

DYNAMIC AIRLINE PRICING AND SEAT AVAILABILITY

By

Kevin R. Williams

August 2017

Updated February 2018

COWLES FOUNDATION DISCUSSION PAPER NO. 3003U



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY

Box 208281

New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

DYNAMIC AIRLINE PRICING AND SEAT AVAILABILITY

Kevin R. Williams
School of Management
Yale University*

February 2018[†]

Abstract

Airfares are determined by both intertemporal price discrimination and dynamic adjustment to stochastic demand. I estimate a model of dynamic airline pricing accounting for both forces with new flight-level data. With model estimates, I disentangle key interactions between the arrival pattern of consumer types and remaining capacity under stochastic demand. I show that the forces are complements in airline markets and lead to significantly higher revenues, as well as increased consumer surplus, compared to a more restrictive pricing regime. Finally, I show that abstracting from stochastic demand leads to a systematic bias in estimating demand elasticities.

JEL: L11, L12, L93

*kevin.williams@yale.edu

[†]This paper is an updated chapter of my dissertation at the University of Minnesota. I thank Tom Holmes, Jim Dana, Aniko Öry, Amil Petrin, Tom Quan and Joel Waldfogel for comments. I thank the seminar participants at the Federal Reserve Bank of Minneapolis, Yale School of Management, University of Chicago - Booth, Georgetown University, University of British Columbia, University of Rochester - Simon, Dartmouth University, Northwestern University - Kellogg, Federal Reserve Board, Federal Reserve Bank of Richmond, Indiana University, Indiana University - Kelley, Marketing Science, and the Stanford Institute for Theoretical Economics (SITE). I also thank the Minnesota Supercomputing Institute (MSI) for providing computational resources.

1 Introduction

Air Asia (2013) on intertemporal price discrimination:

*Want cheap fares, book early. If you book your tickets late, chances are you are desperate to fly and therefore don't mind paying a little more.*¹

easyJet (2003) on dynamic adjustment to stochastic demand:

*Our booking system continually reviews bookings for all future flights and tries to predict how popular each flight is likely to be. If the rate at which seats are selling is higher than normal, then the price would go up. This way we avoid the undesirable situation of selling out popular flights months in advance.*²

Airlines tend to charge high prices to passengers who search for tickets close to their date of travel. The conventional view is that these are business travelers and that airlines capture their high willingness to pay through intertemporal price discrimination. Airlines also adjust prices on a day-to-day basis, as capacity is limited and the future demand for any given flight is uncertain. They may adjust fares upward to avoid selling out flights in advance, or fares may actually fall from one day to the next, after a sequence of low demand realizations.

This paper examines pricing in the airline industry, taking into account both forces: intertemporal price discrimination (fares responding to time) and dynamic adjustment to stochastic demand (fares responding to seats sold). Collectively, I call this dynamic airline pricing. I use a new flight-level data set consisting of daily fares and seat availabilities to estimate a model of dynamic airline pricing in which firms face a stochastic arrival of consumers. The mix of consumer types – business and leisure travelers – is allowed to change over time, and in the estimated model, late-arriving consumers are significantly more price-inelastic than early-arriving consumers. I use the model to quantify the welfare effects of dynamic airline pricing, but also to establish important interactions between intertemporal price discrimination and dynamic adjustment to stochastic demand in airline

¹Accessed through AirAsia.com's Investor Relations page entitled, "What is low cost?"

²Appeared on easyjet.com's FAQs. Accessed through "Low-Cost Carriers and Low Fares" and "Online Marketing: A Customer-Led Approach."

markets. Finally, I show that analyzing both forces together is critical to quantifying the welfare implications of airline pricing.

The existing research documents the importance of intertemporal price discrimination and dynamic adjustment to stochastic demand separately in airline markets, and the central contribution of this paper is to study them jointly and quantify their interactions. Consistent with the idea of market segmentation, Puller, Sengupta, and Wiggins (2015) find that ticket characteristics, such as advance purchase discount (APD) requirements, explain much of the dispersion in fares. Lazarev (2013) quantifies the welfare effects of airline pricing by estimating a model of intertemporal price discrimination but not dynamic adjustment. He finds a substantial role for this force.³ Escobari (2012) and Alderighi, Nicolini, and Piga (2015) find evidence that airlines face stochastic demand and that prices respond to remaining capacity. These results support the theoretical predictions of Gallego and Van Ryzin (1994) and a large branch of research in operations management that studies optimal pricing under uncertain demand, limited capacity, and limited time to sell.⁴ Indeed, this research has been used to inform airline pricing algorithms.

The estimated structural model allows me establish two key points about the interaction between the pricing forces. First, dynamic adjustment complements intertemporal price discrimination in the airline industry. This is due to the fact that price-inelastic consumers (i.e., business travelers) tend to buy tickets close to the departure date. To be able to price discriminate towards these late-arriving consumers, airlines must successfully save seats until close to their departure date. These consumers are then charged high prices. I show that dynamic pricing more efficiently allocates capacity and increases consumer welfare in the monopoly markets I study relative to either uniform pricing or a pricing

³Intertemporal price discrimination can be found in many markets, including video games (Nair 2007), Broadway theater (Leslie 2004), and concerts (Courty and Pagliero 2012). Lambrecht et. al. (2012) provide an overview of empirical work on price discrimination more broadly.

⁴An overview of the dynamic pricing literature can be found in Elmaghraby and Keskinocak (2003) and Talluri and Van Ryzin (2005). Sweeting (2012) analyzes ticket resale markets. Pashigian and Bowen (1991) and Soysal and Krishnamurthi (2012) study clearance sales and seasonal goods, respectively. Zhao and Zheng (2000), Su (2007), and Dilme and Li (2016) discuss extensions to dynamic pricing models, including consumer dynamics. Seat inventory control has also been studied; see Dana (1999).

system in which fares vary with only the purchase time but not with past demand shocks. Second, I show that in order to quantify the welfare effects of price discrimination in airline markets, it is necessary to take stochastic demand into account. When abstracting from stochastic demand, the opportunity cost of selling a seat is the same regardless of the date of purchase. In reality, the opportunity costs tend to increase in airline markets because of the late arrival of price insensitive consumers. As a consequence, I show an empirical procedure abstracting from this feature of the market will systematically overstate consumers' price insensitivity. The bias is large close to the departure date.

In order to investigate dynamic airline pricing, a detailed data set of ticket purchases is required. However, the standard airline data sets used in economic studies (e.g., Goolsbee and Syverson (2008); Gerardi and Shapiro (2009); Berry and Jia (2010)) are either at the monthly or the quarterly level. Recent papers use new data to get to the flight level. McAfee and Te Velde (2006) and Lazarev (2013) create data sets containing high-frequency fares. Other papers have obtained high frequency fares as well as a measure of seats sold. Puller, Sengupta, and Wiggins (2015) capture a fraction of sales through a single computer reservation system used by airlines. Escobari (2012) and Clark and Vincent (2012) collect fare and flight availability data, with the available number of seats derived from publicly available seat maps. I use a new data source that allows me to capture the same information that travel agents see, such as which seats are blocked, occupied, and available. I merge these data with daily fares and flight availability data. In total, I track over 1,300 flights in US monopoly markets over a six-month period. I use the data to provide descriptive evidence of both intertemporal price discrimination and dynamic adjustment to stochastic demand.

While empirical evidence is informative, it cannot be used to disentangle the interactions between the pricing forces. I proceed by estimating a structural model that contains three key ingredients: (i) a monopolist has fixed capacity and finite time to sell; (ii) the firm faces a stochastic arrival of consumers; and (iii) the mix of consumers, correspond-

ing to business and leisure travelers, is allowed to change over time.⁵ The firm solves a stochastic dynamic programming problem, offering a fare to consumers each day before departure. For the demand system, I assume that a stochastic process brings consumers to the market. The consumers that arrive know when they want to travel and solve a static discrete choice problem.⁶ The demand model differs from earlier theoretical work, including Gale and Holmes (1993), and from empirical work such as Lazarev (2013), in which consumers are uncertain, and waiting provides more information. In my model the only reason to wait is to bet on price, and since prices tend to increase, I show that only a small transaction cost is needed to persuade consumers to decide whether to travel in the period they arrive. I also provide empirical evidence that suggests this is a reasonable assumption.⁷ I use the firm's pricing decision, along with the realizations of demand, to separately identify the arrival process of consumers and the demand elasticity over time.

Using the model estimates, I conduct a series of counterfactual exercises to quantify the welfare implications of dynamic airline pricing and to highlight how these results depend on the arrival process of consumers. The general finding is that, while more-restrictive pricing systems reduce the ability of firms to price discriminate, the aggregate consumer welfare gains are limited, or even negative, due to inefficient capacity allocation. That is, while airlines do price discriminate, which we may think only harms consumers, their pricing also more efficiently allocates scarce seats. In fact, under uniform pricing, revenues drop by 2%, but aggregate consumer welfare is unchanged. In this counterfactual, leisure consumers are charged relatively higher prices, and while business consumers are charged relatively lower prices, many flights sell out in advance. Moreover, I find total consumer welfare is lowered when fares are allowed to respond to the date of purchase but not to the scarcity of seats. However, the pricing forces interact differently under alternative arrival

⁵McAfee and Te Velde (2006) note that stochastic demand models that do not incorporate changes in willingness to pay over time do not match the positive trend in airfares as the departure date approaches.

⁶The specification is similar to that in Talluri and Van Ryzin (2004) and Vulcano, van Ryzin, and Chaar (2010); however, my model also allows for the mix of consumers to change across time.

⁷Li, Granados, and Netessine (2014) study dynamic consumer behavior in airline markets. Depending on the specification, they find that between 5% and 20% of consumers are dynamic.

processes. If price-insensitive consumers were the first to arrive, for example, dynamic adjustment no longer complements intertemporal price discrimination, and stochastic demand pricing is used to clear excess capacity. This is consistent with fire sales that are commonly used in retailing.

2 Data

I create an original data set of airfares and seat availabilities with data collected from two popular online travel services. The first is a travel search engine, owned by Kayak.com, from which I collect daily fares at the itinerary level, with itinerary defined as a routing, airline, flight number(s), and departure date(s) combination. I obtain all one-way and round-trip itinerary fares for stays of less than eight days. The fares recorded correspond to the cheapest ticket available for purchase. The fare data include the fare class and fare basis code, which provide some information regarding the restrictions of the ticket, such as advance purchase discount requirements.⁸

The second web service, operated by Expertflyer.com, allows users to look up flight availability. I collect two sources of information. First, I obtain airline seat maps for each flight. By comparing seat maps across time, I obtain daily bookings. Second, I collect the fare availability (sometimes called fare buckets) for each flight. This information provides censored information regarding which fares are available to purchase. For example, Y9 indicates that the Y class fare, which is typically the most expensive coach fare available, has at least nine tickets available for purchase.⁹ In total, the sample contains 1,362 flight departures, with each flight tracked for 60 days. All flights in my data set departed

⁸G21JN5 is an example. This fare basis code indicates a G fare class ticket, and a 21-day advance purchase discount. Lazarev (2013) has a nice discussion of airline pricing and provides additional details on fare restrictions. Like Lazarev (2013), I model only the pricing choice and not the decision to assign restrictions or quantities for each price.

⁹Regarding the example of the fare basis code G21JN5, the fare availability data might say G5. This means five seats remain for that particular fare. Fare availabilities can change because of ticket purchases, but also because of changes input by managers. Both of these potentially result in changes in price, which is the focus of this paper.

between March 2012 and September 2012.

In the following subsections, I discuss route selection (Section 2.1) and the use of seat maps to infer bookings (Section 2.2), and I provide summary statistics and preliminary evidence from the data (Section 2.3).

2.1 Route Selection

I select markets in which to study using the publicly available DB1B tables, which are frequently used to study airline markets. The DB1B tables contain 10% of domestic ticket purchases and are at the quarterly level. The data contain neither the date flown nor the purchase date. I define a market in the DB1B as an origin, destination, quarter. I single out markets where:

- (i) there is only one carrier operating;
- (ii) there is no nearby alternative airport;
- (iii) at least 95% of flight traffic is not connecting to other cities;
- (iv) total quarterly traffic is greater than 3,000 passengers;
- (v) total quarterly traffic is less than 30,000 passengers; and
- (vi) there is high nonstop traffic.

Criteria (i) and (ii) narrow the focus to monopoly markets. One potential complication with using airline seat maps to recover bookings is that it is not clear what fare to assign to an observed change in a seat map. Since airlines offer extensive networks, the disappearance of a single seat could be associated with one of several thousand possible itineraries. This is an important consideration since pricing may be different across routes, and this is especially true with regard to feeder routes versus main routes. This issue is addressed in (iii), which essentially rules out most routes from major hubs. Criterion (iv) corresponds to routes with a less than 75% load factor (seats occupied / capacity) of a daily 50-seat aircraft. This criterion removes routes with irregular service. Criterion (v)

removes most routes with several daily frequencies. I then look for routes with (vi) high nonstop traffic. This criterion is important for establishing the relevant outside option in the demand model. In the data, (vi) is negatively correlated with distance ($\rho = -.5$). Cities with very high nonstop traffic percentages tend to be short-distance flights, and, given such short distances, many consumers may choose to drive instead. At the same time, the DB1B data suggest that longer flights have lower nonstop traffic percentages, suggesting that more consumers choose a connecting flight option.

I select four city pairs, or eight directional routings, given the selection criteria above. All directional routings either originate or end in Boston, MA. The other cities are: San Diego, CA; Austin, TX; Kansas City, MO; and Jacksonville, FL. The selected markets have close to 100% direct traffic, meaning that very few passengers connect to other cities. The percent of nonstop traffic ranges between 40% and 70%. Three of the four city pairs are operated by JetBlue and the other by Delta Air Lines.¹⁰ The selected markets are all low-frequency, with at most two daily frequencies.

Three other features of the data are worth noting. First, JetBlue prices itineraries at the segment level; that is, consumers wishing to purchase round-trip tickets on this carrier purchase two one-way tickets. As a consequence, round-trip fares in these markets are exactly equal to the sum of the corresponding one-way fares. Since fares must be attributed to each seat map change, this feature of the data makes it easier to justify the fare involved. Second, JetBlue does not oversell flights, while most other air carriers do.¹¹ I will use this feature of the data to simplify the pricing problem in the next section. Third, JetBlue does not offer first class on the routes studied. This allows for investigating all sales and also controls for one aspect of versioning (first class versus economy class).¹²

¹⁰At the time of data collection, flights between Kansas City and Boston were operated by regional carriers on behalf of Delta Air Lines. Since Delta Air Lines determines the fares for this market, I collectively call these regional carriers Delta.

¹¹In the legal section of the JetBlue website, under Passenger Service Plan: "JetBlue does not overbook flights. However some situations, such as flight cancellations and reaccommodation, might create a similar situation."

¹²The Delta first class cabins were typically six seats for the market studied.

2.2 Inference and Accuracy of Seat Maps

A seat map is a graphical representation of occupied and unoccupied seats for a given flight at a select point in time before the departure date. Many airlines that have assigned seating present seat maps to consumers during the booking process. When a consumer books a ticket and selects a seat, the seat map changes to reflect that an unoccupied seat is now occupied. The next consumer wishing to purchase a ticket on the flight is offered an updated seat map and can choose one of the remaining unoccupied seats. By differencing seat maps across time – in this case daily – inferences can be made about daily bookings. I use a new data source that allows me to distinguish between various types of occupied and unoccupied seats.

Figure 1 presents a sample seat map: occupied seats are in solid blue, while the unshaded blocks correspond to unoccupied seats. Seats with a "P" are available, but classified as premium. These seats are located toward the front of the aircraft or in exit rows. Some airlines charge a premium to be seated in these rows. Finally, the seat map indicates seats currently blocked by the airline with "X"s. Seats that are blocked are usually not disclosed on airline websites; however, I am able to capture these data through the web service used. Seats may be blocked due to crew rest, weight and balance, because a seat is broken, or because the airline reserves accessible seats for handicapped passengers until the day of departure.¹³ For every seat map collected, I aggregate the number of occupied, unoccupied, and blocked seats. I compare the aggregate counts across days to determine bookings by day prior to departure.

¹³Seat blocking may be used to encourage consumers to purchase tickets or upgrade, as they give the impression that the cabin is closer to capacity. However, the data suggest that airlines predominantly block seats in exit rows and at the front and/or back of the cabin until closer to the departure date or when bookings demand additional seats. 69% of the flights in the sample experience changes in the number of blocked seats, and while I do not model the decision to block/unblock seats, I do take this information into account when determining bookings. Knowing which seats are blocked is important because it allows me to distinguish between consumers' decisions and airlines' adjustment of the supply of seats. For example, if an airline unblocks six seats, without these data, I would erroneously conclude that six passengers canceled their tickets.





Figure 1: An example seat map. The white blocks are unoccupied seats, the blue blocks are occupied, the blocks with the X's are blocked seats, and the blocks with a "P" are premium, unoccupied seats.

Unfortunately, seat maps may not be accurate representations of true flight loads. This is especially problematic if consumers do not select seats at the time of booking. This measurement error would systematically understate sales early on, but then overstate last-minute sales when consumers without existing seat assignments are assigned seats. From a modeling perspective, this measurement error would lead to an overstatement

of the arrival of business consumers. Ideally, the severity of measurement error of my data could be assessed by matching changes in seat maps with bookings, however this is impossible with the publicly available data. While both imperfect, I perform two analyses to gauge the magnitude of the measurement error in using seat maps.

First, I match monthly enplanements using my seat maps aggregated on the the day of departure with actual monthly enplanements reported in the T100 Segment tables. These tables record the total number of monthly enplanements by airline and route. Figure 2 provides a scatter plot that compares the two statistics. All the points closely follow the 45-degree line and I find seat maps differ with true enplanements by less than two percent on average. However, what really matters is the difference between actual loads and seat maps at the flight-day before departure level. Unfortunately, this analysis cannot be performed with the collected data set.


Figure 2: Estimated Seat Map Measurement Error at the Monthly-Level



Note: Measurement error estimated by comparing monthly enplanements using the T100 Tables and aggregating seat maps to the monthly level. The solid line reflects zero measurement error. The black-squares correspond to Delta observations. The two outliers are due to a change in flight numbers which were not picked up by the data collection process in one month.

Second, I create a new data set that allows me to estimate the measurement error at the flight-day before departure level. The mobile website version of United.com allows users to examine seat maps for upcoming flights. In addition, for premium cabins, the airline also reports the number of consumers booked into the cabin. I randomly select flights, departure dates, and search dates, and, in total, I obtain 15,567 observations. With these data I find that seat maps understate reported load factor by 2.3% on average (around 1-2 seats). Figure 3 plots the average measurement error by day before departure ($t=60$ corresponds to the day that flights leave), as well as a polynomial smooth of the data. I find the difference to range between 0% to 5% across days. Two caveats are worth noting. First, I cannot reject that the measurement error is constant across days. Seat maps are most accurate far in advance (60 to 40 days out) and close to the departure date. Second, while these data do not indicate all blocked seats and the main sample does, they are for a different airline, different markets, and for a different cabin.

Figure 3: Estimated Seat Map Measurement Error by Day Before Departure



Note: Measurement error estimated by comparing seat maps with reported load factor using the United Airlines mobile website. The dots correspond to the daily mean, and the line corresponds to fitted values of an orthogonal polynomial regression of the fourth degree. Total sample size is equal to 15,567 with an average load factor of 70.7%.

2.3 Summary Statistics and Preliminary Evidence

Summary statistics for the data sample appear in Table 1. The average one-way ticket in my sample is \$283, whereas the average round-trip fare is \$528. The discrepancy in one-way and round-trip fares can be attributed to the flights operated by Delta since Delta does not price at the segment level, but JetBlue does.

Table 1: Summary Statistics for the Data Sample


Variable	Mean	Std. Dev.	Median	5th pctile	95th pctile
Oneway Fare (\$)	282.778	117.788	272.800	129.800	498.800
Load Factor	0.849	0.077	0.867	0.688	0.933
Daily Booking Rate	0.771	1.600	0.000	0.000	4.000
Daily Fare Change (\$)	3.170	36.398	0.000	-30.000	60.000
Unique Fares (per itin.)	6.793	2.196	7.000	3.000	10.000

Note: Summary statistics for 1,362 flights tracked between 3/2/2012 and 8/24/2012. The total number of observations is 79,856. Load Factor is reported between zero and one the day of departure. The daily booking rate and fare change differences across search days within flights. Finally, unique fares is at the itinerary level and denotes the number of unique price points per flight.

Reported load factor is the number of occupied seats divided by capacity on the day flights leave, and is reported between 0 and 1. In my sample, the average load factor is 85%, ranging from 77% to 89% by market. The booking rate corresponds to the mean difference in occupied seats across consecutive days. I find the average booking rate to be 0.78 seats per day, per flight. At the 5th percentile, zero seats per flight are booked a day, and at the 95th percentile, four seats per flight are booked a day. Airline markets are associated with low daily demand, as 57% of the seat maps in the sample do not change across consecutive days. The fare change rate is an indicator variable equal to one if the itinerary fare changes across consecutive days. I find the daily rate of fare changes to be 20%, so that the itineraries in my sample typically change price 12 times in 60 days. On average, each itinerary reaches 6.8 unique fares, and given that the average itinerary sees

12 fare changes, on average, this implies that fares usually fluctuate up and down usually several times within 60 days. I use the institutional feature that fares are chosen from a discrete set (fare buckets) in the model.

Figure 4: Frequency and Magnitude of Fare Changes by Day Before Departure




Note: The top panel shows the percent of itineraries which see fares increase or decrease by day before departure. The lower panel plots the magnitude of the fare declines and increases by day before departure. The vertical lines correspond to advance-purchase discount periods (fare fences).

Figure 4 shows the frequency and magnitude of fare changes across time. The top panel indicates the fraction of itineraries that experience fare hikes versus fare discounts by day before departure, and the bottom panel indicates the magnitude of these fare changes (i.e., a plot of first differences, conditional on the direction of the fare change). For example, in the left plot, 40 days prior to departure ($t=20$), roughly 10% of fares increase and 10% of fares decrease. The remaining 80% of fares are held constant. Moving to the bottom panel, the magnitude of fare increases and declines 40 days out is roughly \$50. The top panel confirms fares change throughout time, and this is also true for fare declines. Note that well before the departure date, the number of fare hikes and declines is roughly even. The fraction of itineraries that experience fare hikes increases over time.

There are four noticeable jumps in the line, indicating fare hikes. These jumps correspond to crossing three, seven, 14 and 21 days prior to departure, or when the advance purchase discounts placed on many tickets expire.¹⁴ The use of advance purchase discounts (APDs) is consistent with the story of intertemporal price discrimination. Surprisingly, though, the use of APDs is not universal. Only 20% of itineraries experience fare hikes at 21 days, and less than 50% increase at 14 days. Just under 70% of itineraries see an increase in fare when crossing the seven-day APD requirement.

Figure 5: Mean Fare and Load Factor by Day Before Departure



Note: Average fare and load factor by day before departure. Includes all oneway observations in the data sample. The vertical lines correspond to advance-purchase discount periods (fare fences).

Moving to statistics in levels, Figure 5 plots the mean fare and mean load factor (seats occupied/capacity) by day before departure. The plot confirms that the overall trend in prices is positive, with fares increasing from roughly \$225 to over \$425 in sixty days. The noticeable jumps in the fare time series occur when crossing the APD fences noted in Figure 4. At 60 days prior to departure, roughly 40% of seats are already occupied. Consequently, I observe about half the bookings on any given flight. The booking curve

¹⁴Advance purchase discounts are sometimes placed at four, ten, and 30 days prior to departure, but this is not the case for the data I collect.

for flights in the sample is smooth across time, leveling off at 85% roughly three days prior to departure. The fact that fares tend to double but that consumers still purchase tickets is suggestive evidence that consumers of different types purchase tickets towards the date of travel.

Table 2: Dynamic Substitution Regressions

	(1)	(2)	(3)
APD ₃	-0.231*** (0.042)	-0.223*** (0.042)	-0.220*** (0.042)
APD ₇	-0.028 (0.0548)	-0.030 (0.054)	-0.027 (0.054)
APD ₁₄	0.089 (0.051)	0.087 (0.050)	0.086 (0.050)
APD ₂₁	-0.012 (0.066)	-0.022 (0.065)	-0.029 (0.065)
$m(t)$	Yes	Yes	Yes
d.o.w. Flight FE	No	Yes	–
d.o.w. Search FE	No	Yes	Yes
Flight FE	No	No	Yes
Observations	79,856	79,856	79,856
R^2	0.604	0.623	0.860


Note: Flight clustered standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. $m(t)$ is a third-order polynomial in days before departure, d.o.w. stands for day-of-week indicators for the day the flight leaves and the day of search.

While Figure 5 shows that there are noticeable jumps in prices across time, it also shows there are no noticeable jumps in load factor over time. If consumers are aware that fares tend to increase sharply around APD fences, bunching would be expected in the booking curve. This is not the case, and formally testing for bunching (Table 2) reveals insignificant coefficients (except for the three-day APD). I use this finding to simplify the

demand model. That is, I assume that consumers make a one-shot decision to buy or not buy, and then revisit the idea of forward-looking behavior later.

Finally, Figure 6 plots the mean fare response by day before departure. The graph separates out two scenarios: (1) situations in which the firm sees positive sales in the previous period; and (2) those in which there are no sales in the previous period. The graph shows that both pricing forces are at play. Conditional on positive sales, capacity becomes more scarce, and prices increase. This is consistent with stochastic demand pricing. Also consistent with stochastic demand pricing, prices decrease when sales do not occur, reflecting the declining opportunity cost of capacity. However, close to the departure date, regardless of sales, prices increase. This is inconsistent with a model of just stochastic demand and suggests that late-arriving consumers are less price-sensitive. Firms capture this higher willingness to pay through intertemporal price discrimination.

Figure 6: Fare Response to Sales by Day Before Departure



Note: Average fare changes as a response to sales by day before departure. The vertical lines correspond to advance-purchase discount periods (fare fences). The horizontal line indicates no fare response.

3 An Empirical Model of Dynamic Airline Pricing

In this section, I specify a structural model of dynamic airline pricing. Section 3.1 provides an overview. Section 3.2 presents the demand model, and Section 3.3 presents the firm's problem.

3.1 Model Overview

A monopolist airline offers a single flight for sale in series of sequential markets. I assume that the demand and pricing decisions are not correlated across departure dates (multiple flights), so investigating a single date is representative. Time is discrete and the airline pricing problem has a finite horizon. Period 1 corresponds to the first sales period, and period T corresponds to the day the flight leaves. Initial capacity is exogenous, and the firm is not allowed to oversell.

Each period, the airline offers a single price to all customers. Consumers arrive according to a stochastic process. Each consumer is either a business traveler or a leisure traveler; business travelers are less price-sensitive than leisure travelers. The proportion of business versus leisure consumers is allowed to change across time. Upon arrival, consumers choose to purchase a ticket on the flight or choose not to travel, corresponding to the outside option. While the data (Section 2.3) suggest that modeling consumers as being myopic is a reasonable assumption, after estimating the model, I discuss the extension of allowing consumers to delay purchase. If demand exceeds remaining capacity, tickets are randomly rationed, which ensures that the capacity constraint is not violated. I also assume that passengers do not cancel tickets, as the average number of cancellations in the data per flight is fewer than two. Thus, remaining capacity is monotonically decreasing. With an updated capacity constraint, the firm again chooses a fare to offer, and the process repeats until the perishability date.

3.2 Demand

Each day before the flight leaves, a Poisson process (\tilde{M}_t) brings new consumers to the market. Upon entering the market, all uncertainty about travel preferences is resolved. This approach differs from Lazarev (2013) and earlier theoretical work, including Gale and Holmes (1993), in which existing consumer uncertainty can be resolved by delaying purchase. In this model, at date t , consumers arrive, and choose to either purchase a ticket or exit the market.

The demand model is based on the two-consumer type discrete choice model of Berry, Carnall, and Spiller (2006), which is frequently applied to airline data. Consumer i is a business traveler with probability γ_t or a leisure traveler with probability $1 - \gamma_t$. Consumer i has preferences (β_i, α_i) over product characteristics ($x_t \in \mathbb{R}^K$) and price (p_t), respectively. If consumer i chooses to fly, the consumer receives utility $U_{it1} = x_t\beta_i - \alpha_i p_t + \varepsilon_{it1}$. If i chooses to not fly, the consumer receives normalized utility $U_{it0} = \varepsilon_{it0}$. Assume that the idiosyncratic preferences of consumers, $(\varepsilon_{it1}, \varepsilon_{it0})$, are i.i.d. Type-1 Extreme Value (T1EV). Following the discrete choice literature, consumer i chooses to fly if and only if $U_{it1} \geq U_{it0}$.

Define $y_t = (\alpha_i, \beta_i, \varepsilon_{i1}, \varepsilon_{i0})_{i \in 1, \dots, \tilde{M}_t}$ to be the vector of preferences for the consumers that enter the market. Suppressing the notation on product characteristics for the rest of this section, the demand for the flight at t is defined as

$$Q_t(p_t, y_t) := \sum_{i=0}^{\tilde{M}_t} 1[U_{it1} \geq U_{it0}] \in \{0, \dots, \tilde{M}_t\},$$

where $1(\cdot)$ denotes the indicator function. Demand is integer-valued; however, it may be the case that more consumers want to travel than there are seats remaining – i.e., $Q_t(p, y) > s_t$, where s_t is the number of seats remaining at t . Since the firm is not allowed to oversell, in these instances, I assume that remaining capacity is rationed by random selection. Specifically, I assume that the market first allows consumers to enter and choose the product that maximizes utility. After consumers make their decisions, the capacity

constraint is checked. If demand exceeds remaining capacity, consumers who wished to purchase are randomly shuffled. The first s_t are selected, and the rest receive their outside option.

Although the model assumes that consumers arrive and purchase a single one-way ticket, it allows for round-trip ticket purchases in the following way: a consumer arrives looking to travel, leaving on date d and returning on date d' . The consumer receives idiosyncratic preference shocks for each of the available flights in both directions, and chooses which tickets to purchase. Since several airlines, such as JetBlue, price at the segment level, there is no measurement error in this procedure. That is, a consumer pays the same price for for two one-way tickets as the consumer would for a round-trip ticket.¹⁵

The assumptions placed on the consumer problem, along with the assumption that the firm cannot oversell, allow for deriving closed-form expressions to the demand system. As the firm (and econometrician) does not know how many consumers, and how many consumers of each type will arrive, I integrate over the distribution of y_t ,

$$Q_t^e(s_t, p_t) = \int_{y_t} \min(Q_t(p_t, y_t), s_t) dF_t(y_t).$$

To start, note that the T1EV assumption on the consumer preference shocks leads to the frequently used conditional logit model. Since there is only a single product in the choice set,

$$\pi_{t1}^i := \Pr(i \text{ wants to purchase } j \mid \text{type} = i) = \frac{1}{1 + \exp(-x_t \beta_i + \alpha_i p_t)}.$$

The discrete choice literature typically does not model capacity constraints. Because consumers may be forced to the outside option once the capacity constraint binds, the purchase probabilities represent desired purchases instead of realized purchases. Let π_{t0}^i denote the purchase probability of the outside option for consumer type i .

¹⁵Since remaining capacity is rationed when an oversell would otherwise occur, an implication of this modeling choice is that a consumer may be forced to the outside option for either the outbound or the inbound leg of the round-trip.

Let B denote the business type and L denote the leisure type. Recall that the probability of a consumer being type- B is γ_t . Then, $\gamma_t \pi_t^B$ defines the probability that a consumer is of the business type and wants to purchase a ticket. Consider the market share of the outside good at time t . Conditional on $k \in \mathbb{N}$ consumers arriving (suppressing the dependence on p_t, x_t), the probability that all customers want to purchase the outside option is

$$\Pr(Q_t = 0 | k) = \sum_{n=0}^k \binom{k}{n} \left((1 - \gamma_t) \pi_{t0}^L \right)^n \left(\gamma_t \pi_{t0}^B \right)^{k-n}.^{16}$$

Note that

$$\Pr_t(Q_t = 0) = \sum_{k=0}^{\infty} \Pr_t(Q_t = 0, \tilde{M}_t = k) = \sum_{k=0}^{\infty} \Pr_t(Q_t = 0 | \tilde{M}_t = k) \Pr(\tilde{M}_t = k),$$

and since consumers are assumed to enter the market according to a Poisson process,¹⁷ $\Pr(Q_t = 0)$ has the following analytic form:

$$\Pr_t(Q_t = 0) = \sum_{k=0}^{\infty} \frac{\mu_t^k e^{-\mu_t}}{k!} \sum_{n=0}^k \binom{k}{n} \left((1 - \gamma_t) \pi_{t0}^L \right)^n \left(\gamma_t \pi_{t0}^B \right)^{k-n}. \quad (3.1)$$

The expression for selling a positive number of seats is similar, in that the expression goes through all combinations of choosing to fly, or choosing not to fly, and mixes over

¹⁶For example, let $k = 2$. Then, conditional on two consumers arriving, both consumers want to purchase the outside option. Both consumers are leisure travelers; both consumers are business travelers; or one of each type arrive. The first two situations correspond to $n = 2$ and $n = 0$, respectively, since $[(1 - \gamma_t) \pi_{t0}^L]^2$ is the probability of two leisure consumers arriving and wanting to choose the outside option, and $[\gamma_t \pi_{t0}^B]^2$ is the probability two business consumers want to choose the outside option. Lastly, it could be the case that one business and one leisure consumer arrive. There are two possibilities: the first consumer is the business consumer, or vice versa. Hence, $2(1 - \gamma_t) \pi_{t0}^L (\gamma_t \pi_{t0}^B)$ enters the probability.

¹⁷With this assumption, the demand model closely follows Talluri and Van Ryzin (2004) and Vulcano, van Ryzin, and Chahr (2010), except that this model has two consumer types.

the two consumer types. It also is a Poisson-Binomial mixture and can be shown to equal

$$\Pr_t(Q_t = q | 0 < q < s) = \sum_{k=q}^{\infty} \frac{\mu_t^k e^{-\mu_t}}{k!} \binom{k}{q} \sum_{\ell=0}^q \binom{q}{\ell} (\gamma_t \pi_{t1}^B)^\ell ((1 - \gamma_t) \pi_{t1}^L)^{q-\ell} \times \quad (3.2)$$

$$\left[\sum_{n=0}^{k-q} \binom{k-q}{n} ((1 - \gamma_t) \pi_{t0}^L)^n (\gamma_t \pi_{t0}^B)^{k-q-n} \right].$$

Demand is latent in the case of a sellout since it is possible that some consumers are forced to the outside option. These probabilities can be constructed based on the fact that at least s_t seats are demanded. It can be shown that the Binomial-Mixture is equal to

$$\Pr_t(Q_t \geq q | s) = \sum_{q=s}^{\infty} \sum_{k=q}^{\infty} \frac{\mu_t^k e^{-\mu_t}}{k!} \binom{k}{q} \sum_{\ell=0}^q \binom{q}{\ell} (\gamma_t \pi_{t1}^B)^\ell ((1 - \gamma_t) \pi_{t1}^L)^{q-\ell} \times \quad (3.3)$$

$$\left[\sum_{n=0}^{k-q} \binom{k-q}{n} ((1 - \gamma_t) \pi_{t0}^L)^n (\gamma_t \pi_{t0}^B)^{k-q-n} \right].$$

Thus, all the demand possibilities (Equation 3.1-3.3), which will be denoted $f_t(s' | s, p)$, have analytic expressions.

3.3 Monopoly Pricing Problem

The monopolist maximizes expected discounted revenues in a series of sequential markets. In each of the sequential markets, the airline chooses to offer a single price to all consumers. Because of the institutional feature that airfares are discrete, I assume that the firm chooses a price from a discrete set (P). The pricing decision is based on the states of the flight: seats remaining, time left to sell, flight characteristics x_j (notation suppressed), as well as idiosyncratic shocks $\omega_t := \{\omega_{tp} : p_t \in P\} \in \mathbb{R}^P$, which are assumed to be distributed i.i.d T1EV, with scale parameter σ . These shocks are assumed to be additively separable to the remainder of the per-period payoff function, which, in this case, is expected revenues, $R_t^e(s, p) = p_t \cdot Q_t^e(s_t, p_t)$.

The firm's problem can be written as a dynamic discrete choice model. Let $V_t(s_t, \omega_t)$ be the value function given the state (t, s_t, ω_t) . Denoting δ as the discount factor, the dynamic program (DP) of the firm is

$$V_t(s_t, \omega_t) = \max_{p_t \in P} \left(R_t^e(s_t, p_t) + \omega_t p + \delta \int_{\omega_{t+1}, s_{t+1} | s_t, p_t, \omega_t} V_{t+1}(s_{t+1}, \omega_{t+1}) dH_t(\omega_{t+1}, s_{t+1} | s_t, p_t, \omega_t) \right).$$

Because the firm cannot oversell, capacity transitions as $s_{t+1} = s_t - \min\{Q_t(p_t, y_t), s_t\}$. The firm faces two boundary conditions. The first is $V_t(0, \omega_t) = 0$ – that is, once the airline hits the capacity constraint, it can no longer sell seats. The second is $V_T(s_t, \omega_t) = 0$ – that is, unsold seats on the day that the flight leaves represent lost revenue opportunities.

I assume that conditional independence is satisfied, meaning that the transition probabilities can be written as $h_t(s_{t+1}, \omega_{t+1} | s_t, \omega_t, p_t) = g(\omega_{t+1}) f_t(s_{t+1} | s_t, p_t)$.¹⁸ Following Rust (1987), the expected value function can be expressed as

$$EV_t(p_t, s_t) = \int_{s_{t+1}} \left[\sigma \ln \left(\sum_{p_{t+1} \in P} \exp \left(\frac{R_{t+1}^e(s_{t+1}, p_{t+1}) + EV_{t+1}(p_{t+1}, s_{t+1})}{\sigma} \right) \right) \right] f_t(s_{t+1} | s_t, p_t) ds_{t+1} + \sigma \phi,$$

and the conditional choice probabilities also have a closed form and can be computed as

$$CCP_t(s_t, p_t) = \frac{\exp \left\{ \left(R_t^e(p_t, s_t) + EV_t(p_t, s_t) \right) / \sigma \right\}}{\sum_{p'_t \in P} \left[\exp \left\{ R_t^e(p'_t, s_t) + EV_t(p'_t, s_t) \right\} / \sigma \right]}.$$

Given a set of flights (F) each tracked for (T) periods, the likelihood for the data, accounting for the firm's pricing decision, is given by

$$\max_{\theta} \mathcal{L}(\text{data} | \theta) = \max_{\theta} \prod_F \prod_T CCP_t(s_t, p_t) f_t(s_{t+1} | s_t, p_t), \quad (3.4)$$

¹⁸Since common parameters enter both the payoff function and the transition probabilities, I do not estimate the transition probabilities in a first stage, such as in Hotz and Miller (1993) and Bajari, Benkard, and Levin (2007).

where $\theta := (\beta, \alpha, \gamma, \mu, \sigma)$ are the parameters to be estimated.

4 Model Estimates

In this section, I discuss the identification and the estimation procedure (Section 4.1) and then the results and model fit (Section 4.2). Finally, I discuss the model extension of allowing consumers to delay purchase (Section 4.3).

4.1 Identification and Estimation

The key identification challenge of the paper is to separately identify the demand parameters from the arrival process. This is pointed out in Talluri and Van Ryzin (2004), for example. The issue arises because without search data to pin down the arrival process, an increase in arrivals could instead be seen as a change in the mix of consumer types (demand). For example, the sale of two seats could have occurred because two consumers arrived and both purchased or because four consumers arrived and half purchased. Note that this is an argument about the demand elasticity, and I argue that accounting for the firm's pricing decision allows me to disentangle the two.

Figure 6 demonstrates why the pricing information is useful in separating the overall demand elasticity from the arrival process. Given stochastic demand, we would expect prices to rise when demand exceeds expected demand, and to fall over a sequence of low demand realizations. However, close to the departure date and regardless of sales, Figure 6 shows that prices tend to rise. This would only occur in the model presented if there is an overall change in the demand elasticity. That is, consumers who shop late are less price sensitive than those who shop early. By solving the firm's dynamic programming problem, I recover the opportunity costs of seats, and along with price, I recover the demand elasticity. Changes in average prices over time inform the demand elasticity, and responses to variation in sales given seats remaining and time left to sell inform the arrival process.

The identification strategy utilizes the firm’s pricing decision, but this creates a significant computational burden in estimation as the finite-horizon, non-stationary firm problem needs to be solved for each candidate vector of parameters. Thus, while the model is general to include flight characteristics, such as day-of-week effects, holiday indicators, and potential several flight options¹⁹, incorporating these features only increases the computational burden. To keep the problem tractable, I estimate the model by abstracting from flight characteristics other than price, but allow for city-pair-specific parameters. The model fits the pricing data well, as shown in the next section.

I assign the discount factor to be one. Finally, I place additional restrictions on γ_t and μ_t , or the probability on consumer types and arrival rates, respectively. I assign γ_t to be a logit function of two parameters. This forces monotonicity, but not strict monotonicity, in the mixture of business versus leisure customers over time. I assign:

- μ_1 Greater than 21 days prior to departure;
- μ_2 14 to 21 days prior to departure;
- μ_3 7 to 14 days prior to departure; and
- μ_4 Less than 7 days prior to departure

which corresponds to the advance purchase discount periods commonly seen in airline markets. This adds some flexibility in the Poisson arrivals of customers. As consumer types are allowed to change daily, this allows for day-to-day variation in expected demand. Finally, I assume $\beta^L = \beta^B$, resulting in $1 + 2 + 4 + 2 + 1 = 10$ parameters to be estimated per city pair, or 40 parameters to be estimated in total.

To estimate the problem, I maximize the log-likelihood of the firm’s dynamic programming problem found in Figure 3.4 using the Knitro solver. I utilize 100 random starts in the parameter space. The state space is size $T \times S \times Q \times P$ per route, which is the number of periods, times number of potential sales and current seats remaining, times number of prices in the choice seat. This is a large state space, even though many entries are zero due

¹⁹Additional assumptions on the rationing rule are required in this case.

to restrictions implied by the model, i.e., monotonically decreasing capacity. To decrease the size of the problem, I cluster prices using k-means (size 8), and then take the mean of the clusters, for each market. This groups together fares that only differ by sometimes a dollar, and the resulting fares have a geometric fit of 93-98% with the raw data. This reduces the state space by a factor of 2-3, and results in a state space of up to a few million per city pair.

Finally, I use an important institutional feature of airline markets to define the price choice sets of firms. Since airlines commonly assign fare restrictions, such as advance purchase discounts, this creates a restriction that some fares are only used at particular times. For example, a fare with a 21-day advance purchase discount (APD) requirement is not available for purchase close to the departure date. I use the fare restrictions (fences) attached to the fares to define daily choice sets, P_t . Fares with early APD requirements are less expensive than fares that are offered close to the departure date. I demonstrate later how firm pricing varies across time without imposing the fences.²⁰

4.2 Model Estimates, Model Fit, and Model Pricing

Parameter estimates appear in Table 3. All parameters are significant at the 1% level. The parameter estimates suggest leisure consumers are twice as price sensitive as business consumers on average, and business consumers are willing to pay 60% more in order to secure a seat, on average.²¹ Demand elasticities range from two to four depending on market and time until departure. The probabilities on consumer types implies a significant change in the price sensitivity of consumers over time. Figure 7 plots the fitted values using the (γ^1, γ^2) parameters for the markets studied. The estimates imply early on that essentially all arrivals are leisure consumers. Then, a few weeks prior to departure, there is a steep shift towards more price inelastic arrivals. Moving to the parameters governing

²⁰Advance purchase discounts are modeled from the set of daily fares offered, after clustering. This accounts for the fact that high fares are offered close to the departure date and low fares are not.

²¹The model is estimated where prices are measured in hundreds of dollars.

the arrival process, the estimates suggest relatively low Poisson rates, with a minimum rate of nearly two and a maximum rate of nearly 11 persons per day.


Table 3: Parameter Estimates

Variable	Austin	Jacksonville	Kansas City	San Diego
<u>Demand</u>				
β_0	2.500 (0.051 ^{***})	-1.026 (0.036 ^{***})	-0.179 (0.052 ^{***})	0.953 (0.040 ^{***})
α^L	-1.052 (0.009 ^{***})	-1.033 (0.007 ^{***})	-0.785 (0.008 ^{***})	-0.851 (0.005 ^{***})
α^B	-0.722 (0.004 ^{***})	-0.810 (0.011 ^{***})	-0.697 (0.003 ^{***})	-0.349 (0.003 ^{***})
γ^1	-13.397 (0.124 ^{***})	-17.886 (0.092 ^{***})	-5.933 (0.111 ^{***})	-5.927 (0.071 ^{***})
γ^2	0.274 (0.003 ^{***})	0.578 (0.003 ^{***})	0.157 (0.003 ^{***})	0.106 (0.002 ^{***})
<u>Arrival Process</u>				
μ_1	1.912 (0.081 ^{***})	11.001 (0.053 ^{***})	2.906 (0.086 ^{***})	3.470 (0.074 ^{***})
μ_2	2.461 (0.078 ^{***})	13.977 (0.054 ^{***})	3.230 (0.085 ^{***})	3.478 (0.069 ^{***})
μ_3	1.821 (0.078 ^{***})	5.381 (0.058 ^{***})	3.788 (0.090 ^{***})	3.706 (0.064 ^{***})
μ_4	2.357 (0.062 ^{***})	8.050 (0.047 ^{***})	3.088 (0.073 ^{***})	3.020 (0.050 ^{***})
<u>Firm Shock</u>				
ω	0.425 (0.001 ^{***})	1.539 (0.000 ^{***})	0.246 (0.000 ^{***})	0.260 (0.000 ^{***})
LogLike	46,388	91,581	43,476	73,514

Note: Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Figure 8 addresses model fit by comparing mean fares with model fares, by day before departure. Mean data fares are calculated as daily means using the 1,362 flights in the


Figure 7: Visualizing the Arrival Process over Time



Note: Fitted values of the arrival process of business versus leisure customers across the booking horizon. The y-axis is $\Pr(\text{business})$, so $1 - \Pr(\text{business})$ defines $\Pr(\text{Leisure})$.

sample. As the model utilizes shocks that are private to the firm, model fares are calculated by using the empirical distribution of initial capacity seen in the data and simulating one million flights. The plot shows that the model fares are similar to observed fares, with differences of usually less than \$50. The median differences are smaller. Differences are close to zero for the first half of the sample. The model accurately picks up the increasing pattern of fares starting within three weeks of the departure date and peak fares within three days of the departure date. Model fares are smoother than the mean data fares, but that is to be expected given the number of model simulations. Figure 8 also shows a dashed line representing mean model fares where firms are allowed to choose from all available prices each period, i.e. $P := \cup_{t=1}^T P_t$. This line closely follows the restricted fares line for the first 53 days. However, close to the departure date, this model overpredicts discounts to clear remaining capacity, suggesting fare restrictions enforce high prices close to the date of travel.

Figure 8: Model Fit by Day Before Departure




Note: Comparison of mean data fares and mean model fares across the booking horizon. Two versions of model fares are plotted. The solid black line defines per-period price choice sets using fare restrictions in the data. The dashed grey line allows firms to choose from all prices each period. Model fares come from simulating one million flights. Mean data fares are from the 1,362 flights in the sample.

The model estimates imply the pricing functions of the firm react to changes in capacity and time in predictable ways. Figure 9 plots the pricing functions and integrated value functions for the Kansas City - Boston city pair as a function of remaining capacity. The left panel plots prices given remaining capacity for four selected periods, corresponding to 15, 30, 45 and 60 days prior to departure. The right plot indicates the value functions associated with remaining capacities for these four periods. The left panel plots the expected price using the conditional choice probabilities from the model, i.e., $\int p_t CCP_t(s, t)$. There is a general upward trend or negative relationship between capacity and price. For any given period if capacity is high, prices are low. However, for a given capacity, prices are increasing in number of periods remaining. This is because the shadow value of capacity is higher. If a sell out becomes more likely early on, the firm greatly increases prices which saves seats for business consumers who will likely arrive later. The plot also

shows that the pricing functions depend on the period-specific choice sets. For example, 15 days out corresponds to crossing the 21-advance purchase discount date. Now a high fare is available and firm utilizes this fare when capacity is scarce.

Figure 9: Estimated Policy Functions



Note: Model policy functions and value functions for four periods. This figure demonstrates pricing and revenues for the Kansas City-Boston city-pair. The policy functions utilize per-period price choice sets.

The right panel shows that expected revenues are increasing in capacity for a given period, i.e. $\int_{\omega} V_t(s, \omega) dF\omega \leq \int_{\omega} V_t(s + 1, \omega) dF\omega$. It is also true that expected values are increasing in time to sell for a given capacity, i.e. $\int_{\omega} V_t(s, \omega) dF\omega \leq \int_{\omega} V_{t+1}(s, \omega) dF\omega$. These results are consistent with the theory on dynamic pricing found in Gallego and Van Ryzin (1994). Expected revenues flatten out for a given period because the firm cannot capture additional revenue when there is sufficiently high capacity (and the shadow value goes to zero). The plot shows that the probability of selling out is essentially zero if the firm has at least 20 seats remaining with 45 days left to sell. On the other hand, with 30 days remaining, there is excess capacity when at least 15 seats are remaining (25% of capacity).

4.3 Allowing Consumers to Wait

The demand model assumes that consumers do not wait to purchase tickets. This assumption is motivated by the fact that there is no bunching in bookings before advance purchase discount days in the data. In order to investigate the incentive to wait, I change the model in the following way: after consumers arrive, each consumer has the option to either buy a ticket, choose not to travel, or wait one additional day to decide. By choosing to wait, each consumer retains her private valuations (the ε 's) for traveling but may be offered a new price tomorrow. Consumers have rational expectations regarding future prices. However, in order to wait, each consumer has to pay a transaction cost ϕ_i . This cost reflects the disutility consumers incur when needing to return to the market the next period.

I derive a waiting cost $\bar{\phi}$ such that if all consumers have a waiting cost at least as high as $\bar{\phi}$, then no one will wait. I then calculate the transaction costs.

Dropping the i, t, s subscripts, the choice set of a consumer arriving at time t in a model of waiting is

$$\max\{\varepsilon_0, \beta - \alpha p + \varepsilon_1, EU^{\text{wait}} - \phi\},$$

where EU^{wait} is the expected value of waiting one more period. This expected utility can be written as

$$EU^{\text{wait}} = \mathbb{E}\left[\max\{\varepsilon_0, \beta - \alpha p_{t+1} + \varepsilon_1\}\right].$$

To derive $\bar{\phi}$, I first investigate the decision to wait for the marginal consumer, or the consumer such that $\varepsilon_0 = \beta - \alpha p + \varepsilon_1$. This consumer has no incentive to wait if the price tomorrow is at least as high as today. If the price drops, the gain from waiting is

$$\begin{aligned} u_{t+1} - u_t &= (\beta - \alpha p_{t+1} + \varepsilon_1) - (\beta - \alpha p + \varepsilon_1) \\ &= \alpha(p - p_{t+1}). \end{aligned}$$

The expected gains from waiting for this marginal consumer are

$$\Pr(p_{t+1} < p) \mathbb{E}[\alpha(p - p_{t+1}) \mid p_{t+1} < p].$$

Hence, an indifferent consumer will not wait if $\phi_i > \bar{\phi} = \Pr(p_{t+1} < p) \mathbb{E}[\alpha(p - p_{t+1}) \mid p_{t+1} < p]$. With this transaction cost, one can show all consumers will not wait to purchase (see Section A). For consumers who would purchase today, the gains from waiting are equal to $\bar{\phi}$, but there is an additional cost if prices raise. Hence, waiting is not optimal. For consumers who would prefer not to buy, the expected gains of waiting are negative.

In monetary terms, $\bar{\phi}/\alpha = \Pr(p_{t+1} < p) \mathbb{E}[(p - p_{t+1}) \mid p_{t+1} < p]$ defines a transaction cost such that waiting is never optimal. For these costs to be calculated, the information set of consumers needs to be defined. I assume consumers form expectations given current prices and time, but they do not forecast the changes in number of seats remaining across time. This seems reasonable given that meta search engines do not readily display seat maps to consumers. With these assumptions, I find that the median and mean transaction costs to be \$5.69 and \$7.49, respectively. Recall, these costs are based on the most extreme case: the consumer who is indifferent between purchasing today or delaying the decision.

Overall, these results, along with the empirical evidence, suggest that the incentive to wait to purchase, conditional on knowing preferences, is small. While betting on price in airline markets may result in gains, the trajectory of prices is overwhelmingly positive.

5 Analysis of the Estimated Model

In this section, I conduct a series of counterfactuals given the model estimates. I compare firm revenues and consumer welfare under dynamic pricing with several alternative pricing strategies. I investigate: uniform pricing, dynamic pricing with a restriction to the number of price changes over time, and a pricing system where firms commit to a price schedule that depends on the time until departure, but not the scarcity of

seats (intertemporal price discrimination alone). By comparing uniform pricing to this latter counterfactual, which I call the intermediate case, I quantify relative influence of the intertemporal price discrimination and dynamic adjustment to stochastic demand of dynamic airline pricing.

Next, I highlight this relative importance by exploring pricing under alternative arrival processes (Section 5.3). I first reverse the arrival process of consumers so that price-insensitive consumers are the first to shop. This simulates environments such as high-end fashion. I then hold the proportion of business travelers constant across time, which may simulate environments such as movie theaters. Finally, in the last subsection (Section 5.4), I show in order to quantify the effects of price discrimination in airline markets, it is essential to take into account stochastic demand.

For each counterfactual, I use the empirical distribution of remaining capacity 60 days prior to departure as the initial capacity condition. I simulate 250,000 flights for each of the city pairs. Each flight starts with an initial capacity drawn from the empirical distribution. I then calculate the following benchmarks:

- Fare : mean fare, excluding sell-out observations;
- Load Factor : mean load factor the day flights leave;
- Sell outs : percent of flights that sell all initial capacity;
- Revenue : mean flight revenue;
- CS_L^i : mean leisure consumer surplus;
- CS_B^i : mean business consumer surplus;
- Welfare : mean welfare per flight, day (less sunk costs).

I alter the firm problem in the following ways. For the observed arrival process under dynamic pricing and the intermediate case, I impose the observed fare restrictions described in Section 4.1. This results in period-specific choice sets P_t . This allows for a better comparison between observed pricing and the intermediate case; however, the core results are unchanged by removing this restriction. For the uniform case, I allow

firms to choose from all prices seen in the market. This is because a uniform price does not exist for these markets since the intersection of price sets over all periods is empty. The intermediate case covers commitment under changing prices. For the reversed arrival process, I simply flip the arrival process and choice sets defined under the observed arrival process. For the constant arrival process, all prices are available in all periods.

Finally, I remove the firms shocks for the following analysis. I do this in order to single out the effect of prices and capacity on revenues, and not the role of unobservable errors in determining the pricing decision. For example, under uniform pricing, the firm would receive a single error vector and the expected value for all future periods, whereas in the dynamic counterfactual the firm receives per-period error shocks. By removing the unobservable from the firm's problem, quantifying the impact of price discrimination across pricing regimes is salient.

5.1 Uniform Pricing

I start by removing the firm's ability to price discriminate. The firm maximizes expected revenues subject to the constraint that it must charge a uniform price each period. The price is solely dependent on the initial capacity condition. That is, the firm solves

$$\max_p \mathbb{E}_y \left[\sum_{t=1}^T p Q_t(p, y_t) | s_1 \right],$$

such that the firm is not allowed to oversell the flight. Under uniform pricing, a high fare helps ensure seats are saved for business consumers, but doing is likely to price leisure consumers out of the market. At the same time, setting high prices to save seats and capture the high willingness to pay of late-arriving consumers may result in unused capacity. The optimal pricing strategy balances out these effects.

Results for the counterfactual appear in Table & Fig. 4. In the figure, the left panel plots mean fare across time for flights with seats remaining in that period (the solid line tracks


the mean-weighted fare from the initial capacity distribution). The dashed lines show fares resulting from a one standard deviation increase/decrease from the initial mean-weighted capacity. The left figure shows that under uniform pricing, fares are relatively high early on and relatively low close to the date of travel compared to dynamic pricing. This results in a relocation of capacity over time, which is displayed in the right panel. The right figure plots the booking curve, or mean cumulative seats sold, towards the departure date. The uniform pricing booking curve is bowed out as fewer consumers purchase under the relatively high fare early on. Relatively low fares close to the departure date results in a higher booking rate compared to dynamic pricing. However, even with this increase in the booking rate, the overall load factor for flights under uniform pricing is lower than dynamic pricing by 4.3%. The figure also shows that flights that are one standard deviation below the mean have significantly lower load factors (bottom dashed line). That is, the distribution of load factors under uniform pricing become increasingly dispersed, whereas the distributions under dynamic prices are shifted more to the right (towards 100%).

Uniform pricing results in a reallocation of capacity within consumer type. With dynamic pricing, the purchase rate declines as prices increase and this is most notable roughly two weeks prior to departure. Under uniform pricing, the purchase rate within type is constant; however, it also declines for flights that sell out in advance. The firm cannot increase price in response to significant, positive demand shocks, and consequently, many flights sell out in advance (10.3% more). This also impacts business consumers: they pay lower fares, but more do not receive seats.

The top of Table & Fig. 4 reports the benchmarks of the two pricing regimes. Average consumer welfare is 3.2% higher for leisure consumers and 5.2% lower for business consumers under dynamic pricing compared to uniform pricing. In aggregate, the surplus changes essentially cancel out. However, revenues fall 2% under uniform pricing,

Table & Fig. 4: Dynamic to Uniform Pricing

	Fare	Load Factor	Sell Outs	Revenue	CS_L^i	CS_B^i	Welfare
Dynamic	249.1	91.7	23.1	9906.0	90.4	142.2	584.5
Uniform	242.8	87.3	33.4	9704.3	87.5	150.0	580.4
Diff. (%)	2.6	5.0	-30.9	2.1	3.2	-5.2	0.7



Note: Fare: mean fare for flight observations with positive seats remaining; Load factor (LF): mean for the day in which flights leave; Sell Outs: Percentage of flights with zero seats remaining in the last period; Revenue is the average revenue for flights; Consumer surplus (CS_L, CS_B) account for sell outs, and is reported as per-person; Welfare is per-flight-time-period, but excludes fixed costs. Results come from simulating one million flights given the empirical distribution of capacities 60 days prior to departure.

which is significant given the industry²² Overall, this counterfactual suggests that price discrimination in airline markets greatly benefits firms without significantly impacting average consumer welfare. While price discrimination extracts more surplus from business consumers, it also results in more efficient capacity utilization.

²²IATA (2013) reports airline profit margins are around 1.1%, inclusive of ancillary income and debt interest.


The Role of Frequent Price Adjustments

The previous exercise compares the extremes in pricing capabilities of the firm, where prices are allowed to adjust based on time and seats, or prices are held fixed for all periods. Now I allow the firm to use dynamic pricing, with the restriction that prices must be maintained for k days. I conduct six counterfactuals, corresponding to $k = 2, 3, 6, 10, 20, 30$. The idea here is that dynamic pricing is clearly valuable to the firm, but it is not necessarily true that daily price adjustments are needed to obtain the revenues observed under (daily) dynamic pricing. For these counterfactuals, I remove the fare restrictions so all prices are available in all periods.

Figure 10 plots the revenue loss compared to the baseline case of daily price adjustments. For example, uniform pricing reduces revenue by 3.7% compared to (daily) dynamic pricing. This number is greater than the previous uniform pricing counterfactual because firms have the choice of all prices in all periods. As Figure 9 shows, the trajectory of prices without fare restrictions is to offer discounts, relative to the highest fares, close to the travel date. Fares at departure are still higher than fares early on. One way to view the 1.6% difference in profit gains is that it defines the cost of commitment. Airlines would like to offer discounts close to the departure date, but doing so would cause consumers to wait.

Figure 10 shows the incremental gain coming from increased flexibility in pricing. The ability to update prices just once increases revenue by 1.6% over uniform pricing. An additional price adjustment yields another 0.7% gain. Six-day and ten-day adjustments yield similar results, as well as two-day and three-day adjustments. The losses associated with the former are large relative to daily adjustments. The two and three-day results suggest that volatility in demand is not sufficiently large enough that daily adjustments result in significant gains.

Figure 10: The Role of Frequent Price Adjustments



Note: Revenue drop relative to dynamic (daily) pricing for all markets. For example, 3-day corresponds to firms utilizing dynamic pricing, but restricting the number of price updates to 3-day intervals.

5.2 The Use of Intertemporal Price Discrimination Alone

I now quantify the complementarity that results from fares responding to stochastic demand and a change in consumer elasticity over time. With dynamic pricing, fares respond to both forces. Under uniform pricing, fares respond to neither force. Here, fares are allowed to change over time, but not in response to the realizations of demand. By comparing this intermediate case with dynamic pricing, and uniform pricing to dynamic pricing, I quantify the additional gains coming from adjusting prices in response to demand shocks over the gains coming from intertemporal price discrimination.

In this intermediate case, the firm solves


$$\max_{p_t} \mathbb{E}_y \left[\sum_{t=1}^T p_t Q_t(p_t, y_t) \mid s_1 \right],$$

such that the firm is not allowed to oversell. With eight possible prices to choose from each period, the cardinality of the domain is approximately 1.5×10^{54} . To reduce the dimension of the problem, I add the restriction that the firm can only adjust fares on the usual advance

purchase discount days: 3, 7, 14, and 21 days prior to departure. This results in five prices per flight. The prices depend only on the initial capacity condition. I take the union of fares offered within these periods to define the choice sets, P_t .²³

Table & Fig. 5: Dynamic Pricing to Intermediate Case

	Fare	Load Factor	Sell Outs	Revenue	CS_L^i	CS_B^i	Welfare
Dynamic	249.1	91.7	23.1	9906.0	90.4	142.2	584.5
Intermediate	249.2	87.2	28.1	9772.0	88.2	142.1	577.3
Diff. (%)	0.0	5.1	-17.9	1.4	2.4	0.1	1.2



Note: Fare: mean fare for flight observations with positive seats remaining; Load factor (LF): mean for the day in which flights leave; Sell Outs: Percentage of flights with zero seats remaining in the last period; Revenue is the average revenue for flights; Consumer surplus (CS_L, CS_B) account for sell outs, and is reported as per-person; Welfare is per-flight-time-period, but excludes fixed costs. Results come from simulating one million flights given the empirical distribution of capacities 60 days prior to departure.

²³A previous version of this paper utilized simulated annealing on the full problem. I obtained similar findings compared to the setup proposed here.

Results for this counterfactual appear in Table & Fig 5. The time path of prices under intertemporal price discrimination alone are monotonically increasing, with a substantial increase in fares when crossing the last two advance purchase discount periods. This corresponds to the increase in proportion of business consumers for three out of the four markets. The fares offered to leisure consumers well before the departure date are higher compared to dynamic pricing for the same reason highlighted in the uniform pricing counterfactual – there is a strong incentive to save seats for later periods. Load factors are 5.1% higher under dynamic pricing and load factors are more dispersed in the intermediate case.

The additional flexibility in pricing results in lower fares compared to the uniform pricing counterfactual well in advance of the date of travel. This improves leisure consumer surplus, but as prices start to rise, leisure consumers who arrive late face higher prices than under the uniform pricing scenario. Overall I find leisure consumer surplus is increasing as pricing becomes more flexible (dynamic > intermediate > uniform). I find business consumer surplus is highest under uniform pricing and the same under dynamic pricing and this intermediate case. Without the ability to respond to demand shocks, sell outs are also higher under this counterfactual compared to dynamic pricing (5%). Here, average consumer surplus is 1.2% lower compared to dynamic pricing. Just like the previous counterfactual, revenues are higher under dynamic pricing (1.4%).

This counterfactual, along with the uniform pricing counterfactual, shows that there exists a significant complementarity between the pricing channels in airline markets. Allowing firms to adjust fares with time increases revenues from both early-arriving and late-arriving consumers. However, then allowing firms to also respond to demand shocks only enhances this ability to price discriminate. The results together suggest that 34% of the revenue gains associated with dynamic pricing over uniform pricing come from intertemporal price discrimination (the intermediate case). The remainder comes from the ability of being able to respond to the sequence of demand shocks. Note that the waiting

costs to justify these counterfactuals are lower (zero, in fact) as prices do not decline over time. To highlight the significance of these results, in the next subsection, I alter the arrival process of consumers and conduct the same set of counterfactuals.

5.3 Adjustments to the Arrival Process

The complementarity between intertemporal price discrimination and dynamic adjustment is large. I now demonstrate that this complementarity critically depends on the arrival pattern of consumers. I perform two counterfactuals. First, I flip the arrival pattern of consumers so high-valuing consumers arrive first. This reflects some features in markets such as high-end fashion where prices decrease over time. Then, I hold the arrival pattern of consumers fixed over time, which provides an environment to investigate dynamic adjustment without the intertemporal price discrimination force.

Reversal of the Arrival Process

In airline markets, airlines face the challenge of securing seats for business consumers, which is difficult because of the stochastic arrival of consumers. Later periods bring more valuable consumers, and airlines need to be able to capture their high willingness to pay. This makes dynamic adjustment to stochastic demand particularly useful in allocating capacity, as demonstrated in the previous counterfactuals. Firms can increase price early-on if capacity becomes scarce. At the same time, firms can offer reduced prices on flights with excess capacity, as to not leave as many seats unfilled.

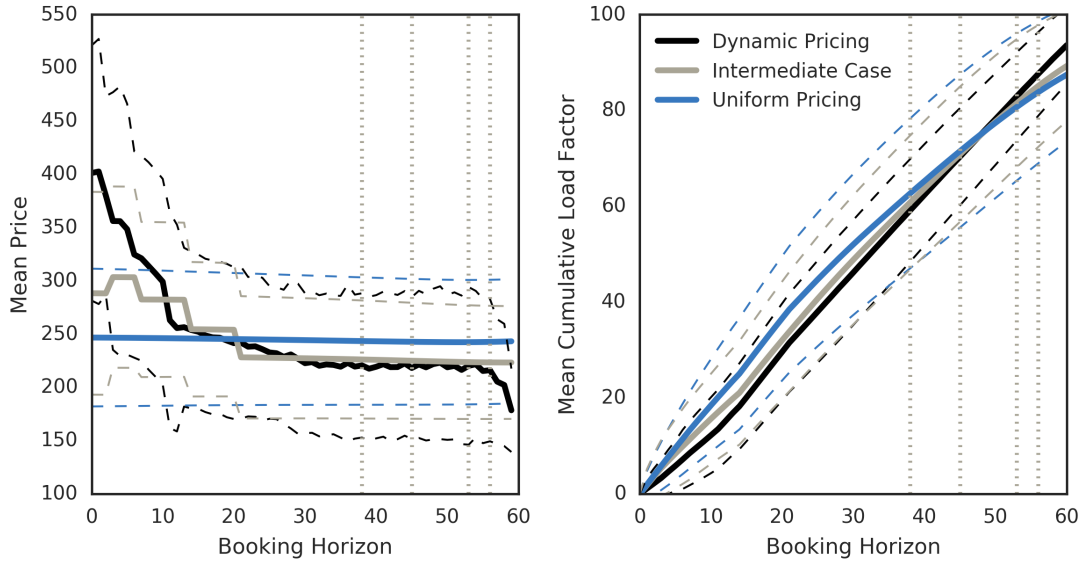
In this counterfactual, I reverse the arrival pattern of consumers and maintain the assumption that consumers are myopic, which I comment on at the end of this subsection. With a reversal of the arrival process, the high-valuing consumers are the first to arrive.²⁴ Table & Fig. 6 presents the counterfactual results of dynamic pricing, the intermediate case, and uniform pricing. Compared to the observed arrival process, the reversed process

²⁴The arrival process is literally flipped, meaning for example γ_T becomes γ_1 , etc. The same is true for μ_t .

brings: lower fares, higher load factors, more sell outs, and increased revenues. Overall, consumer surpluses are also higher under a reversal of the arrival process.

Table & Fig. 6: Reversal of Arrival Process

	Fare	Load Factor	Sell Outs	Revenue	CS_L^i	CS_B^i	Welfare
Dynamic	247.2	93.6	34.5	10019.2	90.6	143.5	588.6
Intermediate	242.2	89.3	34.4	9911.0	88.2	146.9	583.9
Uniform	242.2	87.5	33.8	9717.6	86.6	154.6	582.9



Note: Fare: mean fare for flight observations with positive seats remaining; Load factor (LF): mean for the day in which flights leave; Sell Outs: Percentage of flights with zero seats remaining in the last period; Revenue is the average revenue for flights; Consumer surplus (CS_L, CS_B) account for sell outs, and is reported as per-person; Welfare is per-flight-time-period, but excludes fixed costs. Results come from simulating one million flights given the empirical distribution of capacities 60 days prior to departure.

In this environment, firms have no incentive to hold remaining capacity for the late sequential markets. The early-arriving consumers are the most valuable and these consumers arrive when there is little uncertainty that seats will be available. This is particularly noticeable under intertemporal price discrimination alone, where prices decrease


at a slower rate than they increase under the observed arrival process. Here, firms have the incentive to keep prices high early in the hopes that consumers with high private valuations arrive. Since the first consumers to arrive have the highest willingness to pay, they are easily captured. The price path under dynamic pricing has the same pattern with the key difference being close to the perishability date. Flights that have not already sold out result in fire sales the day before departure as the firm tries to fill any remaining capacity. This results in a reverse ordering of load factors compared to the observed arrival process. Consequently, the percent of sell outs under dynamic pricing goes up by nearly 50%. Under intertemporal price discrimination alone, sell outs are 22% higher.

The most important feature of the reversed arrival process is the role of intertemporal price discrimination. With the observed arrival process, 34% of the revenue gains of using dynamic pricing over uniform pricing can be attributed to intertemporal price discrimination. The remaining can be attributed to dynamic adjustment. However, when the arrival process is reversed, 64% of the revenue gains can be attributed to intertemporal price discrimination. Originally, dynamic adjustment was used to reserve capacity for later periods, when consumers arrive with high valuations. With business consumers arriving early, firms can immediately capture their high willingness to pay and the role of dynamic adjustment is more to ensure that remaining seats are filled.

An important caveat to this analysis is the demand model assumes consumers do not delay their purchase decisions and in this counterfactual there would be a strong incentive to do so. Given the time path of fares, calculating new waiting cost bounds would yield significant higher, unreasonable costs. In fact, we might expect the firm would not be able to price discriminate at all under the reversed arrival process because consumers would just wait to purchase (the capacity constraint would mitigate some waiting). However, that highlights the main results even more. In airline markets, the pricing forces work as complements because of the particular arrival process of consumers. If the arrival process were reversed, either the relative importance of the pricing forces are changed, or they do

not even operate because of the deadline effect and limited capacity.

Constant Arrival Process



Finally, I investigate pricing under a constant arrival process. Demand is still stochastic, but the mixture of business and leisure travelers to be constant over time. Specifically, I take the proportion of business and leisure travels to be equal to the mean of the arrival

process and assign it to all periods. Table & Fig. 7 contains the results for dynamic pricing and uniform pricing. The intermediate case is also plotted, but changes in prices here are not due to elasticity changes over time, but solely due to stochastic demand. Under constant elasticity, the dynamic pricing line is flat, fluctuating around the levels offered under uniform pricing. While there is no intertemporal price discrimination force, if firms commit to pricing in advance, like in the intermediate cases presented earlier, the pattern of prices closely follows dynamic pricing. Close to the perishability date, prices drop to fill capacity.

One of the predictions found in Gallego and Van Ryzin (1994) is that the value of dynamic pricing is lower under constant arrival. My empirical findings support this theory. The revenue gains of dynamic pricing over uniform pricing are 25% higher in the reversed arrival process compared to this counterfactual. Moreover, in both the observed and reversed arrival counterfactuals with dynamic pricing, I utilize period-specific choice sets. This is a constraint on the firm's problem implying the gains over uniform pricing would be even higher.²⁵ The welfare patterns are different yet again in this counterfactual – uniform pricing results in the lowest consumer welfare among all the pricing systems.

5.4 Consequences of Abstracting from Stochastic Demand

The existing work on the welfare effects of price discrimination in airline markets models dynamic consumers but abstracts from stochastic demand (Lazarev 2013). Stochastic demand is an essential feature of airline markets, but here I show it also plays an important role in the price elasticities estimated from data. Stochastic demand, along with a change in the mixture of consumer types over time, results in changing opportunity costs. However, an empirical procedure that abstracts from stochastic demand assigns the opportunity cost to be constant in all periods. This results in a systematic bias in the estimated demand elasticities.

²⁵Revenue Management Overview claims revenue management systems have increased airline revenues by 3-9% over time.

To see this, consider the special limiting case of the model presented in this paper where demand uncertainty is degenerate. A monopolistic faces a series of sequential markets with finite, fixed capacity. Without any demand uncertainty, the firm problem can be written as a static optimization problem. Denoting Q demand and K the capacity constraint, the firm problem is

$$\max_p \sum_{t=1}^T Q_t(p_t)p_t \quad \text{s.t.} \quad \sum_{t=1}^T Q_t(p_t) \leq K.$$

Letting $c(K)$ be the shadow price of capacity, the firm problem can be rewritten as the unconstrained problem,

$$\begin{aligned} & \max_p \sum_{t=1}^T Q_t(p_t)p_t - c(K) \left(\sum_{t=1}^T Q_t(p_t) - K \right) \\ \Leftrightarrow & \max_p \sum_{t=1}^T Q_t(p_t)(p_t - c(K)) + c(K)K. \end{aligned}$$

Letting c be the shadow value at the optimum, Lerner's Index reveals

$$\frac{p_t - c}{p_t} = \frac{1}{e_t^D(p_t)},$$

where $e_t^D(p_t) \in \mathbb{R}_+$ is the elasticity of demand. This can be rearranged to obtain

$$p_t = \frac{e_t^D(p_t)}{e_t^D(p_t) - 1} c.$$

Taking the ratio of prices from the first period ($t = 1$), to any future period, yields

$$\frac{p_t}{p_1} = \frac{\frac{e_t^D(p_t)}{e_t^D(p_t)-1} c}{\frac{e_1^D(p_1)}{e_1^D(p_1)-1} c} = \frac{e_t^D(p_t)}{e_t^D(p_t)-1} \cdot \frac{e_1^D(p_1)-1}{e_1^D(p_1)}. \quad (\text{Abstracting from stochastic demand})$$

Importantly, the shadow values cancel as they do not change across periods. This is convenient because it suggests just by looking at relative prices in the data, we can recover the elasticity ratio in this special limiting case.

However, demand is stochastic in many markets, such is the case in airline markets. With stochastic demand, the shadow values may change over time, and in general we have

$$\frac{p_t}{p_1} = \frac{\frac{\tilde{e}_t^D(p_t)}{\tilde{e}_t^D(p_t)-1} \tilde{c}_t}{\frac{\tilde{e}_1^D(p_1)}{\tilde{e}_1^D(p_1)-1} \tilde{c}_1}. \quad (\text{Accounting for stochastic demand})$$


This suggests that if we were to investigate the welfare effects of airline pricing abstracting from stochastic demand, we would infer the incorrect price elasticities depending on how the \tilde{c} s change over time.

Figure 11 plots the mean and median elasticity ratios for both situations over time coming from one million simulated flights.²⁶ The elasticity ratios accounting for stochastic demand come directly from the estimated model. In Figure 11, the difference in the lines informs the bias. It shows that the empirical procedure that abstracts from stochastic demand will estimate consumers as being too price insensitive as the line is above the true elasticity ratio (the ratio is increasing in consumer price insensitivity). This occurs because the shadow values (\tilde{c} s) tend to rise due to the price insensitivity of late arrivals. Within the last 14 days of departure, the bias grows from 40% in relative terms to nearly 80%.

Like the counterfactual exercises, the bias too depends on the particular arrival process of consumers. If price inelastic consumers were the first to arrive, the direction of the bias is opposite – the procedure that abstracts from stochastic demand will infer late arrivals as too price sensitive. This is because the shadow values tend to decline toward the perishability date. On the other hand, under constant elasticity, the median bias is zero because the \tilde{c} s are constant on average.

²⁶Since means are heavily influenced by prices that price the demand elasticity close to one (the numerator grows without bound), I also plot the median ratios.

Figure 11: Bias When Abstracting from Stochastic Demand



Note: Results come from simulating one million flights according to the model estimates. The bias lines correspond to an elasticity ratio abstracting from stochastic demand. The model elasticity ratios are recovered using the model accounting for stochastic demand.

6 Conclusion

There are two broad rationales for product prices changing over time: segmentation of consumers who differ in their willingness to pay, and changes in scarcity due to stochastic demand. The main contribution of this paper is to study them jointly and quantify their interactions. I do so by examining the pricing decisions of airlines in US monopoly markets with novel data containing high frequency fares and seat availabilities and a structural model that accounts for both forces.

I find that dynamic adjustment to stochastic demand complements intertemporal price discrimination in airline markets. This is because of the particular arrival process of consumers. By having fares respond to demand shocks, airlines are able to secure seats for late-arriving consumers. These consumers are then charged high prices. While airlines utilize sophisticated pricing systems that result in significant price discrimination, these systems also more efficiently ration seats. From a welfare perspective, more restrictive

pricing systems shut down price discrimination, but this does not increase consumer welfare – only revenues decline. I also show that these answers depend critically on the arrival process of consumers. Finally, I show that in order to estimate the welfare effects of dynamic airline pricing, it is essential to control for stochastic demand. An empirical procedure that abstracts from stochastic demand will result in biased demand elasticities.

While this paper is the first to empirically quantify the interactions between dynamic adjustment to stochastic demand and intertemporal price discrimination, both the data and methodology have limitations. In some airline markets, versioning may be important. This complexity is abstracted from here as most of the flights do not have first class. Second, only monopoly markets are studied. An exciting area for future study is to examine the complementarities in the pricing channels in oligopoly. It may be that the ability of airlines to respond to scarcity may be reduced in more competitive markets. Moreover, it has been shown that inventory controls, while typically viewed as a tool to manage demand uncertainty, can also be used to facilitate oligopoly price discrimination (Dana and Williams 2018). Incorporating inventory controls may be necessary in order to extend this analysis to oligopoly.

References

- ALDERIGHI, M., M. NICOLINI, AND C. A. PIGA (2015): “Combined effects of capacity and time on fares: insights from the yield management of a low-cost airline,” *Review of Economics and Statistics*, 97(4), 900–915.
- BAJARI, P., C. L. BENKARD, AND J. LEVIN (2007): “Estimating dynamic models of imperfect competition,” *Econometrica*, 75(5), 1331–1370.
- BERRY, S., M. CARNALL, AND P. SPILLER (2006): “Airline hubs: costs, markups and the implications of customer heterogeneity,” *Advances in airline economics*, 1, 183–214.

- BERRY, S., AND P. JIA (2010): "Tracing the Woes: An Empirical Analysis of the Airline Industry," *American Economic Journal: Microeconomics*, 2(3), 1–43.
- CLARK, R., AND N. VINCENT (2012): "Capacity-contingent pricing and competition in the airline industry," *Journal of Air Transport Management*, 24, 7–11.
- COURTY, P., AND M. PAGLIERO (2012): "The Impact of Price discrimination on Revenue: Evidence from the concert industry," *Review of Economics and Statistics*, 94(1), 359–369.
- DANA, J., AND K. WILLIAMS (2018): "Oligopoly Price Discrimination: The Role of Inventory Controls," *working paper*.
- DANA, J. D. (1999): "Using Yield Management to Shift Demand When the Peak Time is Unknown," *The Rand Journal of Economics*, pp. 456–474.
- DILME, F., AND F. LI (2016): "Revenue management without commitment: dynamic pricing and periodic fire sales," *working paper*.
- ELMAGHRABY, W., AND P. KESKINOCAK (2003): "Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices, and Future Directions," *Management Science*, 49(10), 1287–1309.
- ESCOBARI, D. (2012): "Dynamic Pricing, Advance Sales, and Aggregate Demand Learning in Airlines," *Journal of Industrial Economics*, 60(4), 697–724.
- GALE, I. L., AND T. J. HOLMES (1993): "Advance-Purchase Discounts and Monopoly Allocation of Capacity," *The American Economic Review*, pp. 135–146.
- GALLEGO, G., AND G. VAN RYZIN (1994): "Optimal Dynamic Pricing of Inventories with Stochastic Demand Over Finite Horizons," *Management science*, 40(8), 999–1020.
- GERARDI, K. S., AND A. H. SHAPIRO (2009): "Does competition reduce price dispersion? New evidence from the airline industry," *Journal of Political Economy*, 117(1), 1–37.

- GOOLSBEE, A., AND C. SYVERSON (2008): "How do incumbents respond to the threat of entry? Evidence from the major airlines," *The Quarterly Journal of Economics*, 123(4), 1611–1633.
- HOTZ, V. J., AND R. A. MILLER (1993): "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 60(3), 497–529.
- IATA (2013): "Profitability and the air transport Profitability and the Air Transport Value Chain," Discussion paper, International Air Transport Association.
- LAMBRECHT, A., K. SEIM, N. VILCASSIM, A. CHEEMA, Y. CHEN, G. S. CRAWFORD, K. HOSANAGAR, R. IYENGAR, O. KOENIGSBERG, R. LEE, ET AL. (2012): "Price discrimination in service industries," *Marketing Letters*, 23(2), 423–438.
- LAZAREV, J. (2013): "The Welfare Effects of Intertemporal Price Discrimination: An Empirical Analysis of Airline Pricing in U.S. Monopoly Markets," *working paper*.
- LESLIE, P. (2004): "Price Discrimination in Broadway Theater," *RAND Journal of Economics*, pp. 520–541.
- LI, J., N. GRANADOS, AND S. NETESSINE (2014): "Are consumers strategic? Structural estimation from the air-travel industry," *Management Science*, 60(9), 2114–2137.
- MCAFEE, P., AND V. TE VELDE (2006): "Dynamic Pricing in the Airline Industry," *Handbook on Economics and Information Systems*.
- NAIR, H. (2007): "Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games," *Quantitative Marketing and Economics*, 5(3), 239–292.
- PASHIGIAN, B. P., AND B. BOWEN (1991): "Why are Products Sold on Sale?: Explanations of Pricing Regularities," *The Quarterly Journal of Economics*, 106(4), 1015–1038.
- PULLER, S. L., A. SENGUPTA, AND S. N. WIGGINS (2015): "Does Scarcity Drive Intra-Route Price Dispersion in Airlines?," *working paper*.

- RUST, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, pp. 999–1033.
- SOYSAL, G. P., AND L. KRISHNAMURTHI (2012): "Demand dynamics in the seasonal goods industry: An empirical analysis," *Marketing Science*, 31(2), 293–316.
- SU, X. (2007): "Intertemporal pricing with strategic customer behavior," *Management Science*, 53(5), 726–741.
- SWEETING, A. (2012): "Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets," *Journal of Political Economy*, 120(6), 1133 – 1172.
- TALLURI, K., AND G. VAN RYZIN (2004): "Revenue Management Under a General Discrete Choice Model of Consumer Behavior," *Management Science*, 50(1), 15–33.
- TALLURI, K. T., AND G. J. VAN RYZIN (2005): *The theory and practice of revenue management*, vol. 68. Springer.
- VULCANO, G., G. VAN RYZIN, AND W. CHAAR (2010): "Choice-based revenue management: An empirical study of estimation and optimization," *Manufacturing & Service Operations Management*, 12(3), 371–392.
- ZHAO, W., AND Y.-S. ZHENG (2000): "Optimal Dynamic Pricing for Perishable Assets with Nonhomogeneous Demand," *Management Science*, 46(3), 375–388.

A Proofs

Proposition: With $\bar{\phi} = \Pr(p_{t+1} < p) \mathbb{E}[\alpha(p - p_{t+1}) | p_{t+1} < p]$, then all consumers will choose not to wait.

Proof: Take a consumer who wants to purchase today, i.e., $\varepsilon_0 < \beta - \alpha p + \varepsilon_1$. Then there exists a $\bar{p} > p$ such that $\varepsilon_0 = \beta - \alpha \bar{p} + \varepsilon_1$. The expected gain for this consumer waiting comes from prices dropping below p_t and from price increases up to the indifference point. If prices increase past \bar{p} , then ε_0 is preferred and there is no gain. Hence, the expected gains from waiting are

$$\begin{aligned} & \Pr(p_{t+1} < p) \mathbb{E}[\alpha(p - p_{t+1}) | p_{t+1} < p] \\ & + \Pr(p < p_{t+1} \leq \bar{p}) \mathbb{E}[\alpha(p - p_{t+1}) | p < p_{t+1} \leq \bar{p}] - \bar{\phi}. \end{aligned}$$

The first term above is equal to $\bar{\phi}$, and the second term is less than or equal to zero. Hence, waiting is not optimal for a consumer wishing to buy today.

Next, consider a consumer who prefers not to buy a ticket today, i.e., $\varepsilon_0 > \beta - \alpha p + \varepsilon_1$. Then there exists a $\underline{p} < p$ such that $\varepsilon_0 = \beta - \alpha \underline{p} + \varepsilon_1$. The gains from waiting come from price declines lower than the cutoff, and are equal to

$$\Pr(p_{t+1} < \underline{p}) \mathbb{E}[\beta - \alpha p_{t+1} + \varepsilon_1 - \varepsilon_0 | p_{t+1} < \underline{p}] - \bar{\phi}.$$

Applying the definition of $\bar{\phi}$, this is equivalent to

$$\Pr(p_{t+1} < \underline{p}) \mathbb{E}[\beta - \alpha p_{t+1} + \varepsilon_1 - \varepsilon_0 | p_{t+1} < \underline{p}] - \Pr(p_{t+1} < p) \mathbb{E}[\alpha(p - p_{t+1}) | p_{t+1} < p].$$

Define EG to be the expression above. Since $\underline{p} \leq p$, we have

$$\begin{aligned} \text{EG} & \leq \Pr(p_{t+1} < p) \left(\mathbb{E}[\beta - \alpha p_{t+1} + \varepsilon_1 - \varepsilon_0 | p_{t+1} < \underline{p}] - \mathbb{E}[\alpha(p - p_{t+1}) | p_{t+1} < p] \right) \\ & \leq \Pr(p_{t+1} < p) \left(\mathbb{E}[\beta - \alpha p_{t+1} + \varepsilon_1 - \varepsilon_0 | p_{t+1} < \underline{p}] - \mathbb{E}[\alpha(p - p_{t+1}) | p_{t+1} < \underline{p}] \right). \end{aligned}$$

Moving the expectation operator, the last line above equals

$$\Pr(p_{t+1} < p) \mathbb{E}[\beta - \alpha p_{t+1} + \varepsilon_1 - \varepsilon_0 - \alpha(p - p_{t+1}) | p_{t+1} < \underline{p}],$$

which can be simplified to $\Pr(p_{t+1} < p) \Pr(p_{t+1} < \underline{p}) (\beta - \alpha p + \varepsilon_1 - \varepsilon_0) \leq 0$, since $\beta - \alpha p + \varepsilon_1 - \varepsilon_0 < 0$ by assumption.

Hence, waiting is not optimal for a consumer wishing to not buy today. ■