DYNAMIC AIRLINE PRICING AND SEAT AVAILABILITY

By

Kevin R. Williams

August 2017
Revised May 2020

COWLES FOUNDATION DISCUSSION PAPER NO. 2103R

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

Yale School of Management and NBER*

May 2020†

Abstract

Airfares fluctuate over time due to both demand shocks and intertemporal variation in willingness to pay. I develop and estimate a model of dynamic airline pricing accounting for both forces with new flight-level data. With the model estimates, I disentangle key interactions between the arrival pattern of consumer types and scarcity of remaining capacity due to stochastic demand. I show that dynamic airline pricing expands output by lowering fares charged to early-arriving, price-sensitive customers. It also ensures seats for late-arriving travelers with the highest willingness to pay (e.g. business travelers) who are then charged high prices. I find that dynamic airline pricing increases total welfare relative to a more restrictive pricing regime. Finally, I show that abstracting from stochastic demand results in incorrect inferences regarding the extent to which airlines utilize intertemporal price discrimination.

JEL: L11, L12, L93

*`kevin.williams@yale.edu`

†I thank Judy Chevalier, Jim Dana, Tom Holmes, Olivia Natan, Aniko Öry, Hayden Parsley, Amil Petrin, Tom Quan, Timothy Schwieg, 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, Reserve Bank of Richmond, Indiana University, Indiana University - Kelley, Marketing Science, Stanford Institute for Theoretical Economics (SITE), and the University of Pennsylvania for comments. I also thank the Opportunity & Inclusive Growth Institute at the Federal Reserve Bank of Minneapolis for providing resources that supported this research. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

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]

The airline industry is well known for its complex intertemporal pricing dynamics. Airfares close to the departure date are high. The conventional view is that late shoppers are business travelers, and airlines capture their high willingness to pay through intertemporal price discrimination. In addition, airlines also adjust prices on a day-to-day basis, as capacity is limited and the demand for any given flight is uncertain. They may raise fares to avoid selling out flights in advance, or fares may fall from one day to the next, after a sequence of low demand realizations.

Decomposing the sources of price adjustments in airline markets is critical because they lead to conflicting predictions for welfare. Price adjustments in response to realizations of demand are welfare improving: they increase capacity utilization, and they save seats for business travelers who shop close to the departure date. However, price adjustments respond also to consumer preferences. Having saved seats for these price-insensitive customers, airlines then extract their surplus through high prices. If these adjustments were not possible, the prospect of extracting surplus from late-arriving customers can create the incentive to save an inefficient number of seats and charge an inefficiently high price. Thus, it is an empirical question whether dynamic airline pricing is on net welfare increasing.

In this paper, I examine how dynamic pricing—pricing that depends on both demand shocks and intertemporal variation in willingness to pay—allocates scarce capacity across

[1] Accessed through AirAsia.com's Investor Relations page entitled, "What is low cost?"
[2] Appeared on easyjet.com's FAQs. Accessed through "Low-Cost Carriers and Low Fares" and "Online Marketing: A Customer-Led Approach."

heterogeneous consumers in airline markets. I propose and estimate a model that combines features of stochastic demand and revenue management models from operations research with estimation techniques widely used in empirical economics research. I utilize novel data that track daily prices and seat availabilities for over twelve thousand flights in US monopoly markets. With the model estimates, I disentangle key interactions between the arrival pattern of consumer types and scarcity due to stochastic demand. I find that dynamic pricing increases output by offering discounts to early-arriving, price sensitive consumers, while also ensuring seat availability for late-arriving, price insensitive consumers. I find that dynamic pricing increases welfare in the monopoly markets studied.

Existing research separately examines intertemporal price discrimination and dynamic adjustment to stochastic demand 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) use regression analysis and find that ticket characteristics such as advance-purchase discount (APD) requirements explain much of the dispersion in fares. Lazarev (2013) estimates a model of intertemporal price discrimination and finds a substantial role for this force. Escobari (2012) and Alderighi, Nicolini, and Piga (2015) find evidence that airlines face stochastic demand and that prices respond to remaining capacity.

An investigation of dynamic airline pricing requires a detailed data set of flight-level prices and transactions. However, the standard airline data sets used in economic studies (e.g., Goolsbee and Syverson (2008); Gerardi and Shapiro (2009)) are either monthly or quarterly. Papers have analyzed high-frequency fares—for example, McAfee and Te Velde (2006), or a portion of transactions—for example, Puller, Sengupta, and Wiggins (2015). One of the contributions of this paper is a set of new stylized facts assembled from novel fare (prices) and seat availability (quantities) data. The sample contains thousands of flights in US monopoly markets, where each flight is tracked for up to sixty days. In total, the sample contains over 700,000 observations.

Descriptive evidence provide new insights into the use of both pricing forces. I observe both positive and negative price fluctuations. Fare increases occur after observed bookings. Fares stay constant, or even decline, in the absence of sales. However, the trajectory of fares is overwhelmingly positive. Fares typically double in the sixty days before departure and, regardless of sales, tend to sharply increase close to the departure date. This is consistent with intertemporal price discrimination.

I develop a structural model to estimate both the unobserved arrival process of customers and their preferences. I do so by combining features of dynamic pricing and stochastic demand models commonly used in operations research, including Talluri and Van Ryzin (2004) and Vulcano, van Ryzin, and Chaar (2010), with elements of the discrete unobserved heterogeneity utility specification of Berry, Carnall, and Spiller (2006). Discrete heterogeneity demand models are commonly used in airline studies—for example, in Berry and Jia (2010). Although I tailor the model using institutional features of airline markets, the methodology can be useful for analyzing any perishable goods market with a deadline.

The model 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, corresponding to business and leisure travelers, is allowed to change over time. The model timing is discrete. Each day before departure, the number of business and leisure arrivals is distributed according to independent Poisson distributions with time-dependent arrival rates. Consumers know their preferences and solve a static utility maximization problem. On the supply side, the monopoly solves a finite-horizon, stochastic dynamic programming problem. Within a period, the firm first chooses a price, consumer demand is realized, and then the capacity constraint is updated. Time moves forward, and the process repeats through the perishability date.

This paper proposes explicitly modeling the pricing decision of the firm to address the well-known issue of missing "no purchase" data, or the number of arrivals who opted not to purchase. The identification assumption is that preferences for flights evolve in the same predictable way, but demand shocks can vary. This results in variation in seats sold toward the deadline, and the firm's response to these shocks informs the magnitude of stochastic demand. The model estimates are market specific. They generally suggest that a significant shift in arriving consumer types over time and that stochastic demand is a meaningful driver of the variation in sales.

The estimated model is used to establish two key points about the interaction between the arrival pattern of consumer types and scarcity due to demand shocks. First, through a series of counterfactuals, I decompose the relative importance of intertemporal price discrimination and dynamic adjustment to stochastic demand. I show dynamic adjustment complements intertemporal price discrimination in the airline industry because price-insensitive consumers (i.e., business travelers) tend to buy tickets close to the departure date. Dynamic pricing expands output by lowering fares offered to leisure travelers, but it also ensures seat availability for business travelers who are then charged high prices. Uniform pricing increases aggregate consumer surplus, however, the gains are mitigated because it also results in additional unused capacity. I find that total welfare is higher under dynamic pricing compared to more restrictive pricing regimes in the monopoly markets studied.[3]

Second, I show that managing remaining capacity in airline markets is critical because demand is stochastic and abstracting from its presence affects our understanding of how airlines use intertemporal price discrimination. Over one third of the revenue gains of dynamic pricing over uniform pricing come from the ability to respond to demand shocks. The remaining two thirds come from the ability of the firm to extract surplus through intertemporal price discrimination. According to the model, only 22 percent of the observed flights are projected to be unaffected by scarcity in the initial period. Finally, I show that the presence of stochastic demand and scarcity affects our understanding of the use of intertemporal price discrimination. By abstracting from stochastic demand, the opportunity cost of selling a seat is the same regardless of the date of purchase. In reality, opportunity costs reflect demand shocks and the resolution of uncertainty toward the perishability date. Therefore, it is difficult to identify intertemporal price discrimination without knowing how firms respond to stochastic demand. Empirical procedures that abstract from stochastic demand will systematically overstate consumers' price insensitivity because upward pressure on prices due to scarcity will be inferred as inelastic demand. This bias becomes pronounced close to the departure date.

[3]The proposed model does not fit into existing theoretical work on the welfare effects of monopolistic price discrimination, including Aguirre, Cowan, and Vickers (2010) and Bergemann, Brooks, and Morris (2015), because capacity is constrained and the markets studied are sequential.

1.1 Related Literature

This paper contributes to growing literatures in economics and operations research that study intertemporal price discrimination and revenue management. Intertemporal price discrimination can be found in many markets, including video games (Nair, 2007), Broadway theater (Leslie, 2004), storable goods (Hendel and Nevo, 2013), and concerts (Courty and Pagliero, 2012).[4] Contributions to the study of intertemporal price discrimination in airline markets include Lazarev (2013) and Puller, Sengupta, and Wiggins (2015).

Revenue management (RM) can refer to the dynamic adjustment of either product availability or prices (and sometimes both).[5] Several studies characterize optimal pricing (either analytically or numerically) of RM models with Poisson arrivals (Gallego and Van Ryzin, 1994; McAfee and Te Velde, 2006; Zhao and Zheng, 2000; Talluri and Van Ryzin, 2004; Vulcano, van Ryzin, and Chaar, 2010). Relative to these studies, this paper proposes a model of both time-varying arrivals and multiple consumer types in discrete time. Dana (1999) shows in a theoretical model that business consumers may benefit from RM.

The increasing trajectory of prices observed in airline markets reduces the incentives for consumers to wait to buy, but existing research has shown strategic buyers to be an important consideration. Theoretical contributions include Su (2007), Board and Skrzypacz (2016), Gershkov, Moldovanu, and Strack (2018), and Dilmé and Li (2019). Hendel and Nevo (2006) consider stockpiling and show that dynamic demand impact demand estimates. In the context of major league baseball tickets, Sweeting (2012) estimates a model of strategic delay with search costs. He finds that this leads buyers to sort on participation timing, and he shows dynamic pricing is valuable in this context. Nair (2007) shows that profit losses can be large when firms do not take into account forward-looking behavior. This result is found in an environment where demand becomes more elastic over time; Soysal and Krishnamurthi (2012) study markdowns and show that the incentives to wait decrease because of stock-outs.

Finally, concurrent works provide new insights on the effects of dynamic pricing in

[4]Lambrecht et. al. (2012) provide an overview of empirical work on price discrimination more broadly.

[5]The former is commonly called quantity-based RM; the latter is commonly referred to as price-based RM. Elmaghraby and Keskinocak (2003) and Talluri and Van Ryzin (2005) provide an overview of RM work in operations.

various contexts. Cho et. al. (2018) quantify the gains from dynamic pricing in the hotel industry. They also capture competitive pricing pressures. Chen (2018) examines competitive dynamics in airlines. Aryal, Murry, and Williams (2018) utilize survey data to examine dynamic pricing of different ticket qualities in international airline markets. Finally, D'Haultfœuille et. al. (2018) quantify the effects of revenue management in the French railway system. They also examine the role of demand uncertainty and show that RM results in significant gains relative to uniform pricing.

2 Industry Description

In this section, I provide a short overview of airline pricing practices to motivate my empirical approach. Interested readers can find additional details of revenue management algorithms and practices in McGill and Van Ryzin (1999) and Gallego and Topaloglu (2019).

For a flight, observed prices over time depend on three key inputs: (1) plane capacity, (2) filed fares, and (3) inventory allocation for filed fares. Input (2) means the prices for flights, and input (3) corresponds to the number of seats the airline is willing to sell at prices provided by (2). Each of these decisions is made by separate departments that use different algorithms and methods to determine the most profitable decision, holding the other departments' choices fixed.

A carrier's network-planning group determines which markets are served and the total amount of capacity assigned to them. These decisions occur well in advance of the departure date. Typically, flights are available for purchase over 300 days prior to the departure date; however, adjustments can be made closer to the departure date. This includes entry or exit decisions or a change in gauge of aircraft. In the data collected, I observe that 2.5 percent of flights experience a change in aircraft in the sixty days before departure. Capacity changes overwhelmingly occur close to the departure date: 75 percent of occurrences happen within the two days before departure. Yet these changes do not seem to be associated with flight loads.[6] It is more likely that capacity changes close to the

[6] I cannot reject the null hypothesis that flights that see an upgauge (increase in capacity) have flight loads higher than the average load factor for that particular flight number and vice versa. In the former case, $t = -0.996$; in the latter case, $t = 0.614$. Note that these test statistics go in the opposite direction from what is expected—flights that experience an upgauge actually have load factors slightly lower than the flight average;

departure date occur for operational reasons. I use this finding to motivate my choice to abstract from capacity decisions and focus instead on pricing given remaining capacity.

The pricing department determines filed fares, or a set of fares and associated ticket restrictions that may be offered to consumers. This includes prices of refundable and non-refundable tickets, as well as first-class, economy-class, and basic-economy tickets. A common ticket restriction applied to fares is an advance-purchase discount requirement, a restriction that requires consumers to purchase by a certain day-before-departure. APDs are commonly used three, seven, ten, fourteen, twenty-one, and thirty days before departure, depending on the itinerary and carrier. A fare class (or booking class) is a single- or double-letter code to denote broad ticket characteristics—deeply discounted economy versus full-fare economy, for example. When the additional ticket restrictions are incorporated, this results in a fare basis code.

Fare Basis	Airline	Fare Class	Trip Type	Fare	Adv Purchase Req
LH4OASBN	Alaska	L	One-Way	$174.60	14
LH4OASMN	Alaska	L	One-Way	$189.60	14
QH4OASMN	Alaska	Q	One-Way	$217.60	14
YH0OASMR	Alaska	Y	One-Way	$334.00	–

In this example, there are two L fares filed, one saver economy fare and one economy fare, each with a fourteen-day APD requirement. The two L fares have different fare basis codes. The third fare is a fourteen-day APD Q-class fare. There is a fourth fare, an unrestricted Y-class economy fare.

The pricing group creates a menu of fares for each market. This means the potential number of fares for a particular itinerary is discrete. However, the set may change over time if the pricing department files updated fares. I incorporate this feature in the empirical model by having the firm choose among a discrete set of fares.

Finally, the revenue management department determines fare availability. This process involves setting the number of seats available for purchase for each fare class. In order to dynamically adjust the allocations of fares based on bookings and updated forecasts, the RM department formulates complicated dynamic programming problems and uses techniques developed in operations research, including the well-known ESMR-a and ESMR-b

flights that experience a downgauge have a load factor slightly higher than the flight average.

heuristics (Belobaba, 1987, 1989, 1992; Belobaba and Weatherford, 1996), in order to make them tractable. Phillips (2005) provides an overview of these approaches. Importantly, the allocation decision takes potential fares and forecasts as inputs. Allocations are updated toward the perishability date based on demand realizations. My model of the firm also takes the forecasts and fares as inputs; optimization is assumed to occur daily.

RM systems are designed such that several fares are available at any given point in time. Continuing the example above, airlines would surely be willing to sell all available seats as expensive Y-class fares, but chances are few consumers would purchase at the highest of prices. As a consequence, airlines offer less-expensive fares under different fare classes such as L and Q. An allocation for a at flight a particular point in time may be (Y:10, Q: 2, L: 1). If a consumer purchases the lowest available class (L), the allocation will likely become (Y:9, Q: 1, L: 0), absent inputs from the pricing group or reoptimization by the RM group. The lowest available fare will then become a Q-class ticket.[7]

Combining the decisions of the pricing group and revenue management, this implies that airline prices depend not only on ticket restrictions such as APDs but also on the current seat allocation of each flight, which depends on demand realizations. The data reveal that both restrictions and bookings influence fares—flights with bookings today tend to become more expensive, whereas flights without daily bookings tend to see prices either stay the same or decrease.

Without access to individual-level purchase data and inventory allocations over time (these data are proprietary and available data are top coded, as described in the next section), I do not pursue modeling inventory allocation. Instead, I simplify the problem into a dynamic pricing problem under a number of assumptions, including that consumers purchase only the cheapest available fare. I observe that daily demands for flights are low—less than one seat per flight per day. I argue that this finding removes the need to model the nesting structure of fare buckets, as in the example above.

[7] See Phillips (2005) for additional information on this approach, which is called nesting.

3 Data

I create several original data sets for this study. The data are collected from travel management companies, travel meta-search engines, and airline websites.[8] I collect and merge together three pieces of information.

First, I collect daily prices at the itinerary level, with itinerary defined as an origin-destination, airline, flight number, and departure-date combination. I focus on one-way fares, as for almost all of the sample, round-trip prices are equal to the sum of segment prices. Most analysis concentrates on the cheapest available economy-class ticket for purchase. Whenever possible, I also collect prices for different versions of tickets, such as economy versus first class and restricted economy versus unrestricted economy.

Second, I collect fare class availability information. These data record censored fare class allocations for each flight. For example, on a date prior to departure, I may observe G5. This means the active G-class fare has five available seats. The information is censored in that availabilities are top coded, typically at seven or nine, depending on the carrier. As another example, Y9 means the Y-class fare—almost always the most expensive coach fare available—has at least nine seats available. I utilize this information to confirm when flights are sold out. Whenever possible, I also collect filed fare restrictions. These data record any advance-purchase discount requirements or other restrictions on the ticket. Continuing the previous example, a filed fare in the data is G21JN5. This G-class fare includes a twenty-one-day advance-purchase requirement. The proposed model accommodates the use of APDs.

Third, I collect airline seat maps, which are graphical representations of available and occupied seats flights. By collecting airline seat maps over time and tracking changes to individual seats across consecutive days, I obtain a measure of daily bookings. These data provide quantity information. In Appendix C, I provide evidence two ways that suggests the measurement error associated with using seat maps may be small.

[8]More specifically, the data come from Alaska Airlines, BCD Travel, ExpertFlyer, Fare Compare, JetBlue Airways, United Airlines, and Yapta. The airline websites provide a wealth of information, including seat availabilities, seat maps and fares. ExpertFlyer reports filed fares, seat availabilities, and seat maps; BCD Travel reports seat availabilities; Fare Compare reports filed fares, and Yapta tracks daily fares. Data were collected in 2012 and again in 2019.

In the following subsections, I discuss route selection (Section 3.1) and then document a set of new descriptive facts (Section 3.2).

3.1 Route Selection

Using the publicly available DB1B tables, I select markets in which to study. These data are frequently used to study airline markets. The DB1B tables contain a 10-percent sample of domestic US ticket purchases. The data are at the quarterly level. The data contain neither the date flown nor the purchase date, hence the need to collect data for this study. I define a market in the DB1B as an origin-destination (OD), quarter. With the DB1B data, I single out markets where

(i) there is only one carrier operating nonstop;

(ii) there is no nearby alternative airport serving the same destination;

(iii) total quarterly traffic is greater than 600 passengers;

(iv) total quarterly traffic is less than 45,000 passengers;

(v) a significant portion of traffic is nonstop;

(vi) a significant portion of traffic is not connecting.

Criteria (i) and (ii) narrow the focus to monopoly markets in terms of nonstop flight options. Criteria (iii) and (iv) remove infrequently served markets, and the upper limit on traffic keeps data collection manageable. When I implement these criteria, the resulting markets make up roughly 10 percent of OD traffic in the United States. In addition, quarterly revenues for these markets are roughly $2.5 billion. Criterion (v) is important because it addresses the potential for alternative flight options, such as one-stop connections. Criterion (vi) is equally important because it addresses how fares are assigned to observed changes in remaining capacity.

Criteria (v) and (vi) are negatively correlated, meaning routes with high nonstop traffic percentages typically have low percentages of non-connecting traffic. This is because ODs with very high nonstop traffic percentages tend to be short-distance flights to hubs, after which consumers connect to other flights. Without individual-level data, it is impossible to know the itinerary purchased for each observed booking. Moreover, given that ODs with the highest concentration of nonstop traffic are more than twice as short—comparing above the 95th percentile with below the 95th percentile—it is also possible that alternative modes of transportation, such as taking a bus or train, are valid substitutes to flying.

I collect data on fifty OD pairs that satisfy the selection criteria above. In addition, to compare the descriptive evidence, I select six duopoly markets (Section 3.2).[9] Appendix B presents additional route selection information, market-level statistics, and comparisons with the entire DB1B sample.

All of the routes studied either originate or end at Boston, MA; Portland, OR; or Seattle, WA. Almost all of the data collected study markets operated by either Alaska or JetBlue. Several features of the sample are worth noting. First, both of these carriers price 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. I observe no length-of-stay requirements or Saturday-night stay-overs. 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.[10] I will use this feature of the data to simplify the pricing problem presented in the next section. Third, several of the markets studied feature coach-only flights. This feature allows for investigating all sales and also controls for one aspect of versioning (first class versus economy class). Finally, the sample focuses on airlines that allow consumers to select seats before departure; many carriers now charge fees to choose seats when traveling on restrictive coach tickets.[11]

In contrast with Jetblue, Alaska does offer first class in several of the markets studied—first class appears in 58 percent of the sample, with the average cabin being twelve seats of the plane. I provide some descriptive analysis of first-class pricing, but I do not pursue

[9]Two markets, (Boston, MA - Kansas City, MO) and (Boston, MA - Seattle, WA) are both. The former market originally had nonstop service offered by Delta Air Lines and Frontier Airlines. Frontier exited and Delta became the only carrier flying nonstop. The latter market very briefly was observed to have service offered solely by Alaska prior to the entry of JetBlue.

[10]In 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."

[11]The JetBlue data were collected before the introduction of Blue Basic seats, which feature a fee to select seats. This is also true for Delta Air Lines. Alaska Airlines' restrictive coach tickets are called Saver fares. These fares do allow for limited seat selection in the coach cabin. I observe availability of these seats in 98 percent of seat maps.

versioning in the model. Alaska does allow for overselling, but I note that among the major airlines, Alaska Airlines has an average denied boarding rate.[12]

3.2 Descriptive Evidence

3.2.1 Summary Statistics

I capture fares and seat availabilities for over 12,000 flights, each tracked for the last sixty days before departure. Data collection occurred over two six-month periods (March 2012-August 2012, March 2019-August 2019). In total, I obtain 734,689 observations for analysis, as well as over five million connecting itinerary prices.

Table 1: Summary Statistics for the Data Sample

Variable	Mean	Std. Dev.	Median	5th pctile	95th pctile
Oneway Fare ($)	233.95	139.72	194.00	89.00	504.00
Load Factor (%)	86.95	13.32	91.33	59.21	100.00
Daily Fare Change ($)	3.32	31.94	0.00	-5.00	50.00
Daily Booking Rate	0.73	1.96	0.00	0.00	4.00
Unique Fares (per itin.)	6.98	2.31	7.00	4.00	11.00

Note: Summary statistics for 12,052 flights tracked between 3/2/2012-8/24/2012 and 3/21/2019-8/31/2019. Each flight is tracked for sixty days before departure. The total number of observations is 734,689. Load Factor is reported between zero and 100 the day of departure. The daily booking rate and daily fare change compares consecutive days.

Summary statistics for the data sample appear in Table 1. The average one-way ticket in the sample is $234. This is higher than the average price calculated from the publicly available DB1B tables (Table 12); however, recall that these gathered prices cover the sixty days before departure and also include non-transacted prices.

Reported load factor is the number of occupied seats divided by capacity on the day flights leave and is reported between zero and one hundred. In my sample, the average load factor is 87 percent, ranging from 69 percent to 94 percent, by market. I observe that 9 percent of flights sell out. The median number of daily departures is one and the mean is two.

[12]Source: Air Travel Consumer Report, accessed February 2020.

There is considerable variation in load factor within a market: this supports the presence of flight-level demand shocks. The coefficient of variation (CV) of within-market load factors ranges between 0.04 and 0.27. CVs are higher well in advance of the departure date; the reduction over time is consistent with price adjustments to fill unsold seats. The R^2 of a regression of ending load factor on market-flight number and departure date (subsuming seasonality and day-of-week indicators) fixed effects is only 0.5. This motivates the use of a stochastic demand model of flight-level demand.

The booking rate in Table 1 corresponds to the mean difference in occupied seats across consecutive days. I find the average booking rate to be 0.73 seats per flight-day. 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. This finding shows that airline markets are associated with low daily demand. Fifty-nine percent of the seat maps in the sample do not change across consecutive days. This requires the demand estimation technique to confront the fact that there is a significant number of zero sales.

On average, each itinerary reaches seven unique fares and experiences 10.6 fare changes. This implies that fares fluctuate up and down usually a few times within sixty days. Since the number of fares chosen is small, I will use this institutional feature in the model. Fares will be chosen from a discrete set.

I use data on individual seat assignments in order to gauge the number of passengers per booking—the idea being that adjacent seats becoming occupied likely constitutes a party traveling together.[13] I estimate the average number of passengers per booking to be 1.37. This motivates the unit demand assumption in the consumer demand model.

3.2.2 Dynamic Prices

Figure 1 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

[13]Each row in the data has at most six seats, and I assume whenever more than two seats in row become occupied, this is a party traveling together. This occurs in less than 8 percent of bookings. For rows in which two seats become occupied, I check if the seats are adjacent. Seats with passengers or space in-between are assumed to be two single-passenger bookings. This removes 18 percent of the two-passenger bookings. Thus, as a potential lower bound, I find that 55 percent of passengers, or 75 percent of bookings, are single passenger bookings.

Figure 1: Frequency and Magnitude of Fare Changes by Day Before Departure

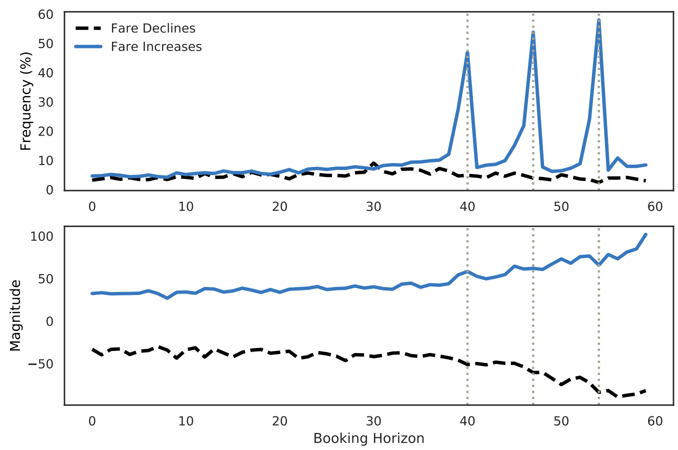

Note: The top panel shows the percentage of itineraries that see fares increase or decrease by day before departure. The lower panel plots the magnitude of the fare declines and increases by day before departure. The vertical lines correspond to advance-purchase discount periods (fare fences).

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 top plot, forty days prior to departure (t=20), roughly 5 percent of fares increase and 5 percent of fares decrease. The remaining 90 percent of fares are held constant. Moving to the bottom panel, the magnitude of fare increases and declines forty days out is roughly \$50. The top panel confirms fares hikes and declines occur throughout time. Note that well before the departure date, the number of fare hikes and the number of fare declines are roughly even.

There are three noticeable jumps in the top panel in Figure 1, indicating fare hikes. These jumps occur seven, fourteen, and twenty-one days before departure, or when the advance-purchase discounts placed on many tickets expire. A few of the markets experience APD restrictions three days before departure.[14] The use of APDs is consistent with the use of intertemporal price discrimination. Surprisingly, the use of APDs is not universal. Just under 50 percent of itineraries experience fare hikes at twenty-one days, and just over 50 percent increase at fourteen days. Nearly 60 percent of itineraries see an increase in fares when crossing the seven-day APD requirement.

[14] Advance purchase discounts are sometimes placed one, four, ten, thirty, and sixty days before departure, but this is not the case for the data I collect.

Figure 2: Mean Fare and Load Factor by Day Before Departure

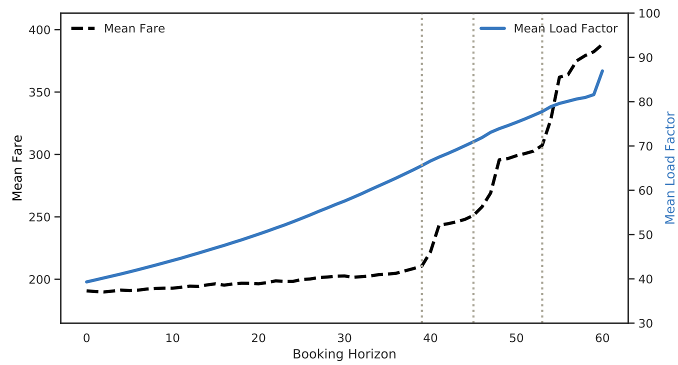

Note: Average fare and load factor by day before departure. The vertical lines correspond to advance-purchase discount periods (fare fences).

Figure 2 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 \$200 to over \$400 in sixty days. The noticeable jumps in the fare time series occur when crossing the APD fences noted in Figure 1. At sixty days before departure, roughly 40 percent of seats are already occupied. The booking curve for flights in the sample is smooth across time and starts to level off around 80 percent a few days before departure. There is a spike in load factor, of around 5 percent the day of departure. This spike could be driven by a combination of measurement error (consumers who were not assigned seats in advance are assigned seats at check-in) and last-minute bookings. I show in Appendix D that on the last day before departure, there is also a sharp decline in economy inventory, which indicates that last-minute bookings do occur.[15]

There is considerable variation in pricing across markets not shown in the aggregate statistic shown in Figure 2. Figure 11 and Figure 12 in Appendix A plot average prices over time as well as the average percent change in prices over time for each market. Levels of fares, the timing of APDs, and the significance of APDs vary by market. These figures suggest the need to include market-specific parameters.

[15]Three markets (Lihue, HI - Portland, OR; Palm Springs, CA - Portland, OR; Santa Barbara, CA - Seattle, WA) are observed to have large seat map changes at the deadline even though the seat maps allow for seat selection regardless of ticket type in over 99 percent of observations. It may be that these spikes represent group bookings or bookings with travel agencies that do not assign seats until check in. For these markets, I assume the spikes reflect bookings made before the collection period as the timing of purchase is not observed.

Figure 3: Fare Response to Sales by Day Before Departure

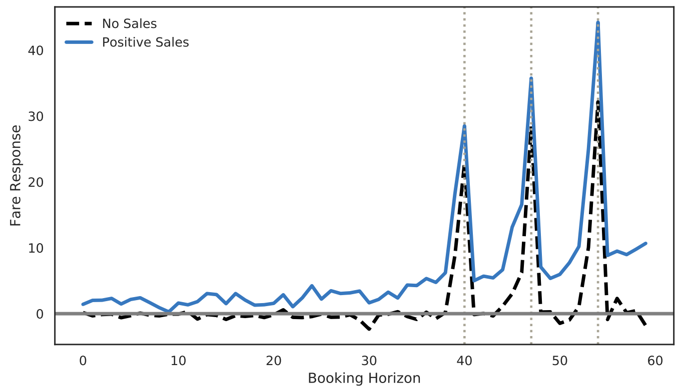

Note: Average fare changes as a response to sales by day before departure. The vertical lines correspond to advance-purchase discount periods (fare fences). The horizontal line indicates no fare response.

Figure 3 establishes a new important link between daily sales and daily price adjustments. The graph separates out two scenarios: (1) a flight experiences positive sales in the previous period; and (2) there are no sales in the previous period. Critically, the graph demonstrates that fares respond to both scarcity and time. It suggests an important interaction between the presence of demand shocks and intertemporal variation in willingness to pay. Conditional on positive sales, capacity becomes more scarce, and prices increase. I find that prices stay constant or decrease when sales do not occur, reflecting the declining opportunity cost of capacity. Both of these price movements are consistent with stochastic demand pricing models. However, close to the departure date and regardless of sales, prices increase. This suggests late-arriving consumers are less price-sensitive and airlines capture their high willingness to pay with intertemporal price discrimination.[16]

These pricing patterns are not limited to economy tickets in monopoly airline markets. In Appendix D, I show that competitive airline markets exhibit similar pricing patterns: fares adjust upward or downward depending on bookings and fares adjust upward regardless of bookings close to the departure date. These patterns are also shown to exist when a carrier offers different ticket qualities, such as first class and economy class. Hence,

Two of these markets also have irregular service.

[16]This was originally pointed out by McAfee and Te Velde (2006). Although, stochastic demand models can result in increasing price paths, they argue that the magnitude of observed price hikes suggest later arrivals are less price sensitive.

the modeling approach and results found in this paper are relevant for these important extensions.[17]

Finally, although fares do occasionally decline, the trajectory is overwhelmingly positive. This greatly reduces the incentive to wait to purchase, conditional on knowing preferences. We may be concerned that consumers strategically time their purchases in order to avoid fare hikes. In Appendix E, I investigate bunching in bookings. Most of the evidence suggests that this is not a concern. Bookings slightly decline (0.08 seats per day) the day after the 7-day APD expires. However, the the booking rate then returns to the same level the following day, when prices are just as high. This motivates Poisson demand.

4 An Empirical Model of Dynamic Airline Pricing

4.1 Model Overview

A monopolist airline offers a flight for sale in a series of sequential markets. More precisely, I will define the markets for a flight on a particular departure date, and I will abstract away from potential correlations in demands across departure dates and other flight options, including connecting flights and other nonstop itineraries. The sales process for every market evolves over a finite and discrete time horizon $t \in \{0, \ldots, T\}$. Period 0 corresponds to the first sales period, and period T corresponds to the day the flight leaves. Initial capacity for the flight is exogenous, and the firm is not allowed to oversell. Unsold capacity on the day of the flight ($t = T$) is scraped with zero value. The only costs modeled are the opportunity costs of remaining capacity, and all other costs are normalized to zero.

Each period t, the airline first offers a single price for the flight, and then consumers arrive according to a stochastic process specified in the next subsection. Each arriving consumer is either a business traveler or a leisure traveler; business travelers are less price sensitive than leisure travelers, and the proportion of each type is allowed to change over time. Note that the terms "business" and "leisure" are used simply to describe a consumer

[17] In Appendix D, I also investigate if demand shocks affect prices of alternative flights, including connections and other nonstop flights when a carrier offers more than a single daily frequency. I present evidence that nonstop bookings do not affect connecting prices in the same way as nonstop prices. If a carrier offers two flights per day, a booking on one flight does not increase the price of the other flight option.

type; they do not identify consumers based on a travel need.[18] Upon entering the market, all uncertainty about travel preferences is resolved. This approach differs from earlier theoretical work such as Gale and Holmes (1993), as well as some empirical work such as Lazarev (2013), 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.

If demand exceeds remaining capacity, tickets are randomly rationed. Consumers who are not selected receive the outside option. This ensures that the capacity constraint is not violated. I also assume that passengers do not cancel tickets, as the average number of cancellations per flight in the data is less than two. Thus, remaining capacity is monotonically decreasing. After tickets are sold in a given period, the capacity constraint is updated, and the firm again chooses a fare to offer. This process repeats until the perishability date. The firm is forward looking and solves the finite horizon, dynamic program.

4.2 Demand

Each day before the flight leaves, $t = 0, 1, ..., T$, a stochastic process brings a discrete number of new consumers to the market. $\tilde{M}_t$ denotes the arrival draw. 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 ($x_{jt} \in \mathbb{R}^K$) and price ($p_{jt} > 0$), respectively.

I assume utility is linear in product characteristics and price. If consumer i chooses to purchase a ticket on flight j, she receives utility $u_{ijt} = x_{jt}\beta_i - \alpha_i p_{jt} + \varepsilon_{ijt}$. If she chooses not to fly, she receives normalized utility $u_{i0t} = \varepsilon_{i0t}$. Arriving consumers solve a straightforward maximization problem: consumer i selects flight j if and only if $u_{ijt} \geq u_{i0t}$.

Define $y_t = \left(\alpha_i, \beta_i, \varepsilon_{ijt}, \varepsilon_{i0t}\right)_{i \in 1,..,\tilde{M}_t}$ to be the vector of preferences for the consumers who enter the market. Suppressing the notation on product characteristics for the rest of this

[18] Booking websites and surveys oftentimes ask the reason for travel. Typically, the two options are for business or for leisure. The model estimates two consumer types that need not coincide with these two rationales for travel.

section, demand for flight j at t is defined as

$$Q_{jt}(p, y_t) := \sum_{i=0}^{\tilde{M}_t} 1\Big[u_{ijt} \geq u_{i0t}\Big] \in \{0, ..., \tilde{M}_t\},$$

where $1(\cdot)$ denotes the indicator function. Demand is integer valued; however, it may be the case that there are more consumers who want to travel than there are seats remaining. That is, $Q_{jt}(p, y) > c_{jt}$, where c_{jt} 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 to fly or not. After consumers make their decisions, the capacity constraint is checked. If demand exceeds remaining capacity for the flight, consumers who chose to travel are randomly shuffled. The first c_{jt} are selected, and the rest receive their outside options. Although this is a restrictive assumption, recall that the average number of seats sold per flight day is less than one.

By abstracting from the ability to oversell and incorporating the rationing rule, expected sales are formed by integrating over the distribution of y_t,

$$Q^e_{jt}(p; c) = \int_{y_t} \min\Big(Q_{jt}(p, y_t), c\Big) dF_t(y_t).$$

Before continuing, note that 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 Alaska and JetBlue price at the segment level, there is no measurement error in this procedure. That is, a consumer pays the same price for two one-way tickets as he or she would for a round-trip ticket.

I incorporate a number of parametric assumptions. First, following McFadden (1973), I assume that the idiosyncratic preferences of consumers are independently and identically distributed according to a Type-1 Extreme Value (T1EV) distribution. This assumption implies that the individual choice probabilities are equal to

$$\pi^i_{jt}\left(p_{jt}\right) = \frac{\exp(x_{jt}\beta_i - \alpha_i p_{jt})}{1 + \exp(x_{jt}\beta_i - \alpha_i p_{jt})}.$$

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, $\gamma_t \pi^B_{jt}$ defines the purchase probability that a consumer is of the business type and wants to purchase a ticket; $(1-\gamma_t)\pi^L_{jt}$ is similarly defined. Hence, integrating over consumer types, product shares is equal to

$$\pi_{jt}\left(p_{jt}\right) = \gamma_t \pi^B_{jt}\left(p_{jt}\right) + (1-\gamma_t)\pi^L_{jt}\left(p_{jt}\right).$$

Next, I assume that consumers arrive according to a Poisson process, $\tilde{M}_t \sim \text{Poisson}_t(\mu_t)$. The arrival rates, μ_t, are also allowed to change over time. Hence, daily demands will depend on both the arrival process as well as preferences of consumers entering the market. Conditional on price, $Q_{jt} \sim \text{Poisson}_t(\mu_t \pi_{jt})$. The probability that q seats are demanded on flight j at time t are equal to

$$Pr_t\left(Q_{jt} = q\,;\, p_{jt}\right) = \frac{\left(\mu_t \pi_{jt}\right)^q \exp\left(-\mu_t \pi_{jt}\right)}{q!}.$$

Finally, using the probability distribution on the number of seats demanded, expected demand can be written as[19]

$$\begin{aligned} Q^e_{jt}(p_{jt}; c_{jt}) &= \sum_{q=0}^{c_{jt}-1} Pr_t\left(Q_{jt} = q\,;\, p_{jt}\right) q + \sum_{q=c_{jt}}^{\infty} Pr_t\left(Q_{jt} = c_{jt}\,;\, p_{jt}\right) c_{jt}. \\ &= \sum_{q=0}^{c_{jt}-1} \frac{\left(\mu_t \pi_{jt}\right)^q \exp\left(-\mu_t \pi_{jt}\right)}{q!} q + \sum_{q=c_{jt}}^{\infty} \frac{\left(\mu_t \pi_{jt}\right)^q \exp\left(-\mu_t \pi_{jt}\right)}{q!} c_{jt}. \end{aligned}$$

[19]This is can be equivalently written as

$$Q^e_{jt}(p_t; c_{jt}) = \sum_{q=0}^{c_{jt}-1} \frac{\left(\mu_t \pi_{jt}\right)^q \exp\left(-\mu_t \pi_{jt}\right)}{q!} q + \left(1 - \sum_{q=0}^{c_{jt}-1} \frac{\left(\mu_t \pi_{jt}\right)^q \exp\left(-\mu_t \pi_{jt}\right)}{q!}\right) c_{jt}$$

because the probability of at least c_{jt} seats being demanded is equivalent to one minus the probability that fewer than c_{jt} seats are demanded.

4.3 Monopoly Pricing Problem

The monopolist maximizes expected revenues of the flight over a series of sequential markets. In each market, the firm chooses to offer a single price before the arrival of customers. Because of the institutional feature that airfares are discrete, I assume that the firm chooses prices from a discrete set, denoted $A(t)$. The set changes over time because of fare restrictions such as advance-purchase discount requirements.[20]

The pricing decision is based on the states of the flight: seats remaining; time left to sell; flight characteristics (notation suppressed); and idiosyncratic shocks $\omega_t \in \mathbb{R}^{A(t)}$, which are assumed to be independently and identically distributed following a Type-1 Extreme Value (T1EV) distribution, with scale parameter σ. These shocks are assumed to be additively separable to the remainder of the per-period payoff function, which are expected revenues (suppressing index j),

$$R_t^e(p_t; c_t) = p_t \cdot Q_t^e(p_t; c_t).$$

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

$$V_t(c_t, \omega_t) = \max_{p \in A(t)} \left(R_t^e(p; c_t) + \omega_{tp} + \delta \int_{\omega_{t+1}, c_{t+1} \,|\, c_t, \omega_t, p} V_{t+1}(c_{t+1}, \omega_{t+1}) dH_t(\omega_{t+1}, c_{t+1} \,|\, \omega_t, p, c_t) \right).$$

Because the firm cannot oversell, capacity transitions as $c_{t+1} = c_t - \min\left\{Q_t, c_t\right\}$, where Q_t is the realized demand draw. The firm faces two boundary conditions. The first is that once the airline hits the capacity constraint, it can no longer sell seats for that flight. The second is that unsold seats are scrapped with zero value.

I follow Rust (1987) and assume that conditional independence is satisfied. This means that the transition probabilities are equal to $h_t(\omega_{t+1}, c_{t+1} \,|\, \omega_t, p_t, c_t) = g(\omega_{t+1}) f_t(c_{t+1} \,|\, p_t, c_t)$.

[20] In principle, the model can be extended to an environment where the monopolist offers multiple flights (J). Two assumptions that can be used so that the model closely follows the exposition here are: (1) consumers do not know remaining capacities when solving the utility maximization problem, (2) when capacity is rationed, consumers not selected receive the outside option. The first assumption addresses that consumers may select less preferred options if the probability of getting a seat is higher. The second assumption implies that conditional on price, Q_{jt} is independent of $Q_{j't}$ for $j' \neq j$ and that $Q_{jt} \sim \text{Poisson}_t(\mu_t \pi_{jt})$. The complexity of the dynamic program increases by $\dim[A(\cdot)]^{(J-1)}$ relative to the complexity of the single-flight problem.

The capacity transitions $f_t(\cdot)$ can be derived from the probability distribution of sales described in the previous section. I return to this momentarily.

By assuming the unobservable is distributed T1EV, along with conditional independence, the expected value function is equal to

$$EV_t(c_t, p_t) = \int_{c_{t+1}} \left[\sigma \ln \left(\sum_{p_{t+1} \in A(t+1)} \exp \left(\frac{R^e_{t+1}(c_{t+1}, p_{t+1}) + EV_{t+1}(p_{t+1}, c_{t+1})}{\sigma} \right) \right) \right] f_t(c_{t+1}|c_t, p_t) + \sigma\phi,$$

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

$$CCP_t(c_t, p_t) = \frac{\exp\left\{ \left(R^e_t(p_t, c_t) + EV_t(p_t, c_t) \right) / \sigma \right\}}{\sum_{p'_t \in A(t)} \exp\left\{ \left(R^e_t(p'_t, c_t) + EV_t(p'_t, c_t) \right) / \sigma \right\}}.$$

Before continuing, I discuss the connections between the notation $Pr_t\left(Q_{jt} = q\,;\, p_{jt}\right)$, which denotes probability masses of sales, and $f_t(c_{t+1}\,|\,c_t, p_t)$, above, which denotes capacity transition probabilities. Consider a two-period model with a single seat. In the first period, expected revenues are formed based on $Pr_t\left(Q_{jt} = q\,;\, p_{jt}\right)$. In this case, expected revenues are simply $Pr_1\left(Q_1 \geq 1\,;\, p_1\right) \cdot 1 \cdot p_1$ because the probability that zero seats are demanded is associated with zero revenues, and with constrained capacity, at most one seat can be sold. The demand probabilities exactly inform the capacity transition probabilities under conditional independence:

$$f_1(c_2\,|\,1, p_1) = \left[\;Pr_1\left(Q_1 \geq 1\,;\, p_1\right) \quad , \quad Pr_1\left(Q_1 = 0\,;\, p_1\right)\; \right].$$

That is, with probability $Pr_1\left(Q_1 \geq 1\,;\, p_1\right)$, the seat sells today and nothing is available for sale tomorrow. On the other hand, with probability $Pr_1\left(Q_1 = 0\,;\, p_1\right)$, the seat is not sold today and is available for purchase tomorrow. The optimal price that affects these probabilities depends on three key inputs: the uncertainty in demand, the share of each consumer type, and the preferences of consumers. Time is a deterministic state. In the general model, with a longer time horizon and additional capacity, many entries in the transition probability matrix are equal to zero. In particular, any entry associated with a probability that $c_{t+1} > c_t$ is equal to zero because capacity is monotonically decreasing.

Given a set of flights (F) each tracked for (T) periods, the likelihood for the data is given by

$$\max_{\theta} \mathcal{L}(\text{data} \mid \theta) = \max_{\theta} \prod_{F} \prod_{T} CCP_t(c_t, p_t) f_t(c_{t+1}|c_t, p_t), \tag{4.1}$$

where $\theta := (\beta, \alpha, \gamma_t, \mu_t, \sigma)$ are the parameters to be estimated. In the next section, I place additional restrictions on the parameters γ_t and μ_t.

5 Model Estimates

In this section, I discuss the identification and the estimation procedure (Section 5.1), model estimates (in Section 5.2), and provide a discussion of model fit and predictions (in Section 5.3).

5.1 Identification and Estimation

The key identification challenge of the paper is to separately identify the demand parameters from the arrival process. This challenge 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 be seen instead 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. This is sometimes called the lack of "no purchase" data. Consumer search data can be used to solve this issue. Unfortunately, these data cannot be collected from public sources.

This paper proposes incorporating the supply-side model in order to separately identify the demand parameters and the arrival process. In particular, I assume that firms optimally price given seats remaining, time left to sell, and their unobservables. Preferences are assumed to evolve in the same predictable way, but demand shocks can vary for each flight toward the deadline. This results in variation in seats sold over time, and the firm's response to these shocks informs the magnitude of stochastic demand. That is, by solving the firm's problem, I recover the opportunity cost of capacity, and along with the pricing decision, I back out the overall demand elasticity. By tracing out price adjustments from variation in seats remaining given time to sell and variation over time given a constant capacity constraint, I separate the incentives to adjust prices in response to demand shocks versus the demand elasticity.

Figure 3 provides graphical evidence of the identification argument. Given stochastic demand, we would expect prices to rise when demand exceeds expectations and fall after a sequence of low demand realizations. This is shown in the figure as the solid (blue) line is above the zero, and the dashed (black) line is at or below zero. However, Figure 3 shows that prices sharply rise close to the departure date and regardless of sales. This sharp rise in prices regardless of the scarcity of seats suggests a change in consumer arrivals over time. That is, consumers who shop late are less price sensitive than those who shop early.

I assign the discount factor to be one. In addition, I place restrictions on γ_t and μ_t, or the probability on consumer types and Poisson arrival rates, respectively. I assign

μ_1 Greater than twenty-one days before departure (21+);

μ_2 Fourteen to twenty-one days before departure (20-14);

μ_3 Seven to fourteen days before departure (13-7); and

μ_4 Within seven days before departure (6-0),

which corresponds to the advance-purchase discount periods commonly seen in airline markets. This adds some flexibility in the Poisson arrivals of customers. In addition, I assign

$$\Pr_t(\text{Business}) = \gamma_t = \frac{\exp\left(\gamma_0 + \gamma_1 t + \gamma_2 t^2\right)}{1 + \exp\left(\gamma_0 + \gamma_1 t + \gamma_2 t^2\right)}, \forall t = 0, ..., T.$$

This parametric specification allows for non-monotonicity in consumer types over time, while keeping the function bounded between zero and one.

I utilize a dynamic discrete choice model because fares are chosen from a pre-determined set—as discussed in Section 2, fares are assigned by the pricing department. The supply model can be interpreted as modeling the decisions of revenue management, conditional on the choices made by other airline departments. In particular, the model takes the initial capacity and observed fares as given. Given the set of fares, identification assumes the pricing choice is optimal. This is perhaps not unreasonable given the sophisticated pricing models used by airlines (McGill and Van Ryzin, 1999). However, airlines operate complex networks and the pricing decision for a single flight may be impacted by forces not accounted for in the model—for example, a persistent, unobserved shock to the network could overstate the role of capacity in the model.

The average number of unique fares observed per flight is less than seven; however, I observe adjustments, sometimes by a single dollar, to fares over time.[21] To account for this without increasing the dimensionality of the problem, I cluster prices using k-means and then take the cluster centers to define prices for each market. I select the minimum in-sample fit threshold of 98.5 percent, which results in choice sets that range from size five to eleven. To preserve to use of APDs, I then assign day-before-departure-specific choice sets based on when the clustered prices appear in the data. That is, the procedure captures the advance-purchase discounts observed in the data, albeit with clustered fares. This approach allows me to utilize the full structure of the model for estimation.[22]

I maximize the log-likelihood of the firm's dynamic programming problem found in Equation 4.1 separately for each market. I group together the directional traffic of the city pairs, which means demand does not vary by direction. Appendix B shows that directional prices are very similar. I do not estimate demand for markets with nonstop competition.

For any candidate solution vector θ, I calculate the Poisson demand functions, expected revenues, and transition probabilities. The firm's problem is finite horizon; thus, with those objects calculated, I solve for the value functions using the recursive structure of the firm's problem. Backward induction allows for computing the conditional choice probabilities (CCP) for any state of the dynamic program. With the transition probabilities and CCPs defined, I calculate the log likelihood given a candidate solution vector θ. I maximize the objective, $\log\left(\mathcal{L}(\text{data}|\theta)\right)$.[23]

[21] One way to incorporate these pricing adjustments is to assume they are known in advance, but doing so greatly increases the complexity of the problem since it requires specifying many dynamic programs (partitioning flights by the unique set of filed fares observed for each departure date).

[22] Other approaches are available. In the hotel setting, Cho et. al. (2018) find the set of prices to be large and they propose using generalized method of moments (GMM) with moment conditions from both the demand and supply side.

[23] Estimation utilizes the interior/direct algorithm using the solver Knitro. The algorithm uses parallel multi-start, selecting at least 100 random initial starting values over a wide set. Using up to sixty-four threads, estimation for a single market takes as few as two hours and at most two days.

5.2 Parameter Estimates

Parameter estimates appear in three tables, Table 6, Table 7, and Table 8, located in Appendix A.[24] Each table contains estimates for eight markets and each table has three subsections.

The first subsection, "Logit Demand," reports consumer preference estimates as well as parameters governing the probability on consumer types over time (γ_t). Consumer preferences are all found to be statistically significant at conventional levels, except for the intercept for the (Palm Springs, CA; Portland, OR) city pair. The parameter estimates suggest that, on average, leisure consumers are twice as price sensitive as business consumers, and business consumers are willing to pay over 68 percent more in order to secure a seat.[25]

The parameters on the probability of consumer types (γ_t) are not easily interpretable so I plot aspects of their distributions below (Figure 4, left panel). The plots depict the average (across markets) business share over time, as well as the interquartile range and the fifth and ninety-fifth percentiles. Most markets exhibit increasing γ_t processes over time; 10 percent of early arrivals are the type labeled "business" and close to 80 percent of late arrivals are the type labeled "business." In early periods, prices are relatively flat and I estimate the average γ_t to be flat. Starting at twenty-one days before departure, I estimate a significant change in the business customer share. This corresponds with the time at which fares start raising rapidly.

Figure 13 in Appendix A plots the fitted values for γ_t for each market separately. The heterogeneity in the arrival estimates is expected given the differences in pricing dynamics across markets (see Figure 11 and Figure 12). I estimate that markets such as (Lihue, HI - Portland, OR) and (Palm Springs, CA - Portland, OR) have mostly flat or non-monotonic γ_t processes. Presumably these markets cater heavily to leisure customers. The Hawaii process decreases in early periods; average fares slightly decrease well in advance of the deadline. Two markets are estimated as having nearly linear γ_t processes; they are (Seattle, WA - Wichita, KS) and (Oklahoma City, OK - Seattle, WA). These markets tend to have

[24]I do not estimate demand for markets that were observed to have more than one nonstop carrier during the data collection period. The excluded markets are: Boston, MA - Kansas City, MO; Boston, MA - Seattle, WA; Boston, MA - Portland, OR; Seattle, WA - Sacramento, CA. Service to/from Lihue, HI and Palm Springs, CA were less than once daily. Estimation uses flights tracked for at least thirty periods.

[25]The mean ratio of price sensitivity is 2.57; the median is 2.31.

Figure 4: Visualizing the Arrival Process over Time

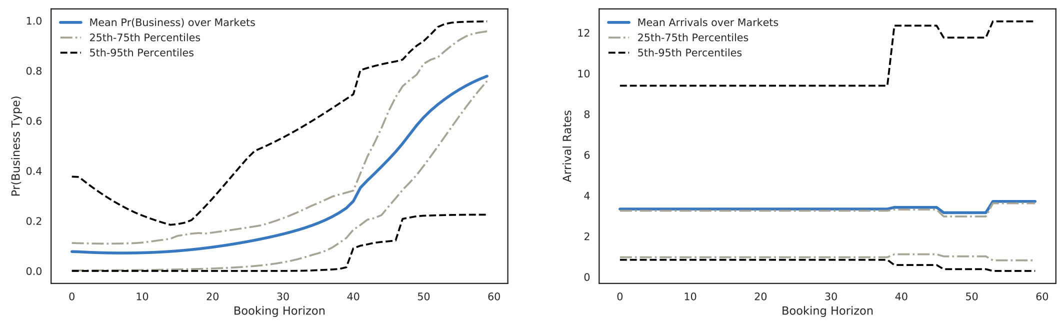

Note: Fitted values of the arrival process of business versus leisure customers across the booking horizon. In the left figure, the y-axis is Pr(business), so $1 - Pr_t$(Business) defines Pr_t(Leisure). In the right figure, the levels of the arrival process are plotted over time.

smaller price increases (in percentage terms), as shown in Figure 12. The first two rows in the table show the general trend: the share of early business arrivals is typically between 0 and 20 percent and the share increases to 60 to 100 percent the day of departure. The shape of the curves correlates with the use of APDs: a larger price increase at the twenty-day APD generally creates a steeper profile. Overall, the estimates establish that a meaningful shift occurs in willingness to pay over time. Average demand elasticities range from 1.5 to 6.3, depending on market and time until departure.

The parametric assumption on consumer types is flexible, as it captures S-shape, almost linear, and non-monotonic arrival paths. However, the model can be restrictive. Two markets, (Helena, MT - Seattle, WA) and (Seattle, WA - Sun Valley, ID) are estimated to shift from one Poisson demand distribution to another (leisure to business) corresponding at the twenty-one and fourteen day APD, respectively. These increases coincide with the prevalence of the routes' APD requirements; however, the γ_t parameters for these two markets are insignificant.

The second set of parameter estimates reports in the tables are labeled "Poisson Rates", which report mean arrival rates for each of the specified time intervals. These functions are also plotted below in Figure 4 (right panel). There is heterogeneity in these estimates as well. Several markets have a general increase in the arrival rates over time while others exhibit a general decreasing trend. The general finding is that the estimated arrival rates are low with only a few potential consumers searching to buy a ticket on a particular flight each day. Combining the γ_t and μ_t parameters, I estimate that 23 percent of arrivals are business travelers. As a point of comparison, Lazarev (2013) estimates 20 percent of consumers are business travelers. All arrival rates are estimated to be statistically significant.

Finally, the last row, "Firm Shock", reports estimates of the scaling parameter. The model is estimated with prices scaled to hundreds of dollars ($100 = 1). All of these parameters are estimated to be less than one and are statistically significant.

5.3 Model Fit and Discussion

Figure 5: Model Fit by Day Before Departure

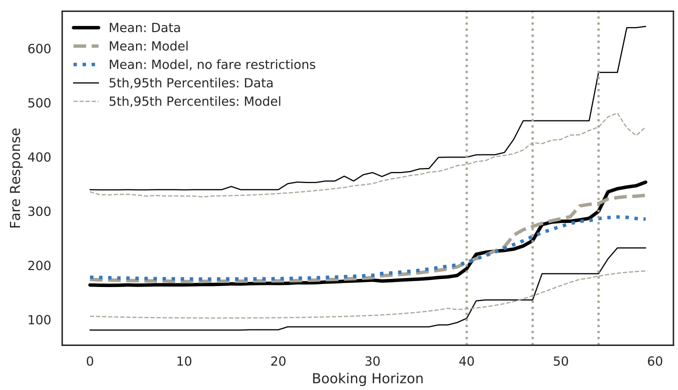

Note: Comparison of mean data fares and mean model fares across the booking horizon. Two versions of model fares are plotted. The solid black line defines per-period price choice sets using fare restrictions in the data. The dashed grey line allows firms to choose from all prices each period.

The model fits the data well. Figure 5 shows within-sample model fit by plotting data and model fares over time. Model fares are shown under the choice set restrictions in the estimated model as well as with the restrictions removed—firms have access to the entire choice set in each period. The figure depicts means as well as the fifth and ninety-fifth percentiles of fares. Model fares closely follow observed fares, with an average difference of $6.35. Differences do vary by day before departure—they are less than $10 for the first half of the sample but the gap increases around APDs. The reason is that the model produces a smoother fare profile that results in fare hikes slightly before the fourteen and seven day APDs. The model accurately captures the use of the twenty-one day APD. The fifth and ninety-fifth percentiles of fares are also aligned, except for close to the departure date, where the top five percent of data fares are higher than what the model assigns.

The dashed line, corresponding to model fares where the firm utilizes the entire choice set, also closely follows the data except close to the deadline. Within the last week before departure, the unrestricted model assigns lower prices. This is because fares are lowered in order to leave fewer seats unsold. One view on this finding is that the utilization of fare restrictions acts as a reputation mechanism that allows firms to commit to high prices close to the date of travel, even for flights with excess capacity.[26]

Figure 6 plots the value functions and policy functions for the Boston, MA - San Diego, CA city pair, focusing on the two state variables, seats remaining and time left to sell. The horizontal axis denotes seats remaining. The four lines cover selected periods (fifteen, thirty, forty-five and sixty days before departure).

Figure 6: Model Policy Functions and Value Functions (Boston, MA - San Diego, CA)

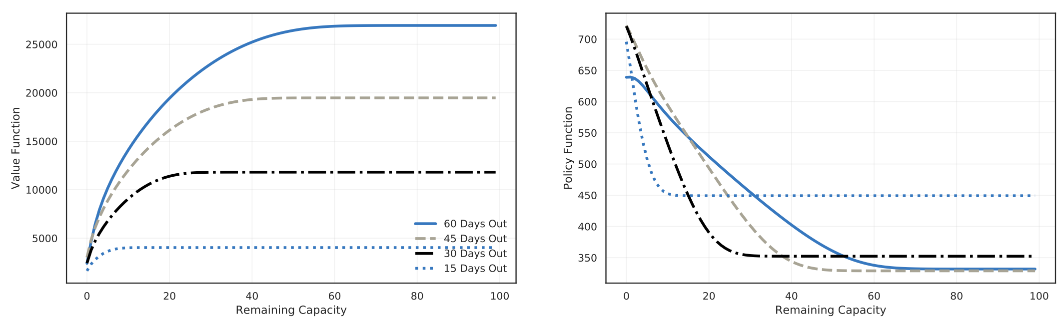

Note: Model policy functions and value functions for four periods. This figure demonstrates pricing and revenues for the Boston, MA - San Diego, CA city pair.

The left panel shows how seats remaining and time left to sell influence expected revenues. Expected revenues are increasing in capacity for a given period, that is, $\int_\omega V_t(c,\omega)dH\omega \leq \int_\omega V_t(c+1,\omega)dH\omega$. However, the curves flatten out as it becomes increasingly unlikely that the firm will sell all remaining capacity. The curves for each time period become completely flat as the probability of sell outs becomes zero. The plot also shows that expected values are increasing in time to sell for a given capacity, that is, $\int_\omega V_t(c,\omega)dH\omega \leq \int_\omega V_{t+1}(c,\omega)dH\omega$. With five seats remaining, expected revenues are very similar with either forty-five or sixty

[26]Dana and Williams (2020) show in a theoretical model that inventory controls can be used to commit to increasing prices in a competitive setting.

days remaining, because remaining capacity is low relative to the sales horizon. Here, the firm is confident that it can sell at units at high prices. As the booking horizon grows, the firm expects to sell more seats, even if it does not expect to sell out. The revenue difference is shown by the gaps between the lines in the right side of the graph.

The right panel depicts the policy functions and demonstrates an important interaction between the estimated arrival process and scarcity. The horizontal axis again shows remaining capacity, and the four lines are separate time periods. The vertical axis shows the expected price using the conditional choice probabilities from the model, that is, $\int p_t CCP_t(c,t)$. Scarcity is shown by the general upward trend or negative relationship between capacity and price. For any given period, if capacity is high, prices are low. This is because the opportunity cost of capacity is decreasing in seats remaining. The highest prices occur with very limited capacity and a long sales horizon. This is because the firm can always decrease price if it does not sell.

The effect of changing preferences is demonstrated by comparing the fifteen-day line with the other lines. The probability of selling forty units is very low, but the optimal price is higher than in other the other periods shown because arriving consumers have higher willingness to pay. In a homogeneous Poisson demand model, such as in Gallego and Van Ryzin (1994), this pricing line curve would be substantially lower because preferences do not cause prices to increase over time. As the plot shows, with fewer than ten units remaining, both scarcity and intertemporal price discrimination are important, and prices increase sharply.

6 Analysis of the Estimated Model

In this section, I conduct a series of counterfactuals given the model estimates. I compare firm revenues and consumer surplus under dynamic pricing with several alternative pricing regimes. In Section 6.1, I investigate uniform pricing and dynamic pricing with a restriction to frequency of price adjustments. In Section 6.2, I solve for a pricing system in which firms commit to a pricing schedule that depends on time until departure but not on demand realizations. By comparing uniform pricing to this latter counterfactual, which I call the intermediate case, I quantify the relative influence of intertemporal price discrimination and dynamic adjustment to stochastic demand in airline markets. I further explore the use of intertemporal price discrimination by studying static pricing problems, where prices reflect per-period preferences. In Section 6.3, I consider hypothetical arrival processes. Finally, in Section 6.4, I show that in order to quantify the effects of intertemporal price discrimination in airline markets, it is essential to account for stochastic demand. Procedures that abstract from stochastic demand will infer that late-arriving consumers are too price insensitive.

For each counterfactual, I use the empirical distribution of remaining capacity sixty days before departure as the initial capacity condition. Note that it may be profitable for firms to adjust capacity if the unmodeled fixed costs are such that the counterfactual pricing systems support a different gauge of aircraft. I explore the role of initial capacity at the end of Section 6.2.

For each counterfactual, I simulate 100,000 flights per market and then combine the results over markets.[27] I report the following benchmarks:

- Fare: mean price;
- Load Factor: mean load factor on day of departure;
- Sell outs: percentage of flights that sell all initial capacity;
- Revenue: mean revenue per flight;
- CS_L^i: mean leisure consumer surplus per passenger;
- CS_B^i: mean business consumer surplus per passenger;
- Welfare: mean daily combined consumer surplus and revenues per flight.

I alter the firm problem in all counterfactuals in two ways. First, I allow firms to use the unrestricted choice set, $A(t) = \cup_{t=0}^{T} A(t)$, in each period. Counterfactuals without this adjustment appear in the appendix. Second, I remove the firm shocks for the following analysis. I do this in order to single out the effect of time remaining and capacity, rather than the role of unobservable errors, in determining the pricing decision. For example, under uniform pricing, the firm would receive a single error vector, whereas in the dynamic

[27] I exclude the three markets observed to have infrequent service and frequent spikes in last-minute bookings (Lihue, HI - Portland, OR; Palm Springs, CA - Portland, OR; Santa Barbara, CA - Seattle, WA). Results are robust to including these markets.

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.

All counterfactuals utilize the important boundary conditions of the initial problem: (1) the firm cannot oversell; (2) unused capacity is scrapped with zero value. Capacity transitions as before.

6.1 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 in each period. The optimal price depends on the initial capacity condition and the distributions of demand over time. The revenue maximization problem is

$$\max_{p} \; \mathbb{E}_y \left[\sum_{t=0}^{T} p \min \{ Q_t(p, y_t), c_t \} \right]$$
$$\text{such that } c_{t+1} = c_t - \min \left\{ Q_t(p, y_t), c_t \right\}, \; c_0 \text{ given.}$$

Under uniform pricing, a high price can be used to target business customers. However, arrivals are sufficiently low that it will result in unused capacity. Lowering the price will allow additional leisure consumers to purchase, thus expanding output, but it will also decrease revenues per seat sold. The optimal price balances out these effects.

Results for the counterfactual appear in Table & Figure 2. In the figure, the left panel plots mean fares over time (the dashed lines show the interquartile distributions). Uniform fares are relatively higher than dynamic fares early on, but then uniform fares are considerably lower than that under dynamic fares close to the departure date. This results in a reallocation of capacity over time, which is shown graphically in the right panel. The right panel plots the booking curve, or mean cumulative seats sold divided by capacity, over time. The uniform pricing booking curve is bowed out as fewer consumers purchase under the relatively high fares early on. Relatively low fares close to the departure date result in a higher booking rate. However, even with this increase in late bookings, total output is lower under uniform pricing (3.1 percent lower load factor). Because the firm cannot respond to demand shocks, load factors under uniform pricing are considerably more dispersed.

Table & Figure 2: Dynamic to Uniform Pricing

	Fare	Load Factor	Sell Outs	Revenue	CS^i_L	CS^i_B	Welfare
Dynamic	242.9	87.8	17.9	11533.5	49.3	161.9	468.0
Uniform	223.7	84.7	28.5	10565.5	46.4	180.2	464.5
Difference (%)	-7.9	-3.1	10.6	-8.4	-6.0	11.3	-0.7

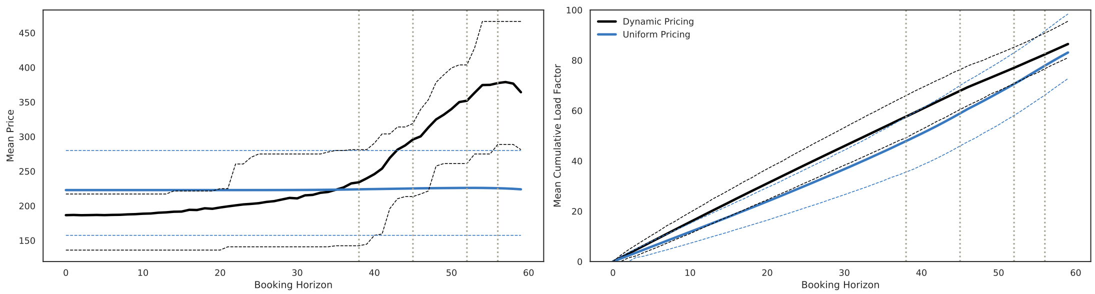

Note: Fare: mean fare for flight observations with positive seats remaining; Load factor (LF): average at departure time; Sell Outs: percentage of flights with zero seats remaining in the last period; Revenue: mean flight revenue; Consumer surplus (CS^i_L, CS^i_B): surplus per-person; Welfare: daily mean revenues plus consumer surplus, excluding fixed costs. Results come from simulating 100,000 flights per market given the empirical distribution of remaining capacity sixty days before departure.

The reallocation of capacity over time necessarily means a reallocation of capacity across consumer types. The top of Table & Figure 2 reports key benchmarks for the two pricing regimes. All values are in levels, except for load factor and sell outs, which are reported as percentages. The difference row is relative to dynamic pricing—a negative number means the value is higher under dynamic pricing. The table shows that leisure consumers are harmed by uniform pricing. This is because fares are higher in early periods. I estimate leisure consumer surplus declines 6.0 percent under uniform pricing. On the other hand, business consumers benefit from considerably lower fares under uniform pricing—flights are up to \$150 less expensive. I estimate business consumer surplus increases by 11.3 percent under uniform pricing. Although this is a significant gain, it is mitigated somewhat by the fact that the firm cannot ensure all business consumers are served under uniform pricing. Sell outs increase by 10.6 percent because the firm cannot respond to demand shocks, and 5.3 percent of arriving business customers cannot obtain a seat. Under dynamic pricing, only 1.2 percent of business travelers cannot obtain a seat. Combining the effects across the two consumer types, aggregate consumer surplus increases by 4.6 percent when the firm cannot use dynamic pricing.

Both the uncertainty about demand and the change in preferences of consumers over time impact the welfare effects of dynamic pricing. Although the relative importance of these forces varies across markets, two findings broadly apply. First, uniform pricing largely reallocates consumer surplus from leisure travelers to business customers. Because business customers have significantly higher willingness to pay, I find the uniform pricing increases consumer surplus. Second, uniform pricing results in significantly reduced revenues; I estimate the decline to be 8.4 percent.[28] Combining both the revenue and consumer surplus effects, I find that uniform pricing lowers total welfare by 0.7 percent in the monopoly markets studied.

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 remaining, 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.

Figure 7 plots the revenue loss compared with the baseline case of daily price adjustments. The bottom bar shows the revenue loss under the uniform pricing scenario just discussed. The ability to update prices just once reduces the revenue loss compared uniform pricing by half (30-day adjustments). An additional price adjustment yields another 2.2 percent gain. I find that two-day, three-day, and six-day adjustments yield fairly similar results. This suggests that several demand shocks can be observed before requiring a price adjustment for revenues to be similar.

[28]IATA (2013) reports airline profit margins are around 1.1%, inclusive of ancillary income and debt interest. Revenue Management Overview claims revenue management systems have increased airline revenues by 3-9% over time.

Figure 7: The Role of Frequent Price Adjustments

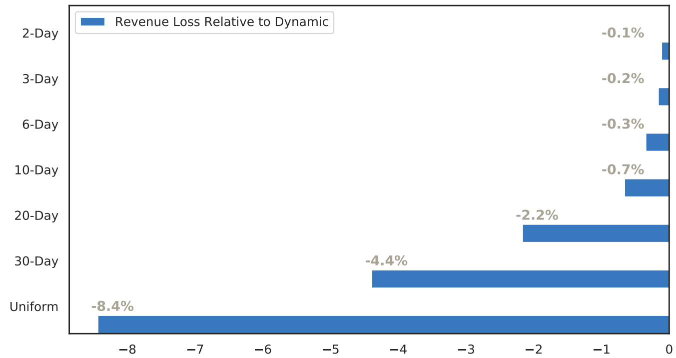

Note: Revenue drop relative to dynamic (daily) pricing for all markets. For example, 3-day corresponds to firms utilizing dynamic pricing, but restricting the number of price updates to 3-day intervals.

6.2 The Use of Intertemporal Price Discrimination Alone

I now solve for optimal pricing where the firm commits to a price schedule that varies across time but not on the realizations of demand. I call this "the use of intertemporal price discrimination alone", or "the intermediate case." The reason for this label is that dynamic pricing varies by seats and time remaining, $p^*(c_t, t)$, and uniform pricing does not vary along either of those two dimensions, p^*. In the intermediate case, prices solely depend on time remaining, $p^*(t)$. Therefore, the following inequalities hold:

$$\underbrace{\sum_{t=0}^{T} R_t^e\Big(p^*; c_t\Big)}_{\text{Uniform Pricing}} \le \underbrace{\sum_{t=0}^{T} R_t^e\Big(p^*(t); c_t\Big)}_{\text{Intermediate Case}} \le \underbrace{\sum_{t=0}^{T} R_t^e\Big(p^*(t, c_t); c_t\Big)}_{\text{Dynamic Pricing}}.$$

By comparing dynamic pricing with uniform pricing, and dynamic pricing with the intermediate case, I quantify the relative importance of having fares respond to preferences versus fares respond to both preferences and demand shocks.

In the intermediate case, the revenue maximization problem is

$$\max_{p_0,\ldots,p_T} \mathbb{E}_y\left[\sum_{t=0}^{T} p_t \min\left\{Q_t(p_t, y_t), c_t\right\}\right]$$
$$\text{such that } c_{t+1} = c_t - \min\Big\{Q_t(p_t, y_t), c_t\Big\},\ c_0 \text{ given.}$$

This is a large dimensional problem—an exhaustive search involves evaluating the objective over dim$(A)^T$ possible price vectors. At a minimum, the problem contains approximately 8.6e41 possibilities. To reduce the dimensionality of the problem, I add the restriction that the firm can adjust fares only on the usual advance purchase discount days: three, seven, fourteen, and twenty-one days before departure. This results in five prices per flight.

Table & Figure 3: Dynamic Pricing to Intermediate Case

	Fare	Load Factor	Sell Outs	Revenue	CS_L^i	CS_B^i	Welfare
Dynamic	242.9	87.8	17.9	11533.5	49.3	161.9	468.0
Intermediate	244.4	83.7	21.4	11212.4	48.1	158.9	456.4
Difference (%)	0.6	-4.1	3.6	-2.8	-2.5	-1.8	-2.5

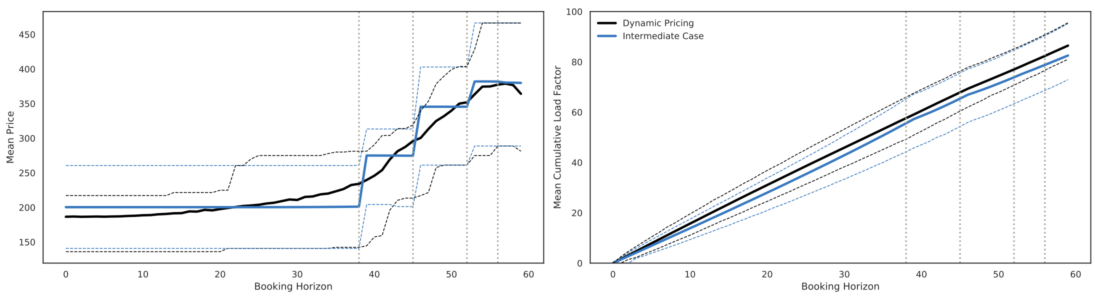

Note: Fare: mean fare for flight observations with positive seats remaining; Load factor (LF): average at departure time; Sell Outs: percentage of flights with zero seats remaining in the last period; Revenue: mean flight revenue; Consumer surplus (CS_L^i, CS_B^i): surplus per-person; Welfare: daily mean revenues plus consumer surplus, excluding fixed costs. Results come from simulating 100,000 flights per market given the empirical distribution of remaining capacity sixty days before departure.

Results for this counterfactual appear in Table & Figure 3. Fares in the intermediate case are monotonically increasing, with a substantial increase in fares when crossing the first three APDs. This largely reflects changes in willingness to pay over time. Because the firm cannot respond to demand shocks, load factors are 4.1 lower and more dispersed than under dynamic pricing.

Just like dynamic pricing, the use of intertemporal price discrimination alone increases leisure consumer surplus relative to uniform pricing due to lower initial fares. As time advances, arriving leisure consumers face higher prices due to the increased presence of business consumers. I find that leisure consumer surplus is increasing as pricing becomes more flexible (dynamic > intermediate > uniform). In contrast, I find business consumer surplus is highest under uniform pricing and lowest in the intermediate case. In the intermediate case, the inability to respond to demand shocks results in high prices, even for flights with low realized demand. Moreover, sellouts increase by 3.6 percent, and an additional 4.2 percent of arriving business customers are not served. In addition, revenues are 2.8 percent lower when the firm cannot respond to demand shocks, hence, total welfare declines in the intermediate case.

Combining both counterfactuals, 66 percent 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 to respond to demand shocks. Dynamic pricing increases revenues two ways. First, for flights realized to have low demand, prices are adjusted downward in order to leave fewer seats unsold. Second, for flights realized to have high demand, output is constrained, but prices adjust upward in order to save seats for business consumers who are then charged high prices.

Static Pricing and Responding to Consumer Types

In the previous counterfactuals, the firm takes expectations over future demand when committing to prices in the initial period. Prices in the intermediate case adjust solely based on time remaining, but they also reflect aggregate demand uncertainty. To understand how increasing prices reflect intertemporal price discrimination isolated from scarcity, I examine static revenue maximization problems. When the firm is not forward looking, the optimal price solely reflects the demand it faces in each period.

I consider two static pricing problems. In the first scenario, in each period, the firm solves

$$\max_{p_t} \ \mathbb{E}_{y_t}\left[p_t \min\left\{Q_t(p_t, y_t), c_t\right\}\right].$$

Here, the optimal price in each period depends on the current arrival process and the current capacity constraint. When capacity becomes sufficiently low relative to the expected per-period demand, prices will reflect that scarcity. This problem is equivalent to the original dynamic pricing problem with the discount factor set equal to zero.

To completely remove the impact of remaining capacity on prices, I consider an alternative environment in which the firm has unlimited capacity. In each period, the firm solves

$$\max_{p_t} \ \mathbb{E}_{y_t}\left[p_t Q_t(p_t, y_t)\right]$$

Since demand is independent over time, this model is equivalent to the original dynamic pricing model with an infinite capacity constraint.

The results of these counterfactuals appear in Table & Figure 4. In the figure, the three curves represent dynamic pricing and the two static models described above. The interquartile range lines are removed for readability. I refer to "static pricing" (or "static model") when referring to the model with capacity constraints and "static$_\infty$ pricing" (or "static$_\infty$ model") when referring to the model without a capacity limit. I use the term "static models" (plural form) to refer to both simultaneously. In the table, only numbers for the static model are reported.

Table & Figure 4: Static Pricing

	Fare	Load Factor	Sell Outs	Revenue	CS_L^i	CS_B^i	Welfare
Dynamic	242.9	87.8	17.9	11533.5	49.3	161.9	468.0
Static	190.3	91.6	60.8	9617.5	56.5	141.6	421.8
Difference (%)	-21.7	3.8	43.0	-16.6	14.5	-12.5	-9.9

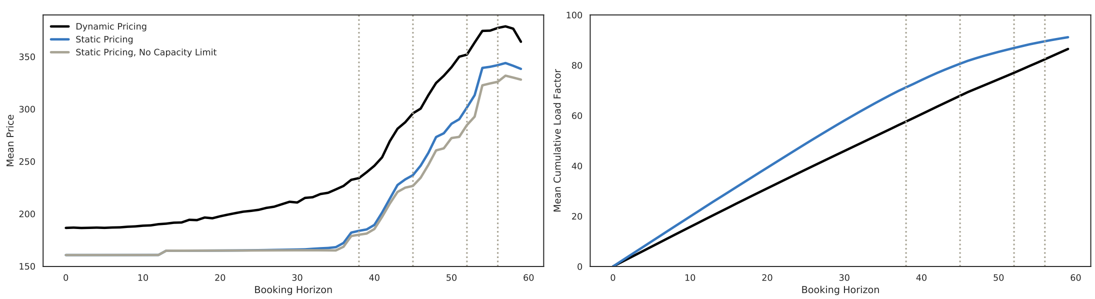

Note: Fare: mean fare for flight observations with positive seats remaining; Load factor (LF): average at departure time; Sell Outs: percentage of flights with zero seats remaining in the last period; Revenue: mean flight revenue; Consumer surplus (CS_L^i, CS_B^i): surplus per-person; Welfare: daily mean revenues plus consumer surplus, excluding fixed costs. Results come from simulating 100,000 flights per market given the empirical distribution of remaining capacity sixty days before departure.

The figure shows three notable findings. First, optimal prices for the static models are well below the baseline dynamic pricing model. When the firm internalizes the scarcity of seats, opportunity costs raise prices. The plot shows capacity plays a significant role in pricing, raising prices close to $50 relative to the static models. Second, both of the static pricing models naturally create APDs because preferences change systematically over time and prices are chosen from a discrete set. As the share of business consumers increases, if capacity remains, the optimal static price increases discreetly. Third, static$_\infty$ pricing closely follows static pricing, except when capacity becomes sufficiently scarce. For example, optimal prices differ when there is one seat left and the demand rate is sufficiently high, but prices coincide when fifty seats remain.

The results in Table & Figure 4 show stark contrasts between static pricing and dynamic pricing. Static pricing results in a 16.6 percent decline in revenues, nearly double the drop estimated under the uniform pricing counterfactual. Output is the highest among all the models considered, which makes sense since the firm does not value holding capacity for future periods. Leisure consumers greatly benefit from static pricing—they arrive first and enjoy low prices. However, business consumer surplus under static pricing is estimated to be the lowest among all the models considered. Fares are lower than under dynamic pricing, but availability is limited and this force dominates. This is shown in Figure 8, which plots the cumulative distribution of sell outs (left) and cumulative consumer surplus relative to dynamic pricing (right).

Figure 8: Cumulative Distributions of Sell Outs and Consumer Surplus

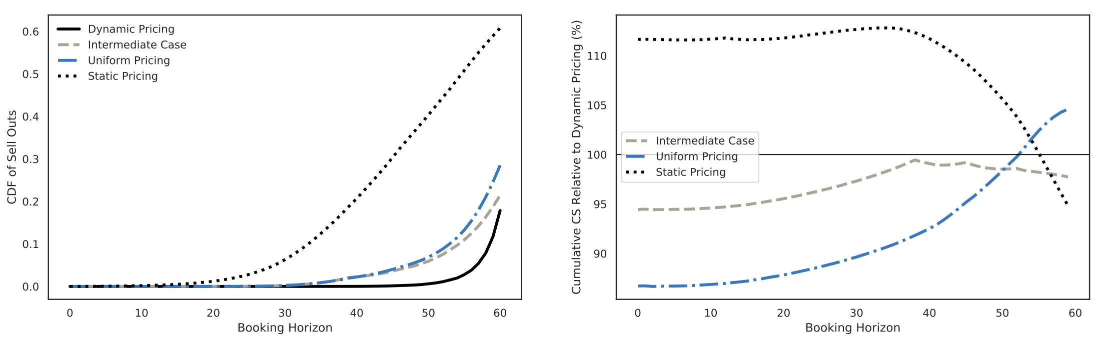

Note: Revenue drop relative to dynamic (daily) pricing for all markets. For example, 3-day corresponds to firms utilizing dynamic pricing, but restricting the number of price updates to 3-day intervals.

The Role of Initial Capacity

In the static pricing counterfactuals, static prices are substantially lower than dynamic prices. This is because the firm does not internalize that selling a seat today forgoes the increased revenue opportunity of selling the seat to a business consumer tomorrow. Dynamic pricing helps ensure that scarce capacity is allocated efficiently and results in substantially higher revenues. However, if the firm had *enough* initial capacity, it could sell a seat to all arriving customers.

To investigate this further, I analyze how the initial capacity constraint affects dynamic prices. Specifically, I investigate how large initial capacity has to be in order for static pricing to be a reasonable approximation of the environment. To accomplish this, I compute optimal dynamic prices and simulate outcomes for a wide range of initial capacity values.

Figure 9 shows the counterfactual for Fresno, CA = Portland, OR. In the left panel, the horizontal axis is the initial capacity condition. The left vertical axis is the percentage change in sales from increasing the initial capacity constraint by one. The steepness of the curve in early values reflects the fact that the firm will almost surely sell the additional unit. The curve starts to flatten and eventually approaches zero as additional capacity will not be sold. The right vertical axis plots total expected revenues by initial capacity. The (black) vertical line depicts the average observed initial capacity. The (grey) square denotes revenues with six fewer seats than the average (a row of a plane); the (blue) triangle denotes the minimum initial capacity such that prices follow a non-stochastic model.[29]

Figure 9: Initial Capacity Counterfactual (Fresno, CA - Portland, OR)

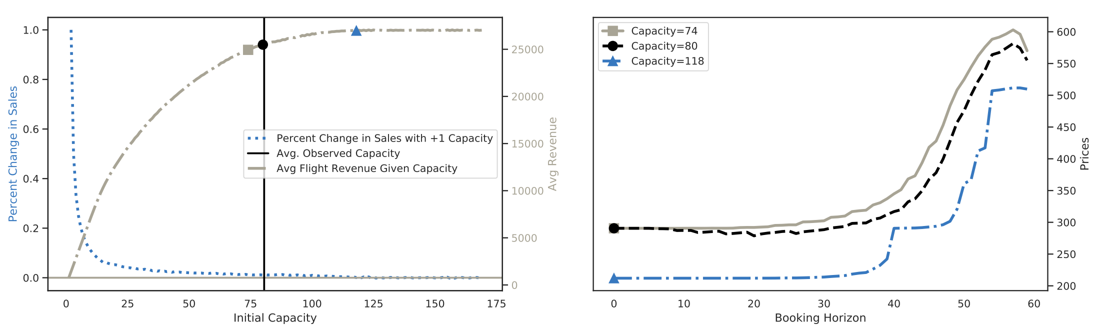

Note: The left panel shows the percentage change in quantity sold by increasing the initial capacity constraint by one (dotted blue). Also show are expected revenues by initial capacity constraint (dashed grey). The black vertical line shows the (weighted) average initial capacity observed in the data. The black dot shows expected revenues under this capacity. The grey square shows expected revenues with six fewer seats. The blue triangle shows expected revenues in the first instance when the percentage change in quantity sold is less than 0.1%. The right panel shows average prices over time for those three scenarios (average less ten, average, and the limiting case).

[29] I calculate this point is the first instance in which the percentage change in expected sales from increasing capacity by one is less than 0.1 percent.

The right panel plots average prices over time for the three initial capacities just described. The dashed blue (triangle) line shows the limiting case, where dynamic prices correspond to static prices. If the firm starts with fewer initial seats, realizations of demand impact prices.

I repeat this exercise for all markets in order to calculate the percentage of observed flights where optimal prices can be approximated using a static model. To do this, I calculate the capacity threshold for each market and then compare this threshold to the observed distribution of initial capacity. I find that 22 percent of the observed flights meet this threshold, meaning that stochastic demand and scarcity play an important role for a large majority of the flights studied.

6.3 Highlighting the Interaction under an Alternative Arrival Process

The importance of responding to demand shocks depends on the particular arrival process of consumers. I demonstrate this by considering a hypothetical environment in which preferences do not change over time, thereby removing the incentive to utilize intertemporal price discrimination. More specifically, I take the arrival rates to be equal to the average over time, by market. I also assign the proportion of business and leisure travelers to be equal to the average over time, by market. Thus, demand remains stochastic, but preferences are held fixed.

Table & Figure 5 contains the results for dynamic pricing and uniform pricing. The intermediate case is also plotted. The results show that average prices can increase (for much of the booking horizon) even when preferences do not change over time.[30] The reason why average prices increase is that the purchase rate is sufficiently low relative to remaining capacity and the deadline. The firm does not expect to sell many (any) seats in a given period, so when a sale does occur, this drives up the opportunity cost of capacity and price. There is a natural asymmetry: firms can choose to cut prices in later periods but cannot recover sold seats after a positive demand shock. A change in the arrival rate, holding all else constant, changes the shape of average prices. For example, with the arrival process multiplied by two, prices are mostly flat, except for a drop at the deadline. When the arrival

[30] Similar simulated optimal price paths can be found in Gallego and Van Ryzin (1994) and McAfee and Te Velde (2006), for example.

Table & Figure 5: Constant Arrival Process

	Fare	Load Factor	Sell Outs	Revenue	CS_L^i	CS_B^i	Welfare
Dynamic	223.3	90.6	25.2	10948.1	49.1	202.6	479.2
Intermediate	226.3	86.3	27.1	10735.5	47.9	197.5	468.6
Uniform	224.9	84.4	27.6	10625.3	47.1	196.5	464.3

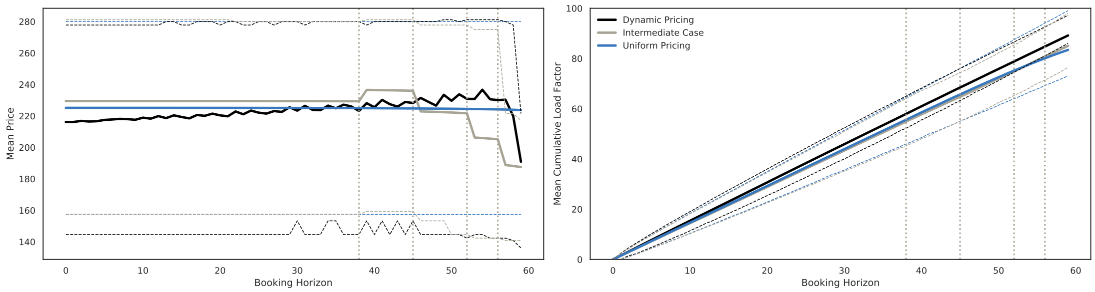

Note: Fare: mean fare for flight observations with positive seats remaining; Load factor (LF): average at departure time; Sell Outs: percentage of flights with zero seats remaining in the last period; Revenue: mean flight revenue; Consumer surplus (CS_L^i, CS_B^i): surplus per-person; Welfare: daily mean revenues plus consumer surplus, excluding fixed costs. Results come from simulating 100,000 flights per market given the empirical distribution of remaining capacity sixty days before departure.

process is multiplied by four, average prices decline over time. This is because the firm expects to sell seats each period, but when it receives a low demand shock, opportunity costs decline and prices are lowered.[31] Thus, stochastic demand alone can affect the shape of average prices over time. In the next subsection, I show the importance of accounting for stochastic demand when identifying intertemporal price discrimination.

Relative to the estimated arrival process, this counterfactual brings less dispersion in prices, sell outs, revenues, and consumer surpluses. One of the predictions 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 comparative revenue gains of dynamic pricing over uniform pricing are almost three times higher under the observed arrival process compared with this counterfactual (8.4 percent versus 3.0 percent).

6.4 Demand Elasticities Abstracting from Stochastic Demand

This paper demonstrates that stochastic demand is an essential feature of airline markets, and it quantifies how stochastic demand interacts with a changing mix of consumer types

[31] This intuition focuses on the arrival rate but the variance of arrivals also affects prices in a similar way. The Poisson distribution is equidispersed. I do not isolate the effect driven by the variance.

toward the perishability date. If demand is stochastic and capacity is scarce, but we abstract from this feature of the market, we will infer the incorrect demand elasticites. To understand why, note that empirical procedures that abstract from stochastic demand assign a constant opportunity cost of remaining capacity over time. In reality, opportunity costs adjust according to the realizations of demand. I demonstrate that the bias in estimated elasticities can be significant, and the direction of the bias depends on the arrival process.

Consider the special limiting case of the model presented in this paper where demand uncertainty is degenerate. A monopolist 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

$$\max_p \sum_{t=0}^{T} Q_t(p_t)p_t \quad \text{s.t.} \quad \sum_{t=0}^{T} Q_t(p_t) \leq K.$$

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

$$\begin{aligned} &\max_p \sum_{t=0}^{T} Q_t(p_t)p_t - c(K)\left(\sum_{t=0}^{T} Q_t(p_t) - K\right) \\ \Leftrightarrow &\max_p \sum_{t=0}^{T} Q_t(p_t)\big(p_t - c(K)\big) + c(K)K. \end{aligned}$$

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 ($t = 0$), to any future period, yields

$$\frac{p_t}{p_0} = \frac{\frac{e_t^D(p_t)}{e_t^D(p_t)-1}c}{\frac{e_0^D(p_0)}{e_0^D(p_0)-1}c} = \frac{\frac{e_t^D(p_t)}{e_t^D(p_t)-1}}{\frac{e_0^D(p_0)}{e_0^D(p_0)-1}}. \qquad \text{(Model without stochastic demand)}$$

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

When we account for stochastic demand, the shadow values may change over time. In general, we have

$$\frac{p_t}{p_0} = \frac{\frac{\tilde{e}_t^D(p_t)}{\tilde{e}_t^D(p_t)-1}\tilde{c}_t}{\frac{\tilde{e}_0^D(p_0)}{\tilde{e}_0^D(p_0)-1}\tilde{c}_0}. \qquad \text{(Model with stochastic demand)}$$

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

Figure 10 plots the median as well as the interquartile range of elasticity ratios for both situations. The plot contains 100,000 simulated flights per market. The elasticity ratio accounting for stochastic demand come directly from the estimated model, and the elasticity ratio abstracting from stochastic demand are derived from relative prices. In Figure 10, 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 over time due to the arrival of price insensitive customers. Within the last three weeks before departure, prices rise sharply and the bias increases significantly.[32]

[32] The direction of the bias depends on the particular arrival process of consumers. If price inelastic consumers were the first to arrive, the direction of the bias would be 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 bias tends to be close to zero.

Figure 10: Bias When Abstracting from Stochastic Demand

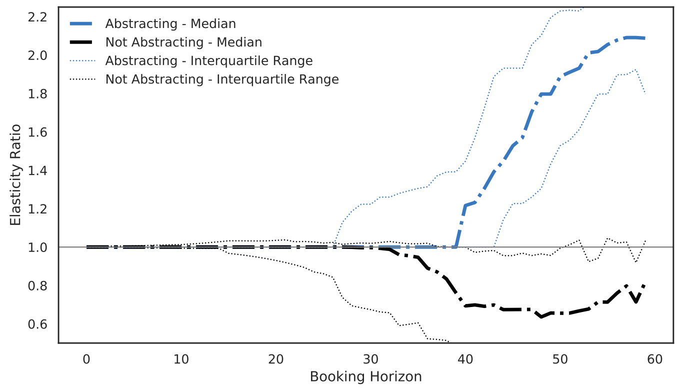

Note: Results come from simulating 100,000 per market. The bias lines correspond to an elasticity ratio abstracting from stochastic demand. The model elasticity ratios are recovered using the model accounting for stochastic demand.

7 Conclusion

Airfares fluctuate over time due to both demand shocks and intertemporal variation in willingness to pay. The main contribution of this paper is to jointly study intertemporal price discrimination and dynamic adjustment in response to stochastic demand 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 offer discounts to early-arriving, price sensitive consumers. It also allows them to secure seats for late-arriving consumers who are then charged high prices. I show that dynamic airline pricing results in more capacity utilization and significantly higher revenues. From a welfare perspective, I show that uniform pricing results in higher consumer surplus in nearly all of the markets studied, but total welfare tends to be higher under dynamic pricing. This is because the revenue gains can exceed the consumer surplus losses, particularly because price discrimination increases leisure consumer surplus and decreases business consumer surplus. Finally, I show that in order to study intertemporal price discrimination in airline markets, it is essential to account for stochastic demand.

References

Aguirre, I., S. Cowan, and J. Vickers (2010): "Monopoly price discrimination and demand curvature," *American Economic Review*, 100(4), 1601–15.

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.

Aryal, G., C. Murry, and J. W. Williams (2018): "Price discrimination in international airline markets," *Available at SSRN 3288276*.

Belobaba, P. (1987): "Air travel demand and airline seat inventory management," Ph.D. thesis, Massachusetts Institute of Technology.

Belobaba, P. P. (1989): "OR practice-application of a probabilistic decision model to airline seat inventory control," *Operations Research*, 37(2), 183–197.

——— (1992): "Optimal vs. heuristic methods for nested seat allocation," in *Presentation at ORSA/TIMS Joint National Meeting*.

Belobaba, P. P., and L. R. Weatherford (1996): "Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations," *Decision Sciences*, 27(2), 343–363.

Bergemann, D., B. Brooks, and S. Morris (2015): "The limits of price discrimination," *American Economic Review*, 105(3), 921–57.

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.

Board, S., and A. Skrzypacz (2016): "Revenue management with forward-looking buyers," *Journal of Political Economy*, 124(4), 1046–1087.

Bulow, J. I. (1982): "Durable-goods monopolists," *Journal of Political Economy*, 90(2), 314–332.

Chen, N. (2018): "Perishable Good Dynamic Pricing Under Competition: An Empirical Study in the Airline Markets," *Available at SSRN 3228392*.

Cho, S., G. Lee, J. Rust, and M. Yu (2018): "Optimal dynamic hotel pricing," .

Conlisk, J., E. Gerstner, and J. Sobel (1984): "Cyclic pricing by a durable goods monopolist," *The Quarterly Journal of Economics*, 99(3), 489–505.

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 (2020): "Intertemporal Price Discrimination in Sequential Quantity-Price Games," *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.

Dilmé, F., and F. Li (2019): "Revenue management without commitment: Dynamic pricing and periodic flash sales," *The Review of Economic Studies*, 86(5), 1999–2034.

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 H. Topaloglu (2019): *Revenue management and pricing analytics*, vol. 209. Springer.

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.

Gershkov, A., B. Moldovanu, and P. Strack (2018): "Revenue-maximizing mechanisms with strategic customers and unknown, markovian demand," *Management Science*, 64(5), 2031–2046.

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.

Hendel, I., and A. Nevo (2006): "Measuring the implications of sales and consumer inventory behavior," *Econometrica*, 74(6), 1637–1673.

——— (2013): "Intertemporal price discrimination in storable goods markets," *American Economic Review*, 103(7), 2722–51.

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*.

McFadden, D. (1973): "Conditional Logit Analysis of Qualitative Choice Behavior," *Frontiers in Econometrics*, pp. 105–142.

McGill, J. I., and G. J. Van Ryzin (1999): "Revenue management: Research overview and prospects," *Transportation science*, 33(2), 233–256.

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.

Öry, A. (2016): "Consumers on a Leash: Advertised Sales and Intertemporal Price Discrimination," .

Phillips, R. L. (2005): *Pricing and revenue optimization*. Stanford University Press.

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.

Siegert, C., and R. Ulbricht (2019): "Dynamic oligopoly pricing: Evidence from the airline industry," *working paper*.

Sobel, J. (1991): "Durable goods monopoly with entry of new consumers," *Econometrica: Journal of the Econometric Society*, pp. 1455–1485.

Soysal, G. P., and L. Krishnamurthi (2012): "Demand dynamics in the seasonal goods industry: An empirical analysis," *Marketing Science*, 31(2), 293–316.

Stokey, N. L. (1979): "Intertemporal price discrimination," *The Quarterly Journal of Economics*, pp. 355–371.

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.

——— (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 Additional Figures and Tables

Figure 11: Average Fares over Time by Market

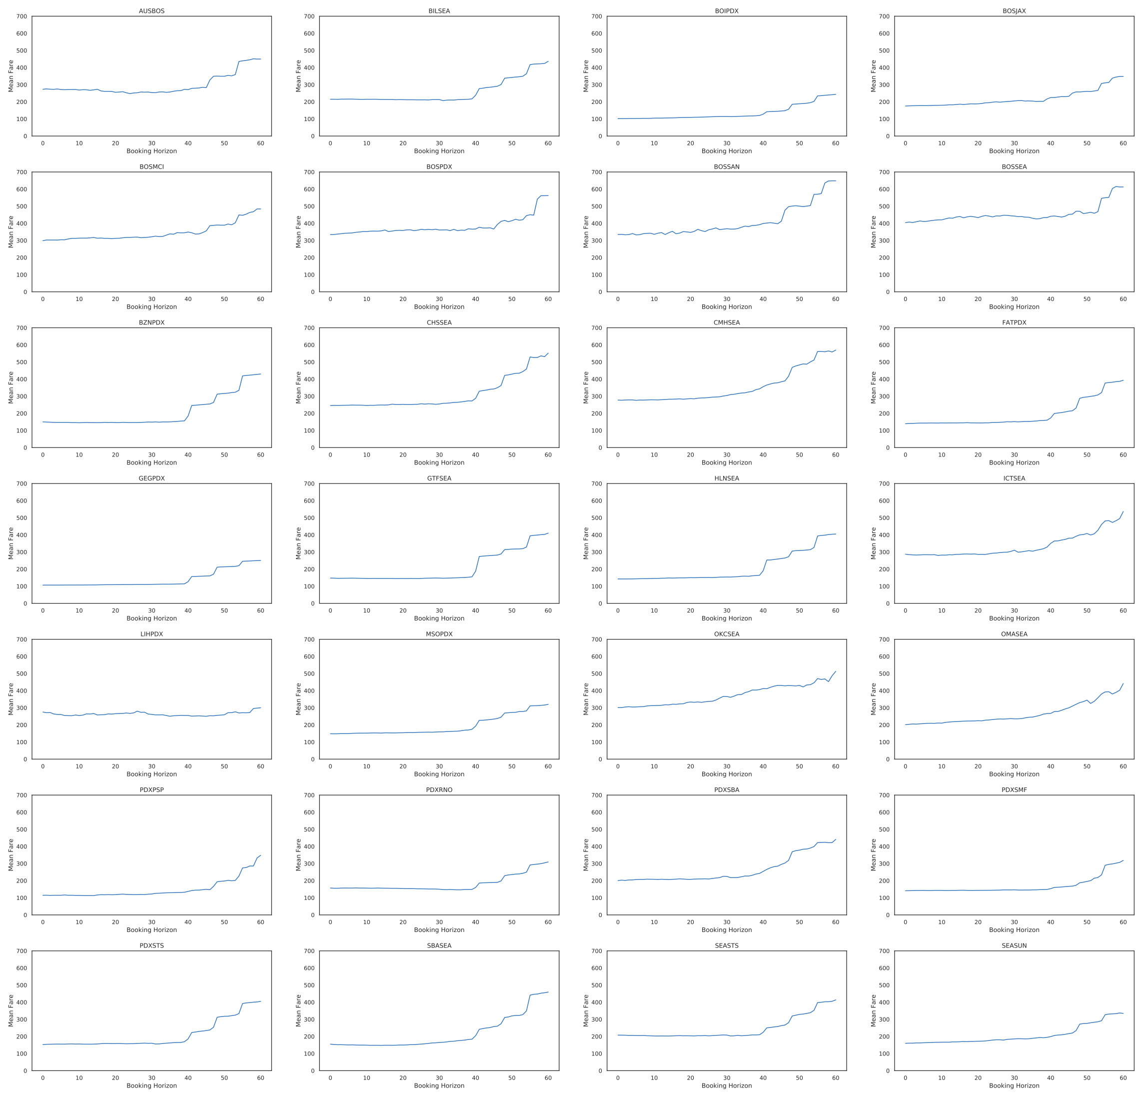

Note: Average fares over time for each market separately. This analysis combines origin-destination and destination-origin fares. Both axes are common across all plots.

Figure 12: Percent Change in Fares over Time by Market

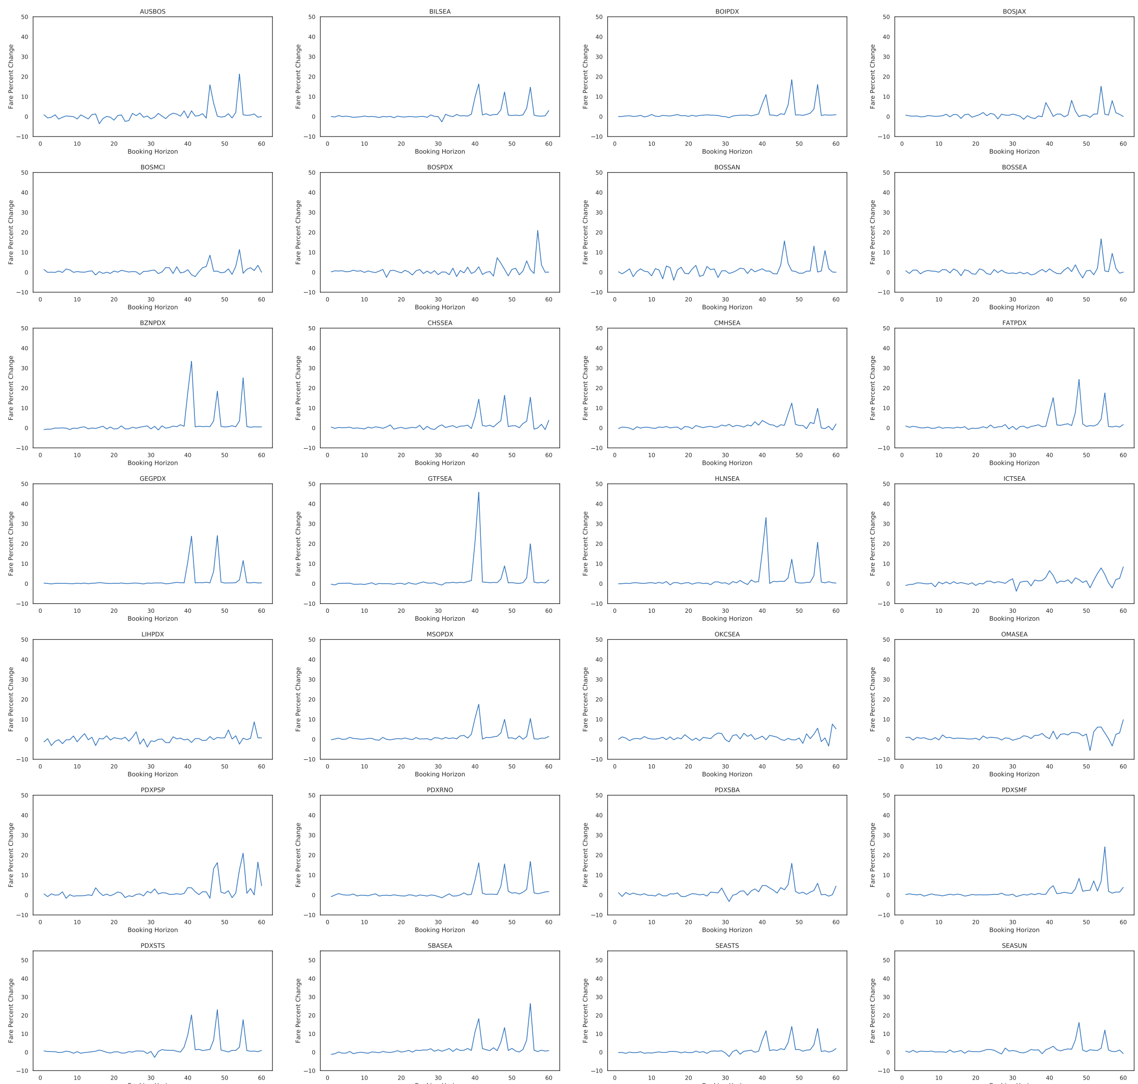

Note: Percent change in average fares over time for each market separately. This analysis combines origin-destination and destination-origin fares. Both axes are common across all plots.

Table 6: Parameter Estimates

Variable		AUSBOS	BILSEA	BOIPDX	BOSJAX	BOSSAN	BZNPDX	CHSSEA	CMHSEA
Logit Demand									
Intercept	β_0	5.629	0.672	6.470	0.929	10.532	4.472	2.644	1.796
		(0.104***)	(0.135***)	(0.060***)	(0.153***)	(0.131***)	(0.074***)	(0.092***)	(0.120***)
Leisure Price Sens.	α^L	-3.503	-0.851	-4.471	-1.709	-3.347	-2.615	-1.155	-0.888
		(0.059***)	(0.030***)	(0.036***)	(0.040***)	(0.040***)	(0.034***)	(0.018***)	(0.024***)
Business Price Sens.	α^B	-1.567	-0.338	-2.344	-1.171	-1.948	-0.982	-0.482	-0.319
		(0.021***)	(0.014***)	(0.020***)	(0.017***)	(0.021***)	(0.015***)	(0.007***)	(0.010***)
Pr(Business) Constant	γ^1	-1.980	-4.574	-5.999	-5.318	-0.081	-30.701	-20.000	-4.463
		(0.058***)	(1.396***)	(0.638***)	(0.452***)	(0.041**)	(4.891***)	(2.324***)	(0.538***)
Pr(Business) Slope	γ^2	0.013	0.005	0.014	0.280	-0.035	1.082	0.610	0.050
		(0.005**)	(0.055)	(0.031)	(0.030***)	(0.004***)	(0.239***)	(0.101***)	(0.023**)
Pr(Business) Quadratic	γ^3	0.001	0.002	0.002	-0.003	0.002	-0.009	-0.004	0.001
		(0.000***)	(0.001***)	(0.000***)	(0.000***)	(0.000***)	(0.003***)	(0.001***)	(0.000***)
Poisson Rates									
Greater than 21 days	μ^1	3.508	2.652	0.749	5.542	1.486	0.925	3.233	3.659
		(0.136***)	(0.173***)	(0.004***)	(0.475***)	(0.013***)	(0.009***)	(0.085***)	(0.149***)
14 to 21 days	μ^2	2.836	3.347	1.092	4.301	1.039	0.646	3.295	4.185
		(0.088***)	(0.215***)	(0.009***)	(0.389***)	(0.015***)	(0.012***)	(0.099***)	(0.160***)
7 to 14 days	μ^3	2.514	3.686	1.070	3.349	0.828	0.451	2.417	3.524
		(0.067***)	(0.235***)	(0.009***)	(0.319***)	(0.014***)	(0.016***)	(0.075***)	(0.127***)
Less than 7 days	μ^4	1.294	3.699	0.951	3.589	0.644	0.575	3.083	4.412
		(0.034***)	(0.242***)	(0.009***)	(0.354***)	(0.016***)	(0.023***)	(0.142***)	(0.167***)
Firm Shock									
	σ	0.991	0.166	0.298	0.645	0.865	0.204	0.836	0.339
		(0.009***)	(0.001***)	(0.001***)	(0.005***)	(0.009***)	(0.001***)	(0.007***)	(0.003***)
Log Likelihood		47,036	76,247	278,381	95,160	61,173	54,782	34,965	39,524
Num. Flights		232	423	1,626	464	312	333	173	210
Num. Dep. Dates		116	106	106	116	116	106	87	106
Num. Obs.		13,389	25,205	96,807	25,843	18,011	19,846	10,322	12,407

Note: Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ The $\Pr_t$Business function is specified as a logistic function with form $\gamma_0 + \gamma_1 t + \gamma_2 t^2$. Direction markets are combined for each city pair. Num. flights reports the total number of unique flights tracked, num. dep. dates reports the number of departure dates observed. Prices are scaled to $100.

Table 7: Parameter Estimates

Variable		FATPDX	GEGPDX	GTFSEA	HLNSEA	ICTSEA	LIHPDX	MSOPDX	OKCSEA
Logit Demand									
Intercept	β_0	-0.811	3.846	0.938	6.500	0.695	9.878	9.177	4.620
		(0.341**)	(0.034***)	(0.188***)	(0.106***)	(0.210***)	(0.402***)	(0.223***)	(0.151***)
Leisure Price Sens.	α^L	-1.411	-3.264	-1.645	-3.617	-0.833	-3.878	-5.496	-1.642
		(0.037***)	(0.021***)	(0.030***)	(0.051***)	(0.033***)	(0.145***)	(0.116***)	(0.044***)
Business Price Sens.	α^B	-0.360	-1.435	-0.470	-1.706	-0.165	-1.787	-2.993	-0.531
		(0.013***)	(0.011***)	(0.019***)	(0.026***)	(0.017***)	(0.061***)	(0.065***)	(0.028***)
Pr(Business) Constant	γ^1	-3.758	-1.353	-38.548	-108.573	-3.267	-0.376	-2.078	-2.027
		(0.198***)	(0.020***)	(6.556***)	(5.1e4)	(0.156***)	(0.095***)	(0.111***)	(0.066***)
Pr(Business) Slope	γ^2	0.010	-0.055	1.398	1.598	0.045	-0.104	-0.008	0.042
		(0.008)	(0.001***)	(0.321***)	(255.451)	(0.005***)	(0.008***)	(0.009)	(0.003***)
Pr(Business) Quadratic	γ^3	0.001	0.002	-0.012	0.027	-0.000	0.001	0.002	-0.000
		(0.000***)	(0.000***)	(0.004***)	(3.152)	(0.000)	(0.000***)	(0.000***)	(0.000***)
Poisson Rates									
Greater than 21 days	μ^1	10.075	0.834	3.300	0.871	3.220	0.880	0.973	1.399
		(2.794***)	(0.005***)	(0.393***)	(0.012***)	(0.339***)	(0.018***)	(0.012***)	(0.023***)
14 to 21 days	μ^2	13.766	1.118	3.116	0.463	3.645	0.719	0.573	1.682
		(3.660***)	(0.007***)	(0.308***)	(0.011***)	(0.384***)	(0.034***)	(0.015***)	(0.036***)
7 to 14 days	μ^3	13.070	1.072	2.040	0.335	4.327	0.660	0.368	1.725
		(3.380***)	(0.007***)	(0.196***)	(0.007***)	(0.448***)	(0.028***)	(0.011***)	(0.031***)
Less than 7 days	μ^4	13.594	0.565	2.268	0.230	6.674	0.863	0.244	2.285
		(3.487***)	(0.003***)	(0.215***)	(0.006***)	(0.687***)	(0.026***)	(0.007***)	(0.036***)
Firm Shock									
	σ	0.173	0.043	0.143	0.259	0.289	0.644	0.181	0.312
		(0.001***)	(0.000***)	(0.001***)	(0.002***)	(0.007***)	(0.010***)	(0.002***)	(0.007***)
Log Likelihood		75,797	234,198	51,490	30,208	37,534	14,846	30,032	33,982
Num. Flights		424	1,718	412	212	212	96	208	212
Num. Dep. Dates		106	106	106	106	106	48	106	106
Num. Obs.		25,273	102,292	24,531	12,644	12,608	4,552	12,390	12,581

Note: Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ The $\Pr_t$Business function is specified as a logistic function with form $\gamma_0 + \gamma_1 t + \gamma_2 t^2$. Direction markets are combined for each city pair. Num. flights reports the total number of unique flights tracked, num. dep. dates reports the number of departure dates observed. Prices are scaled to \$100.

Table 8: Parameter Estimates

Variable		OMASEA	PDXPSP	PDXRNO	PDXSBA	PDXSTS	SBASEA	SEASTS	SEASUN
Logit Demand									
Intercept	β_0	5.899	0.110	4.359	1.392	3.162	3.664	3.859	1.286
		(0.139***)	(0.272)	(0.066***)	(0.155***)	(0.119***)	(0.124***)	(0.123***)	(0.099***)
Leisure Price Sens.	α^L	-2.161	-4.722	-2.460	-1.260	-1.760	-1.930	-1.814	-1.434
		(0.045***)	(0.431***)	(0.030***)	(0.040***)	(0.047***)	(0.046***)	(0.045***)	(0.020***)
Business Price Sens.	α^B	-0.770	-1.198	-1.198	-0.489	-0.745	-0.880	-0.902	-0.693
		(0.026***)	(0.019***)	(0.016***)	(0.016***)	(0.022***)	(0.022***)	(0.023***)	(0.021***)
Pr(Business) Constant	γ^1	-4.146	-2.477	-2.751	-2.062	-5.633	-6.645	-2.848	-72.852
		(0.186***)	(0.252***)	(0.107***)	(0.093***)	(1.552***)	(1.220***)	(0.283***)	(118.611)
Pr(Business) Slope	γ^2	0.076	0.049	-0.105	-0.016	0.006	0.099	-0.118	1.995
		(0.008***)	(0.005***)	(0.006***)	(0.005***)	(0.070)	(0.064)	(0.013***)	(5.023)
Pr(Business) Quadratic	γ^3	-0.000	-0.000	0.003	0.001	0.002	0.001	0.004	-0.010
		(0.000)	(0.000***)	(0.000***)	(0.000***)	(0.001***)	(0.001)	(0.000***)	(0.053)
Poisson Rates									
Greater than 21 days	μ^1	2.068	24.957	0.951	2.403	1.195	2.039	1.232	1.896
		(0.020***)	(9.569***)	(0.008***)	(0.147***)	(0.023***)	(0.037***)	(0.021***)	(0.094***)
14 to 21 days	μ^2	2.207	22.238	1.277	2.643	2.062	2.670	1.664	1.440
		(0.033***)	(8.484***)	(0.014***)	(0.145***)	(0.042***)	(0.051***)	(0.037***)	(0.071***)
7 to 14 days	μ^3	2.081	22.002	1.251	2.844	1.772	1.986	1.550	0.798
		(0.032***)	(8.417***)	(0.013***)	(0.149***)	(0.037***)	(0.049***)	(0.032***)	(0.042***)
Less than 7 days	μ^4	5.025	29.475	1.555	3.484	1.397	1.224	1.247	0.665
		(0.046***)	(11.384***)	(0.015***)	(0.174***)	(0.031***)	(0.030***)	(0.023***)	(0.038***)
Firm Shock									
	σ	0.988	0.594	0.235	0.207	0.242	0.464	0.406	0.175
		(0.013***)	(0.006***)	(0.001***)	(0.002***)	(0.002***)	(0.004***)	(0.003***)	(0.002***)
Log Likelihood		51,934	40,800	130,491	39,725	58,173	38,817	52,236	26,458
Num. Flights		212	216	721	212	384	212	306	168
Num. Dep. Dates		106	48	106	106	106	106	106	85
Num. Obs.		12,633	10,465	42,997	12,626	22,867	12,377	18,205	10,076

Note: Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ The $\Pr_t$Business function is specified as a logistic function with form $\gamma_0 + \gamma_1 t + \gamma_2 t^2$. Direction markets are combined for each city pair. Num. flights reports the total number of unique flights tracked, num. dep. dates reports the number of departure dates observed. Prices are scaled to \$100.

Figure 13: Fitted Values of γ_t over Time for Each Market

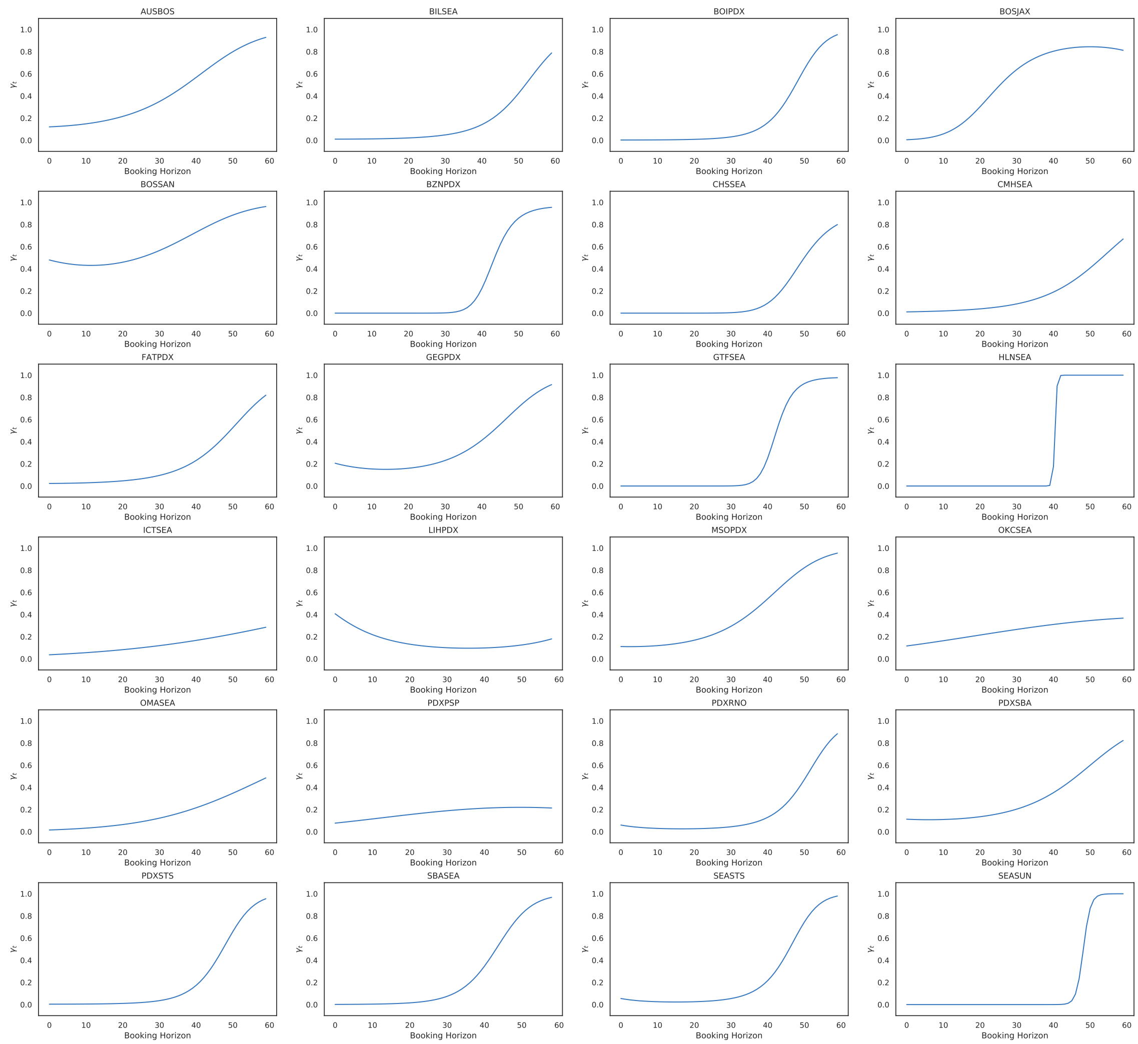

Note: Estimates of γ_t for each market separately. The function γ_t is defined as $\exp(\gamma_1 + \gamma_1 t + \gamma_1 t^2)/(1 + \exp(\gamma_1 + \gamma_1 t + \gamma_1 t^2))$.

Counterfactuals with Restricted Pricing Choice Set

Table & Figure 9: Dynamic to Uniform Pricing (Restricted Choice Set)

	Fare	Load Factor	Sell Outs	Revenue	CS_L^i	CS_B^i	Welfare
Dynamic	243.7	87.6	16.5	11475.0	49.4	161.5	466.7
Uniform	223.7	84.7	28.5	10565.5	46.4	180.2	464.5
Difference (%)	-8.2	-2.9	12.0	-7.9	-6.1	11.6	-0.5

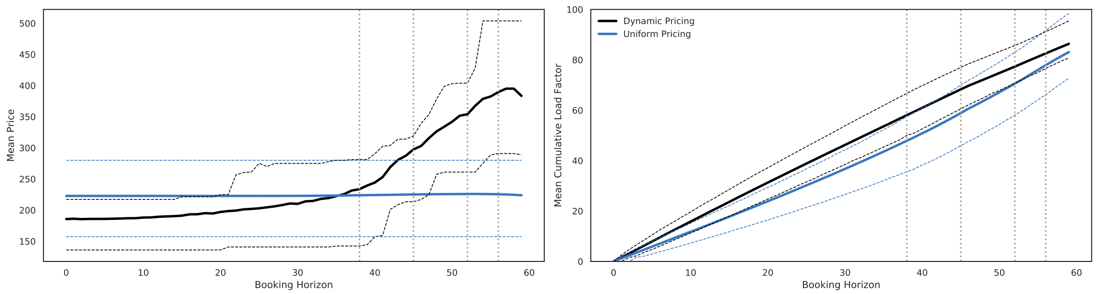

Table & Figure 10: Dynamic Pricing to Intermediate Case (Restricted Choice Set)

	Fare	Load Factor	Sell Outs	Revenue	CS_L^i	CS_B^i	Welfare
Dynamic	243.7	87.6	16.5	11475.0	49.4	161.5	466.7
Intermediate	244.6	83.7	21.0	11188.0	48.1	158.8	456.0
Difference (%)	0.4	-4.0	4.5	-2.5	-2.5	-1.7	-2.3

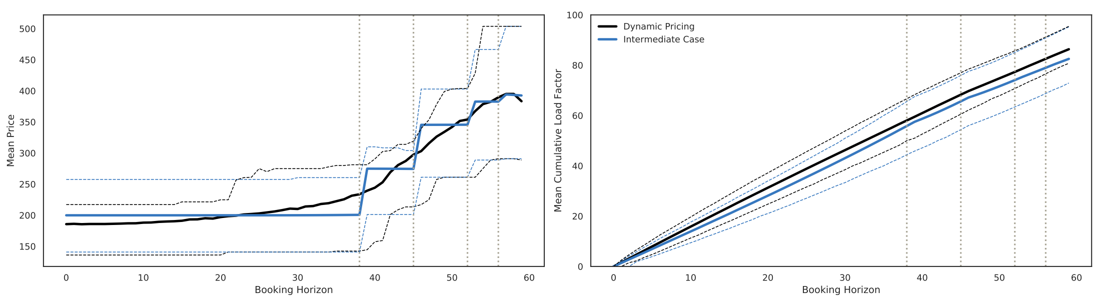

B Route Selection

Using the publicly available DB1B data, I select origin-destination pairs to study. These data contain a 10 percent sample of bookings and are at the quarterly level. The data contain neither the date of travel nor the date of purchase.

I first combine traffic from all airports in which there exists a nearby airport within sixty miles. This combines, for example, Laguardia (LGA), John F. Kennedy (JFK), and Newark (EWR).[33] Next, I focus on ODs with a nonstop option; this reduces the number of potential markets studied from 80,000 to 10,000. Over 70 percent of these markets have a single carrier providing nonstop service and they make up a total of 10 percent of OD traffic in the United States. Out of these 10,000, I implement the following cleaning criteria: (1) total quarterly traffic, including connecting traffic with up to four stops, exceeds 600 passengers;[34] (2) a single carrier operates nonstop on the OD leg. This reduces the number of potential markets by over half, to roughly 4,000.

Next, I calculate the following statistics: (1) OD nonstop traffic; (2) OD total traffic (including one-stop connections, all the way up to four-stop connections); (3) passenger traffic connecting to OD or connecting from OD, which again is allowed to have at most five legs. The fraction (1)/(2) calculates the percentage of traffic flying nonstop. The fraction

[33]This creates the following groupings: (DAB, MCO, SFB); (OGD, SLC); (EWN, OAJ); (KOA, MUE); (SBP, SMX); (AZA, PHX); (BRO, HRL, MFE); (CMI, DEC); (PIE, SRQ, TPA); (MHT, PSM); (BUR, LAX, LGB, ONT, SNA); (BTV, PBG); (BFM, MOB); (HHH, SAV); (DAL, DFW); (EVV, OWB); (MSS, OGS); (BQN, MAZ); (PSG, WRG); (HOU, IAH); (ORF, PHF); (FAT, VIS); (ATW, GRB); (PAE, SEA); (LNS, MDT); (CLT, USA); (OAK, SFO, SJC); (AOO, JST, LBE); (BLV, STL); (CPX, SPB, STT, VQS); (LWS, PUW); (BGM, ELM, ITH); (BGR, BHB); (ACK, EWB, HYA, MVY, PVC, PVD, BOS); (BWI, DCA, IAD); (CLD, SAN); (CHO, SHD); (ASE, EGE); (SCM, VAK); (GYY, MDW, ORD); (BUF, IAG); (CMH, LCK); (PHL, TTN); (PGD, RSW); (FLL, MIA); (HNM, JHM, LNY, LUP, MKK, OGG); (MCE, MOD, SCK); (LEB, RUT); (CKB, MGW); (GLV, WMO); (EWR, HPN, HVN, ISP, JFK, LGA, SWF).

[34]This is calculated as half a fifty-seat plane, offering at least weekend service (eight monthly flights), for the quarter, e.g. .5*50*8*3 = 600. This level of the criterion is not critical, but a minimum passenger threshold of 10 (scaling 1 passenger up to 10, as it is a 10% sample) is important because it removes erroneous entries in the DB1B. For example, in 2012, United Airlines did not operate nonstop between Lehigh Valley International Airport (ABE) and Nashville (BNA). Another method is to look at scheduled service in the T100 segment tables.

(1)/[(1) + (3)] calculates the percentage of traffic not connecting. Shown another way,

$$\begin{aligned}\text{FracNonstop} &:= \frac{\text{Passengers OD Nonstop}}{(\text{Passengers OD Nonstop}) + (\text{Passengers OD } \geq \text{ 1 Stops})}\\ &:= \frac{(O \to D)}{(O \to D) + (O \to C \to D)},\end{aligned}$$

where C denotes potential connections for passengers flying on OD. Using similar notation,

$$\text{FracNotConnecting} := \frac{(O \to D)}{(O \to D) + (C \to O \to D) + (O \to D \to C)},$$

which is simply the fraction of passengers on planes flying OD that are not connecting on either end.

Single carrier markets have percent nonstop and percent non-connecting means of 74 percent and 55 percent as compared with 81 percent and 59 percent for competitive markets (medians of 80 percent, 52 percent, 86 percent, 59 percent, respectively). I limit myself to markets with at most 15,000 monthly passengers. This is to keep the data collection process manageable.

The two fractions are negatively correlated ($\rho = -0.34$), each is correlated with distance. The correlation between percentage non-connecting and distance is 0.23; nearby ODs have higher connecting traffic. The correlation between percentage nonstop and distance is -0.53; nearby ODs have a higher percentage of nonstop traffic. Both make intuitive sense. Passengers at small spoke airports connect to a nearby hub and then travel to their destinations. Hence, non-connecting traffic decreases with distance. On the other hand, nonstop traffic percentage is strongly correlated with distance: if a passenger wants to fly from Madison, WI to Chicago, IL, it makes little sense to connect in Detroit, MI. ODs in with a nonstop percentage above 95 percent, are ODs with an average distance of 400 miles; ODs below 95 percent have an average distance of over 1,000 miles.

It may make sense to maximize both metrics in selecting routes. There is one important caveat. Although a high nonstop percentage implies consumers are not connecting to reach their destination via a flight, it does not mean alternative transportation options for these consumers (such as trains) do not exist. In 2019, there exist 661 ODs that meet the arbitrary threshold of 95 percent in both metrics. Of those ODs, 607 are operated by low-cost carriers Allegiant Air and Spirit Airlines. Unfortunately, both airlines charge for a seat assignment; thus, utilizing seat maps to determine bookings will likely be inaccurate. The next two carriers that commonly meet threshold criteria are JetBlue Airways and Alaska Airlines.

I select fifty ODs and concentrate data collection on two carriers, JetBlue Airways and Alaska Airlines. In addition, for a comparison in the descriptive analysis, I collect data on six duopoly markets.[35] Figure 14 maps the markets and Table 11 provides a dictionary for the airport codes. The data were collected in two phases: The data on markets operated by JetBlue were collected in 2012, and the data for Alaska Air Lines were collected in 2019. Prices for data collected in 2012 are adjusted for inflation.

Figure 14: Markets of Study

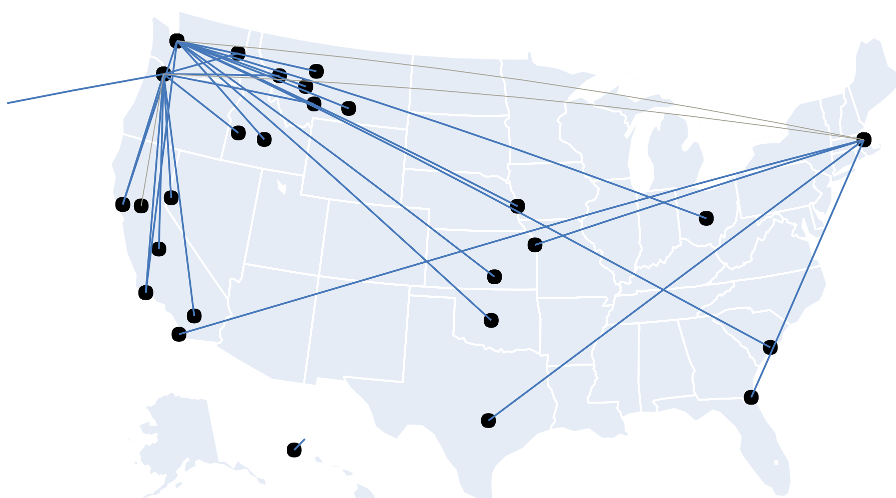

Note: Map of the markets selected for study. All of the markets either start or end at Seattle, WI; Portland, OR; and Boston, MA.

Figure 15 depicts all OD pairs in the DB1B data that meet the thresholds stated above. Each dot corresponds to an OD pair. The vertical axis reports the percentage (0-100) of non-connecting traffic. The horizontal axis reports the percentage of nonstop traffic. The left panel (a) includes all markets, and the right panel (b) removes Allegiant and Spirit

[35] The city pair Boston, MA - Kansas City, MO was a duopoly market, with nonstop offered by both Delta Air Lines and Frontier Airlines in 2012. Frontier then exiting the market.

because of the fee charged to select seats. These 607 ODs removed in (b) lie mostly along the top of the graph, corresponding to markets with 100 percent non-connecting traffic. The red squares show the markets selected for data collection and analysis. The dashed grey lines show the mean of each statistic and the solid black line depicts the fit of a linear simple regression.

The graphs show the negative correlation between the two statistics previously mentioned, with a large cluster of ODs having close to 100 percent nonstop traffic but also very high levels of connecting traffic. For this study, “ideal" markets arguably lie in the upper right of the graph. These are markets in which most consumers travel nonstop (versus one-stop) and do not connect to other flights. Note that this region is less dense compared with other areas in the graph. The graphs show that all but nine (panel a) or five (panel b) of the selected markets appear above the regression line, and most lie in the upper-right region of the graph.

Figure 15: Nonstop and Non-Connecting Traffic in the DB1B

(a) Full Sample (b) Removing Allegiant & Spirit

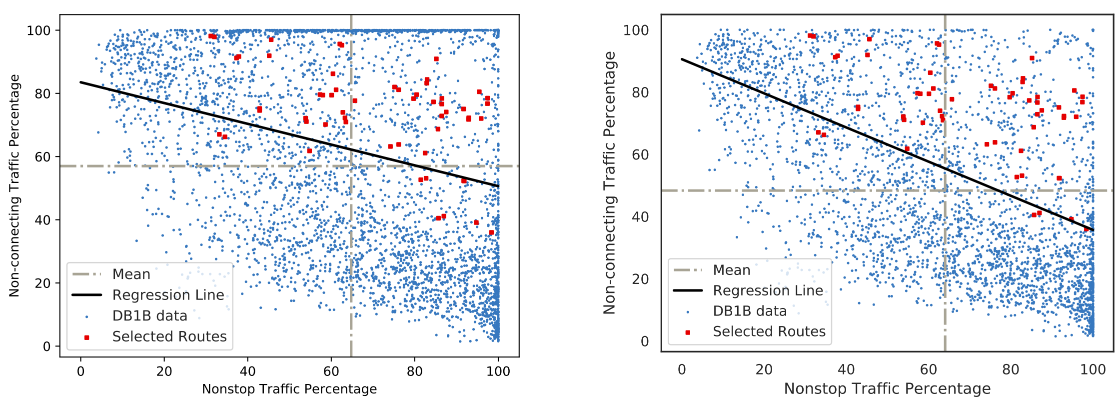

Note: (a) Percentage nonstop traffic and percentage non-connecting traffic for markets that meet selection criteria in the DB1B data. (b) Repeat of (a), excluding markets operated by Allegiant and Spirit.

Table 12 provides traffic and price statistics in the DB1B for each OD in the sample. Note that OD fares are very similar to DO (the reverse) fares in the DB1B, and I use this finding in order to aggregate observations in estimation. Finally, one-stop fares are not necessarily cheaper than nonstop options. For example, nonstop fares from Billings, MT to Seattle, WA are cheaper than one-stop connections.

Table 11: Airport Code Lookup

Airport Code	City	Airport Code	City
AUS	Austin, TX	JAX	Jacksonville, FL
BIL	Billings, MT	LIH	Lihue, HI
BOI	Boise, ID	MSO	Missoula, MT
BOS	Boston, MA	OKC	Oklahoma, OK
BZN	Bozeman, MT	OMA	Omaha, NE
CHS	Charleston, SC	PDX	Portland, OR
CMH	Columbus, OH	PSP	Palm Springs, CA
FAT	Fresno, CA	RNO	Reno, NV
GEG	Spokane, WA	SAN	San Diego, CA
GTF	Great Falls, MT	SBA	Santa Barbara, CA
HLN	Helena, MT	STS	Santa Rosa, CA
ICT	Wichita, KS	SUN	Sun Valley, ID

Table 12: Markets Summary

Origin	Dest.	Nonstop_{OD}	Nonstop_{DO}	Total_{OD}	Total_{DO}	Connect._{OD}	Connect._{DO}	Fare_{OD}	Fare_{DO}	1-stop Fare
AUS	BOS	8,105	7,860	8,277	8,005	17,425	17,687	274	262	237
BIL	SEA	5,213	5,356	9,743	9,946	1,153	1,066	202	201	246
BOI	PDX	25,963	26,423	35,940	36,840	1,966	2,006	136	137	134
BOS	JAX	12,655	12,837	13,222	13,470	7,660	7,622	188	188	193
BOS	MCI	7,847	7,952	8,607	8,220	9,372	9,280	229	227	248
BOS	PDX	13,015	13,140	16,250	16,497	9,642	9,480	270	273	275
BOS	SAN	15,915	15,770	17,435	17,335	26,242	26,495	302	301	262
BOS	SEA	28,055	27,705	34,935	35,192	17,555	18,240	297	299	283
BZN	PDX	4,130	4,180	6,480	6,436	1,356	1,190	156	153	190
CHS	SEA	6,440	6,390	8,986	8,823	5,363	5,390	257	248	261
CMH	SEA	8,730	8,745	11,630	11,735	11,755	11,810	257	262	224
FAT	PDX	7,563	7,576	10,326	10,196	1,110	1,036	161	162	194
GEG	PDX	20,926	21,030	40,003	40,203	1,850	1,883	142	142	133
GTF	SEA	2,733	2,940	7,670	7,696	46	156	177	177	238
HLN	SEA	2,190	2,150	5,320	5,303	326	356	159	163	266
ICT	SEA	4,053	4,233	5,460	5,843	2,533	2,526	216	225	229
LIH	PDX	6,680	7,085	10,760	10,110	5,300	4,870	245	236	288
MSO	PDX	3,736	3,493	5,426	5,696	623	736	156	158	198
OKC	SEA	3,940	3,940	5,946	5,873	7,990	8,330	252	258	192
OMA	SEA	9,336	9,863	13,063	12,646	5,510	5,120	184	177	196
PDX	PSP	21,430	14,906	23,480	16,496	3,715	4,133	142	147	177
PDX	RNO	12,710	12,633	16,423	16,046	2,306	1,950	152	153	136
PDX	SBA	4,530	4,620	5,500	5,683	1,526	1,510	195	195	175
PDX	SMF	47,886	47,573	60,966	62,033	1,260	1,266	143	142	180
PDX	STS	5,433	5,483	6,710	7,506	213	196	167	171	218
SBA	SEA	8,580	8,560	10,933	10,746	2,220	2,133	168	168	179
SEA	STS	5,843	5,566	7,803	7,310	483	743	180	179	192
SEA	SUN	3,253	3,076	3,856	3,696	490	460	153	156	212
All Markets			22,237		41,105		4,708		181	235

Note: Origin-destination (OD) statistics for the markets included in the sample. Summary statistics calculated as the quarterly average in the year in which data was collected. 2019 data is summarized with the three available quarters. 2012 fare data is adjusted for inflation. The All Markets row summarizes all ODs in the DB1B data, including all observations with less than five connections. The subscript *OD* denotes corresponds to the Origin and Destination columns; *DO* denotes the reverse routing. (PDX,SMF); (SEA,BOS); and (BOS,PDX) are oligopoly markets. (BOS,MCI) was a duopoly market prior to the exit of Frontier. All markets operated by Alaska, except (BOS,MCI) [Delta], (AUS,BOS) [JetBlue], (BOS,JAX) [JetBlue], and (BOS,SAN) [JetBlue].

C Inference and Accuracy of Seat Maps

I utilize airline seat maps to recover bookings for flights. 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 accordingly. 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.

Unfortunately, seat maps may not accurately represent 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 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 are imperfect, I perform two analyses to gauge the magnitude of the measurement error in using seat maps.

Figure 16: Estimated Seat Map Measurement Error at the Monthly Level

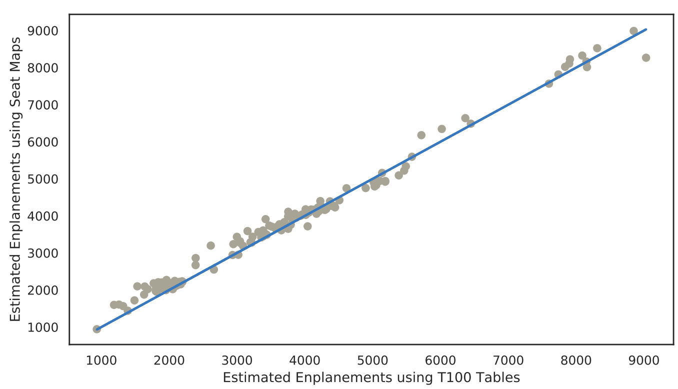

Note: Measurement error estimated by comparing monthly enplanements, using the T100 Tables and aggregating seat maps to the monthly level. The solid line reflects zero measurement error.

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 16 provides a scatter plot that compares the two statistics. All the points closely follow the 45-degree line, and I find seat maps overstate recorded enplanements, with the median difference being 2.57 percent. Since seat maps are not obtainable after departure, some of this difference could be driven by last-minute cancellations.

Second, I create a new data set that allows me to estimate seat-map measurement error for each day before departure. The mobile version of United.com allows users to examine seat maps for upcoming flights. In addition, for premium cabins, the airline reports the number of consumers booked into the cabin. I randomly select flights, departure dates, and search dates. In total, I obtain 15,567 observations. With these data, I find that seat maps understate reported load factor by 2.3 percent, or around one to two seats on average.

Figure 17: Estimated Seat Map Measurement Error by Day Before Departure

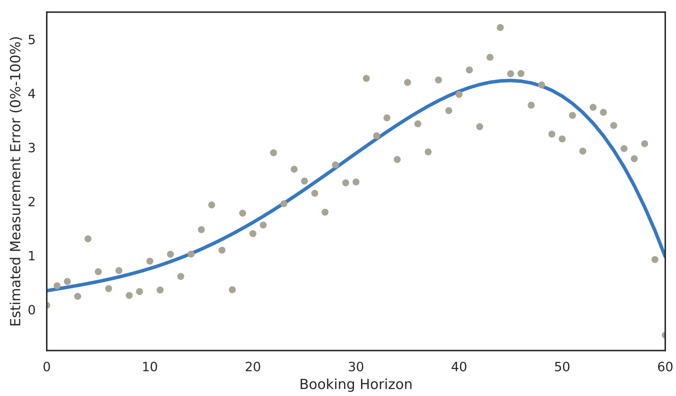

Note: Measurement error estimated by comparing seat maps with reported load factor using the United Airlines mobile website. The dots correspond to the daily mean, and the line corresponds to fitted values of an orthogonal polynomial regression of the fourth degree. Total sample size is equal to 15,567,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

Figure 17 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 ranges between 0 to 5 percent across days, or at most four seats. This suggests seat maps are useful for recovering bookings as the departure date approaches. However, 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 (sixty to forty days out) and close to the departure date. Second, compared to the main sample, these data cover a different airline, different markets, and a different cabin.

D Additional Descriptive Analysis

The data show that the patterns documented in monopoly markets are also present in competitive airline markets (in Appendix D.1). I then show that first class tickets also respond to bookings, and that the trajectory in prices is overwhelmingly positive (in Appendix D.2). Finally, I discuss how alternative flight options respond to nonstop bookings (in Appendix D.3).

D.1 Competitive Pricing Dynamics

Figures 18-(a) and 18-(b) recreate Figure 3 and Figure 1, respectively, for markets with competition. Although more noisy—the sample is one-fifth the size—the general patterns are present: (i) fares tend to increase when sales occur; (ii) fares tend to decrease when sales do not occur; (iii) fares increase, regardless of sales, close to the departure date. However, Figure 3 also suggests, at least for the markets selected, that competition may decrease the magnitude of systematic fare increases and increase the number of fare decreases.

Figure 18: Fare Dynamics in Competitive Markets

(a) Frequency and Magnitude of Fare Changes by Day Before Departure

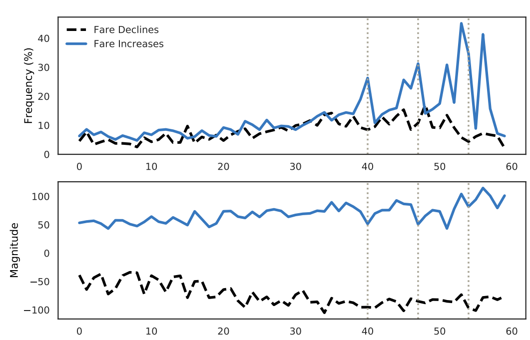

(b) Fare Response to Sales by Day Before Departure

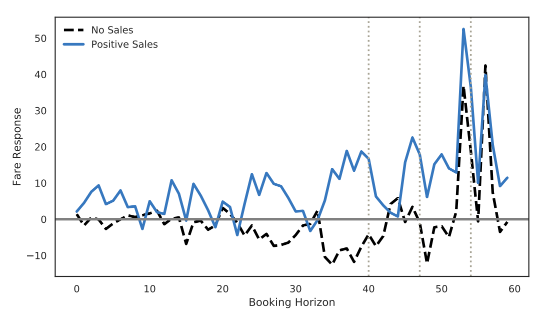

Note: (a) Reproduction of Figure 1 for routes with competition. (b) Reproduction of Figure 3 for routes with competition.

These descriptive facts support findings in Siegert and Ulbricht (2019), use fare data to show that competition is correlated with a flattening of prices over time. Dana and Williams (2020) show in a theoretical model of sequential quantity-pricing games that strong competitive effects work to equalize prices across periods and that inventory controls can facilitate intertemporal price discrimination in oligopoly. If the role of intertemporal price discrimination is reduced in competitive markets, this may suggest the efficiency aspect of dynamic pricing may be higher, and the ability to extract surplus from business customers lower, compared to the markets studied in this paper.

D.2 Static Forms of Price Discrimination and Dynamic Versioning

Figure 19 shows the average prices and availability of different fare categories over time. I use the term "fare categories" to describe versioning broadly, different qualities of both economy tickets as well as first-class ones. Panel (a) shows five different versions of tickets. Full-fare (refundable) tickets for both economy and first class are flat over time. Average prices for saver-economy, economy, and first-class tickets raises over time, with first class and economy class following similar patterns. However, the gap in prices between saver-economy and economy prices grows as the departure date approaches. Panel (b) recreates Figure 3 for first class. There are two noticeable differences: the presence of APDs is diminished in first class; and fare increases are more pronounced throughout the booking horizon. It is unclear why first-class prices decline forty-five days before departure, regardless of bookings.[36]

Figure 19(c) shows the availability of saver-economy, economy, and first-class tickets over time (full-fare availability closely follows the two respective non-full fare versions). Panel (c) establishes two results: First, there is a general downward trend in availability as flights start selling out. Close to departure, economy fare availability abruptly drops suggesting that the spike in load factor shown in Figure 2 captures last-minute bookings. Second, the availability of the lowest-quality economy tickets (saver economy) substantially decreases over time. The large declines coincide with APD restrictions. This implies that average fares go up over time for two reasons: economy fares become more expensive, and

[36] The systematic decrease in first-class prices occurs in several, but not all, markets. I find no associated spike in bookings.

Figure 19: Fare Category Pricing Dynamics

(a) Mean Fare Category Pricing over Time

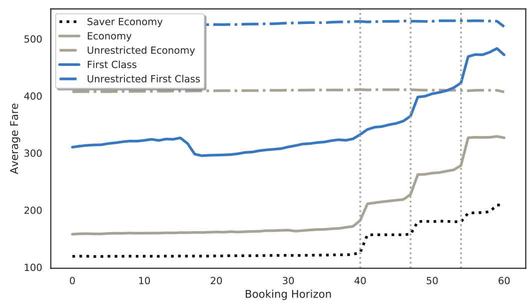

(b) Fare Response to First Class Bookings over Time

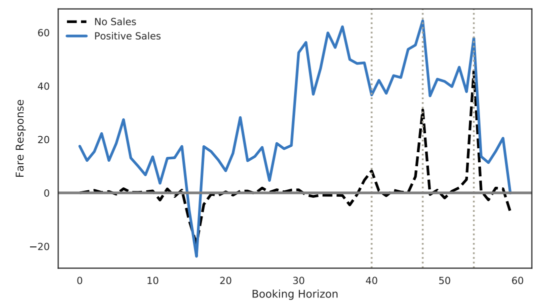

(c) Fare Category Availability over Time

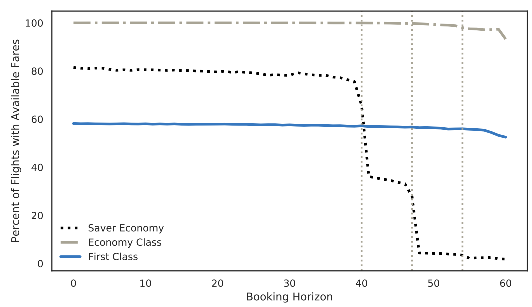

Note: (a) Mean prices of different fare categories over time. (b) Recreation of Figure 3, but separately showing price effects for two fare types. (c) Percentage of flights that offer individual fare categories over time. First-class denominator is the number of flights with first class, not the number of flights in the sample.

saver fares are no longer offered.

D.3 Flight Substitutability: Connecting and Nonstop Options

The data show that connecting flight fares do not respond to nonstop bookings in the same way as nonstop bookings on nonstop fares. Figure 20(a) recomputes Figure 3; the outcome is the fare response on connecting fares instead of own-flight fares. The connecting fare is the average fare among connecting flight options for the same carrier corresponding to the same departure date and day before departure.[37] There are two noticeable differences between the figures. First, greater than twenty-one days from departure, the two lines coincide. This means that connecting fares are unaffected by nonstop bookings. Second, closer to the departure date, unlike Figure 3, I find that connecting fares always increase in

[37]Specifically, I take the average of connecting flights at the origin-destination-departure date-day before departure date level and then merge it to the nonstop OD data at the same level.

both cases (no sales and positive sales).

Figure 20: Pricing Effects on Other Itineraries

(a) Connecting Fare Response

(b) Multiple Nonstop Prices and Load Factors

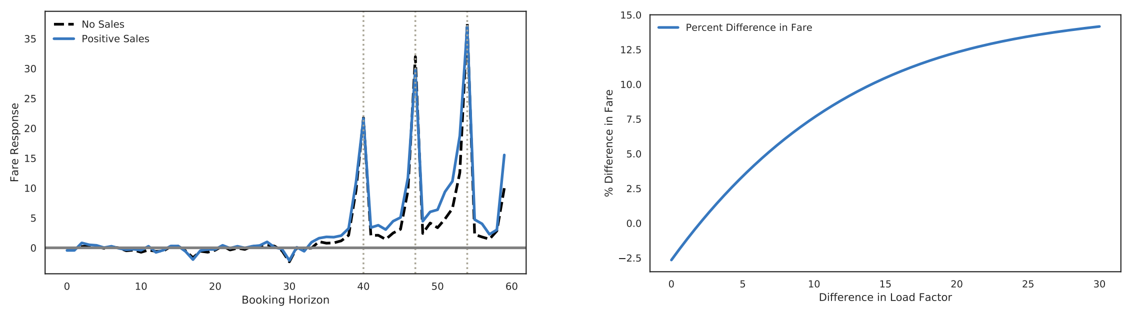

Note: (a) Recreation of Figure 3, but with connecting fares. Average connecting fare changes as a response to sales by day before departure. The vertical lines correspond to advance-purchase discount periods (fare fences). The horizontal line indicates no fare response. (b) Fourth order polynomial fit of a regression of the percent difference in fares on the percent difference in load factor when a carrier operates two flights a day.

I also use the data to document pricing differences when a carrier offers twice-daily service. I first calculate the differences over time in load factors and fares across the two flights. When flights have the same load factor, I find that the average difference in fares is -.07 percent. Figure 20(b) plots fitted values of a fourth order polynomial fit of percent difference in fares on the percent difference in load factors. The line is upward sloping, meaning that the flight with the higher load factor is, on average, more expensive. Examining fare responses, I find that if one flight sells a seat but the other one does not, the flight that sells a seat becomes more expensive by \$17.78, on average, but the other flight option becomes \$1.40 cheaper.[38] We might have expected both prices to increase, owing to an overall reduction in remaining capacity if airlines internalized the substitution effect. This outcome is consistent with a model in which the flights are independently priced. It is also consistent with a learning model in which a demand shock on one flight informs the airline that the other flight is less preferred.

[38]The test statistics for these two scenarios are $t = 62$ and $t = -7$.

E Dynamic Demand in Airline Markets

There are noticeable jumps in prices over time, however, the booking curve for flights is smooth. If consumers are aware that fares tend to increase sharply around APD requirements, and they can strategically enter into the market, we should expect to see bunching in sales before APDs expire and few sales after expiration. At the same time, if enough consumers behaved in this way, having APD on set days before departure is not optimal. Two possibilities are that (a sufficient number of) consumers do not strategically enter—indeed, Li, Granados, and Netessine (2014) estimate that between 5 and 20 percent of consumers dynamically substitute across days—or firms are not correctly responding to the strategic timing of purchasing tickets.

Theoretical work has considered strategic consumers in durable, storable, and consumption goods markets.[39] For the airline context, a relevant paper is Dilmé and Li (2019), who consider strategic consumers, capacity, and a deadline. They show that in response to waiting, firms may offer flash sales to reduce inventory, and doing so increases the later arrivals' willingness to pay. I do not model this strategic interaction and note that this assumption can impact willingness to pay or timing of participation (Hendel and Nevo, 2013; Sweeting, 2012). However, unlike environments where prices decrease (Nair, 2007), the strong upward trajectory of prices greatly diminishes the incentives to delay, conditional on consumers knowing their preferences.

I investigate bunching (strategic purchasing timing) by modeling the booking curve as a function of time and include dummy variables for the day-before-departure (DFD) times immediately before APDs expires. Table 13 reports regression results under three fixed effects specifications. I find insignificant bunching at the fourteen-day APD period, although the signs are positive. I find negative bunching at the three-day and twenty-one-day APD period, meaning sales are lower prior to the price increases. Finally, I find a positive and significant coefficient for the seven-day APD; that is, sales are higher before the usual seven-day fare increase. It may be that at least some consumers anticipate price hikes and time their purchases accordingly.

[39]These include Stokey (1979); Bulow (1982); Conlisk, Gerstner, and Sobel (1984); Sobel (1991); Board and Skrzypacz (2016); Öry (2016).

Table 13: Consumer Bunching Regressions

	(1)	(2)	(3)
APD_3	-1.095***	-1.093***	-1.091***
	(0.144)	(0.146)	(0.146)
APD_7	0.389**	0.388**	0.387**
	(0.0737)	(0.0742)	(0.0747)
APD_{14}	0.176	0.172	0.173
	(0.0959)	(0.0925)	(0.0919)
APD_{21}	-0.0984	-0.101	-0.101
	(0.0545)	(0.0503)	(0.0500)
$m(t)$	Yes	Yes	Yes
OD FE	Yes	Yes	–
Month FE	Yes	Yes	–
D.o.W. Search FE	No	Yes	Yes
D.o.W. Departure FE	No	Yes	–
Flight FE	No	No	Yes
Observations	734,297	734,297	734,297
R^2	0.593	0.602	0.907

Note: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. $m(t)$ is a sixth-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. OD-Month clustered standard errors in parentheses.

The collected data make it impossible to discern if a ticket was purchased by a waiting consumer or a consumer that just arrived. However, if consumers strategically time purchases, we would expect that they not buy immediately after the APD expires—on average fares only continue to increase. However, I find little evidence for this: the booking rate after APDs expire is the same as either the booking rate the day before or the booking rate two days before. This is also the case when the 7-day APD expires. Here, the booking rate drops by 0.08 for one day and then immediately returns to the same level before the APD expiration date. This finding supports modeling consumers under Poisson arrivals, as bookings occur throughout time even though prices rise.

I also investigate the incentive to wait by changing the estimated model in the following way: after consumers arrive, each consumer has the option to 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 in the next period.

I derive a waiting cost $\overline{\phi}$ such that if all consumers have a waiting cost at least as high as $\overline{\phi}$, 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

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

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

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

To derive $\overline{\phi}$, 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

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

For this marginal consumer, the expected gains from waiting are

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

Hence, an indifferent consumer will not wait if $\phi_i > \overline{\phi} = \Pr\Big(p_{t+1} < p\Big)\mathbb{E}\Big[\alpha(p - p_{t+1})\,|\,p_{t+1} < p\Big]$. This leads to the following proposition.

Proposition: With $\overline{\phi} = \Pr\Big(p_{t+1} < p\Big)\mathbb{E}\Big[\alpha(p - p_{t+1})\,|\,p_{t+1} < p\Big]$, 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 $\overline{p} > p$ such that $\varepsilon_0 = \beta - \alpha\overline{p} + \varepsilon_1$. 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 $\overline{p}$, then ε_0 is preferred and there is no gain. Hence, the expected gains from waiting are

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

The first term above is equal to $\overline{\phi}$, 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

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

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

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

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

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

Moving the expectation operator, the last line above equals

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

which can be simplified to $\Pr\Big(p_{t+1} < p\Big)\Pr\Big(p_{t+1} < \underline{p}\Big)\Big(\beta - \alpha p + \varepsilon_1 - \varepsilon_0\Big) \leq 0$, since $\beta - \alpha p + \varepsilon_1 - \varepsilon_0 < 0$ by assumption. Hence, waiting is not optimal for a consumer wishing to not buy today. ■

For consumers who would purchase today, the gains from waiting are equal to $\overline{\phi}$, but there is an additional cost if prices rise. Hence, waiting is not optimal. For consumers who would prefer not to buy, the expected gains of waiting are negative.

In monetary terms, $\overline{\phi}/\alpha = \Pr\Big(p_{t+1} < p\Big)\mathbb{E}\Big[(p - p_{t+1})\Big| p_{t+1} < p\Big]$ 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 remaining capacity is not reported to consumers. With these assumptions, I find the median and mean transaction costs to be \$2.14 and \$2.20, respectively. Recall that these costs are based on the most extreme case: the consumer who is indifferent between purchasing today or delaying the decision.

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