket was purchased 15 or 5 days in advance. However, we do know the relevant market (Chicago to Seattle), the operating carrier (United) and that the flight is direct.

The DB1B has been used extensively to analyse different market and product features of the US domestic industry as it is freely available and standardised by aggregating prices at the market or product levels, respectively. Airline alliances on the US domestic market have also been analysed using the DB1B, with some studies focusing on their market impact as in Gayle (2008) and some - on the different kinds of alliance products as in Ito and Lee (2007). We discuss this work with more detail in the next section.⁹

Domestic alliances in the literature

A large portion of the literature on domestic airline alliances uses the DB1B prices aggregated at the market level when estimating the impact of alliances. This means

⁷The *flight* is characterised by the hour and day of the week of take-off for the outbound and inbound flights.

⁸*Advance purchase* provides the number of days between the date of purchase and the date of take-off.

⁹A different branch of literature is concerned with the variation of prices, or price dispersion, that is partially due to the unobserved flight features and advance purchase in the DB1B data set. This variation is generally seen as evidence of the airline to price discriminating between customers based on flight characteristics and inter-temporally. The ability to price discriminate may be related to a number of factors, such as the level of competition or stochastic demand. The price dispersion literature typically uses the Gini coefficient as a measure of the variation, following the early work by Borenstein and Rose (1994). More recently, data with flight-level and advance purchase characteristics has become publicly available for collection through online marketing of tickets by the airlines themselves and by online travel agents. Escobari (2012) uses data from Expedia to study the dynamic pricing of inventories with uncertain demand and over a finite horizon. Similarly, Lazarev (2012) and Williams (2013) examine the issue of dynamic pricing in a structural framework. The structural approach applied to data collected from online sources has great potential for discovering more about airline dynamic pricing. However, the data collection process should be carefully considered for external validity. Moreover, structural models using this kind of data so far are limited to the monopoly case because of the high complexity of modelling competition in this framework. This is the case in both Williams (2013) and Lazarev (2012). Additionally, there is the issue of the ownership of the data and airlines may not be willing to grant access to it, as experienced by McAfee and te Velde (2006) with American Airlines.

the impact of an alliance is the difference in prices between markets with and without an alliance, before and after the alliance entry or cross-sectionally. One of the first studies of alliances in the US domestic market is Bamberger, Carlton and Neumann (2004), who use a before-after approach to compare fares and traffic in connecting markets where the alliances operate relative to non-alliance *interline*¹⁰ markets. Similarly, Armantier and Richard (2006) use panel data methods to estimate the effect of the presence of alliance products on markets when the markets are segmented based on whether at least one carrier already offers a non-stop service. Lastly, Gayle (2008) focuses on own (alliance members only) and market price and traffic effects of the Continental-Northwest-Delta alliance, again using a before-after approach. These early studies find higher fares in markets where the alliance partners operate, but they also find increase in traffic volume in many markets. Armantier and Richard (2006) attribute the increased prices to alliance partners being able to price discriminate and manage the stochastic demand more profitably.

Another strand of the literature emphasizes the type of cooperation between alliance partners as a feature of the offered product. Ito and Lee (2007) analyse different itineraries¹¹ within a market, and distinguish between *virtual* code-shared products (where one partner operates the flight, and the other can sell tickets on that flight) and *traditional* code-shared products (where both partners are involved in the operation of the flight and both can sell tickets).¹² Their conclusion is that alliance products are seen as inferior by consumers in comparison pure *online* flights (flights operated and marketed by the same airline), and used by airlines to price-discriminate between consumers with different willingness to pay. Gayle (2008) performs a somewhat similar exercise, but he focuses on the market price effect of the presence of traditional and virtual code-share flights. More recently, Urdanoz and Sampaio (2012) use panel data extracted from the DB1B database to explore the persistence of gaps in the fares of different kinds of alliance and non-alliance products over a period of six years and find that these gaps are diminishing over time.

While the empirical evidence on the competitive effect of alliances is rich, it is difficult to draw unilateral conclusions. For example, Gayle (2008) and Ito and Lee (2007) both look at the same time period, but the effects they estimate are different. The first examine the impact that each cooperation type has on the market-level fares and traffic. The second, on the other hand, focus on cooperation type as a product characteristic.

¹⁰*Interline* products or markets are such that the service is performed by two carriers with no agreement, hence the service of each carrier is independently provided.

¹¹Their definition of itinerary is what we call here products, or prices identified by a combination of market and operating and marketing carrier.

¹²It is also interesting to note that traditional code-sharing, which takes advantage of the complementarity of the networks of alliance partners, is not as widespread as some might believe. In fact, Ito and Lee (2007) report that 85 percent of their products are virtual code-shares and 70 percent exhibit overlap of service between alliance partners.

Another difference is that Gayle (2008) performs a before and after estimation, while Ito and Lee (2007) look at a cross-sectional difference in price between different type of cooperation products. As a result, comparison between empirical studies can be difficult. While Gayle (2008) finds a positive price effect of the presence of virtual code-shared products in a market (on both the market fare and the alliance carrier fare), Ito and Lee (2007) show that virtual code-shared products are sold at a discount compared to online products.

The degree of competition in the airline industry

The US airline industry was regulated until 1978, and the major carriers of the period (the so-called *legacy* carriers) had exclusive rights to operation in certain regions of the US market and prices were set by the Civilian Aeronautics Board. During the regulated period carriers put emphasis on continuously improving the quality of service. After deregulation, the market was open to entry and price competition. Legacy carriers started forming alliances with international carriers in the early 90's, and these type of agreements were soon extended to other domestic legacy carriers, the first being the agreement between Northwest (NW) and Continental (CO) in 1998. While the majority of early research work on the industry focused on competition and network formation in the post-deregulation period, recently the industry has experienced a number of other developments. The entry and increased market share of low cost carriers (*LCCs*), the economic crises, advances in aviation technology, higher oil costs, and transparent internet ticket sales have all put pressure on the market players, and the industry has witnessed quite a few bankruptcies in the 2000's. These changes have been explored in the literature that we summarize below.

Intense competitive pressure in the airline industry has led to lowered emphasis on quality and increased product homogenisation. In their recent study of determinants of airline profitability in the years between 1999 and 2006, Berry and Jia (2010) find that legacy carriers have reduced their services and started to compete more intensively on prices. Brueckner, Lee and Singer (2012) suggest that service homogenisation is linked to competition from of LCCs, who have significantly expanded their operations in the last decade. An additional contributor to this effect, noted by Borenstein and Rose (2013), is the elimination of search costs by internet ticket sales. The different kinds of tools that have recently become available on the internet (e.g. Kayak, Expedia, Orbitz, Google Flights) facilitate search based on ticket parameters, and the comparison of prices and characteristics. This has reduced the demand for full-fare unrestricted tickets and induced the "unbundling" of services that are now available at additional fees (e.g. meals and check-in luggage).

Price elasticity has also increased considerably in the recent period. Berry and Jia (2010) find an increase in aggregate price elasticity by 8 percent between 1999 and 2006, with both business and tourist-specific elasticities rising. The authors credit the dot-com bust and internet ticket sales for this trend, and demonstrate that these changes in demand explain roughly 50 percent of the observed reduction in legacy carrier profits in the examined period. In their study of LCCs and adjacent airports, Brueckner, Lee and Singer (2012) similarly discuss how in recent years internet ticket sales have increased price transparency and business travellers have become more prudent, resulting in the increase of price sensitivity of customers.

Lastly, there has been higher provision of direct flights resulting from a combination of increased consumer sensitivity to connections, higher fuel price and aviation technology improvements. Berry and Jia (2010) find a 13 percent increase in direct passengers and a 23 decrease in connecting passengers over the period between 1999 and 2006. The semi-elasticity of connection has increased by 17 percent, as consumers can more easily compare and assess the relative merit of direct and connecting flights using online tools. Additionally, the cost advantages of connecting flights have disappeared with more efficient smaller aircraft coming in use and higher fuel prices that make multiple take-offs and landings too costly. This has likely brought about the recent trend in de-hubbing by legacy carriers, for example Delta de-hubbing at Memphis and Southwest at Atlanta (Berry and Jia (2010)).

Overall, the current economic conditions of the US domestic market seems to have put a rather strong competitive pressure on legacy airlines, through increased price sensitivity, homogenisation of the service, and low demand for connecting flights. These trends are important in the recent strategic environment of the airline industry and they motivate our assumptions and data selection. Our model emphasises price competition between carriers offering homogeneous service, and we focus on direct markets.

The next section presents the structural model we propose, and how it will use the DB1B data in a new way to estimate the effect of airline alliances on both the mean and standard deviation of prices. Next, in Section 3 we introduce the data and covariates used in the estimation. Lastly, Section 4 discusses the estimation results and robustness checks.

2 Bidding for travellers

We now provide an overview of the model, briefly describe the derivation of the maximum likelihood estimator, and then discuss how we use the DB1B data set to estimate the distribution of reservation costs and the effect of airline alliances.

Overview

As we already discussed in the introduction, an important consideration about working with the DB1B data set is that flight level and advance purchase information are not observed. For this reason, most work aggregates (by averaging) the observed fares at the product or market level. This kind of approach may be potentially problematic, and it ignores the issue of price variability in airline markets. In addition, working with average product or market prices weighs each market equally in the estimation of a given effect, but the correct weight might be at the ticket sale level to account for different levels of demand.¹³

We propose a different framework that we believe is appropriate in the current economic environment faced by airlines. The recent trends in the industry, specifically service homogenisation and high price elasticity, motivate our assumption that firms compete by setting prices. For example, consider two airlines that offer a flight from New York to Chicago leaving on Monday morning, returning Wednesday afternoon; the consumer interested in these particular flight characteristics will buy the cheapest ticket on offer.¹⁴

Assumption 1. *The observed transaction price (ticket fare) is the winning price in a reverse English auction between airlines.*

This assumption is based on the following argument. Given the particular travel characteristics he is interested in, the consumer chooses the cheapest ticket on offer. The

¹³To the best of our knowledge, the only study which uses the DB1B data at the flight level is Armantier and Richard (2008)'s discrete choice model. A potential problem in working with the data at the flight level may come from not acknowledging the difference in the set of alternatives (and prices) faced by each consumer. Armantier and Richard (2008) warn against using the average product fare as a proxy for the unobserved alternative prices because that would create measurement error. Their proposed solution to this problem is innovative yet rather challenging as it requires an auxiliary data set and a significant number of modelling assumptions.

¹⁴We let aside entry and exit issues, which are studied by Berry (1992). However, we control for the potential bias that we could face by including market fixed effects.

competitive pressure on the price causes airlines to undercut each other's price until it is no longer profitable to do so. This set-up describes Bertrand competition but also the reverse English auction, as the two games are strategically and outcome equivalent (Vives (2002)).¹⁵ The dominant strategy in this kind of game for each airline is to lower their price up to the point where it reaches their reservation cost, below which the airline would not be willing to make the sale. Take, for example, two airlines with reservation costs which are two random draws from a probability distribution. The equilibrium strategy implies that the airlines will undercut each other's prices until the one with the lower reservation cost (the player with the competitive advantage) slightly undercuts the price of its rival. The airline with the higher cost is not willing to lower its price further. Hence, the airline with the lower reservation cost wins the game and makes the sale at a price equal to the higher reservation cost.

Our model is in line with the widely spread yield management method of bid price control. *Yield* or *revenue management* is a variable pricing strategy that allows airlines to increase revenues in an environment with fixed capacity with an expiration date (the take-off of a plane) and uncertain demand. In practice, there are a number of techniques that the airline can use to achieve this, some of which involve the estimation of a marginal cost of each seat on a plane. One such method is bid price control, where this marginal cost is used as a bid - an optimum cut-off required to accept a booking. These bids correspond to reservation costs in our model, below which airlines are unwilling to sell tickets. Bid prices are dynamically adjusted over time to reflect the changing reservation costs under dynamic demand.¹⁶

In our framework, the reservation cost for a ticket sale is determined by the operating cost of providing the service, but also by its opportunity cost. The *operating cost* can be explained by the distance, and the economies of scale and scope, which are constant for all flights in a given market. Operating costs on a market need not be constant - multiple flights scheduled on the same market may have different operational costs depending on the departure time or day of the week. For example, flights at more busy hours may be more costly both in terms of the airline's own aircraft and personnel allocation and in terms of the airport-related costs. The reservation cost is also based on the ticket's *opportunity cost*, which is the value an airline places on selling the ticket now relative to selling it closer to the departure date with uncertainty but for a higher

¹⁵A reverse auction is when the auctioneer is a buyer and the participants are sellers who compete by offering the prices (their bids) at which they are willing to provide the service. An open auction of this kind, known as an English auction, is when competitors can observe each other's bids and react to them, as is the case when firms compete Bertrand-style. Note that in a sealed-bid auction the players cannot revise their bids after observing the bids of their competitors, even if they would like to. The English auction is strategically and outcome equivalent to the famous Vickrey auction (a second-price sealed-bid auction) in the case when players' valuations are independent (private).

¹⁶This practice has been analysed in the operations management literature by Talluri and Van Ryzin (1998) and Adelman (2007), among others.

price. For a given flight, the opportunity cost varies over time because it is related to the remaining capacity on the flight and on competing flights at the current moment, and the expectation of demand in future moments before take-off. The opportunity cost also incorporates the ability of an airline to price discriminate over time. In the present version of our empirical model it is not possible to estimate separately the opportunity and operating costs because of data limitations, however it is important to make clear the concept of the reservation cost and its underlying components.^{17,18}

Given our assumption that airlines compete on price, we develop an empirical model based on the observational equivalence between Bertrand competition and the reverse English auction. This observational equivalence between the two models allows us to apply an MLE method developed by Paarsch (1997) for estimating the parameters in auction models. In particular, we find this method appealing because we can estimate the parametric effect of available covariates on the mean and standard deviation of the reservation costs and transaction prices (the observed fares). We treat our sample of ticket sales as realisations of repeated auction games. In each round of this game, we assume that the airlines offer a product with similar flight features and compete on price, and the consumer chooses the cheapest ticket for the service. Each observation in the data represents the result of a different game, where the reservation costs have been re-drawn.

The model

In this section, we outline the equilibrium bidding strategy in the reverse English auction in the symmetric independent private values paradigm. We begin by assuming that the competition in this market can be modelled as a non-cooperative game. We identify the players, characterize the information each player has, describe the strategies available to the players, describe how each player is rewarded, and characterize the equilibrium.

¹⁷The proposed cost concepts are very close to the systematic and stochastic peak-load pricing described in the seminal paper on airline ticket price dispersion by Borenstein and Rose (1994). The authors distinguish between systematic peak-load pricing (the mean variation in the expected shadow cost of capacity that is known at the scheduling of the flight) and stochastic peak-load pricing (which is related to demand uncertainty and the pricing flexibility of a firm after capacity is set). The concept of *systematic peak – load pricing* is close to our operating costs component which depends on flight characteristics such as time and day of the week of take-off. Similarly, the concept of *stochastic peak – load pricing* is not unlike our definition of opportunity cost that depends on current market conditions, demand expectation and the competitive structure of the market. Borenstein and Rose (1994) note the difficulty in identifying these two effects using the DB1B data.

¹⁸A complete model of the operating costs would be possible if we had information available on the flight characteristics of each sale (day and hour of take-off of both inbound and outbound flights). To identify the dynamic opportunity cost, we would need to know the time of sale of each ticket, the remaining capacity of the carriers, and any information related to the anticipated level of demand.

We consider auctions at which a single buyer wishes to purchase one ticket from two sellers (a duopoly). Each player i is assumed to know his cost of providing the ticket c_i but not those of his opponents. The heterogeneity in costs is resulting from cost C being a continuous random variable with probability density function $f_C(c)$ and cumulative distribution $F_C(c)$. The costs of players are assumed to be independent draws from the same distribution $f_C(c)$. Together, the above assumptions constitute an independent, private costs auction game where the players are also symmetric.

Assumption 2. *Reservation costs are private, independent and ex-ante symmetric.*

The strategies available to the buyers are their bids as a function of the cost: $p_i = \sigma(c_i)$. Bidding at a reverse English auction was modelled by Milgrom and Weber (1982). We describe the proceedings of the Bertrand game as follows: given that the consumer wants to buy the cheapest ticket, the players will react to each other's prices and try to undercut the price of the competitor. The equilibrium strategy is to lower one's price up to $p_i = c_i$, and stop after. The winner is the player with the lowest reservation cost, and he is paid the price is the cost of his opponent (the player with the highest reservation cost).

Let us take a closer look at the symmetric private costs framework, and how it is motivated by some observations on airline markets and ticket sales. In the private values framework, the c_i of each player i is private information and it also has private relevance, in the sense that if it is revealed it would not make other players re-evaluate their cost estimates.¹⁹ The private values framework does allow for the inclusion of common, deterministic reservation costs components. For example, characteristics related to the operating costs like the market distance are clearly known to the airlines operating in the market and they affect them in the same way.²⁰ Many opportunity cost determinants are also common and deterministic, for example the time left before the take-off date but also how much capacity remains on own and competitor planes. We base this assumption on the observation that airlines have access to ample amounts of historical and real time information concerning their own and their competitor's sales and operations, and they use very sophisticated techniques to assess this information.²¹

¹⁹This is different from a common cost framework, where the cost draw for each player carries information for the final, interdependent cost.

²⁰Such costs, for example the cost of distance, are the same to all carriers to be consistent with the symmetric framework.

²¹For example, Sabre Airline Solutions, a major airline software developer and consultancy, states that their fare management tools have the following capabilities: "Our pricing solution equips airlines with the ability to manage fares in a competitive and timely manner. The systems powerful data query tools help pricing analysts examine relative market data including competitors changes to make the right decisions at the right time." Source: <http://www.sabreairlinesolutions.com/>.

Hence, at any given moment the reservation cost of an airline is determined by some common and observed components and a random component, which is private.²²

Given the equilibrium strategy of the game, the observed winning price P is a function of the reservation cost C - it is the highest out of two realisations of the random variable C . A ranked draw of a random variable is itself a random variable, known as an order statistic.²³ The distribution $F_P(p)$ can be derived from the underlying reservation cost distribution $F_C(c)$ as follows:

$$F_P(p) = Pr[P \leq p] = \prod_{i=1}^2 Pr[C_i \leq p] = F_C^2(p).^{24} \quad (1)$$

Then, we find the probability density $f_C(c)$ by taking its derivative:

$$f_P(p) = 2F_C(p)f_C(p). \quad (2)$$

The distribution of prices and underlying costs are linked as above under the assumptions of our model. In the next sub-section, we describe how we will use this relationship to estimate the model using the method of maximum likelihood, and how we will obtain the effect of a number of cost covariates on both the cost and the price distribution.

Specification

We specify a parametric model of the type $F_C(c) = F_C(c; \theta)$, where θ is the vector of parameters to be estimated. In our baseline specification, we assume the distribution

²²Flight level characteristics and time of purchase, which are not observed in our data set, represent unobserved heterogeneity that cannot be identified from the distribution of the random component (see Athey and Haile (2002)). However, there is no reason to believe they are correlated with the presence of an alliance on a market, our covariate of interest, and their omission should not bias our estimates.

²³Order statistics and their distribution are important components of auction models, where the winner is chosen based on a ranking of the bids and the bids are monotonic functions of the underlying random costs or valuations. For more information and derivations, see Paarsch and Hong (2006) and particularly Appendices 1 and 2.

²⁴Consider two independent and identically distributed random variables, C_1 and C_2 with probability cumulative distribution $F_C(c)$ and probability density $f_C(c)$. Let P be the largest of the two: $P = \max(C_1, C_2)$. P has a cumulative distribution $F_P(p) = Pr(P \leq p)$. The probability that the random variable P is less than some p requires the event $(C_1 \leq p) \cap (C_2 \leq p)$. When each C_i is independent, $(C_i \leq p)$ are also independent events and their joint probability is the multiplication of the individual probabilities. Thus, we are able to construct $F_P(p) = Pr[C_1 \leq p]Pr[C_2 \leq p] = F_C^2(p)$.

of the costs C is log-normal, or the logged costs $\ln(C)$ are distributed as $N(\mu, \sigma)$, with $\theta = (\mu, \sigma)$.²⁵ The maximum likelihood estimation method allows us to estimate the parameters of the distribution θ by maximising the joint probability of the sample with respect to the parameters. The likelihood of observing a sample of prices of size T is the joint probability of the individual price observations, $f_P(p)$. We assume the sample is coming from a set of independent auction events, the joint probability is simply the multiplied individual probabilities.

$$L = \prod_{t=1}^T f_P(p_t; \theta) = \prod_{t=1}^T 2F_C(p_t; \theta) f_C(p_t; \theta). \quad (3)$$

The log-likelihood is used in the maximisation to transform the multiplicative relationships into additive ones and thus make maximisation easier:

$$l = N \ln(2) + \sum_{t=1}^T \ln \left(\Phi \left(\frac{\ln p_t - \mu}{\sigma} \right) \right) + \sum_{t=1}^T \ln \left(\frac{1}{\sigma p_t} \phi \left(\frac{\ln p_t - \mu}{\sigma} \right) \right). \quad (4)$$

The set of auctions in the data are not identical - different markets have different distance and population demographics, for example, that would naturally affect the reservation costs distribution. To account for these deterministic differences in the estimation procedure, we specify the parameters of the reservation cost $F_C(c, \theta)$ distribution as linear functions of covariates:

$$\mu = \alpha W \quad \text{and} \quad \sigma = \beta W, \quad (5)$$

where α and β are the coefficient vectors for the covariates W . They represent the marginal effects of covariates on the mean and standard deviation, respectively. It is important to note that the structural model will estimate the parameters θ of the underlying reservation cost distribution $F_C(c, \theta)$, but we are also interested how these parameters affect the distribution of transaction prices actually paid by consumers. Let us denote the parameters of the price distribution $\gamma = (m, s)$ and the parametrised

²⁵In an alternative specification, we assume the reservation costs to be distributed according to the Beta distribution. The Beta is defined on a compact set and bounded, which is appealing for modelling reservation costs as they should be positive. Additionally, the Beta can accommodate skewness and allows the dependent variable to be heteroskedastic in the sense that the dispersion may vary for different costs of the mean. The results obtained using the Beta are similar to those from the estimation with the Normal, and are available upon request.

distribution $F_P(p, \gamma)$. Then, using theoretical derivations from the field of statistics²⁶, we can express the parameters of the price distribution γ in terms of the parameters of the reservation cost distribution θ . The mean of the price m is a combination of the the mean μ and standard deviation σ of the reservation cost distribution. The price standard deviation s is simply the scaled σ of the reservation cost. The marginal effects of the covariates W on m and s can then be simply calculated using the marginal effects α and β of the reservation cost distribution.

$$m = \mu + \frac{\sigma}{\sqrt{\pi}} = \left(\alpha + \frac{\beta}{\sqrt{\pi}} \right) W = aW. \quad (6)$$

$$s = \sigma \sqrt{\frac{\pi - 1}{\pi}} = \beta \sqrt{\frac{\pi - 1}{\pi}} W = bW. \quad (7)$$

Note that since $\sqrt{\frac{\pi-1}{\pi}} < 1$, the marginal effect of the covariates W on s (the b) is lower than that on σ (the β). The distribution of prices becomes more compact, because the prices are reservation costs that are selected away from small values. What happens to the mean depends on the signs and relative magnitude of α and β . If they have the same sign, then the marginal effect on the mean a is augmented. If they have different signs, the overall effect depends on their relative magnitude and significance. For example, it is possible that both α and β are significant but of opposite sign and a is insignificant.

Next, we define a *market* as a directional (origin to destination) city pair as in Ito and Lee (2007). Markets are heterogeneous in terms of characteristics W , which can be defined at the market, origin or destination level. Let the market be indexed by jk , where j represents the city of origin and k represents the destination city. Let X_{jk} be covariates relevant to the market, Y_j relevant to the origin and Z_k to the destination. Then the mean μ_{jk} on a market jk and standard deviation σ_{jk} can be written the following way:

$$\mu_{jk} = \alpha_A Alliance_{jk} + \alpha_X X_{jk} + \alpha_Y Y_j + \alpha_Z Z_k \quad (8)$$

$$\sigma_{jk} = \beta_A Alliance_{jk} + \beta_X X_{jk} + \beta_Y Y_j + \beta_Z Z_k \quad (9)$$

²⁶Our source is Nadarajah and Kotz (2008), although the exact forms of the different moments of the distribution of order statistics have been known for a while and are available in many good reference books such as David and Nagaraja (1970).

The alliance effect is at the market level - the market can either be an alliance market ($Alliance_{jk} = 1$), or a non-alliance market ($Alliance_{jk} = 0$). Our analysis focuses on direct markets and direct products (non-stop flights offered on these markets). The literature on alliances has approached the estimation of their effect in two different ways. Ito and Lee (2007) and Urdanoz and Sampaio (2012) look at the price of different types of alliance products on the same market. Gayle (2008) and Bamberger, Carlton and Neumann (2004), on the other hand, look at the effect of an alliance product being offered in a given market. Our methodology similarly estimates how the presence of an alliance agreement between the carriers in a market affects the prices in that market, but we compare markets cross-sectionally rather than before and after the agreement. We prefer comparing cross-sectionally because the before-after comparison is only possible around the date of the alliance agreement, and this means we should focus on an earlier period that may no longer be relevant given the new market conditions.

An important assumption in our approach is that markets where an alliance is present must be comparable to those without. In other words, there is no selection on markets based on unobservables such that, for example, the strategic choice of allied carriers is to enter in markets where price discrimination is more profitable and this would cause the alliance coefficient to be biased. To avoid this type of problem, Brueckner (2003) uses a model with entry. Another more direct approach that is used by Brueckner and Whalen (2000) and Ito and Lee (2007) is to introduce fixed effects, which we will consider in our sensitivity analysis.

3 Data and variables

The standard airline industry data base DB1B is published by the US Department of Transportation and represents a 10 percent sample of all tickets sold within the US domestic market. We use data for the third quarter (Q3) of 2008 and we work with all direct service duopoly markets. Direct ticket sales are a significant portion of all sales - they represent 36 percent of all ticket sales in 2008. Duopoly markets, in turn, represent 33 percent of competitive (non-monopoly) sales.

We focus on the analysis of direct products for the following reason. We model competition between carriers offering identical products to the consumer. Connecting flights have characteristics (connecting airport, time of layover, etc) which differentiate them in a number of dimensions and accounting for that is not feasible in our framework. Still, we believe direct markets are representative and increasingly relevant as demonstrated by the high demand sensitivity to connecting flights and provision of direct service in the recent years.

Furthermore, our analysis concerns only duopoly markets where major legacy carriers operate. To be consistent with the symmetry assumption, we do not consider LCCs and markets where they operate. Indeed LCC's have rather different cost structures than the legacy carriers in addition to the fact that they do not enter alliances.²⁷ Operating in 2008 time period are American (AA), Alaska (AS), Continental (CO), Delta (DL), Midwest (YX), Northwest (NW), United (UA), and US Air (US). Sales in markets where these eight compete against each other represent 30 percent of all duopoly sales for Q3 2008. The representation of each airline in the sample is found in Table 1 below. Additional information on the cleaning of the data set is presented in Appendix A.

Table 1: Ticket sales in selected direct duopoly markets.

Carrier	Freq.	Percent
American Airlines (AA)	13,669	31.16
Alaska (AS)	4,001	9.12
Continental (CO)	1,366	3.11
Delta (DL)	3,948	9.00
Northwest (NW)	5,576	12.71
United (UA)	7,685	17.52
US Air (US)	7,166	16.33
Midwest (YX)	458	1.04
Total	43,869	100

Alliance presence

The original and main objective of an alliance is to sell tickets on another carrier's network, the so-called code-sharing which allows a given carrier to expand the range of its services. The level of cooperation between alliance carriers can go further than simple code-sharing by the sharing of frequent flyer programs, coordinating schedules, and sharing of airport facilities (lounges, operating and maintenance facilities, even staff). In our definition of an alliance, we follow Ito and Lee(2007): "carriers are alliance partners if passengers on one of the alliance carriers can earn elite-qualifying frequent flyer miles on flights marketed or operated by the other alliance partner and vice versa". The complete list of alliances, and the number of markets where they operate in 2008, is provided in Table 2.

²⁷Regional-legacy carrier agreements are not considered to be alliances but rather an integrated service. As is standard in the literature, we re-code tickets sold by regional carriers as the legacy partner. More details on this can be found in Appendix A.

Table 2: Domestic alliances in 2008

	Markets	Tickets
Alaska & American	2	1,177
Alaska & Continental	4	489
Alaska & Delta	2	154
Continental & Delta	3	358
Continental & Northwest	4	1,088
Delta & Northwest	8	1,125
Northwest & Midwest	2	346
United & US Airways	26	3,392
All alliance	51	8,129
No alliance	177	35,740
Total	228	43,869

Our *Alliance* covariate is an indicator variable, which is equal to 1 in markets where both carriers are in an alliance (e. g., markets where CO and NW both operate). The covariate is 0 when the carriers operating on a market are not in an alliance together (e. g., markets where CO and US operate). In this sense, we are not concerned with the exact type of cooperation products or level of coordination that takes place in alliance markets. The alliance variable can affect both the mean and the standard deviation of the reservation cost, and therefore prices.

Ex ante, it is not obvious how the alliance presence will affect the level and variability of the reservation cost on a market. On the one hand, alliances are allowed to share certain operating costs such as personnel and airport facilities, which could reduce operating costs and make them more stable (lower their variability). On the other hand, the ability to sell tickets on a competitor's flights can make price discrimination more profitable, hence it will alter the opportunity cost of a ticket and how this opportunity cost varies over time before take-off. The alliance covariate thus measures an overall effect of the alliance on the reservation cost's mean and variance.

Covariates

Covariates W are the characteristics of the market, the origin and the destination that shift the distribution mean and standard deviation deterministically. The covariates are available in the DB1B or calculated using information contained in the data set.

To control for demographic characteristics, we added information on the population and income at the origin and destination of the markets obtained from the US Census Bureau.

Distance is a market-level covariate, measured in number of miles between the origin and destination cities. The distance between two cities relates to the level of operating costs, as longer distances require more fuel to reach. The effect on the standard deviation of the reservation cost, however, is undefined ex-ante. Consider the substitutability between air and land travel transport modes, which is quite high for short-distance trips. Demand for last-minute trips, however, might be quite inelastic - business travellers might find flights more convenient than driving when there is little time to plan ahead. A steeper price path might be optimal because early ticket sales are highly substitutable with land travel (and need to be prices low) while later ticket sales are rather inelastic. This kind of effect would show up in the data as a negative effect of distance on price variance.

The *Market volume* variable was constructed as a measure of traffic or scale on the market: it is the number of ticket sales observed in that market. A higher volume is correlated with higher number of flights and/or larger plane size, which would allow the airlines to have lower average costs. Alternatively, each additional flight is a substitute for the other flights the airline offers, and a larger number of flights may lead to smaller difference between their prices, hence lower standard deviation.²⁸

We construct the variables measuring the scale of operations - *Origin volume* and *Destination volume* - as the total number of ticket sales at the origin and destination of the market. We also control for the location of the origin and destination in the transport network, or how central is the market. To quantify this centrality we use *Origin connections* and *Destination connections*. Origin connections counts the number of cities directly accessible from the origin, while destination connections counts the number of cities from which one can fly to the destination. The centrality in a network affects operating costs through the alternative use of resources (planes, personnel) and through the opportunity cost changing with the stochastic demand in adjacent markets.

Our demographic measures - *Population* and *Income* - are measured at the origin and destination cities. The effect of income is related to the opportunity cost of providing the service. Higher income cities have both richer leisure travellers who do not need to plan too much in advance and more business travellers who book tickets in the last days before departure. The airline has a higher opportunity cost of a ticket sale, and possibly a steeper price path is more profitable. Population, on the other hand, is another measure of scale similar to market volume and it would be associated with

²⁸We prefer to use volume and connections because "scale" and "scope" applies to operating costs, and here we could also have opportunity cost as contributing to the reservation cost.

lower reservation costs as larger-scale operations are more efficient.

Summary statistics

Below we present the summary statistics for the sub-samples for which *Alliance* = 1 and *Alliance* = 0. The units of the variables are as follows. Price (airfare) and income are measured in US dollars, distance is in miles, population is in number of people, volumes are in number of tickets, and connections are number of origin/destination cities. We notice that prices are higher for the alliance sub-sample, but so is distance which naturally increases costs. Price has a higher standard deviation in the alliance subsample.

Table 3: Summary statistics

Variable	Alliance=0		Alliance=1	
	Obs: Mean	35,740 Std.	Obs: Mean	8,129 Std.
Price	425	234	547	322
Distance	1,773	992	2,560	1486
Origin population	5,329,810	5,262,221	3,362,796	1,925,202
Destination population	6,121,631	4,360,113	4,239,602	3,629,085
Origin income	53,215	8,273	59,450	6,922
Destination income	52,019	4,938	51,671	6,821
Market volume	455	399	285	172
Origin volume	19,692	13,621	16050	9,116
Destination volume	16,929	11,783	14076	7,964
Origin connecions	74	36	67	27
Destination connections	70	34	64	26

Although the variables are presented here in levels, we consider their logarithmic values as it is common in the airline literature (see, for example, Ito and Lee (2007)). Together with the assumption of log-Normality of the reservation cost, this allows us to interpret the effect of the coefficients on the mean μ (the α 's) and standard deviation σ (the β 's) as elasticities. Hence, each estimated coefficient represents the percentage change in μ or σ given one percent change in the relevant covariate. The alliance coefficient, on the other hand, is a semi-elasticity and is interpreted slightly different: to get the percentage chance of μ or σ when *Alliance* = 1, we multiply the estimated coefficient by 100. This interpretation is also relevant for m and s , the mean and standard deviation of prices. The effects of the covariates on m and s , a and b respectively, will be constructed using

Equations (6) and (7).

As mentioned previously, there could be a potential selection bias on unobservables - alliance partners entering into markets based on unobserved (to the econometrician) factors. To evaluate the potential impact of this selection, we perform sensitivity tests of our results by also running the regressions with fixed effects at the origin and destination levels.

4 Empirical results

The MLE methodology for auction data by Paarsch (1997) allows us to use the observed distribution of fares to infer the parametric effect of different covariates on the underlying distribution of reservation costs, itself comprised of an operative cost and an opportunity cost. Through the distributional relationship between the reservation cost and the transaction price, we are able to derive how these covariates affect the level and standard deviation of the observed fares.

Parameter estimates

We have assumed a log-Normal distribution for reservation costs, with mean $\mu = \alpha W$ and standard deviation $\sigma = \beta W$. Our baseline model includes all discussed covariates as mean and standard deviation shifters. Table 5 has the results of the main estimation.²⁹ We present the estimated effects on the parameters of the reservation value (our MLE results), and the effects on the parameters of the price that are constructed using Equations (6) and (7).

We start the discussion of the results by looking at our main covariate of interest, the *Alliance* dummy. The results show a significant positive relationship between the existence of alliance on a given market and both the mean and standard deviation of the reservation cost. It is unlikely that the positive mean shift of 9.8 percent would be related to higher operating costs for alliance members, when in fact alliance members are able to share facilities and coordinate schedules in a way that should reduce operating expenses. Rather, this positive coefficient could be due to the alliance improving the ability of an airline to price discriminate by allowing the selling of the partner's tickets. Moreover, a higher standard deviation of 5.2 percent for alliance markets supports the

²⁹As customary, different significance levels are denoted by an asterisk. We have the following: *** for significance at 1 percent, ** for significance at 5 percent, and * for significance at 10 percent.

hypothesis that alliance members can use the ability to sell each other's tickets to segment demand more effectively. As a result, we have prices that are 12.7 percent higher and 4.3 percent more spread relative to non-alliance markets.

The estimated alliance coefficient is comparable to other results in the domestic airline literature. For example, Armantier and Richard (2006) find an increase of 10.7 percent for ticket fares in markets through which CO and NW code-share using a before-after approach and data for 1998-2001. Gayle (2008) also finds a positive effect of an alliance operating in a given market using a similar before-after approach for the DL/CO/NW alliance and data for 2002-2003, but of only 1.8 percent. In contrast, Bamberger, Carleton and Neumann (2004) find that the CO/HP and NW/Alaska alliances formed in the mid-90's are associated with 7.5 and 3.9 lower fares, respectively, in connecting markets. One major difference between our results and previous analysis is that we work with all eight alliances operating in 2008.

The effect of *Distance* on reservation costs is positive (0.193 percent), which is not surprising since distance is associated with higher operating costs. A larger distance between the origin and destination cities is associated with a lower price variation (-0.047 percent). This is possibly due to the substitution of shorter distance trips with land transportation. For example, on shorter distances, the airlines may offer a combination of low priced tickets which would be competitive with respect to travelling by car, and high priced last-minute tickets for urgent business travels. Naturally, this causes higher variation of prices for flights of shorter distance. The overall effect of distance on the price mean is lower (0.166 percent) compared to that on the reservation value, because the effects on μ and σ go in the opposite direction.

Moving on to the demographic variables, *Population* has a significant negative effect on the mean reservation value (-0.057 percent for population at the origin and -0.036 percent at the destination) but does not affect the standard deviation. A higher population may lead to lower costs through higher demand and potentially more efficient levels of operation. Given that the effect of population on the standard deviation is not significant, the effect of population on μ and m almost identical.

Origin and destination *Income*, on the other hand, have a significant and positive effect on the mean reservation cost (0.145 and 0.292 percent, respectively) and the standard deviation (0.207 and 0.072 percent, respectively). Higher income at the origin can be associated with richer leisure travellers, that would increase the opportunity cost of a ticket. Rich leisure travellers may also be able to afford to not plan their trips too long in advance, causing the opportunity cost of sales to increase over time and raising the observed dispersion. Higher income (both at the origin and the destination) is associated with more business travel, that would also increase the opportunity cost of a ticket. Business travel could make the opportunity cost more variable over the course

of selling tickets as it is uncertain and usually concentrated in the last moment before take-off. This causes capacity to be more filled in a more random manner, causing the opportunity cost of a ticket to have more variation over time. The mean elasticity of the price is higher than that of the reservation cost (0.261 for the origin and 0.331 for the destination).

We have three different measures of *Volume* - at the market level, and the origin and at the destination. All three lower the standard deviation of the reservation cost (by 0.018 percent, 0.027 percent, and 0.023 percent, respectively). A higher scale of operation is usually associated with more optimised costs and more stable costs. An alternative and complementary explanation for this effect could be that more flights on a market would decrease ability to price discriminate because of the substitutability between them, hence leading to lower variability in the opportunity cost of the sale. The negative effect on σ translates to a negative effect on s . Additionally the destination volume and market volume achieve a negative effect on price levels of 0.041 and 0.016 percent respectively.

Lastly, origin and destination *Connections* measure the centrality of the origin or destination to the travel network. Both are associated with higher reservation cost (by 0.05 percent for the origin, and 0.07 percent for the destination), and higher standard deviation (0.051 and 0.037 percent, respectively). A market with more central origin and destination airports may mean that the cost to the airline to service that market must be weighed against using its resources to service other adjacent markets, affecting the operating costs. Alternatively, there may be dynamic demand spill-overs where more connections increase the variability of the stochastic demand, and make the opportunity cost more volatile.

Sensitivity analysis

To test the robustness of our results, we estimate the model with fixed effects at the origin and destination levels. Naturally, the covariates specific to the origin and destination are collinear with the fixed effects and they are excluded from this estimation. The covariates which remain are *Distance*, *Marketvolume* and *Alliance*. The results of the fixed effects estimation, displayed in Table 6, demonstrate that the significance of the alliance indicator remains when controlling for unobserved factors. An *Alliance* is associated with a 11 percent increase in reservation cost mean, and a 4.8 increase in reservation cost standard deviation, both effects being comparable to the magnitudes we have in the non-fixed effects regression. The corresponding semi-elasticities to alliance presence for the price are 13.6 percent for the mean and 3.9 percent for the standard deviation. The effect of *Distance* is also consistent with what was estimated

Table 4: Estimation results with covariates

	Reservation cost		Price	
	μ	σ	m	s
	Coef./Std.	Coef./Std.	Coef./Std.	Coef./Std.
Alliance	0.098*** (0.008)	0.052*** (0.006)	0.127*** (0.007)	0.043*** (0.005)
Distance	0.193*** (0.006)	-0.047*** (0.005)	0.166*** (0.006)	-0.039*** (0.004)
Origin population	-0.057*** (0.004)	0.004 (0.004)	-0.055*** (0.004)	0.004 (0.003)
Destination population	-0.036*** (0.004)	-0.003 (0.003)	-0.037*** (0.004)	-0.003 (0.003)
Origin income	0.145*** (0.023)	0.207*** (0.020)	0.261*** (0.023)	0.171*** (0.016)
Destination income	0.290*** (0.030)	0.072** (0.024)	0.331*** (0.028)	0.06*** (0.020)
Market volume	-0.031*** (0.005)	-0.018*** (0.004)	-0.041*** (0.005)	-0.015*** (0.004)
Origin volume	0.020* (0.009)	-0.027*** (0.008)	0.005 (0.009)	-0.023*** (0.006)
Destination volume	-0.003 (0.009)	-0.023** (0.008)	-0.016* (0.009)	-0.019*** (0.006)
Origin connections	0.050*** (0.010)	0.051*** (0.008)	0.078*** (0.010)	0.042*** (0.007)
Destination connections	0.070*** (0.010)	0.037*** (0.008)	0.091*** (0.010)	0.03*** (0.007)
Constant	0.373 (0.431)	-1.892*** (0.364)	-0.695* (0.415)	-1.562*** (0.301)
Likelihood	-109,780			

previously. The *Marketvolume* coefficient, on the other hand, changes sign and this indicates that there is bias in the main regression that affects this particular variable.

Table 5: Estimation results with fixed effects

	Reservation cost		Price	
	μ	σ	m	s
	Coef./Std.	Coef./Std.	Coef./Std.	Coef./Std.
Alliance	0.110*** (0.014)	0.048*** (0.012)	0.136*** (0.013)	0.039*** (0.010)
Distance	0.139*** (0.010)	-0.020* (0.008)	0.127*** (0.009)	-0.016** (0.007)
Market volume	0.055*** (0.008)	-0.028*** (0.007)	0.039*** (0.008)	-0.023*** (0.006)
Constant	4.530*** (0.097)	1.067*** (0.083)	5.132*** (0.094)	0.881*** (0.068)
Likelihood	-106,756			

We tested the sensitivity of our results by also looking at data for other years, notably for 2007 and 2009. A potential problem with 2009 is that many of the major alliances (UA-US and CO-NW-DL) ended by the end of that year as Continental merged with United in 2010 and Northwest merged with Delta also in 2010. It is not known how these two events, known in advance, may have affected the carriers. The estimated reservation cost coefficients for years 2007, 2008 and 2009 are available in Appendix B. The year 2009, just before the mergers and possibly at the height of the economic crisis, is somewhat inconsistent in the significance and signs of a number of coefficients, while years 2007 and 2008 are mostly consistent. The *Alliance* effect on μ is significant and of similar magnitude in all estimations: 10 percent in 2007, 9.8 in 2008, and 13.4 in 2009. The effect of alliance on σ is actually negative (-1.6 percent) in 2007, indicating a lower variance on alliance markets. In 2008 and 2009, it is significant, positive and of similar magnitude (5.2 and 4.1, respectively). The effect on the price mean ranges from 9.1 (for 2007) to 15.7 (for 2009), while the effect on price standard deviation - from -1.3 percent (2007) to 4.3 percent (2009). Overall, the sensitivity analyses with respect to the chosen time period produce results that are consistent with what we have obtained in the main regression.

5 Conclusion

Our objective in this paper has been the following: to propose a new framework in regard to how airlines set prices for their tickets, and to use it to estimate the effect of domestic alliances in the US market. We are motivated by the increased competitiveness of the airline industry that has resulted in higher price sensitivity and product homogenisation, which are well documented by academics who study this industry. Our model is thus based on price competition between the airlines, also known as Bertrand competition. The observational equivalence between Bertrand competition and the reverse English auction allows us to employ an MLE method from the auction literature. We fit our model on a subsample of flight-level price observations extracted from the DB1B data set, and we estimate the effect of covariates on both the mean and the variance of prices. Our result from the main regression indicates that the presence of an alliance in the market is associated with prices higher by 12.7 percent, and with a 4.3 percent increase in the price standard deviation. We believe this result indicates an improved ability to price discriminate when alliance partners.

For simplicity, our analysis assumed an auction environment of private values and symmetric players. A richer model should take into account of the difference in market presence of each competitor - their relative position in the market, or the origin and destination of the market - and how this affects the operating costs, and the incentives and ability to price discriminate. Incorporating asymmetry in reservation cost distribution is a direction that we would like to explore in the future, as from the auction theory literature we know that asymmetric private values models are identified as long as the transaction price and the identity of the winner are observed. Generalising the private values framework by assuming either affiliated values or common values can be problematic, especially if one is to use the DB1B data set. Indeed, identification and testing in such more general environments is not straightforward and can be achieved only by employing more detailed data (Athey and Haile (2002)).

Our work is not without limitations, and this provides scope for future work on modelling competition in the airline industry. For example, we are not able to identify whether the effect of the different covariates on the reservation cost distribution comes from the operating cost, the opportunity cost or both. To do that, one would need to construct a more complex structural model of competition that takes into account dynamic pricing under demand uncertainty. As mentioned earlier, this type of models are of high complexity, and have typically been analysed only in the monopoly case. Moreover, estimation would require a more detailed data set than the DB1B with information on flight characteristics and time of purchase, as well as on capacity levels and sales at each point in time. We believe the framework we propose - an open competitive environment where consumers make their choice based on the price and competitors

differ in their private costs - would be an appropriate starting point for a more complex model.

References

- Adelman, D. (2007). Dynamic bid prices in revenue management. *Operations Research*, 55(4), 647-661.
- Armantier, O., and Richard, O. (2003). Exchanges of cost information in the airline industry. *RAND Journal of Economics*, 461-477.
- Armantier, O., and Richard, O. (2006). Evidence on pricing from the Continental Airlines and Northwest Airlines code-share agreement. *Advances in airline economics*, 1, 91-108.
- Armantier, O., and Richard, O. (2008). Domestic airline alliances and consumer welfare. *The RAND Journal of Economics*, 39(3), 875-904.
- Athey, S., and Haile, P. A. (2002). Identification of standard auction models. *Econometrica*, 70(6), 2107-2140.
- Bamberger, G., D. Carlton, and L. Neumann. (2004). An Empirical Investigation of the Competitive Effects of Domestic Airline Alliances. *Journal of Law and Economics*, 47, 195-222.
- Berry, S. T. (1992). Estimation of a Model of Entry in the Airline Industry. *Econometrica: Journal of the Econometric Society*, 889-917.
- Berry, S., and Jia, P. (2010). Tracing the woes: An empirical analysis of the airline industry. *American Economic Journal: Microeconomics*.
- Borenstein, S., and Rose, N. L. (1994). Competition and Price Dispersion in the US Airline Industry. *Journal of Political Economy*, 102(4), 653-683. Chicago
- Borenstein, S., and Rose, N. L. (2013). How airline markets work or do they? Regulatory reform in the airline industry. In *Economic Regulation and Its Reform: What Have We Learned?*. *University of Chicago Press*.
- Brueckner, J. K., Lee, D., and Singer, E. S. (2012). Airline competition and domestic US airfares: A comprehensive reappraisal. *Economics of Transportation*.
- David, H. A., and Nagaraja, H. N. (1970). *Order statistics*. John Wiley and Sons, Inc.

- Escobari, D. (2012). Dynamic Pricing, Advance Sales and Aggregate Demand Learning in Airlines. *The Journal of Industrial Economics*, 60(4), 697-724.
- Evans, W. N., and Kessides, I. N. (1993). Localized market power in the US airline industry. *The Review of Economics and Statistics*, 66-75.
- Evans, W. N., and Kessides, I. N. (1994). Living by the" Golden Rule": Multimarket contact in the US airline industry. *The Quarterly Journal of Economics*, 341-366.
- Gayle, P. G. (2008). An Empirical Analysis of the Competitive Effects of the Delta /Continental /Northwest Code-Share Alliance. *Journal of Law and Economics*, 51(4), 743-766.
- Ito, H., and Lee, D. (2007). Domestic code sharing, alliances, and airfares in the US airline industry. *Journal of Law and Economics*, 50(2), 355-380.
- Lazarev, J. (2012). The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in US monopoly markets. *Working paper*.
- Nadarajah, S., and Kotz, S. (2008). Exact distribution of the max/min of two gaussian random variables. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 16(2), 210-212.
- Paarsch, H. J. (1997). Deriving an estimate of the optimal reserve price: an application to British Columbian timber sales. *Journal of Econometrics*, 78(1), 333-357.
- Sampaio, R., and Miguel Urdanoz. (2011). Airlines' Cooperation in the US Domestic Market: Measuring the Evolution of Price Gaps. *Working Paper*.
- Talluri, K., and Van Ryzin, G. (1998). An analysis of bid-price controls for network revenue management. *Management Science*, 44(11-part-1), 1577-1593.
- Vives, X. (2002). Private information, strategic behavior, and efficiency in Cournot markets. *RAND Journal of Economics*, 361-376.
- Williams, K. R. (2013). Dynamic Airline Pricing and Seat Availability. *mimeo*.

Appendix

A: Data cleaning

Regional "feeder" or "commuter" carriers are recoded as their major carrier partner. See Table 4 in Sampaio and Urdanoz (2012). Strong cooperation between such commuter carriers and the major carriers started in the 1980's and in some cases this has resulted in vertical integration, for example American Eagle (AM) and American Airlines (AA).

Carriers with less than 15 passengers were deleted, since these probably reflect coding errors. We also removed tickets with cost less than 50 USD and more than 3000 USD. The majority of these happen to be tickets at 0 USD, representing frequent flyer purchases. We also focus on markets with more than 9 passengers per quarter, as that is equivalent to one passenger per day given that the sample represents 10 percent of ticket sales (Sampaio and Urdanoz (2012)).

One other modification of the data set consists of grouping airports in the same metropolitan area. The six groups of airports are: Dallas-Fort Worth International and Love Field in Dallas, TX; Baltimore/ Washington International, Dulles, and National in Washington, DC; Midway and OHare in Chicago, IL; Kennedy, LaGuardia, and Newark in New York, NY; Los Angeles, Burbank, and Long Beach in Los Angeles, CA; San Francisco, Oakland, and San Jose in San Francisco, CA. For example, Chicago Midway and Chicago O'Hare International will represent the same market. Again, this is a standard treatment in the literature that allows more properly to look at metropolitan areas with multiple airports (Berry and Jia (2010, Sampaio and Urdanoz (2012))).

Lastly, following Evan and Kessides (1993, 1994), we count carriers as operating in a given market if their sales represent at least 1 percent of observations in the data, equivalently 1 percent of total sales.

B: Additional results

Table 6: Estimate results for reservation value coefficients 2007, 2008 and 2009

	2007			2008			2009		
	μ	Coef./Std.	σ	μ	Coef./Std.	σ	μ	Coef./Std.	σ
Alliance	0.100*** (0.008)	-0.016** (0.006)	0.030*** (0.004)	0.098*** (0.008)	0.052*** (0.006)	0.004 (0.004)	0.134*** (0.007)	0.041*** (0.006)	0.000 (0.003)
Distance	0.174*** (0.006)	-0.020*** (0.005)	-0.015*** (0.003)	0.193*** (0.006)	-0.047*** (0.005)	-0.036*** (0.004)	0.134*** (0.006)	0.008 (0.005)	0.008 (0.000)
Origin population	-0.062*** (0.005)	0.030*** (0.004)	0.030*** (0.004)	-0.057*** (0.004)	0.004 (0.004)	0.004 (0.004)	-0.040*** (0.004)	0.000 (0.003)	0.000 (0.003)
Destination population	-0.035*** (0.004)	-0.015*** (0.003)	-0.015*** (0.003)	-0.036*** (0.004)	-0.003 (0.003)	-0.003 (0.003)	0.010** (0.003)	0.005* (0.003)	0.005* (0.003)
Origin income	0.150*** (0.019)	0.186*** (0.016)	0.186*** (0.016)	0.145*** (0.023)	0.207*** (0.020)	0.207*** (0.020)	-0.007 (0.022)	0.189*** (0.019)	0.189*** (0.019)
Destination income	0.242*** (0.032)	0.218*** (0.027)	0.218*** (0.027)	0.290*** (0.030)	0.072** (0.024)	0.072** (0.024)	-0.037 (0.036)	0.161*** (0.028)	0.161*** (0.028)
Market volume	-0.029*** (0.005)	-0.029*** (0.005)	-0.029*** (0.005)	-0.031*** (0.005)	-0.018*** (0.004)	-0.018*** (0.004)	0.094*** (0.006)	-0.020*** (0.005)	-0.020*** (0.005)
Origin volume	0.004 (0.009)	-0.009 (0.007)	-0.009 (0.007)	0.020* (0.009)	-0.027*** (0.008)	-0.027*** (0.008)	-0.004 (0.008)	-0.022** (0.007)	-0.022** (0.007)
Dest volume	-0.010 (0.009)	-0.002 (0.007)	-0.002 (0.007)	-0.003 (0.009)	-0.023** (0.008)	-0.023** (0.008)	-0.011 (0.009)	-0.036*** (0.007)	-0.036*** (0.007)
Origin connections	0.074*** (0.010)	0.051*** (0.008)	0.051*** (0.008)	0.050*** (0.010)	0.051*** (0.008)	0.051*** (0.008)	0.004 (0.011)	0.032*** (0.009)	0.032*** (0.009)
Destination connections	0.074*** (0.010)	0.032*** (0.007)	0.032*** (0.007)	0.070*** (0.010)	0.037*** (0.008)	0.037*** (0.008)	0.012 (0.011)	0.043*** (0.009)	0.043*** (0.009)
Constant	1.092** (0.423)	-3.922*** (0.368)	-3.922*** (0.368)	0.373 (0.431)	-1.892*** (0.364)	-1.892*** (0.364)	4.952*** (0.476)	-3.023*** (0.372)	-3.023*** (0.372)
Likelihood	- 108,118			- 109,780			- 94,749		