Transportation Research Part A 66 (2014) 196–212

Contents lists available at ScienceDirect

Transportation Research Part A

journal homepage: www.elsevier.com/locate/tra

Estimating flight-level price elasticities using online airline data: A first step toward integrating pricing, demand, and revenue optimization

Stacey Mumbower [a], Laurie A. Garrow [a,*], Matthew J. Higgins [b,c]

[a] Georgia Institute of Technology, School of Civil and Environmental Engineering, 790 Atlantic Drive, Atlanta, GA 30332-0355, United States
[b] Georgia Institute of Technology, Ernest Scheller Jr. College of Business, 800 West Peachtree NW, Atlanta, GA 30308, United States
[c] National Bureau of Economic Research, 1050 Massachusetts Ave., Cambridge, MA 02138, United States

ARTICLE INFO

Article history:
Received 17 August 2013
Received in revised form 17 February 2014
Accepted 9 May 2014
Available online 13 June 2014

Keywords:
Air travel demand
Price elasticity
Price endogeneity

ABSTRACT

We estimate flight-level price elasticities using a database of online prices and seat map displays. In contrast to market-level and route-level elasticities reported in the literature, flight-level elasticities can forecast responses in demand due to day-to-day price fluctuations. Knowing how elasticities vary by flight and booking characteristics and in response to competitors' pricing actions allows airlines to design better promotions. It also allows policy makers the ability to evaluate the impacts of proposed tax increases or time-of-day congestion pricing policies. Our elasticity results show how airlines can design optimal promotions by considering not only which departure dates should be targeted, but also which days of the week customers should be allowed to purchase. Additionally, we show how elasticities can be used by carriers to strategically match a subset of their competitors' sale fares. Methodologically, we use an approach that corrects for price endogeneity; failure to do so results in biased estimates and incorrect pricing recommendations. Using an instrumental variable approach to address this problem we find a set of valid instruments that can be used in future studies of air travel demand. We conclude by describing how our approach contributes to the literature, by offering an approach to estimate flight-level demand elasticities that the research community needs as an input to more advanced optimization models that integrate demand forecasting, price optimization, and revenue optimization models.

1. Introduction and motivation

Within the airline industry, there is growing interest in understanding how prices influence demand. As strange as this may seem, current airline revenue management (RM) systems do not forecast demand as a function of price. Instead, these systems forecast demand for a particular booking class. To generate booking class forecasts, all prices sold in the market (which can exceed more than a hundred for a single flight) are mapped into a smaller – and more manageable – number of booking classes. RM systems optimize revenue by using information about the historical demand and average fares associated with each booking class.

* Corresponding author. Tel.: +1 (404) 385 6634.
E-mail addresses: stacey.mumbower@gatech.edu (S. Mumbower), laurie.garrow@ce.gatech.edu (L.A. Garrow), matt.higgins@scheller.gatech.edu (M.J. Higgins).

http://dx.doi.org/10.1016/j.tra.2014.05.003
0965-8564/

These RM systems worked well in the era after deregulation because fare restrictions (such as advance purchase, minimum stay, and Saturday night stay requirements) made it relatively straight-forward to segment customers and map them to distinct booking classes with monotonically-increasing average fares. However, these systems are currently struggling because the market today is fundamentally different than it was after deregulation when these first-generation RM systems were built. Overall, the market has become more competitive. The Internet has become a more significant distribution channel and low cost carriers (LCC) have increased market penetration. For example, in 2012, U.S. business and leisure travelers are estimated to have spent $85.7 billion online for airline tickets (Harteveldt, 2012). In 2009 Southwest Airlines was the largest U.S. domestic carrier, carrying over 101 million passengers; 81% of these passengers made their bookings via www.southwest.com (Southwest Airlines, 2009, 2010).

These factors have increased price and flight transparency making it easier for consumers to compare prices across multiple competitors and tailor travel plans to take advantage of lower fares. Airlines have responded to this increased competition by investing in automated price response systems. These systems help airlines identify when competitors have introduced new fares into the market, and provide recommendations as to how to respond to these changes (*e.g.*, match the fares of United, or price fares $50 higher than Spirit). This automation is essential, as there is no way to manually manage the process. To put this in perspective, at any given time, there are more than 100 million fares in the world (ATPCo, 2013b). On a given day, there may be more than one million fare changes (Vinod, 2010). In the U.S., domestic fares are updated through the Airline Tariff Publishing Company (ATPCo) up to four times a day and international fares can be updated hourly (ATPCo, 2013a). Clearly, with such a dynamic environment, it is challenging for airlines to maintain accurate inputs into their RM systems and map customer bookings and their associated fares into smaller sets of booking classes with monotonically-increasing average fares. However, as noted by Vinod (2010), "although frequently overlooked, addressing the fare class misalignment problem is *mandatory* for revenue management to produce positive results."

These and other challenges have spurred interest in developing the next generation of RM systems that better represent how customers make decisions in today's online environments. The development of these choice-based RM systems requires information about the prices (or choices) viewed by customers at the time of booking – both on the carrier of interest and, potentially, across several different competitors. The ultimate goal of these new RM systems is to forecast demand as a function of price and maximize revenue by jointly determining what prices to offer in the market, as well as how many seats to sell at each price. In turn, this means that airlines will need to develop methods for estimating demand elasticities that take into account day-to-day fluctuations in prices.

The majority of extant work has estimated air travel demand elasticities at a high level of data aggregation, typically at the market or route level. This is due to researchers only having access to highly aggregated datasets, most notably the U.S. Department of Transportation's Origin and Destination Survey Databank 1A/1B, which provides a 10% sample of route-level prices over an entire quarter. Although these measures are useful for such things as forecasting the impacts of mergers, the entry of a low cost carrier into a market, or the imposition of system-wide fuel surcharges and passenger taxes, they provide limited insight on how to design promotions and when (and how) to respond to competitors' fare changes. Moreover, these more aggregate measures provide limited insights into questions that policy makers need to address, such as the potential impact of time-of-day congestion pricing. To fully answer these types of questions, researchers need flight-level price elasticities that can be used to forecast responses in demand due to day-to-day fluctuations in prices.

In this paper, we show how flight-level price elasticities can be estimated using publically-available online data. Importantly, we use an instrumental variable approach to correct for price endogeneity. This is critical since failure to correct for endogeneity in these types of models leads to biased estimates and incorrect pricing recommendations. Our results indicate that price elasticities vary as a function of advance booking, departure day of week, departure time of day, booking day of week, and promotional sales dates of a competitor. We use these elasticities to show how they can be used to support airline pricing decisions.

The remaining sections are organized as follows. Section 2 describes the data and discusses potential sources of selection bias. Section 3 presents our methodology, with a particular focus on how we addressed missing data and price endogeneity. Empirical results are presented in Section 4 while robustness and study limitations are discussed in Section 5. We provide specific examples of how our results can help support airline pricing decisions in Section 6. We conclude by highlighting how our model contributes to the literature by offering an approach to estimate flight-level demand elasticities that will allow the research community to move one step closer toward its ultimate goal of developing advanced optimization models that integrate demand forecasting, price optimization, and revenue optimization models.[1]

2. Data

This section describes the data and variables used in the study.[2] It highlights information about the data that is relevant for interpreting results. For additional information on the pricing data, readers are referred to Mumbower and Garrow (2014) and to other papers that have used this data for pricing and revenue management applications (*e.g.*, Newman et al., 2013; Mumbower et al., 2013).

[1] Because airlines do not store information about the prices actually paid by consumers within their RM systems, airlines would need to use data sources such as the ones we use in this study to develop prototypes that will help them justify multi-million dollar investments required for them to store and access data required to successfully implement choice-based RM systems.

[2] The pricing data is available upon request and is described more fully in Mumbower and Garrow (2014).

Table 1
JetBlue descriptive statistics.

Market[a]	Nonstop competitors[b]	Flight number	DTOD[c]	Total bookings	Min price	Mean price	Max price
BOSLAX	AA, B6, UA, VX	473	8	844	$114	$205	$586
		483	18	927	$114	$191	$466
JFKLAS	AA, B6, DL, VX	187	7	451	$129	$254	$463
		191	18	405	$129	$231	$463
		197	10	458	$129	$282	$586
		199	21	313	$129	$225	$463
		711	14	481	$129	$251	$526
JFKLAX	AA, B6, DL, UA, VX	671	11	691	$129	$257	$586
		673	16	671	$129	$244	$586
		675	7	974	$129	$205	$526
		677	19	747	$129	$223	$466
JFKSFO	AA, B6, DL, UA, VX	641	8	339	$129	$300	$586
		647	17	221	$129	$287	$586
	Totals/Averages:			7522	$114	$232	$586

[a] Airport codes: BOS = Boston; LAX = Los Angeles; JFK = JFK, New York; LAS = Las Vegas; SFO = San Francisco.
[b] Airline codes: AA = American; B6 = JetBlue; DL = Delta; UA = United; VX = Virgin America.
[c] DTOD is defined as "flight departure time of day", in local military time. For example, a DTOD of 18 means the flights departed between 6:00 PM and 6:59 PM.

2.1. Overview and descriptive statistics

We predict demand for JetBlue flights in four transcontinental markets. Automated web client robots (or webbots) collected detailed flight, fare, and seat map information for 21 departure dates (September 2, 2010 to September 22, 2010) over a 28-day booking horizon. Our data collection methods are consistent with those used in prior studies (*e.g.*, Bilotkach, 2006; Bilotkach and Pejcinovska, 2012; Bilotkach et al., 2010; Button and Vega, 2006, 2007; Horner et al., 2006; Li et al., 2011; McAfee and Vera, 2007; Mentzer, 2000; Pels and Rietveld, 2004; Pitfield, 2008).

The main difference between our study and prior work is that we collect seat maps in addition to pricing data; all other studies focused exclusively on pricing data. By tracking seat maps over the booking horizon, we are able to calculate a measure of demand for JetBlue flights. We define the daily number of JetBlue bookings on a flight as the number of seats that switched from being "available" one day to "reserved" the next day. Descriptive statistics for the four markets are provided in Table 1.

We observe between 0 and 16 bookings per flight per day, with a mean (median) of 1.9 (1) bookings.[3] One-way flight prices range between $114 and $586, with a mean (median) of $232 ($199). A total of 7522 bookings are observed for JetBlue. As seen in Fig. 1, both average prices and average demand increase as the flight departure date approaches (with correlation coefficients of −0.76 and −0.56, respectively). Most importantly, Fig. 1 highlights the underlying price endogeneity that must to be addressed as part of any empirical analysis.

2.2. Selection bias

Predicting demand for JetBlue flights enables us to control for potential sources of selection bias. We expect the number of occupied seats to be highly correlated with JetBlue's actual nonstop demand. This is because, unlike the majority of U.S. airlines, JetBlue does not overbook its flights (*i.e.*, JetBlue does not sell more tickets than the actual number of seats on each flight). This means that all customers have the option to select a regular coach seat for free at the time of booking. Further, we expect the majority of customers will select seats at the time they book, as they are prompted to do so during the booking process.

We also expect online fares shown to consumers to be strongly correlated with the actual fares paid by consumers. Many airline websites offer multiple prices per flight, which makes it impossible to know how much a customer actually paid when observing seat maps and screen displays. However, during our data collection period, JetBlue only showed a single one-way fare for each nonstop flight. Further, the majority of JetBlue's customers purchased a nonstop fare and were local passengers. Between 96.4% and 99.6% of JetBlue passengers in these markets traveled on nonstop flights (U.S. DOT, 2010). Further, across these markets between 88.0% and 94.9% of passengers are "local" passengers who do not originate from or connect to another airport (U.S. DOT, 2010).[4]

[3] Throughout this paper, we use the terms "number of bookings" and "demand" interchangeably, although we realize that the two measures are not exactly the same. JetBlue's flights rarely sellout, so in general, there is not more demand for flights than we can observe from the actual bookings.

[4] We acknowledge that there is another selection bias present for which we are unable to correct for given current data limitations. Our results are conditional on a customer having selected JetBlue from a choice set that includes other airlines. In order to correctly solve this potential bias we would need detailed data about the customers, which does not currently exist. (Such data could be obtained, for example, via survey.) If such data were available, a two-stage model could be used where the first-stage focused on the customer characteristics that led to the selection of a particular airline.

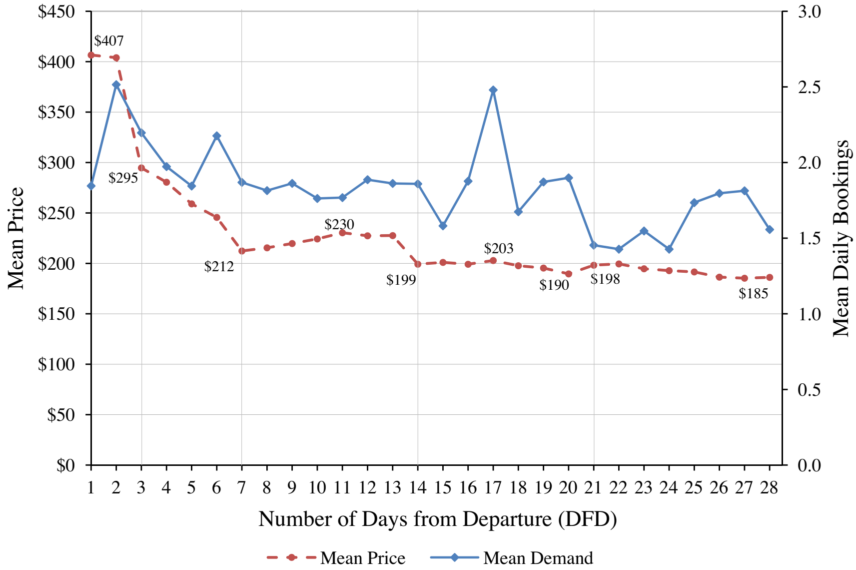

Fig. 1. Average daily demand and prices as a function of days from departure.

Using the BOSLAX market as an example, a nonstop passenger is defined as anyone who traveled between the market origin (BOS) and market destination (LAX) on a nonstop versus connecting flight. However, these "nonstop" BOS to LAX passengers may have had a flight before BOS or continued on after LAX. Local passengers are those who traveled just from BOS to LAX, *i.e.*, they boarded in BOS and disembarked in LAX. Thus, the percent of "nonstop" passengers in combination with the percent local passengers provide information about the number of customers who (just) purchased a flight from BOS to LAX. For these four JetBlue markets, the percentages of nonstop and local passengers are very high, so we simplify the model by ignoring network-level effects in our model.[5]

2.3. Holidays and promotions

During the time of data collection, the Labor Day holiday was observed on Monday, September 6, 2010. We control for this holiday with dummy variables for bookings made for flights departing the day before, the day of, and the day after Labor Day.

Also during the data collection period, a competing airline, Virgin America had three sales that were promoted on its website and directly to consumers via emails from Travelzoo®. One of these sales, promoting Virgin America's plans to launch new service to Los Cabos and Cancun, Mexico (Virgin America, 2010), increased demand for Virgin America flights. It was noted that "....significant online buzz circulating about the promotion helped make it the fifth highest sales day in Virgin America's history..." (Arrington, 2010).

The presence of competitor sales during the data collection period is fortuitous, as it provides us with a unique opportunity to examine how JetBlue's bookings were impacted by the sales and whether JetBlue responded to the sales by also lowering its prices. We controlled for the influence of Virgin America's sales by including a dummy variable for those booking dates and departure dates in which Virgin America was offering a promotion.

2.4. Other variables

Table 2 defines and describes the dependent variable, independent variables, and controls utilized in our analysis. Instrumental variables, more fully discussed below, are also defined.

[5] In datasets with airlines and markets that have more connecting traffic than those used in the dataset used in this paper, it will likely be important to take into account network-level effects. Using nonstop fares for all passengers could potentially bias price elasticity results, but it is difficult to determine in which direction. Airlines use different strategies for deciding whether to sell a seat on a particular flight to a nonstop passenger or a connecting passenger. When priority is given to a connecting passenger, the pro-rated amount a connecting passenger would pay for a seat on flight A is likely less than the price a nonstop passenger would pay for flight A. This would potentially lead to price elasticity estimates that over-predict changes in quantity demand associated with a price change. Conversely, when priority is given to a nonstop passenger, the pro-rated amount a nonstop passenger would pay for a seat on flight A is likely less than the pro-rated amount a connecting passenger would pay for a seat on flight A. This would potentially lead to price elasticity estimates that under-predict changes in quantity demand associated with a price change.

Table 2
Variables and descriptions.

Variable	Variable description
numbookings	The total number of daily bookings for a flight (dependent variable)
price	Price of the flight (JetBlue's oneway price)
vxsale	Indicates a date that Virgin America was offering promotional sales
travelsep5, ..., travelsep7	Indicates bookings made for travel on and around Labor Day holiday (Sep. 6, 2010)
dtod7, dtod8, ... dtod16+	Indicates flight departure is 7–7:59 AM, 8–8:59 AM, 4–9:59 PM[a]
dfd1, dfd2, ... dfd28	Indicates a booking made 1, 2, ..., 28 days from flight departure
ddow1, ..., ddow7	Indicates flight departs on a Sun, Mon, ..., Sat
bdow1, ..., bdow7	Indicates flight was booked on a Sun, Mon, ..., Sat
Market dummies	Dummy variable for each market
IV1	Instrumental variable: JetBlue's mean prices in other markets (Hausman-type price instrument)
IV2	Instrumental variable: the average number of Virgin America's nonstop flights in a market (Stern-type competition instrument)

[a] In the data no flights depart between 9–9:59 AM, 12–1:59 PM, 3–3:59 PM, 8–8:59 PM and 10:00 PM–6:59 AM.

3. Methodology

Consistent with the extant literature, we use a linear model to predict air travel demand (*e.g.*, Bhadra, 2003; Granados et al., 2012). Specifically, we use linear regression methods to estimate the number of bookings for flight i with departure date j in market m that are made t days in advance. We observe bookings made 1–28 days in advance of flight departures for 13 flights in 4 markets across 21 departure dates. We first present results obtained from an ordinary least squares (OLS) regression model; however, in the presence of endogeneity the estimates will be biased. To correct for this we implement and present results from a two-stage least squares (2SLS) instrumental variable model (Greene, 2003).

There were two additional key methodological challenges we need to address as part of our analysis. The first relates to missing data and the second relates to finding a set of valid instruments to correct for endogeneity.

3.1. Missing data

When collecting data using webbots, it is common to have missing data.[6] Our dataset is approximately 73% percent complete. The data is more complete for flights that are closer to their departure dates. The pattern of missing data is related to the number of queries that were executed over time. For example, on September 1, 2010, pricing data for all 21 departure dates needed to be collected. However, on the last day of the data collection, only information for a single remaining departure date needed to be collected. The structure of the webbot queries over time results in data that is not missing completely at random. That is, missing data are related to the collection date; however, conditional on the collection date, the data are missing at random.

There are many methods that can be used to account for missing data. In the complete case analysis method, only the observations with complete data for every variable are included in the analysis. However, this method will typically lead to invalid and biased results (Carpenter et al., 2006). One way to correct for bias introduced by missing data is to use inverse probability weights with complete case analysis method. First, the probability of an observation being missing is modeled using binary logistic regression and all fully observed explanatory variables. Then, the inverse of the predicted probability of an observation being missing is used as a weight on the complete cases that are observed. Observations that have a small probability of being observed are given larger weights to compensate for the similar missing observations (Bartlett, 2012). We used this approach to correct for missing data.[7]

3.2. Price endogeneity

Many prior studies of airline demand have failed to properly address price endogeneity and have assumed that prices are exogenous. Results from these studies are therefore biased. Endogeneity occurs when correlation exists between an explanatory variable and the error term (or unobserved factors) in a model. This correlation means that the conditional expectation

[6] We have seen similar patterns of missing data in webbot data we collected, as well as in pricing data compiled by firms that specialize in collecting online pricing data and selling this data to airlines. The missing data patterns we describe here are likely to be encountered in similar research and industry contexts.

[7] When missing observations occur (due to webbots not completing on certain data collection dates), we are missing prices and seat maps (demand). However, due to the date-based structure of the data collection, we know which queries did not complete (*i.e.*, we know which observations are missing) for each market, departure date and query date combination. For example, on query date X, a webbot did not capture data for flights that depart on date Y in market Z. In this case, the response variable for the binary logit model would equal zero; all non-price explanatory variables in Table 2 would be known and used in the binary logit model as explanatory variables. Forming the weights in this way creates values that are specific to each unique market, departure date, and query date combination. The weights are not flight-specific because within each webbot query, either all flights on a webpage were captured or none of them were.

of the error term on the endogenous explanatory variable will not equal zero, which violates a main assumption required to ensure estimator consistency for most models (Greene, 2003).

In demand models, prices are endogenous because they are influenced by demand, which is influenced by prices (often referred to as simultaneity of supply and demand). For excellent comprehensive reviews of endogeneity, see Guevara-Cue (2010) and Train (2009). Price endogeneity is well documented in the economics and management literatures. Many empirical demand studies have shown that price coefficients are underestimated if endogeneity is not corrected, including recent studies that estimate: demand for high speed rail travel (Pekgün et al., 2013), household choice of television reception options (Goolsbee and Petrin, 2004; Petrin and Train, 2010), household choice of residential location (Guevara and Ben-Akiva, 2006; Guevara-Cue, 2010), choice of yogurt and ketchup brands (Villas-Boas and Winer, 1999), consumer-level choice of and aggregate product demand for the make and model of a new vehicle (Berry et al., 1995, 2004; Train and Winston, 2007), and brand-level demand for hypertension drugs in the U.S. (Branstetter et al., 2011).

There are multiple methods that can be used to correct for price endogeneity, including two-stage least squares (2SLS) regression that accounts for endogeneity using instruments. An instrument is a variable that does not belong in the demand equation, but is correlated with the endogenous price variable. Instruments that satisfy the following two conditions will generate consistent estimates of the parameters, subject to the model being correctly specified: (1) instruments should be correlated with the endogenous variable, and (2) they should be independent of the error term in the model (Rivers and Vuong, 1988; Villas-Boas and Winer, 1999). Therefore, we need to find instruments that are correlated with airfares but not correlated with a customer's purchase or choice of a flight. Validity tests are used to statistically determine whether the instruments are correlated with airfares, but not correlated with the error term of the demand model (*i.e.*, customers' purchase or choice of a flight).

The first-stage of our 2SLS model is an OLS regression that uses price as the dependent variable. Explanatory variables include the set of instruments, along with all other exogenous regressors. The predicted price (*predprice*) from the first stage regression is used in place of the price variable in the second-stage OLS regression. The first- and second-stage regressions are formulated as follows:

Stage 1:

$$\begin{aligned}\text{price}_{ijmt} &= \alpha_0 + \beta_1\text{IV1}_{jmt} + \beta_2\text{IV2}_{jm} + \beta_3\text{vxsale}_{jmt} + \beta_4\text{travelsep5}_j + \ldots + \beta_7\text{dtod7}_{ij} + \ldots + \beta_{12}\text{dfd1}_t + \ldots \\ &\quad + \beta_{27}\text{ddow1}_j + \ldots + \beta_{33}\text{bdow1}_{jt} + \ldots + \beta_{39}\text{market1}_m + \ldots + \beta_{41}\text{market3}_m + u_{ijmt}\end{aligned} \tag{1}$$

Stage 2:

$$\begin{aligned}\text{NumBook}_{ijmt} &= \tau_0 + \gamma_1\text{predprice}_{ijmt} + \gamma_2\,\textit{vxsale}_{jmt} + \gamma_3\text{travelsep5}_j + \ldots + \gamma_6\text{dtod7}_{ij} + \ldots + \gamma_{11}\text{dfd1}_t + \ldots \\ &\quad + \gamma_{26}\text{ddow1}_j + \ldots + \gamma_{32}\text{bdow1}_{jt} + \ldots + \gamma_{38}\text{market1}_m + \ldots + \gamma_{40}\text{market3}_m + \varepsilon_{ijmt}\end{aligned} \tag{2}$$

where we estimate the number of bookings for flight i with departure date j in market m that are made t days in advance.

To estimate the 2SLS model, we used the statistical software Stata 10 (StataCorp, 2007) and the enhanced estimation routine *ivreg2* (Baum et al., 2007, 2010a). *Ivreg2* performs the two stages within one estimation command, automatically making the necessary corrections to the standard errors (SE) of the second stage. It should be noted that the general standard errors (SE) of IV estimates are inconsistent when heterogeneity is present. This does not impact the consistency of the IV coefficient estimates, but the general forms of diagnostic tests will be inconsistent (Baum et al., 2003). We control for market heteroskedasticity by using a dataset of similar markets (nonstop transcontinental flights in markets where JetBlue and Virgin America compete head-to-head). Additionally, we use cluster-robust SE to control for heteroskedasticity and intra-group correlation (Greene, 2003). We cluster by itinerary screen (*i.e.*, a cluster includes all flights for sale on a particular day for a unique market and departure date), which relaxes the assumption of independent observations within each set of flights on an itinerary screen.

3.3. Instruments

Many researchers have noted the difficulty in finding valid instruments. Moreover, there exists a debate as to the theoretical soundness of certain types of instruments (*e.g.*, Bresnahan, 1997). Table 3 summarizes the types of instruments that have been used in the literature and offers examples that have been used (or could possibly be used) in air travel demand models.

3.3.1. Cost-shifting variables as instruments

Variables that shift cost and are uncorrelated with demand have been used in many applications of aggregate-level demand estimation. For example, Hausman (1996) estimates aggregate brand choice models for ready-to-eat cereal and uses instruments that shift the cost of cereal (such as ingredients, packaging, and labor). Within the airline industry, Hsiao (2008) uses route distance multiplied by unit jet fuel cost as an instrument for discrete choice models of aggregate quarterly air travel demand. Route distance and unit jet fuel cost can be thought of as cost-shifters because they are expected to impact the price of tickets.

Theoretically, these are good instruments if one believes that route distance and unit jet fuel cost are correlated with ticket prices, but not with a customer's decision to travel. Unfortunately, these variables are unable to capture day-to-day

Table 3
Summary of instrument types and examples of instruments in the airline context.

Instrument Type (with reference to authors who introduced it)	Instrument description	Examples of instruments in the airline context
Cost-shifting instruments	Variables that impact a product's cost but that are uncorrelated with demand shocks	Hsiao (2008) uses route distance and unit jet fuel costs Berry and Jia (2010) and Granados et al. (2012) use a hub indicator Granados et al. (2012) use distance
Hausman-type price instruments Hausman et al. (1994), Hausman (1996)	Prices of the same brand in other geographic contexts are used as instruments of the brand in the market of interest	Gayle (2004) uses an airline's average prices in all other markets with similar length of haul (also used in this paper)
Measures of competition and market power Stern (1996)	Measures of the level of market power by multiproduct firms, and measures of the level of competition	Berry and Jia (2010) use the number of all carriers offering service on a route Granados et al. (2012) use the degree of market concentration, calculated as the Herfindahl index Number of daily nonstop flights in the market operated by competitor airlines (also used in this paper)
Measures of non-price characteristics of other products Berry et al. (1995)	Average non-price characteristics of the other products supplied by the same firm in the same market	Average flight capacity of other flights operated by the airline of interest in the same market
	Average non-price characteristics of the other products supplied by other firms in the same market	Berry and Jia (2010) use the percentage of rival routes that offer direct flights, the average distance of rival routes, and the number of rival routes

fluctuations in airfares, which are more likely to be driven by revenue management practices and competitor price matching. In a disaggregate model of air travel demand, most cost-shifting variables will not be good instruments due to this lack of day-to-day variation.

3.3.2. Hausman-type price instruments

For the disaggregate models of brand choice in the ready-to-eat cereal industry, Hausman (1996) could not use cost-side instruments. He could, however, observe price in several different cities (markets). He therefore defined a price instrument for the city of interest using prices of the same brand in other cities (many researchers now refer to this type of instrument as "Hausman-type price instruments"). These instruments are based on the economic theory that a firm's price in one city (market) are a function of the average marginal costs of a product plus a markup amount that the firm is able to charge due to customers' differing willingness to pay for that product in that specific city or market. Hausman assumes that a city-specific value and demand for a brand is independent across cities or markets. The basic idea is that after eliminating city-specific and brand-specific effects (by including fixed-effects in the model), the price of a brand in city j will be correlated with the prices of the brand in other cities due to common marginal costs. However, the price of a brand in city j will (ideally) be uncorrelated with city-specific valuation and demand in other cities.[8]

Many studies have used these Hausman-type price instruments (*e.g.*, Guevara and Ben-Akiva, 2006; Guevara-Cue, 2010; Nevo, 2000b, 2001; Petrin and Train, 2010). Within our airline context, this translates into a price instrument for market m that is defined as the average price in all other markets with a similar length of haul.

3.3.3. Measures of competition and market power as instruments

In contrast to Hausman, Stern (1996) uses measures of market power by multiproduct firms and measures of competition as instruments. Specifically, he notes (Stern, 1996, p.18) that "Unless consumers value products sold by a particular firm *because* it is a multiproduct firm, measures of multiproduct ownership will be correlated with price and advertising, but be uncorrelated with unobserved quality." Levels of market power focus on the number of products in the market and also the time since a product (and/or firm) was introduced into the market. In the context of pharmaceutical drugs, he measures the level of market

[8] See Hausman (1996) and Nevo (2000a,b) for a more formal discussion.

power by multiproduct firms as the number of products produced within a drug category by a firm that produces product *j*, and the sum of the time since entry over each of all other products (excluding product *j*). Additionally, Stern (1996, p.18) also notes that "…measures of the level of competition in the market, such as the number and characteristics of other products, will also affect price but, under the assumption that entry is exogenous, be uncorrelated with unobserved quality". For example, one instrument Stern uses to capture the degree of competition facing product *j* is the number of manufacturers in the market.

Translating Stern's approach into an airline context, the number of competitors' flights in a market or the number of carriers in a market could be used as potential instruments. This approach has been followed in two airline-specific studies. First, Berry and Jia (2010) use the number of carriers offering service on a route as an instrument. Second, Granados et al. (2012) use the degree of market concentration (the Herfindahl index) as an instrument.

3.3.4. Non-price product characteristics of other products as instruments

Berry et al. (1995), commonly referred to as "BLP", derive a set of instruments using observed exogenous product characteristics that excludes price. Their suggested instruments include: (1) observed product characteristics for a firm; (2) the sums, if any, of the values of the same product characteristics of *other products* offered by that firm; and, (3) the sums of the values, if any, of the same characteristics of the same products offered by other firms.

Instruments of this type have been used in many applications, including choice of an automobile (*e.g.*, Berry et al., 1995, 2004; Train and Winston, 2007). Nevo (2000a, p. 535) provides a clear description of how these instruments have been used within the automobile industry: "Suppose the product has two characteristics: horsepower (HP) and size (*S*), and assume there are two firms producing three products each. Then we have six instrumental variables: the values of HP and *S* for each product, the sum of HP and *S* for the firm's other two products, and the sum of HP and *S* for the three products produced by the competition."

In the airline context, the average flight capacity of other flights operated by the airline of interest in the same market could be used as instruments. One airline study utilizes similar types of instruments; notably the percentage of rival routes that offer direct flights, the average distance of rival routes, and the number of rival routes (Berry and Jia, 2010). For the disaggregate flight-level models of air travel demand discussed in this paper, we found that this type of instrument suffers from a problem similar to that of the cost-shifting instruments. Instruments of this type generally lack day-to-day variation, and (by themselves) cannot fully explain day-to-day fluctuations observed in airfares of highly disaggregate data. The dataset discussed in this paper is for a rather homogeneous set of nonstop flights in four markets. For a larger, more heterogeneous set of markets and routes (that includes both nonstop and connecting flights), including a variable of this type in the set of instruments may be helpful.

4. Model results

The instruments we use to correct for endogeneity are based on Hausman-type price and Stern-type competition instruments. Specifically, we construct the Hausman-type instrument by using JetBlue's equivalent one-way price from the online travel agency (OTA) website (round-trip prices divided by two). These instruments have day-to-day variation as JetBlue changes prices. For the Stern-type competition instrument, we use the number of daily flights offered by a competitor (Virgin America) as a proxy for the degree of competition facing JetBlue. This instrument has variation as Virgin America operates a different number of nonstop flights across markets and departure days of the week.

We performed three diagnostic tests: an endogeneity test of endogenous regressors[9], a test for instrument validity[10], and a test for instrument relevance.[11] The results of these diagnostic tests indicate that price should be treated as endogenous and that the set of instruments are valid and relevant. Additionally, as a robustness check, we compared different assumptions on the SE, including models with and without cluster-robust SE. We also tested model specifications with and without the inverse probability weight correction for missing values. Finally, we compared the 2SLS model to a generalized method of moments (GMM) model specification.[12] Importantly, the model results (and corresponding price elasticities) were consistent with our main findings.

4.1. Comparison of OLS and 2SLS estimates

Table 4 compares the results obtained from an ordinary least squares (OLS) regression model to those obtained from a two-stage least squares (2SLS) instrumental variable model, keeping in mind that the estimates from OLS are biased. The

[9] The *ivreg2* option *endog* tests the null hypothesis that price can be treated as an exogenous regressor. The null hypothesis was rejected, with Chi-square(1) *p*-value = 0.0394, indicating that price should be treated as endogenous.

[10] The Hansen *J* test of over-identifying restrictions tests the joint null hypothesis that the set of instruments are valid (uncorrelated with the error term) and correctly excluded from the demand model. The null hypothesis was not rejected, with Chi-square(1) *p*-value = 0.9098, indicating no evidence of invalid instruments. Hansen's *J* statistic is consistent in the presence of heteroskedasticity, whereas some other test statistics are not (Baum et al., 2010b)

[11] The first stage regression on price has an adjusted *R*-square of 0.48 and a Kleibergen-Paap Wald *F*-statistic of 10.034. The Kleibergen-Paap Wald *F*-statistic is used in place of the Cragg-Donald *F*-statistic when errors are not assumed to be independent and identically distributed, as is the case when robust standard errors are used (Stock and Yogo, 2005; Baum et al., 2003, 2010b). Using a critical value of 8.75 (5% significance level test that the maximum size distortion is no more than 20%), we reject the null hypothesis that the instruments are weak (Stock and Yogo, 2005: Table of Critical Values for the Weak Instrument Test Based on TSLS Size).

[12] Two-stage least squares is a special case of the generalized method of moments (GMM) model that was introduced in Hansen (1982).

Table 4
OLS and 2SLS regression results.

	OLS		2SLS	
price	−0.0056***	(0.000)	−0.0151***	(0.005)
vxsale	−0.2417***	(0.089)	−0.2522***	(0.096)
Departure time of day (reference: evening departures 4pm or later)				
7–7:59 AM	0.2741***	(0.105)	0.2817***	(0.105)
8–8:59 AM	0.0404	(0.129)	0.1836	(0.152)
10–10:59 AM	0.4507***	(0.158)	0.8719***	(0.277)
11–11:59 AM	0.1492	(0.138)	0.3831**	(0.184)
2–2:59 PM	0.1982*	(0.118)	0.3363**	(0.138)
Number of days from flight departure (reference: dfd22–28)				
dfd1	1.3997***	(0.218)	3.2804***	(1.047)
dfd2	2.0543***	(0.209)	3.9633***	(1.039)
dfd3	1.2128***	(0.165)	2.0728***	(0.508)
dfd4	0.8894***	(0.163)	1.5955***	(0.420)
dfd5	0.6015***	(0.143)	1.0796***	(0.302)
dfd6	0.8464***	(0.204)	1.2213***	(0.290)
dfd7	0.4608***	(0.137)	0.5556***	(0.163)
dfd8	0.4199***	(0.154)	0.5618***	(0.181)
dfd9	0.5152***	(0.155)	0.706***	(0.197)
dfd10	0.3232**	(0.158)	0.5613***	(0.204)
dfd11	0.3687**	(0.145)	0.6502***	(0.221)
dfd12	0.4302**	(0.178)	0.6647***	(0.211)
dfd13	0.446**	(0.210)	0.698***	(0.265)
dfd14	0.2697*	(0.150)	0.2426	(0.160)
dfd15_21	0.2105**	(0.096)	0.2268**	(0.101)
Departure day of week variables (reference: Saturday departure)				
ddow1 (Sunday)	0.3971***	(0.133)	0.8898***	(0.293)
ddow2 (Monday)	0.5039***	(0.128)	1.0985***	(0.337)
ddow3 (Tuesday)	0.2493**	(0.122)	0.3719**	(0.151)
ddow4 (Wednesday)	0.4306***	(0.123)	0.5255***	(0.142)
ddow5 (Thursday)	0.2193*	(0.123)	0.6797**	(0.267)
ddow6 (Friday)	0.3198**	(0.125)	0.486***	(0.162)
Booking day of week variables (reference: Friday booking)				
bdow1 (Sunday)	−0.7744***	(0.103)	−0.6539***	(0.138)
bdow2 (Monday)	0.2502*	(0.130)	0.477***	(0.180)
bdow3 (Tuesday)	0.4097***	(0.124)	0.4618***	(0.132)
bdow4 (Wednesday)	0.3611***	(0.114)	0.276**	(0.126)
bdow5 (Thursday)	0.3218***	(0.122)	0.3978***	(0.130)
bdow7 (Saturday)	−0.7647***	(0.100)	−0.689***	(0.122)

Cluster robust standard errors in parenthesis.
Note 1: Both models include dummies for departure dates on and around Labor Day, market dummies and a constant term, which are not reported.
Note 2: Both models use inverse probability weights to account for missing data.
Note 3: R-squared for OLS is 0.137.
* $p < 0.10$.
** $p < 0.05$.
*** $p < 0.01$.

price coefficient in the 2SLS regression becomes more negative, a trend which is consistent with prior findings in the literature (*e.g.*, Berry et al., 1995, 2004; Branstetter et al., 2011; Goolsbee and Petrin, 2004; Guevara and Ben-Akiva, 2006; Guevara-Cue, 2010; Train and Winston, 2007; Pekgün et al., 2013; Petrin and Train, 2010; Villas-Boas and Winer, 1999).

4.2. Average price elasticities

Price elasticity of demand estimates are commonly used in policy decisions and to justify increases in airport fees and taxes, as well as national departure and emissions taxes (International Air Transport Association, 2008). Therefore, calculating reliable price elasticities is an essential part of ensuring effective air transport policy (International Air Transport Association, 2008).

Due to their importance for decision-making within the airline industry, many studies have estimated price elasticities. Estimated elasticities have varied widely depending on the data used, the modeling methodology, and the markets and time period included. Some studies have corrected for price endogeneity, and others have not. Most studies have used aggregate data to estimate price elasticity. Price elasticities have been found to differ across many dimensions of air travel, including: business versus leisure travel, length of haul, level of data aggregation, and booking channel. Business travelers are generally found to be more inelastic (less price sensitive) than leisure travelers, as people traveling for business have less flexibility to

Table 5
Comparison of price elasticity estimates across studies.

Study	Level of aggregation	Elasticity estimate	Data source
Gillen et al. (2002)	Market	−0.198 to −1.743[1]	Meta-study
InterVistas (2007)	Route/Market	−1.40 to −1.54[2]	DB1B (quarterly fares)
	National	−0.80 to −0.88[2]	
	Pan-National	−0.60 to −0.66[2]	
Hsiao (2008)	Market	−1.05 to −2.66	DB1B (quarterly fares)
	Route	−1.76 to −2.97	
Granados et al. (2012)	Booking channel:		Booking data
	Leisure travel	−1.33 to −2.28	
	Business travel	−0.34 to −1.29	
This study	Flight	−1.32 to −1.97[3]	Daily online pricing and seat map data

[1] These elasticities represent a wide range of markets. Long haul domestic elasticities range from −0.79 to −1.43.
[2] Elasticities of −1.40, −0.80 and −0.60 represent long haul markets. Elasticities of −1.54, −0.88, and −0.66 represent short haul markets.
[3] These numbers represent median and mean elasticities. However, elasticities also vary as a function of flight characteristics, booking characteristics and competitor promotions and are found to be between −0.57 and −3.21.

Table 6
OLS and 2SLS price elasticity results (at the mean and median of price).

Model	Price elasticity	
	At median price $199	At mean price $232
OLS	−0.58	−0.75
2SLS	−1.32	−1.97

Note: Price elasticities are calculated over the means of all non-price variables.

postpone or cancel their trip. Travelers in short-haul markets are generally more elastic (more price sensitive) due to availability of inter-modal substitutes, such as driving or taking a bus.

Table 5 compares elasticity estimates from prior studies to those obtained from our model. A meta-study by Gillen et al. (2002) found that market-level price elasticities in the literature have ranged from −0.198 (*i.e.*, very inelastic or non-price sensitive) in long-haul international business markets to −1.743 (*i.e.*, very elastic or price sensitive) in short-haul leisure markets. InterVISTAS (2007) corrects for price endogeneity and finds average elasticity estimates that vary based on the length of haul and level of data aggregation, for long haul markets: −1.40 at the route/market-level, −0.80 at the national-level, and −0.60 at the pan-national level. Hsiao (2008) estimates discrete choice models of aggregate quarterly air passenger demand at the market-level and route-level and finds price elasticity estimates between −1.05 and −2.66, and −1.76 to −2.97, respectively. Finally, using disaggregate bookings data and correcting for price endogeneity, Granados et al. (2012) find that leisure travel booked through offline channels, transparent online travel agents, and opaque online travel agents have price elasticity estimates of −1.33, −1.56, and −2.28, respectively.[13]

As is clearly demonstrated by the InterVISTAS (2007) study, price elasticities become more negative (more elastic or price sensitive) as the level of aggregation becomes more refined. Intuitively, this is because there are more substitutes available at lower levels of aggregation. In other words, overall market-level demand for air travel will not be as sensitive to average market-level prices as flight-level demand. In this context, the results of our study are consistent with those reported in earlier studies, and show relatively high price sensitivity at the flight-level.

Table 6 shows the comparison between the price elasticities of demand estimated from our OLS and 2SLS regression models. For the OLS regression model, the estimated price elasticity of demand evaluated at the sample median is −0.58, which represents inelastic demand. After correcting for endogeneity using 2SLS, the estimated price elasticity of demand is −1.32, which represents elastic demand. An estimated price elasticity of −1.32 is interpreted in the following way: a 10% increase in JetBlue's fares leads to a 13.2% decrease in demand. Another way to interpret this is: a 10% decrease in JetBlue's fares leads to a 13.2% increase in demand. Evaluating the price elasticities at the sample mean gives similar results, with elasticity estimates of −0.75 and −1.97, respectively. This difference is important, as pricing recommendations differ for inelastic and elastic models. Specifically, inelastic models suggest that prices should be raised whereas elastic models suggest prices should be lowered. This underscores the importance of correcting for endogeneity in our models: failure to correct for endogeneity results in biased results and leads to incorrect pricing recommendations.

[13] Price elasticities of less than −1 are referred to as elastic (or price sensitive) while price elasticities greater than −1 are inelastic (or price insensitive). Price elasticities equal to −1 are referred to as unit elastic. In the case of the Hsiao (2008) and Granados et al. (2012) for leisure travel, all the price elasticities were less than −1 or elastic.

Table 7
2SLS price elasticity results as a function of flight characteristics.

	Price = $199 (median)	Price = $232 (mean)
Departure time of day[a]		
7–7:59 AM	−1.30	−1.93
8–8:59 AM	−1.26	−1.85
10–10:59 AM	−1.11	−1.58
11–11:59 AM	−1.10	−1.57
2–2:59 PM	−1.37	−2.07
4–4:59 PM	−1.34	−2.01
5–5:59 PM	−1.17	−1.70
6–6:59 PM	−1.72	−2.81
7–7:59 PM	−1.33	−1.99
9–9:59 PM	−1.47	−2.27
Departure day of week		
Sunday	−1.29	−1.91
Monday	−1.16	−1.67
Tuesday	−1.41	−2.14
Wednesday	−1.45	−2.22
Thursday	−1.21	−1.77
Friday	−1.26	−1.86
Saturday	−1.72	−2.80

[a] In the data no flights depart between 9–9:59 AM, 12–1:59 PM, 3–3:59 PM, 8–8:59 PM and 10 PM–6:59 AM.

4.3. Price elasticity estimates as a function of flight, booking, and sale characteristics

Our model differs from others reported in the literature in that it predicts flight-level demand using daily measures of bookings and prices. Our ability to capture day-to-day fluctuations in demand and prices at the flight-level provides us with the ability to understand how price elasticities differ across different flight and booking characteristics, as well as in response to competitor pricing actions. To the best of our knowledge, this is the first time that price elasticity estimates have been reported in the literature at this level of disaggregation.

4.3.1. Price elasticities for flight characteristics

Average daily demand and average prices are also observed to differ by a flight's departure time of day and departure day of week, as shown in Table 7. Customers who book flights that depart on Saturdays are the most price sensitive, and those who book flights that depart on Mondays are the least price sensitive. This is intuitive, as many leisure travelers travel on Saturdays, whereas many business travelers travel on Mondays.

Customers who book flights for morning departures between 10 and 11:59 AM are the least price sensitive, whereas customers who book flights that depart between 6 and 6:59 PM are the most price sensitive. These differences in price elasticities are important when considering the potential impact of time-of-day congestion pricing policies. If taxes associated with using a slot controlled airport during peak hours of the day were to be increased, JetBlue may be able to pass these increases on to the less price sensitive customers who book flights departing between 10 and 11:59 AM. However, JetBlue could expect a greater loss in demand for price increases associated with flights departing between 6 and 6:59 PM.

These results are specific to transcontinental markets and reflect a preference for departure times that depart and arrive mid-day (versus very early in the morning or very late in the evening). Price elasticities across departure times of day will likely vary for different airlines and will depend on the mix of customers flying in the market (*e.g.*, business passengers may prefer earlier morning departures). However, the example highlights how, at a given airport, particular departure times may be highly desirable in one market, and highly undesirable in another market. In this case, mid-morning departures are desirable in transcontinental markets, but are likely undesirable for short-haul markets that operate in the same time zone (such as Boston to Atlanta). Consequently, airlines' ability to pass on time-of-day congestion pricing policies will differ based on which markets are served during the airport's peak period.

4.3.2. Price elasticities for booking characteristics

Price elasticities were calculated from the 2SLS model as a function of number of days from flight departure. Table 8 provides the price elasticities of demand at both the median and mean of price. The table shows that JetBlue's customers are less price sensitive closer to flight departure. This is intuitive as leisure passengers generally book further in advance of departure and business passengers often book closer to departure. In fact, customers who book only 1 or 2 days before flight departure are estimated to be demand inelastic (i.e., price insensitive), whereas customers over all other advance purchase ranges are estimated to be demand elastic (i.e., price sensitive).

Price elasticities are also shown to vary as a function of booking day of week. Customers who book on Saturdays and Sundays are significantly more price sensitive than those customers who book on weekdays. We hypothesize this is because the type of consumers searching on weekends are more likely to be leisure customers who are more price sensitive and have

Table 8
2SLS price elasticity results as a function of booking characteristics.

	Price = \$199 (median)	Price = \$232 (mean)
Advance booking: days from departure		
1–2 days	−0.57	−0.73
3–7 days	−1.03	−1.45
8–14 days	−1.36	−2.04
15–21 days	−1.58	−2.50
22–28 days	−1.89	−3.21
Booking day of week		
Sunday	−1.81	−3.01
Monday	−1.12	−1.60
Tuesday	−1.09	−1.55
Wednesday	−1.20	−1.75
Thursday	−1.20	−1.75
Friday	−1.34	−2.00
Saturday	−1.85	−3.11

Table 9
2SLS price elasticities during competitor promotions vs. all other dates.

	Number of days from departure		
	7–14	15–21	22–28
Not during competitor's promotional sales	−1.36	−1.20	−1.39
During competitor's promotional sales	−1.39	−1.33	−1.59

Note: Price elasticities during competitor sales and not during competitor sales are calculated at the median price for each DFD range. For DFD 7–14, median prices are \$199. For DFD 15–21 and DFD 22–28, median prices are both \$179.

lower search costs than business customers. Leisure customers who have lower search costs will spend more time looking for (and finding) lower fares. The price elasticities show that customers who book on Mondays and Tuesdays are the least price sensitive. This is a particularly interesting result, as it suggests that the success of an airline's promotion depends on which day of the week customers learn about the sale. An email promotion sent on a Monday will stimulate less of a demand response than an email promotion sent on a Saturday (assuming the number of potential customers who read the email is identical across the 2 days). Similarly, differences in price sensitivities will impact demand when a promotion ends. If a promotion ends on the weekend and prices are increased, more demand will be lost than if the promotion ends on a Monday (assuming the same number of potential customers are searching for flights).

4.3.3. Price elasticities during competitor sales and promotions

Price elasticities were also calculated during dates where Virgin America was offering promotional sales. *A priori*, we had two competing theories about how JetBlue's customers would react to competitor sales. On one hand, during Virgin America's promotional sales, JetBlue's most price sensitive customers could choose to book on Virgin America, leaving the more price insensitive (and more brand loyal) customers with JetBlue. In the model, this would show that customers were *less price sensitive* during Virgin America's promotional sales. On the other hand, Virgin America's promotion may have stimulated a significant number of price sensitive customers in the market, some of whom chose to purchase low fares available from JetBlue. In the model, this would show that customers became *more price sensitive* during Virgin America's promotional sale.

Our demand model shows a significant and negative coefficient on the indicator variable for Virgin America's promotional sales dates. Price elasticities were calculated during these dates versus all other dates, as a function of days before departure (DFD). Virgin America did not offer promotions for any flights that were departing in less than 7 days. Therefore, price elasticities were calculated for three DFD ranges: 7–14, 15–21, and 22–28. Table 9 shows JetBlue's price elasticities during Virgin America's promotional sales dates, as compared to all other dates, for each DFD range. The table shows that JetBlue's customers are *more price sensitive* during Virgin America's promotional sales (consistent with our second theory). For example, for DFD 22–28, price elasticities are −1.59 during Virgin America's promotional sales and –1.39 during all other dates. This is interpreted in the following way: a 10% increase in JetBlue's fares during Virgin America's promotional sales dates would lead to a 15.9% decrease in demand. However, during all other dates, a 10% increase in JetBlue's fares would lead to a 13.9% decrease in demand. The table further shows that JetBlue's customers are less price sensitive to Virgin America's promotional sales for flights that are closer to departing, as seen for price elasticities for DFD in the range of 7–14 days.

Athough we can observe that JetBlue demand decreased during Virgin America's promotional sales, we are not able to observe whether these passengers decided to book on Virgin America, or whether they just shopped more strategically by waiting for JetBlue to lower prices at some point in the future. At any rate, there does appear to be evidence that customers in the market during Virgin America's promotion were more price sensitive.

5. Study limitations and robustness of results

All studies have limitations and ours is no different. Our analysis database is relatively small and includes only four markets across 21 departure dates. However, a comparison of our price elasticity estimates with those reported in prior studies suggests that our estimates are reasonable. The magnitudes of our price elasticity estimates are also consistent with those reported in Granados et al. (2012), which used booking data. Thus, while it would be desirable to repeat our study on a larger dataset, we expect that the directional results will hold.[14]

A second limitation is that more than 25% of our observations are missing both price and demand information. We used an inverse probability weighting methodology to correct for the missing observations and conducted a sensitivity analysis by comparing estimates of models that performed regressions on the complete observations without weighting. The results were quite similar and all interpretations on the price elasticity estimates as a function of the explanatory variables remained the same. Instrumental variable methods allow consistent parameter estimation in the presence of an endogenous variable, given that the model is correctly specified. Due to missing data in our dataset, our parameter estimates may not reach full consistency. However, the price elasticity estimates provide directional results and are robust to the method we used to account for missing data.

A modeling limitation is heterogeneity that we cannot account for; we do not observe individual customer information and cannot therefore identify the trip purpose of each booking. Within the data and model we expect that there is heterogeneity between business and leisure travel that we cannot completely account for. However, including advance purchase variables helps control for this to some extent, as business passengers often book much later than leisure travelers. There may also be heterogeneity present across bookings made through different channels (such as online vs. offline bookings). This issue is less important in the context of our study, as the majority of JetBlue bookings (77%) are made through its website; an additional 13% is made through global distribution systems and the remaining 10% is made through JetBlue's call center (JetBlue, 2009). However, modeling price elasticities by distribution channels will be an important consideration for researchers and airlines that rely more heavily on multiple distribution channels.

Despite these limitations, the dataset is unique in that it provides daily airline prices and demand data at the flight-level, which is a level of detail that is not available in public datasets. Consequently, the approach we present here for calculating flight-level price elasticities can be used in many new decision-support contexts.

6. Example applications

Our data provides the first insights into price elasticities as a function of advance booking, departure day of week, departure time of day, booking day of week, and during promotional sales of a low cost competitor. These new insights into price elasticities should be of interest to airline pricing and flight scheduling analysts, policymakers at airports, and to researchers from economics, marketing, and revenue management. In this section, we illustrate two ways in which airlines can use our model to support their pricing decisions.

6.1. When to launch a promotion

Our results show that price elasticities vary not only as a function of which days customers want to travel (defined as the flight departure day of week), but which day of the week that customers purchase their tickets (defined as the booking day of week). Customers who book on Saturdays and Sundays are more price sensitive than customers who book their flights on weekdays. As noted earlier, this is a particularly interesting result, as it suggests the success of an airline promotion may depend on which day of the week the promotion is launched, and which days of the week customers are allowed to purchase tickets.

Some airlines prefer to launch promotions on weekends because they can evaluate the success of the promotion on days in which the total number of customers in the market is smaller. This is a conservative philosophy, and one that is designed to "test the market" at a time the promotion will cause the "least potential damage" on profits if customers' responses to the sale are stronger than predicted. However, as our model results show, this may not be the best strategy, as customers purchasing on weekends are not representative of the total market. Evaluating the success of a promotion based on weekend sales may result in over-estimates of stimulated demand, particularly if the promotion is continued beyond the weekend and available for customers to purchase on weekdays.

Understanding how the number of potential customers and their price sensitivities vary across different booking days is critical to deciding when to launch a promotion. In some situations, it may be optimal to allow customers to purchase the promotion on all days of the week. In other situations, it may be optimal to only allow customers to purchase the promotional fares on the weekends. Fig. 2 shows how the percent change in JetBlue's revenues associated with a 5% decrease in

[14] Although the estimated elasticities are similar to those reported in similar applications, it would be ideal to perform other validation checks by transforming our disaggregate dataset into an aggregate dataset and then estimating price elasticities on the aggregate data. We explored market-level aggregations of the price elasticities during our data analysis, but were unable to estimate acceptable aggregate elasticities due to the small sample sizes that resulted after the data was aggregated. Future researchers may want to consider using a larger sample of markets and departure dates in order to estimate aggregate elasticities as additional validation checks.

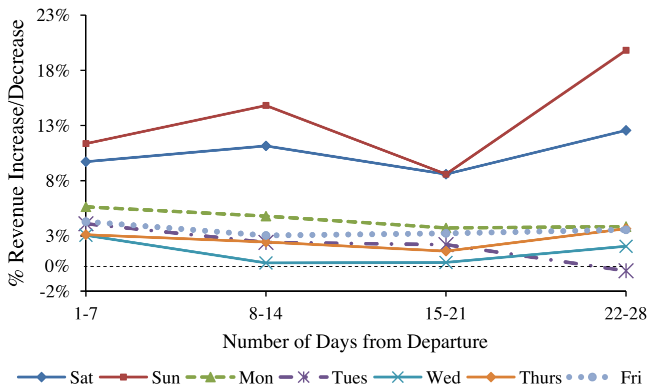

Fig. 2. Percentage change in JetBlue revenue by booking day of week for a 5% price decrease.

its fares varies by booking day of week and advance purchase dates. The model predicts that JetBlue could increase revenues by decreasing prices across all days from departure and booking days of the week, except for customers purchasing tickets on Tuesdays for flights that depart more than 21 days in advance.

This figure does not take into account competitive pricing responses and capacity constraints. Determining the "optimal" promotional design is a much more complex problem that should account for these and other factors. Our intent is not to suggest that JetBlue's revenues would increase if it simply lowered all of its fares. Our intent is to illustrate how the elasticity estimates derived from flight-level data can be used as inputs to more complex optimization models that can jointly optimize pricing and allocation decisions.

6.2. Determining when to match a competitor's promotional sale fares

Earlier, we noted that Virgin America's promotions appear to have stimulated a significant number of price sensitive customers in the market. Our model results showed that customers purchasing tickets during Virgin America's promotional sales were more price sensitive.

Fig. 3 shows the percent change in revenue for JetBlue when it lowers its average fares by 10%. The percent change in JetBlue's revenue differs as a function of days from departure and whether Virgin America is offering a promotion. For fares within 14 days from departure, our model predicts that JetBlue could increase its revenues if it lowered its fares by 10%; however, these revenues are diminished slightly if JetBlue lowers its fares when Virgin America is offering a promotion. However, the opposite result is observed for fares associated with flights departing 15–21 days from departure. In this case, our model predicts that JetBlue could have substantially increased its revenues by lowering its fares when Virgin America offered its promotion – more so than if JetBlue had simply lowered its fares during periods in which Virgin America was not offering a sale. During Virgin America's promotional sales for flights departing 15–21 days from departure, our model predicts that a 10% decrease in JetBlue prices would lead to a revenue increase for JetBlue of 10.8%, whereas the same decrease in price during other dates will lead to a revenue increase of only 4.8%. It appears that Virgin America's sales

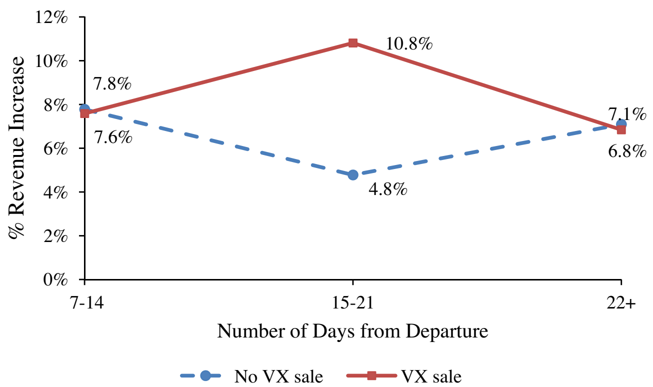

Fig. 3. Percentage change in revenue when JetBlue decreases price by 10% during a competitor's sale.

and promotional expenditures stimulated demand. By lowering its prices during this period, JetBlue could have benefitted from Virgin's promotional expenditures and captured part of this stimulated demand.

Of course, JetBlue's decision to match (or not match) Virgin America's fares will be influenced by many other factors, most notably its forecasts of how many high-paying customers it expects to make purchases close to the flight departure. Our example highlights another way in which our model, which predicts price elasticities as a function of booking, flight, and competitor promotions, can be used as inputs into optimization models that simultaneously recommend pricing, demand, and product allocation decisions.

7. Conclusions and future research directions

This study uses a dataset of online prices and seat maps to build disaggregate models of flight-level demand for JetBlue air travel. An instrumental variable approach is used to control for price endogeneity, and a set of unique instruments were found that passed the necessary diagnostic tests. These instruments can be used in future disaggregate air travel models of demand. OLS and 2SLS models are compared and demonstrate the importance of correcting for endogeneity. After correcting for endogeneity the price coefficient is found to be 2.7 times more negative than the price coefficient of an uncorrected model. Additionally, without correcting for endogeneity, the estimated mean price elasticity of demand is -0.75, which represents inelastic demand. After correcting for endogeneity, the estimated price elasticity of demand is -1.97, which represents elastic demand. We also find that price elasticity estimates vary as a function of advance booking, departure day of week, departure time of day, booking day of week, and promotional sales dates of a competitor.

We show how these detailed price elasticity estimates can be used by airlines to better design promotions by considering not only which departure dates should be offered for sale, but also what days of the week the promotion should be available for purchase. We also show how these detailed price elasticity estimates can be used by airlines to decide whether to match a competitor's promotional sale fares and/or which subset of a competitor's fares to match. Our approach is not restricted to airline applications, and can help support evaluation of proposed airport fees and taxes, national departure and emission taxes, landing fees, and congestion pricing policies.

To the best of our knowledge, this is the first time that flight-level elasticities have been estimated using online data. Although the data used in this study represents only four markets and 21 departure dates, the elasticity estimates are intuitive and consistent with (meaning more negative than) route-level and market-level elasticities reported in the literature. This is encouraging from an implementation perspective, as it suggests that airlines and policy makers may be able to accurately estimate price elasticities for network-level demand. Historically, one of the key challenges associated with forecasting network-level demand is that many of the origin–destination paths observe very few bookings. However, in our application, we were able to obtain robust results using less than 7500 bookings by grouping four similar transcontinental markets. We expect that by grouping origin–destination pairs with thin demand that share similar characteristics (such as the same origin and destination time zones and/or similar booking profiles), it will be possible to extend our methodology to estimate prices for markets that are served by connecting flights.

Our paper provides a framework that other researchers can use to estimate flight-level elasticities, by illustrating how methodological challenges related to missing data and price endogeneity can be addressed. The ability to estimate flight-level elasticities allows researchers to forecast demand as a function of day-to-day variation in price. This is a critical first step toward the research community's goal of developing the next-generation of revenue management systems that seek to forecast demand as a function of price, determine which prices to offer in the market, and how many seats to sell at each price.

Importantly, our paper provides some of the first evidence that when developing models to forecast demand as a function of price, prices for the carrier of interest as well as for competitors should be considered. That is, this is the first paper we are aware of that shows how a carrier's bookings are impacted by a competitor's promotional sale, demonstrating that customers are more price sensitive during a competitor promotion. Controlling for competitor promotions when using historical data to estimate price elasticities is important in order to obtain accurate price elasticity estimates.

References

Arrington, M., 2010. Virgin America rides Loopt taco truck special to fifth largest revenue day ever. TechCrunch. <http://techcrunch.com/2010/09/02/virgin-america-rides-loopt-taco-truck-special-to-fifth-largest-revenue-day-ever/> (accessed 07.05.11).

ATPCo, 2013a. Data distribution and subscriptions. <http://www.atpco.net/atpco/products/prodserv_dd.shtml> (accessed 07.08.13).

ATPCo, 2013b. Life cycle of the fare. <http://www.atpco.net/atpco/products/lfcFare.shtml> (accessed 07.28.13).

Bartlett, J., 2012. Handling missing data in Stata – a whirlwind tour. 2012 Italian Stata Users Group Meeting. <http://www.stata.com/meeting/italy12/abstracts/materials/it12_bartlett.pdf> (accessed 07.05.11).

Baum, C.F., Schaffer, M.E., Stillman, S., 2003. Instrumental variables and GMM: estimation and testing. Stata J. 1(3), 1–31.<http://www.stata-journal.com/article.html?article=st0030> (accessed 01.02.14).

Baum, C.F., Schaffer, M.E., Stillman, S., 2007. Enhanced routines for instrumental variables/GMM estimation and testing. Stata J. 4(7), 465–506. <http://www.stata-journal.com/sjpdf.html?articlenum=st0030_3> (accessed 01.02.14).

Baum, C.F., Schaffer, M.E., Stillman, S., 2010a. Ivreg2: Stata module for extended instrumental variables/2SLS, GMM and AC/HAC, LIML and k-class regression. A software component provided by Boston College, Department of Economics, Statistical Software Components series, S425401. <http://ideas.repec.org/c/boc/bocode/s425401.html> (accessed 01.02.14.

Baum, C.F., Schaffer, M.E., Stillman, S., 2010b. Help documentation for ivreg2. <http://repec.org/bocode/i/ivreg2.html> (accessed 01.02.14).

Berry, S., Jia, P., 2010. Tracing the woes: an empirical analysis of the airline industry. Am. Econ. J. Microecon. 2 (3), 1–43.

Berry, S., Levinsohn, J., Pakes, A., 1995. Automobile prices in market equilibrium. Econometrica 63 (4), 841–890.
Berry, S., Levinsohn, J., Pakes, A., 2004. Differentiated products demand systems from a combination of micro and macro data: the new car market. J. Polit. Econ. 112 (1), 68–105.
Bhadra, D., 2003. Demand for air travel in the United States: bottom-up econometric estimation and implications for forecasts by origin and destination pairs. J. Air Transport. 8 (2), 19–56.
Bilotkach, V., 2006. Understanding price dispersion in the airline industry: capacity constraints and consumer heterogeneity. In: Lee, Darin (Ed.), Advances in Airline Economics, Competition Policy and Antitrust, vol. 1. Elsevier Science, New York, pp. 329–345.
Bilotkach, V., Pejcinovska, M., 2012. Distribution of airline tickets: a tale of two market structures. In: Peoples, James H., Jr. (Ed.), Advances in Airline Economics, Pricing Behavior and Non-Price Characteristics in the Airline Industry, vol. 3. Emerald Group Publishing Limited, Bingley, West Yorkshire, England, pp. 107–138.
Bilotkach, V., Talavera, O., Gorodnichenko, Y., Zubenko, I., 2010. Are airlines' price setting strategies different? J. Air Transp. Manage. 16 (1), 1–6.
Branstetter, L., Chatterjee, C., Higgins, M.J., 2011. Regulation and welfare: evidence from paragraph-IV generic entry in the pharmaceutical industry. Working paper, Carnegie Mellon University and Georgia Institute of Technology.
Bresnahan, T., 1997. The Apple-Cinnamon Cheerios war: valuing new goods, identifying market power, and economic measurement. Unpublished paper, Department of Economics, Stanford University. <http://www.stanford.edu/~tbres/Unpublished_Papers/hausman%20recomment.pdf> (accessed 05.04.13).
Button, K.J., Vega, H., 2006. Airlines competing with themselves: a note on the temporal pattern of fare setting prior to departure. Int. J. Transp. Econ. 31 (3), 341–350.
Button, K.J., Vega, H., 2007. The uses of the "temporal-fares-offered curve" in air transportation. J. Transport. Res. Forum 46, 83–99.
Carpenter, J.R., Kenward, M.G., Vansteelandt, S., 2006. A comparison of multiple imputation and inverse probability weighting for analyses with missing data. J. R. Stat. Soc. Ser. A Stat. Soc. 169, 571–584.
Gayle, P.G., 2004. Does price matter? Price and non-price competition in the airline industry. No. 163, Econometric Society 2004 North American Summer Meetings, Econometric Society. http://EconPapers.repec.org/RePEc:ecm:nasm04:163.
Gillen, D.W., Morrison, W.G., Stewart, C., 2002. Air travel demand elasticities: concepts, issues and measurement. Final Report, Department of Finance, Canada. <http://www.fin.gc.ca/consultresp/airtravel/airtravstdy_-eng.asp> (accessed 05.20.13).
Goolsbee, A., Petrin, A., 2004. The consumer gains from direct broadcast satellites and the competition with cable TV. Econometrica 72 (2), 351–381.
Granados, N., Gupta, A., Kauffman, R.J., 2012. Online and offline demand and price elasticities: evidence from the air travel industry. Inform. Syst. Res. INFORMS 23 (1), 164–181.
Greene, W.H., 2003. Econometric Analysis Rod Banister, 5th ed. Prentice Hall, Upper Saddle River, New Jersey.
Guevara-Cue, C.A., 2010. Endogeneity and sampling of alternatives in spatial choice models. Dissertation for Doctor of Philosophy, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology.
Guevara, C.A., Ben-Akiva, M., 2006. Endogeneity in residential location choice models. Transport. Res. Rec. J. Transport. Res. Board 1977, 60–66.
Hansen, L.P., 1982. Large sample properties of generalized method of moments estimators. Econometrica 50 (4), 1029–1054.
Harteveldt, H.H., 2012. The future of airline distribution: a look ahead to 2017. Special Report, Atmosphere Research Group LLC. <www.iata.org/whatwedo/stb/Documents/future-airline-distribution-report.pdf> (accessed 07.20.13).
Hausman, J.A., 1996. Valuation of new goods under perfect and imperfect competition. In: Gordon, Robert J., Bresnahan, Timothy F. (Eds.), The Economics of New Goods. University of Chicago Press, Chicago, pp. 207–248.
Hausman, J., Leonard, G., Zona, J.D., 1994. Competitive analysis with differentiated products. Ann. Econ. Stat. 34, 159–180.
Horner, M.W., Grubesic, T.H., Zook, M.A., Leinbach, T.R., 2006. Global distribution systems and U.S. commercial air industry: gathering real-time airline flight and fare information for spatial and economic analysis. 85th Annual Meeting of the Transportation Research Board Compendium of Papers CR-ROM, Washington, DC.
Hsiao, C-.Y., 2008. Passenger demand for air transportation in a hub-and-spoke network. Dissertation for Doctor of Philosophy, Civil and Environmental Engineering, University of California, Berkeley. <http://www.nextor.org/pubs/HsiaoDissertation2008.pdf> (accessed 05.18.11.
International Air Transport Association, 2008. Economics briefing no. 9: air travel demand <http://www.iata.org/whatwedo/Documents/economics/air_travel_demand.pdf> (accessed 07.15.13).
InterVISTAS, 2007. Estimating air travel demand elasticities: final report. <http://www.iata.org/whatwedo/Documents/economics/Intervistas_Elasticity_Study_2007.pdf> (accessed 05.20.13).
JetBlue, 2009. JetBlue Airways 10-K Filing on February 13, 2009. <http://www.wikinvest.com/stock/JetBlue_Airways_%28JBLU%29/Marketing%20Distribution> (accessed 07.27.13).
Li, J., Granados, N., Netessine, S., 2011. Are consumers strategic? Structural estimation from the air-travel industry. INSEAD Working Paper Collection. <https://flora.insead.edu/fichiersti_wp/inseadwp2011/2011-104.pdf> (accessed 07.15.13).
McAfee, R.P., Vera, V., 2007. Dynamic pricing in the airline industry. In: Hendershott, T.J. (Ed.), Economics and Information Systems, Handbooks in Information Systems, vol. 1. Emerald Group Publishing Limited, Bingley, West Yorkshire, England, pp. 527–567.
Mentzer, M.S., 2000. The impact of discount airlines on domestic fares in Canada. Transport. J. 39 (4), 35–42.
Mumbower, S., Garrow, L.A., 2014. Online pricing data for multiple U.S. carriers. Manufact. Serv. Operat. Manage. Art. Adv., 1–6.
Mumbower, S., Garrow, L.A., Newman, J., 2013. Investigating airline customers' premium coach seat purchases and implications for optimal pricing strategies. Working paper, Georgia Institute of Technology.
Newman, J.P., Ferguson, M.E., Garrow, L.A., 2013. Estimating GEV models with censored data. Transport. Res. Part B Methodol. 58, 170–184.
Nevo, A., 2000a. A practitioner's guide to estimation of random-coefficients logit models of demand. J. Econ. Manage. Strat. 9 (4), 513–548.
Nevo, A., 2000b. Mergers with differentiated products: the case of the ready-to-eat cereal industry. Rand J. Econ. 31 (3), 395–421.
Nevo, A., 2001. Measuring market power in the ready-to-eat cereal industry. Econometrica 69 (2), 307–342.
Pekgün, P., Griffin, P.M., Keskinocak, P., 2013. An empirical study for estimating price elasticities in the travel industry. Working Paper, University of South Carolina.
Pels, E., Rietveld, P., 2004. Airline pricing behaviour in the London–Paris market. J. Air Transp. Manage. 10 (4), 277–281.
Petrin, A., Train, K., 2010. A control function approach to endogeneity in consumer choice models. J. Mark. Res. 47 (1), 3–13.
Pitfield, D.E., 2008. Some insights into competition between low-cost airlines. Res. Transport. Econ. 24 (1), 5–14.
Rivers, D., Vuong, Q.H., 1988. Limited information estimators and exogeneity tests for simultaneous probit models. J. Econometr. 39, 347–366.
Southwest Airlines, 2009. Southwest Airlines 10-k Filing, Part 2, Item 6.
Southwest Airlines, 2010. Southwest Airlines Fun Facts. <http://www.southwest.com/about_swa/press/factsheet.html#Financial%20Statistics> (accessed 05.17.10).
StataCorp, 2007. Stata Statistical Software: Release 10. College Station, StataCorp LP, TX.
Stern, S., 1996. Market definition and the returns to innovation: Substitution patterns in pharmaceutical markets. Working paper, Sloan School of Management, Massachusetts Institute of Technology.
Stock, J.H., Yogo, M., 2005. Testing for weak instruments in linear IV regression. identification and inference for econometric models: essays in Honor of Thomas Rothenberg. In: Andrews, D.W.K., Stock, J.H. (Eds.), Cambridge University Press, New York, NY, pp. 80–108. Working paper version: <http://scholar.harvard.edu/files/stock/files/testing_for_weak_instruments_in_linear_iv_regression.pdf> (accessed 01.02.14).
Train, K., 2009. Endogeneity. In: Discrete Choice Methods with Simulation. Cambridge University Press, New York, NY, pp. 315–346, Chapter 13.
Train, K.E., Winston, C., 2007. Vehicle choice behavior and the declining market share of U.S. automakers. Int. Econ. Rev. 48 (4), 1469–1496.
Villas-Boas, J.M., Winer, R.S., 1999. Endogeneity in brand choice models. Manage. Sci. 45 (10), 1324–1338.

Vinod, B., 2010. The complexities and challenges of the airline fare management process and alignment with revenue management. J. Reven. Pric. Manage. 9 (1–2), 137–151.
Virgin America, 2010. Virgin America launches new service to Los Cabos and Cancun. Press Release. <http://www.virginamerica.com/press-release/2010/virgin-america-launches-new-service-to-los-cabos-and-cancun.html> (accessed 07.05.11).
U.S. Department of Transportation, Bureau of Transportation Statistics, 2010. Origin and Destination Survey Databank 1A/1B. <www.transtats.bts.gov> (accessed 05.16.11).