

NBER WORKING PAPER SERIES

THE WELFARE EFFECTS OF DYNAMIC PRICING:
EVIDENCE FROM AIRLINE MARKETS

Kevin R. Williams

Working Paper 28989
<http://www.nber.org/papers/w28989>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2021

This paper was previously titled "Dynamic Airline Pricing and Seat Availability." I thank Steve Berry, Judy Chevalier, Jim Dana, Drew DiPrinzio, Tom Holmes, Shih-Hsuan Hsu, Olivia Natan, Aniko Oery, 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. Finally, I thank the Yale School of Management for providing the computational resources used in this study. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Kevin R. Williams. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Welfare Effects of Dynamic Pricing: Evidence from Airline Markets

Kevin R. Williams

NBER Working Paper No. 28989

July 2021

JEL No. L11,L12,L93

ABSTRACT

Airfares fluctuate due to demand shocks and intertemporal variation in willingness to pay. I estimate a model of dynamic airline pricing accounting for both sources of price adjustments using novel flight-level data. I use the model estimates to evaluate the welfare effects of dynamic airline pricing. Relative to uniform pricing, dynamic pricing benefits early-arriving, leisure consumers at the expense of late-arriving, business travelers. Although dynamic pricing ensures seat availability for business travelers, these consumers are then charged higher prices. When aggregated over markets, welfare is higher under dynamic pricing than under uniform pricing. The directionality of the welfare effect at the market level depends on whether dynamic price adjustments are mainly driven by demand shocks or by changes in the overall demand elasticity.

Kevin R. Williams

Yale School of Management

165 Whitney Avenue

New Haven, CT 06520

and NBER

kevin.williams@yale.edu

1 Introduction

The airline industry is well known for employing complex intertemporal pricing strategies. In principle, dynamic pricing could be welfare-reducing or welfare-increasing, depending on the sources of airline price adjustments. Fare adjustments may arise in part because aggregate demand shocks change the opportunity cost of selling a seat. Airlines 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. These price adjustments are welfare-improving as they increase capacity utilization. However, fare adjustments may also reflect changes in the aggregate demand elasticity. If late shoppers are business travelers, airlines will raise prices over time to capture these consumers' high willingness to pay through intertemporal price discrimination. This allows airlines to extract more surplus, but it could also lower welfare if seats remain empty more frequently. Existing theoretical frameworks on the welfare effects of price discrimination, including Aguirre, Cowan, and Vickers (2010) and Bergemann, Brooks, and Morris (2015), do not consider sequential markets with limited capacity. However, these works establish that the welfare predictions are ambiguous—depending on demand elasticities and information structure—in the static setting, therefore, it is likely also true when considering dynamic prices. This suggests it is an empirical question whether dynamic airline pricing is on net welfare increasing.

In this paper, I estimate the welfare effects of dynamic pricing in the airline industry, and in doing so, examine the sources of price adjustments over time. I develop a dynamic pricing model that combines features of stochastic demand and revenue management models from operations research with estimation techniques widely used in empirical economics research. I estimate this model using novel data that track daily prices and seat availabilities for over 12,000 flights in US monopoly markets. With the model estimates, I disentangle key interactions between changes

in willingness to pay over time and changes in opportunity costs arising from stochastic demand and limited capacity. The results establish that if intertemporal price adjustments were not possible, the prospect of extracting surplus from late-arriving customers creates the incentive for airlines to save an inefficient number of seats and charge an inefficiently high price. Allowing fares to respond to changing aggregate price responsiveness and demand shock realizations expands output, but also results in a significant reallocation of capacity across time. Leisure consumers benefit, and business consumers are made significantly worse off.

I begin by describing airline pricing practices (Section 2) and documenting stylized facts from novel airline data (Section 3) to motivate my empirical approach. With a sample containing over 700,000 observations, I show that price adjustments are consistent with standard dynamic pricing models: fare increases are common after bookings, and fares stay constant, or even decline, in the absence of sales. In addition, the trajectory of fares is overwhelmingly positive. Fares typically double in the sixty days before departure and, regardless of bookings, tend to sharply increase close to the departure date. This is consistent with intertemporal price discrimination. I show that these pricing patterns also occur in competitive markets and when considering tickets of different qualities (the pricing of economy versus basic economy, and economy versus first class). This indicates that the forces and trade-offs explored in this paper are relevant for these important extensions. Finally, I show there is significant heterogeneity in dynamic pricing patterns across markets: the frequency of fare increases and decreases varies, as well as the presence and depth of advance purchase discount opportunities. This motivates an empirical design that includes route-specific parameters.

In Section 4, I develop a structural model that allows for aggregate demand uncertainty and a changing composition of arriving consumers over time. I do so by combining features of dynamic pricing and stochastic demand models commonly used in operations research, including Gallego and Van Ryzin (1994), Zhao

and Zheng (2000), Talluri and Van Ryzin (2004) and Su (2007), 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), where demand is comprised of "business" travelers and "leisure" travelers. Although I tailor the model to reflect 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- and day-of-week-dependent arrival rates. Consumers know their preferences and solve a static utility maximization problem. On the supply side, the monopolist 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 or until the plane is full.

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 (Vulcano, van Ryzin, and Chahr, 2010). The identification assumption is that preferences for flights evolve in the same predictable way, but demand shocks can vary (Section 5). 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 route specific, with flexibility to account for day of the week variation in demand. The estimates generally suggest a significant shift in arriving consumer types over time and that demand shocks

are a meaningful driver of the variation in sales (Section 6). Variation in demand across days of the week matches travel patterns documented with data provided by the Transportation Security Administration (TSA).

I use the model estimates to quantify the welfare effects of dynamic airline pricing and to examine the drivers of dynamic price adjustments (Section 7). I show that relative to uniform pricing, dynamic pricing expands output (by 2.7 percent), primarily through lower fares offered to leisure travelers. Dynamic pricing also ensures seat availability for business travelers; however, these consumers are then charged significantly higher fares. This latter effect is sufficiently strong that total consumer welfare is 6.3 percent lower under dynamic pricing compared to uniform pricing. Increased revenues more than offset this decline, and I estimate total welfare to be one percent higher under dynamic pricing compared to uniform pricing.

Dynamic pricing increases welfare in most—but not all—of the monopoly markets studied. I show that the directionality of the overall welfare effect depends on which sources of price adjustments drive revenues. Welfare declines under dynamic pricing when price changes are mainly in response to changes in willingness to pay and not in response to demand shocks. Intertemporal price discrimination explains the strong upward trajectory in prices and accounts for two thirds of the revenue gains of dynamic pricing over uniform pricing. The remaining one third comes from responses to demand shocks that occur greater than 21 days before departure, when aggregate price responsiveness is stable but overall demand uncertainty is at its highest. In such instances, price adjustments ensure seat availability for later arriving, price insensitive customers. If airlines did not react to demand shocks, price changes would occur one third as frequently in the markets studied.

1.1 Related Literature

This paper contributes to growing literatures in economics, marketing, and operations research that study intertemporal pricing dynamics. Intertemporal price discrimination can be found in many markets, including video games (Nair, 2007), Broadway theater (Leslie, 2004), storable goods (Hendel and Nevo, 2006, 2013; Gowrisankaran and Rysman, 2012), and concerts (Courty and Pagliero, 2012).¹ Importantly, this paper focuses on third degree intertemporal price discrimination resulting from time-varying arrivals of different consumer types, instead of second degree intertemporal price discrimination as a result of screening (Stokey, 1979; Bulow, 1982; Conlisk, Gerstner, and Sobel, 1984; Sobel, 1991; Su, 2007; Board and Skrzypacz, 2016; Öry, 2016; Gershkov, Moldovanu, and Strack, 2018; Dilmé and Li, 2019). This large theoretical literature focuses on forward-looking buyer behavior, but abstracts from a changing composition of arriving customers over time. McAfee and Te Velde (2006) argue that a change in the elasticity of demand is required to rationalize airfare pricing patterns.

The strong upward trajectory in prices observed in the airline setting greatly reduces buyers' incentive to wait to purchase. Imposing the assumptions that buyers will not wait to buy and arrivals are exogenous to price, however, can affect welfare estimates, via estimates of willingness to pay (Hendel and Nevo, 2006) or from sorting and the timing of market participation (Sweeting, 2012). Nair (2007) shows that abstracting from forward-looking consumers can lead to profit losses when demand becomes more elastic over time.

However, the incentive to wait to purchase decreases when capacity constraints are modeled, even in environments where prices typically fall (Soysal and Krishnamurthi, 2012).² Gale and Holmes (1993) and Dana (1998) show that firms may

¹Lambrecht et. al. (2012) provide an overview of empirical work on price discrimination more broadly.

²Aguirregabiria (1999) also considers a model with markdowns and studies how pricing varies with remaining inventory.

offer advance purchase discounts when consumers are heterogeneous and learn their preferences over time. Although Gale and Holmes show that learning implies increasing prices even without aggregate uncertainty, Dana (1998) emphasizes that aggregate uncertainty increases airlines' costs, and that even in competitive markets these costs are passed on to consumers who learn their preferences later. This causes prices to rise over time, which implies that once consumers know their preferences, they no longer have an incentive to wait to purchase.

This paper advances consumer demand models that have been extensively used in operations research to characterize optimal dynamic prices (Gallego and Van Ryzin, 1994; Zhao and Zheng, 2000; Talluri and Van Ryzin, 2004; McAfee and Te Velde, 2006).³ My approach differs from those used in operations research in that I empirically estimate preferences and arrivals; I extend baseline Poisson demand models to include time-varying arrivals, censored demand, and discrete random coefficients, where the share of consumer types changes over time.

This paper also complements recent studies on the airline industry, including Escobari (2012), Alderighi, Nicolini, and Piga (2015), and Puller, Sengupta, and Wiggins (2015). In closely related work, Lazarev (2013) estimates the welfare effects of intertemporal price discrimination in airline markets by modeling how changes in willingness to pay over time affect the firm's choice of the distribution of fares to offer, prior to the realization of demand shocks. In this project, I investigate the firm's responses to demand shocks over time (the "revenue management" problem), however, I abstract away from the set of fares chosen (see Section 2). Chen (2018) extends the methodology presented here to investigate competitive dynamics. Aryal, Murry, and Williams (2018) utilize survey data to examine dynamic pricing in international airline markets where seats have different qualities.

Finally, concurrent works provide new insights on the effects of dynamic pricing in other industries. Cho et. al. (2018) quantify the gains from dynamic pricing in the

³Elmaghraby and Keskinocak (2003) and Talluri and Van Ryzin (2005) provide an overview of revenue management work in operations.

hotel industry. They also capture competitive pricing pressures. D’Haultfoeuille 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 revenue management results in significant gains relative to uniform pricing.

2 Motivating Facts and Industry Pricing Practices

In this section, I provide a short overview on airline pricing practices to motivate my empirical approach. Additional details on airline revenue management algorithms and practices can be found in McGill and Van Ryzin (1999) and Gallego and Topaloglu (2019).

Flight prices depend on three key inputs: (1) plane capacity, (2) filed fares, and (3) inventory allocation for filed fares, or revenue management. Filed fares (input 2) are the pre-set price levels at which the airline is willing to sell tickets for a flight, and inventory allocation (input 3) is the number of tickets allocated to each fare level. Each of these decisions is made by separate departments, holding the other departments’ choices fixed. This paper focuses on modeling dynamic prices arising from (3).

A carrier’s network-planning department determines which markets are served, assigns capacity, and flight frequencies. These decisions typically occur well in advance of the departure date. Exceptions include entry or exit decisions or a change in size of aircraft. While I do observe aircraft substitutions in the collected data, they are not correlated with flight loads.⁴ It is more likely that these gauge adjustments occur for operational reasons. This motivates my assumption that initial

⁴I observe that 3.0% of flights experience a change in aircraft in the sixty days before departure. 79% of occurrences happen within the two days before departure. Yet these changes do not seem to be associated with flight loads. I cannot reject the null hypothesis that flights which see an upgauge (increase in capacity) have flight loads higher than the average load factor for that route and vice versa. In the former case, $p = 0.999$; in the latter case, $p = 0.197$. Flights which see an upgauge actually have lower load factors than the route average.

capacity is exogenous.⁵

The pricing department determines filed fares, or prices and associated ticket restrictions, that consumers may face. Examples include refundable and non-refundable tickets, as well as first-class, economy-class, and basic-economy tickets. A common ticket restriction is an advance-purchase (AP) requirement, or a restriction that requires consumers to purchase by a deadline. AP requirements are commonly observed three, seven, ten, 14, 31, and 30 days before departure. 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 what is called a fare basis code (the fare class, price, and restrictions).

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-class fares filed, one saver economy fare and one economy fare, each with a 14-day AP requirement. The two L-class fares have different fare basis codes. The third fare is a 14-day AP Q-class fare, and the fourth is an unrestricted Y-class economy fare. I incorporate this feature in the empirical model by having firms choose among a discrete set of fares. Lazarev (2013) models this distribution of fares chosen by airlines. Because I do not observe airline ticketing data, I assume consumers only consider the lowest available economy fare.

Finally, the revenue management department dynamically determines fare availability, among the fare classes set by the pricing department. This process involves setting the number of seats available for purchase for each fare class over time. Allocations are determined using techniques developed in operations re-

⁵I explore the effects of initial capacity on dynamic pricing in Online Appendix E.

search, including the well-known ESMR-b heuristic, in order to make them tractable (Belobaba, 1987, 1989, 1992; Belobaba and Weatherford, 1996). Phillips (2005) provides an overview of these approaches. Importantly, the allocation decision takes fares and forecasts as inputs, which are also the inputs I consider in my model. Although I do not model inventory allocations explicitly, I note that the average number of seats booked per day is less than one. This means customers are unlikely to face intra-day price dispersion due to fare classes closing.⁶

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.⁷ I collect and merge together three pieces of information. First, I collect daily prices for thousands of flight itineraries. I focus on one-way fares, since 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. Second, I collect censored fare class allocations for each flight. These show available fare classes and can be used to determine when flights are sold out.⁸ Third, I collect airline seat maps, which are graphical representations of available and occupied seats. By tracking changes to individual seats across consecutive days, I obtain a measure of daily bookings. I show in Online Appendix C that the measurement error in using seat maps to proxy bookings may be small.

⁶Many RM systems are designed such that several fares are available at any given point in time, which is called nesting (Phillips, 2005). Without access to individual-level purchase data and inventory allocations (these data are proprietary and available data are censored), I cannot pursue modeling inventory allocation.

⁷The 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.

⁸For example, G5 means the active G-class fare has five available seats, however, airlines censor these data at seven or none, depending on the carrier.

In the following subsections, I discuss route selection (Section 3.1) and document a set of new descriptive facts on dynamic pricing in the airline industry (Section 3.2).

3.1 Route Selection

I use the publicly available the Department of Transportation DB1B tables to select markets to study.⁹ I define a market in the DB1B as an origin-destination (OD), quarter, year. With the DB1B data, I filter based on the following criteria:

- (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 number make up roughly 14 percent of OD traffic in the United States. In addition, quarterly revenues for these markets are roughly \$2.3 billion. Criterion (v) addresses the potential for alternative flight options, including connecting flights for OD. Criterion (vi) is equally important because it addresses how fares are assigned to observed changes in remaining capacity. For this analysis, it is important to find markets that have relatively homogeneous itineraries, primarily non-connecting, nonstop trips.

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

⁹The DB1B tables contain a 10-percent sample of domestic US ticket purchases. The data are at the quarterly level and contain neither the date flown nor the purchase date.

flights to hubs, but overwhelmingly, most consumers on these flights connect to other destinations. Without individual-level data, it is impossible to know the itinerary 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 which satisfy the selection criteria above. In addition, to compare the descriptive evidence, I select six duopoly markets.¹⁰ Online 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. Most of the sample covers markets served by Alaska Air Lines (JetBlue and Delta are the other carriers studied).

Several features of the sample are worth noting. First, Alaska and JetBlue price itineraries at the segment level; that is, consumers wishing to purchase round-trip tickets on these carriers 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 a fare 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.¹¹ I will use this feature of the data to simplify the pricing problem presented in the next section. Third, several selected markets offer 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

¹⁰Two markets, (Boston, MA - Kansas City, MO) and (Boston, MA - Seattle, WA) were both monopoly and duopoly markets. The former market originally had nonstop service offered by Delta and Frontier. Frontier exited early on in the sample and Delta became the only carrier flying nonstop. The latter market was very briefly served by just Alaska, prior to the entry of JetBlue.

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

on airlines that allow consumers to select seats before departure. Many carriers now charge fees to choose seats when traveling on restrictive coach tickets.¹²

In contrast with Jetblue, Alaska and Delta offer first class in several of the markets studied—first class appears in 58 percent of the sample, with the average cabin size being twelve seats of the plane. I provide some descriptive analysis of first-class pricing (see Online Appendix A), but I do not pursue 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 (overselling).¹³

3.2 Descriptive Evidence

3.2.1 Summary Statistics

The sample contains over 12,000 flights, each tracked for the last sixty days before departure. The sample contains 738,625 observations, as well as over five million connecting fares. Data collection occurred over two six-month periods (March 2012-August 2012, March 2019-August 2019).

Summary statistics appear in Table 1. The average one-way ticket price in the sample is \$233.¹⁴ Load factor is the number of occupied seats divided by capacity on the day of departure. Average load factor is 89 percent, ranging from 70 percent to 98 percent, by market. I observe that 15.7 percent of flights sell out. There is considerable variation in load factor within a market, which shows evidence for 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

¹²The JetBlue data were collected before the introduction of Blue Basic seats, which feature a fee to select seats. This is also true for Delta. Alaska's 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.

¹³Source: Air Travel Consumer Report, accessed February 2020.

¹⁴This is higher than the average price calculated from the publicly available DB1B tables (Table 8); however, recall that these gathered prices cover the sixty days before departure and also include non-transacted prices.

Table 1: Summary Statistics for the Data Sample

Variable	Mean	Std. Dev.	Median	5th Pctile.	95th Pctile.
Oneway Fare (\$)	232.60	139.33	190.18	89.00	504.00
Load Factor	88.76	13.52	93.42	59.21	100.00
Daily Booking Rate	0.68	1.94	-0.00	-0.00	4.00
Daily Fare Change (\$)	3.43	31.25	0.00	0.00	46.00
Unique Fares (per itin.)	6.97	2.16	7.00	4.00	11.00

Note: Summary statistics for 12,119 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 738,625. Load Factor is reported between zero and 100 the day of departure. The daily booking rate and daily fare change compares consecutive days.

to fill unsold seats. The R^2 of a regression of load factor on origin-destination-flight number and departure date (subsuming seasonality and day-of-week indicators) fixed effects is only 0.56, which suggests that demand shocks may be significant in explaining flight loads.

The booking rate in Table 1 corresponds to the mean difference in occupied seats across consecutive days. The average booking rate is 0.68, with the 5th and 95th percentiles of zero and four seats per flight, respectively. This finding shows that airline markets are associated with low daily demand. 61% seat maps 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 10.4 fare changes. This implies that fares fluctuate up and down, usually a few times, and that the number of realized prices is relatively small. For the markets studied, the median number of daily departures is one and the mean is two. Finally, I use individual seat assignments to estimate the number of passengers per booking.¹⁵ I

¹⁵Each 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 eight 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.

estimate this to be 1.37, which motivates the unit demand assumption in the model.

There are a few differences across the data samples worth noting. Relative to the data collected in 2012, the 2019 data contain substantially lower fares (-\$170); the booking rate and daily price increase are slightly lower (-0.2; -\$0.3); slightly fewer fares are offered per flight (-0.5); and load factors are lower (-7%). These differences cannot be attributed to a single factor, as the carriers and markets differ across the samples. In Section 7, I highlight how the welfare estimates vary across markets.

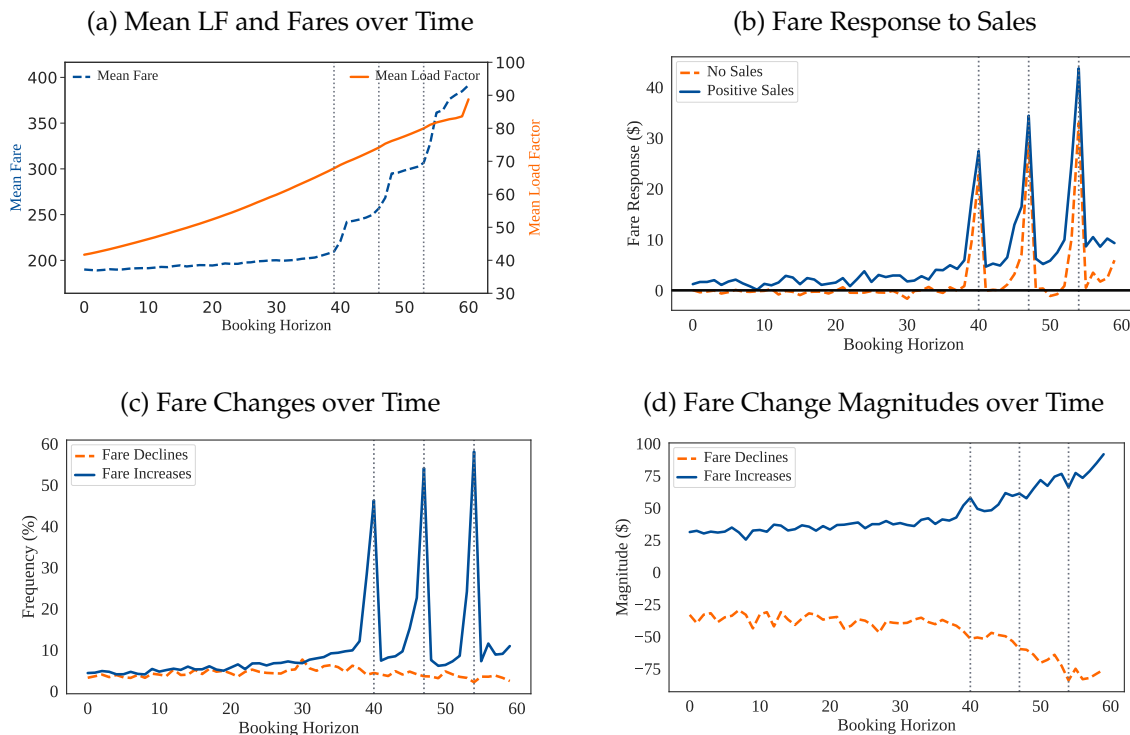
3.2.2 Dynamic Prices

Figure 1 documents patterns on dynamic price adjustments in airline markets. Panel (a) plots the mean fare and mean load factor by day before departure. The plot confirms that the overall trend in prices is positive, with fares increasing from roughly \$200 to nearly \$400 in sixty days. The noticeable jumps in the time series occur when crossing advance purchase (AP) restrictions. At sixty days before departure, roughly 42 percent of seats are already occupied. The booking curve for flights in the sample is smooth over 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. On the last day before departure, there is also a sharp decline in available economy tickets, which suggests that last-minute bookings do occur (see Online Appendix A).

There is considerable variation in pricing across markets. Figure 7 and Figure 8 in Appendix A plot average fares over time as well as the average percentage change in fares over time for each route studied. Price levels, the timing of AP restrictions, and the depth of AP discounts vary by market.

Figure 1-(b) establishes an important link between bookings and price adjust-

Figure 1: Average Fares, Load Factors, and Fare Response to Sales



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

ments. 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. First, fares respond to demand shocks as predicted by standard dynamic pricing models: When bookings occur, prices tend to rise; when bookings do not occur, prices stay the same, or fall. However, close to the departure date and regardless of bookings, prices increase. This suggests late-arriving consumers are less price-sensitive and airlines capture their high willingness to pay with intertemporal price discrimination.¹⁶

¹⁶This 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 suggests later arrivals are less price sensitive.

Figure 1-(c) shows the frequency of fare increases and decreases over time. Well before the departure date, the number of fare hikes and the number of fare declines are roughly even. The spikes correspond to advance purchase requirements. Note their use is not universal. Less than 60 percent of flights experience a price increase at the 7-day AP requirement. Finally, Figure 1-(d) shows the magnitude of fare adjustments over time; conditional on fares changing, the magnitude of the change is increasing over time.

The pricing patterns documented here also occur in competitive markets and when considering tickets of different qualities, i.e., first class and basic economy). I highlight two findings.¹⁷ First, for the few competitive markets in the sample,, the magnitude of systematic fare increases is lower, and the number of systematic fare decreases is higher. This may suggest the role of intertemporal price discrimination is reduced in competitive markets.¹⁸ Second, all ticket qualities respond to AP restrictions, the gap between economy and basic economy grows over time, and the availability of basic economy fares decreases. Therefore, economy cabin fares rise over time for two reasons: regular economy fares become more expensive and basic economy fares are no longer offered.

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,

¹⁷Additional analysis can be found in Figure 9, Figure 10, and Figure 11 in Online Appendix A.

¹⁸This finding complements the work of Siegert and Ulbricht (2020), who use fare data to show that competition is correlated with a flattening of prices over time. Dana and Williams (2021) show in a theoretical model 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 compared to the markets studied in this paper.

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 $t \in \{0, \dots, T\}$. Period 0 corresponds to the first sales period, and period T corresponds to the flight departure date. 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 scrapped 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 type; they do not identify consumers based on a travel need.¹⁹ Upon entering the market, all uncertainty about travel preferences is resolved.²⁰ Arriving consumers 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. Consumers do not cancel seats so remaining capacity is monotonically decreasing.²¹ This process repeats each sales period; the firm is forward looking and solves the finite horizon, dynamic program.

¹⁹Booking 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.

²⁰This 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. This assumption is motivated by the fact that I do not find significant bunching in bookings before the expiration of AP fares (see Online Appendix D).

²¹The average number of cancellations per flight in the data is less than two.

4.2 Demand

Each day before the flight leaves, $t = 0, 1, \dots, T$, 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 $y_t = (\alpha_i, \beta_i, \varepsilon_{ijt}, \varepsilon_{i0t})_{i \in 1, \dots, \tilde{M}_t}$ to be the vector of consumer preferences. Suppressing the notation on product characteristics for the rest of this section, demand for flight j at t is defined as $Q_{jt}(p, y_t) := \sum_{i=0}^{\tilde{M}_t} 1[u_{ijt} \geq u_{i0t}] \in \{0, \dots, \tilde{M}_t\}$, 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, consumers arrive and choose to fly or not. The capacity constraint is then checked. If demand exceeds remaining capacity, c_{jt} consumers are randomly selected from the set of consumers who chose to travel, and the rest receive their outside options. Although this assumption may appear restrictive, the daily booking rate is less than one.²²

²²Although 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

Without the ability to oversell and incorporating the rationing rule, expected sales are formed by integrating over the distribution of y_t ,

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

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

$$\pi_{jt}^i(p_{jt}) = \frac{\exp(x_{jt}\beta_i - \alpha_i p_{jt})}{1 + \exp(x_{jt}\beta_i - \alpha_i p_{jt})}.$$

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, $\gamma_t \pi_{jt}^B$ defines the purchase probability that a consumer is of the business type and wants to purchase a ticket; $(1 - \gamma_t) \pi_{jt}^L$ is similarly defined. Hence, integrating over consumer types, product shares is equal to $\pi_{jt}(p_{jt}) = \gamma_t \pi_{jt}^B(p_{jt}) + (1 - \gamma_t) \pi_{jt}^L(p_{jt})$. Next, I assume that consumers arrive according to a Poisson distribution, $\tilde{M}_t \sim \text{Poisson}_t(\mu_t)$. 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, it follows that demand is also Poisson, $Q_{jt} \sim \text{Poisson}_t(\mu_t \pi_{jt})$. The probability that q seats are demanded on flight j at time t are equal to

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

With these probabilities defined and noting that demand is censored at remain-

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.

ing capacity, expected sales is equal to²³

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

4.3 Monopoly Pricing Problem

The monopolist maximizes expected revenues of flight j (subscript suppressed) over a series of sequential markets. Each day before departure, the firm chooses to offer a single price before the arrival of customers. Using the institutional features discussed in Section 2, I assume the firm chooses a price from a discrete set, denoted $A(t)$. The set may change over time due to advance purchase restrictions.²⁴

The pricing decision is based on the states of the flight: seats remaining; time left to sell; flight characteristics; 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 $\sigma > 0$. These shocks are assumed to be additively separable to the remainder of the per-period payoff function, which are expected revenues, $R_t^e(p_t; c_t) = p_t \cdot Q_t^e(p_t; c_t)$.

The firm's problem can be written as a dynamic discrete choice model. Let

²³This can be equivalently written as

$$Q_{jt}^e(p_{jt}; c_{jt}) = \sum_{q=0}^{c_{jt}-1} \frac{(\mu_t \pi_{jt})^q \exp(-\mu_t \pi_{jt})}{q!} q + \left(1 - \sum_{q=0}^{c_{jt}-1} \frac{(\mu_t \pi_{jt})^q \exp(-\mu_t \pi_{jt})}{q!} \right) c_{jt}.$$

²⁴In 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.

$V_t(c_t, \omega_t)$ 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

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

Because the firm cannot oversell, capacity transitions as $c_{t+1} = c_t - \min\{Q_t, c_t\}$, 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 $h_t(\omega_{t+1}, c_{t+1} | \omega_t, p_t, c_t) = g(\omega_{t+1})f_t(c_{t+1} | p_t, c_t)$. 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 conditional value function is equal to

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

where ϕ is Euler's constant. The conditional choice probabilities also have a closed form and are computed as

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

Before continuing, I discuss the connection between the notation $\Pr_t(Q_{jt} = q; p_{jt})$ and $f_t(c_{t+1} | c_t, p_t)$. Consider a two-period model with a single seat. In the first period, expected revenues are simply $\Pr_1(Q_1 \geq 1; p_1) \cdot 1 \cdot p_1$ because at most one seat can be sold. The demand probabilities exactly inform the capacity transition probabilities under conditional independence, that is, $f_1(c_2 | 1, p_1) =$

$[\Pr_1(Q_1 \geq 1; p_1), \Pr_1(Q_1 = 0; p_1)]$. With probability $\Pr_1(Q_1 \geq 1; p_1)$, the seat sells today and nothing is available for sale tomorrow, and with probability $\Pr_1(Q_1 = 0; p_1)$, the seat is not sold today and is available for purchase tomorrow. The optimal price that affects these probabilities depends on the arrival process and product shares. Time is a deterministic state. Note, in the general model, any transition probability where $c_{t+1} > c_t$ is equal to zero because capacity is monotonically decreasing.

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 that 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.

Another potential limitation of the model is that consumers are assumed to make a one-shot decision upon entering the market, and market participation is exogenous. This can impact estimated demand elasticities (Hendel and Nevo, 2006). If increasing prices are also used to shape consumer expectations, my estimates may overstate the proportion of business travelers and understate their price sensitivity. In addition, if consumers learn about their preferences toward the deadline, this will cause opportunity costs to rise over time (Dana, 1998), which may act to reinforce this potential overstatement.

5 Estimation

I assign the discount factor to be one. Arrival rates, μ_t^d , vary by day before departure (t) and departure date (d) in the following way. Over the booking horizon, I let arrival rates vary corresponding to observed advance-purchase discount intervals, which are then scaled according to the day-of-week of the departure date,

$$\mu_t^d = \begin{cases} \mu^{\text{dow}} \cdot \mu_1, & \text{Greater than twenty-one days before departure (21+);} \\ \mu^{\text{dow}} \cdot \mu_2, & \text{Fourteen to twenty-one days before departure (20-14);} \\ \mu^{\text{dow}} \cdot \mu_3, & \text{Seven to fourteen days before departure (13-7); and} \\ \mu^{\text{dow}} \cdot \mu_4, & \text{Within seven days before departure (6-0).} \end{cases}$$

Here, μ^{dow} is a day-of-week shifter for each departure date. Mondays are normalized to one, and parameters are estimated for Tuesday through Sunday. In total, there are ten arrival rate parameters per market.

I introduce flexibility in the composition of consumer types by assuming

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

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

Finally, I assume consumer utility is of the form

$$u_{ijt} = \beta^{\text{dow}_j} - \alpha_i p_{jt} + \varepsilon_{ijt},$$

where β^{dow_j} is a day-of-week preference for the departure date.²⁵ There are nine preference parameters per market.

To reduce computational burden, I construct a single pricing menu for each

²⁵This control explains more of the variation in bookings than day-of-week of purchase.

route by reducing the dimensionality of observed prices. The average number of unique fares observed per flight is less than seven, however, I observe price differences across departure dates within a route, sometimes by a single dollar. To avoid constructing likelihoods for each flight individually, I first cluster all observed prices for a given route using k-means with a minimum in-sample fit threshold of 99 percent. This results in pricing choice sets that range in size from five to eleven. I then map each observed fare to its clustered fare, creating pricing menus that only vary by route and day before departure. Because lower-priced fares are typically not offered close to the departure date, this procedure preserves advance-purchase discounts, albeit with clustered fares.²⁶

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

$$\max_{(\beta, \alpha, \gamma, \mu, \sigma)} \sum_F \sum_T \log(\text{CCP}_t(p_t; c_t)) + \log(f_t(c_{t+1}|c_t, p_t)).$$

I maximize this objective separately for each market. To increase sample sizes, I group together the directional traffic of the city pairs, which means demand does not vary by direction. Online Appendix B shows that directional prices are very similar. For any candidate parameter vector, I calculate the censored-Poisson demand functions, expected revenues, and transition probabilities. I then solve for the value functions using the recursive structure of the firm's problem, which defines the conditional choice probabilities (CCP).²⁷

²⁶Other 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.

²⁷Estimation utilizes analytical gradients computed via the module JAX using GPUs (set to 64-bit) and the solver Knitro. I select the Sequential Quadratic Programming (SQP) algorithm. I first use parallel multi-start, selecting 200 random initial starting values, using relaxed parameter bounds. I then rerun the estimation script with tighter bounds centered around the first solution, using 42 random random starts.

5.1 Identification

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 proprietary search data to pin down the arrival process, an increase in arrivals could instead be inferred as inelastic demand. For example, the sale of two seats could have occurred because two consumers arrived and both purchased, or because twenty consumers arrived and a tenth purchased. This is sometimes called the lack of "no purchase" data.

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 1-(b) 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 (orange) line is at or below zero. However, Figure 1-(b) shows that prices sharply rise close to the departure date and regardless of bookings. This sharp rise in prices, regardless of the scarcity of seats, suggests a change in demand elasticity. That is, consumers who shop late are less price sensitive than those who shop early.

6 Empirical Results

Complete parameter estimates span three tables, Table 4–Table 6 in the appendix.²⁸ Each table reports results for a set of markets and has three sections. The first section, "Logit Demand," reports day-of-week preferences, price sensitivities, and the parameters governing the probability on consumer types over time (γ_t). The second section, "Poisson Rates", reports mean arrival rates for each of the specified time intervals for Monday departures. The rows labeled "DoW Effect" contain the multiplicative factor for Tuesday through Sunday departures. Finally, the last row, "Firm Shock", reports estimates of the scaling parameter. A summary of the demand estimates is shown in Table 2 below.

Almost all consumer preferences are all found to be statistically significant at conventional levels.²⁹ The parameter estimates suggest that, on average, leisure consumers are over twice as price sensitive as business consumers, and business consumers are willing to pay up to 125% more in order to secure a seat.³⁰ I estimate meaningful differences in demand across departure dates due to day-of-week effects. In Figure 2-(a), I plot the average willingness to pay for the days of the week, relative to the minimum estimated day-of-week preference. The histogram is over markets. I estimate that willingness to pay is highest for flights departing on Sunday, Friday, Thursday, Monday (in that order; highest to lowest). Saturday, Tuesday, and Wednesday are estimated to be the most off-peak days. These values closely match day-of-week patterns found using security checkpoint data from the Transportation Security Administration (TSA).³¹

²⁸I do not estimate demand in competitive routes or routes with infrequent service. The excluded routes are: Boston, MA - Seattle, WA; Boston, MA - Portland, OR; Portland, OR - Sacramento, CA; Portland, OR - Lihue, HI; and Portland, OR - Palm Springs, CA. In addition, Omaha, NE - Seattle, WA is excluded from the analysis due to numerical stability issues and resource constraints.

²⁹The exception being Oklahoma City, OK - Seattle, WA. All random starts converge to the same maximum; however, several parameters are estimated to be insignificant.

³⁰The mean ratio of price sensitivity across markets is 3.34; the median is 2.25.

³¹In 2019, the busiest to least busy travel days in the United States were Friday (2.44 mil.), Sunday (2.38 mil.), Thursday (2.37 mil.), Monday (2.36 mil.), Wednesday (2.15 mil.),

Table 2: Demand Results Summary Table

Parameter	Mean	Std. Dev.	Median	25th Pctile.	75th Pctile.
DoW Preferences	5.77	4.67	4.72	3.27	6.99
Leisure Price Sensitivity	-3.48	3.46	-2.37	-3.57	-1.71
Business Price Sensitivity	-1.61	1.72	-1.24	-1.72	-0.74
Prob(Business)	0.26	0.28	0.15	0.03	0.38
DoW Arrival Rates	2.02	2.36	1.35	0.93	2.01

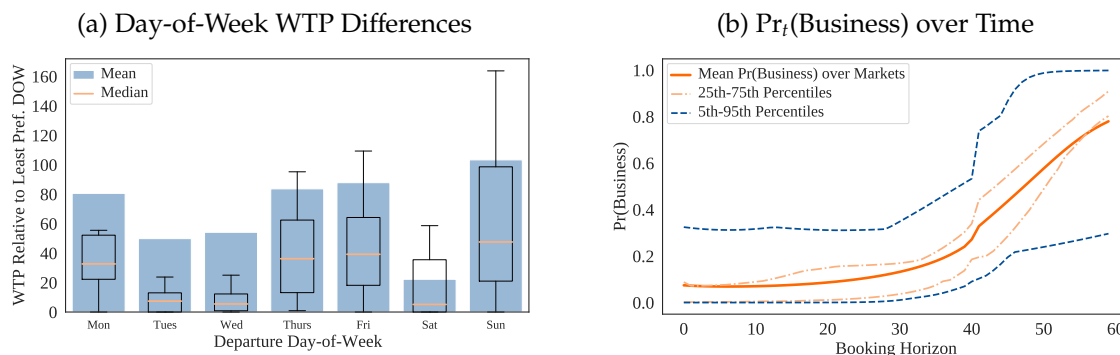
Note: Summary of demand estimates. See Table 4–Table 6 for all parameter estimates. DoW preference statistics are computed using all β_r^d parameters. Leisure and Business price sensitivity statistics are computed using all $\alpha_{L,r}$ and $\alpha_{B,r}$ parameters. Probability of business uses the predicted values of the Logit specification at the $\gamma_{t,r}$ level. DoW Arrival Rates are computed using all $\mu_{t,d}^d$ parameters.

Fitted values of the probability that a customer is of the business type (γ_t) is shown in Figure 2-(b). The plots depict the average (across routes) business share over time, as well as the interquartile range and the fifth and ninety-fifth percentiles over markets. Most routes exhibit increasing γ_t processes over time. On average, 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 21 days before departure, I estimate a significant change in the business customer share. This corresponds with the time at which fares start raising rapidly.

There is substantial heterogeneity in the fitted values for γ_t across routes, as shown in Figure 12 in Online Appendix A. The heterogeneity reflects the differences in pricing dynamics across markets (see Figure 7 and Figure 8 in Online Appendix A). In general, the shape of the curves correlates with the use of AP restrictions: a larger price increase at the 21-day AP requirement generally creates a steeper profile. The share of business arrivals well before the departure date is typically between 0 and 20 percent and increases to between 60 to 80 percent the day of

Tuesday (2.06 mil.), and Saturday (2.01 mil.), respectively (daily average). Compiled from <https://www.tsa.gov/foia/readingroom>.

Figure 2: Day-of-week Preferences and Consumer Types over Time



Note: (a) Average willingness to pay for the days of the week, relative to the minimum estimated day-of-week effect for each market. The plot shows an average over markets. (b) Fitted values of the arrival process of business versus leisure customers across the booking horizon. The y-axis is $Pr(\text{business})$, so $1 - Pr_t(\text{Business})$ defines $Pr_t(\text{Leisure})$.

departure.

The parametric assumption on consumer types is flexible, as it captures S-shape, almost linear, and convex arrival paths. It can also be restrictive. One market is estimated to shift from one Poisson demand distribution to another (leisure to business) corresponding to the 21-day AP requirement.

All arrival rates are estimated to be statistically significant.³² There are three levels of heterogeneity in these estimates. First, across markets, the average number of arrivals ranges from around one to up to ten. Second, in some markets, the arrival rates rise over time, whereas in most of the estimates, the rates stay low. Finally, there is variation in which days of the week experience the largest market sizes across routes. Monday and Sunday are estimated to have the largest market sizes in forty percent of markets, followed by Thursday and Friday. I estimate that 24.6 percent of arrivals are business travelers. As a point of comparison, Lazarev (2013) estimates 20 percent of consumers are business travelers.

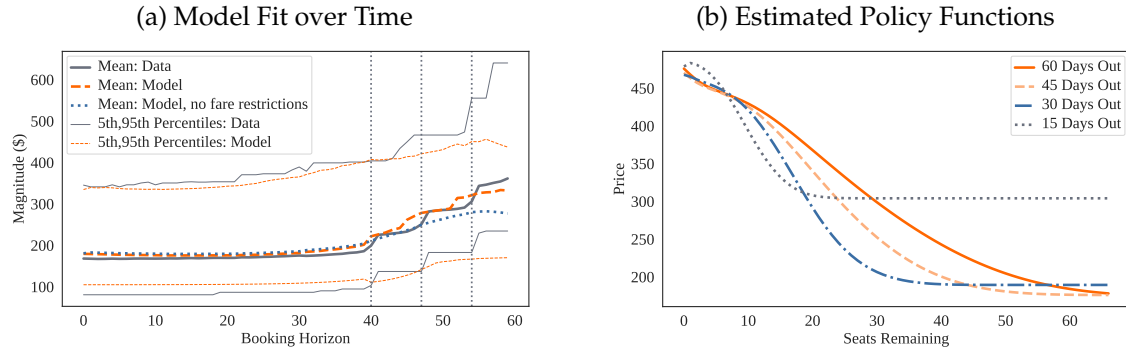
Overall, the demand estimates establish that a meaningful shift occurs in willingness to pay over time. Demand elasticities range from -9.38 to -1.46, depending

³²The exception being Oklahoma City, OK - Seattle, WA, where both the DoW preferences and blocked arrival rates are found to be insignificant.

on route and time until departure. I estimate the average price elasticity to be -3.31.

6.1 Model Fit and Discussion

Figure 3: Model Fit and Optimal Pricing



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 3-(a) 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—the firm has 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 \$7.50. Differences do vary by day before departure—they are less than \$11 for the first half of the sample but the gap increases around AP requirements. The reason is that the model produces a smoother fare profile that results in fare hikes slightly before the 14 and 7-day AP requirements. 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, where the unrestricted model assigns lower prices. 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.

Price adjustments occur because of the time-varying composition of customers and in response to demand shocks. Figure 3-(b) highlights how remaining capacity affects pricing within a period for a sample route. Plotted are the firm's policy functions. Each line corresponds to a different time period. With fewer seats remaining (moving toward the origin on the x-axis), fares increase. The plot also demonstrates for a given amount of seats remaining, opportunity costs are increasing in time left to sell. For example, the price of having forty seats remaining sixty days out is higher than forty seats remaining thirty days out. However, close to departure, fares are higher regardless of remaining capacity due to demand being more inelastic. Also, consistent with Dana (1998), aggregate demand uncertainty results in unused capacity that raises opportunity costs over time.

7 The Welfare Effects of Dynamic Airline Pricing

In this section, I estimate the welfare effects of dynamic pricing through a series of counterfactual exercises. In Section 7.1, I study uniform pricing, where the firm is not able to respond to demand shocks nor changes in the demand elasticity, and intertemporal price discrimination (IPD), where the firm is not able to respond to demand shocks. In Section 7.2, I examine the the sources of price adjustments and show how both forces explain airline pricing patterns.

To set up all counterfactuals, I use the empirical distribution of remaining capacity sixty days before departure as the initial capacity condition.³³ All coun-

³³Note 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. This is explored further in Online Appendix E, where I consider dynamic pricing under a large range of initial capacities. I do not bound the fixed costs, but I do show that capacity would need to be significantly higher for demand shocks to not affect airline prices.

terfactuals utilize the important boundary conditions of the initial problem: (1) the firm cannot oversell, and capacity transitions as $c_{t+1} = c_t - \min\{Q_t(p, y_t), c_t\}$; (2) unused capacity is scrapped with zero value. I simulate 100,000 flights per market using the distribution of initial observed capacities. I then combine the results over markets. Route-level heterogeneity is then explored.

For all counterfactual analysis, I make two changes to the estimated model. First, I allow firms to use the unrestricted choice set, $A(t) = \cup_{t=0}^T A(t)$, in each period, in order to streamline the counterfactuals, e.g., under uniform pricing, the firm may wish to charge a low fare that is not available close to departure. Second, I remove the firm shocks (ω) in order to single out the effects of the demand elasticity and scarcity (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 counterfactual, the firm receives per-period error shocks.

7.1 Uniform Pricing and a Model of Intertemporal Price Discrimination

With uniform pricing, the firm sets a single price for each flight by integrating over future demands in the initial period. The revenue maximization problem under uniform pricing is

$$\max_p \mathbb{E}_y \left[\sum_{t=0}^T p \min \{Q_t(p, y_t), c_t\} \right]$$

such that $c_{t+1} = c_t - \min \{Q_t(p, y_t), c_t\}$, c_0 given.

With a constant price, the firm cannot respond to both demand shocks and changes in willingness to pay.

In the model of intertemporal price discrimination, I assume the firm sets a sequence of prices before sales begin, at $t = 0$. Price changes over time reflect changes in willingness to pay—intertemporal price discrimination as third degree price dis-

crimination (as opposed to screening with second degree price discrimination). The revenue maximization problem is therefore

$$\max_{p_0, \dots, p_T} \mathbb{E}_y \left[\sum_{t=0}^T p_t \min \{Q_t(p_t, y_t), c_t\} \right]$$

such that $c_{t+1} = c_t - \min \{Q_t(p_t, y_t), c_t\}$, c_0 given.

Since prices cannot depend on the remaining capacity, they cannot react to changes in the opportunity cost of a seat.³⁴

The counterfactuals are nested such that as the pricing strategy becomes more flexible, expected revenues are necessarily increasing. This is because under dynamic pricing, prices are defined by $p^*(c_t, t)$, whereas in the model of intertemporal price discrimination, prices are time-dependent, $p^*(t)$. Finally, under uniform pricing, prices do not vary with both seats and time remaining, p^* . Therefore, expected revenues are increasing in pricing flexibility, that is,

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

Note that if capacity were sufficiently large, then the outcomes of the IPD and dynamic pricing models would coincide. The extent to which they differ suggests that responding to demand shocks is particularly important in the airline context.³⁵

Note that aggregate demand uncertainty affects prices in all scenarios, but in distinct ways. With dynamic pricing, prices are state-dependent and the firm reacts directly to demand shock realizations. However, with uniform pricing and in the model of IPD, the pricing decision reflects the integral of all future demands

³⁴Note that because demand becomes more inelastic over time, there is little to no role for Coasian forces (consumer waiting for fare declines). In Online Appendix D, I provide a bound on the waiting costs so that no consumer would choose to wait to buy under dynamic pricing.

³⁵See Online Appendix E for additional counterfactual analysis on initial capacities.

(for given prices) before any uncertainty is resolved. The magnitude of demand uncertainty affects both the overall price level, but also the incentive to set different prices over time, irrespective of changes in the demand elasticity. Therefore, in order to completely separate the effects of demand uncertainty changes in willingness to pay on intertemporal price adjustments, I consider an alternative model of intertemporal price discrimination in Section 7.2.

The firm’s objective function the model of IPD is 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 when the advance purchase requirements typically expire (days 3, 7, 14, and 21). This results in five prices per flight.

Table 3: Welfare Effects of Dynamic Airline Pricing

	<u>Fare</u>	<u>Load Factor</u>	<u>Sell Outs</u>	<u>Revenue</u>	<u>CS_L</u>	<u>CS_B</u>	<u>CS</u>	<u>Welfare</u>
Dynamic	243.3	87.6	18.7	10.7	2.2	6.5	8.7	19.5
IPD	243.6	83.6	22.2	10.4	2.1	6.4	8.4	18.9
Uniform	219.9	84.9	29.4	9.9	2.0	7.4	9.3	19.2

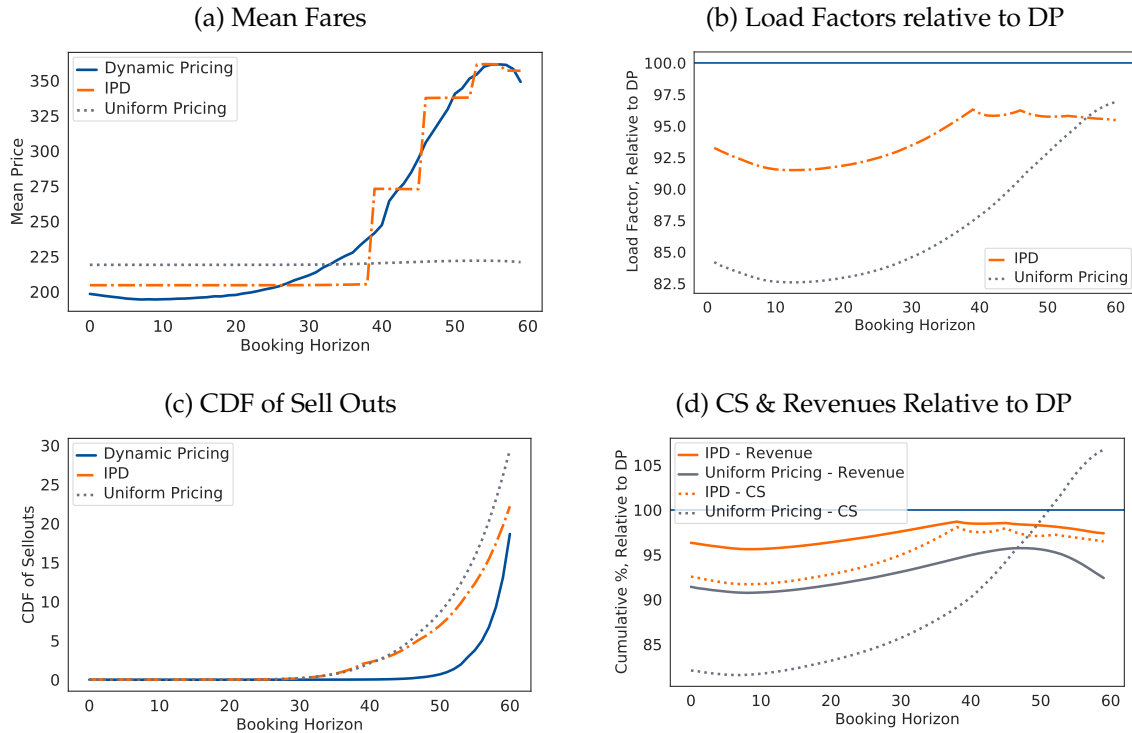
Fare: mean fare for flight observations with positive seats remaining; Load factor (LF): average at departure time; Sell Outs: percentage of flights with zero seats remaining in the last period; Revenue: mean flight revenue; Consumer surplus (CS_L, CS_B): surplus per flight; Welfare: daily mean revenues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

Welfare estimates for the baseline dynamic pricing model and the two counterfactuals are shown in Table 3. All values are in levels, except for load factor and sell outs, which are reported as percentages. A visual summary of the intertemporal dynamics are shown in Figure 4.

I find that average fares are over ten percent lower under uniform pricing compared to the other pricing models. However, this does not lead to an increase in output—load factors are 2.7 percent higher under dynamic pricing. This occurs

because although uniform pricing prohibits price discrimination, it creates the incentive for the firm to save an inefficient number of seats for later arrivals. As shown in Figure 4-(a), uniform pricing results in fares that are relatively high early on, but that are relatively low close to the departure date.

Figure 4: Counterfactual Results over Time



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 primary driver of market expansion under dynamic pricing is that leisure consumers are offered lower fares. This is also true in the model of IPD; however, in this model, the firm cannot adjust fares downward when flights are realized to have low demand. This is shown in Figure 4-(b), which depicts load factors relative to dynamic pricing over time. There are two notable findings. First, output remains highest under dynamic pricing throughout the entire booking horizon. Second, although the relative booking rate increases over time under both uniform

pricing and in the model of IPD (the lines move closer to 100%), both curves level off. This occurs because without the ability to respond to demand shocks, more flights sell out in advance. Figure 4-(c) highlights this result, which shows the cumulative distributions of sell outs for all scenarios. Uniform pricing results in not only more sell outs, but sell outs also occur much earlier than under dynamic pricing. Relative to dynamic pricing, I find sell outs are 3.5 percent higher under IPD, and 10.7 percent higher under uniform pricing.

Dynamic pricing leads to substantial changes in how capacity is allocated across consumer types (and time). In general, leisure consumers benefit under more flexible pricing systems (by construction, as does the firm). I estimate leisure consumer surplus would decline by 12.4 percent under uniform pricing. For business customers, there are two opposing forces. On one hand, dynamic pricing allows for increased price targeting, which lowers surplus relative to uniform pricing. On the other hand, dynamic pricing also acts to ensure seat availability, thus potentially allowing more business consumers to purchase. The model estimates show that the former negative effect on business consumer surplus dominates. Optimal uniform prices result in flights that are up to \$125 less expensive than under dynamic pricing, which more than offsets the increased probability that later arrivals may not be able to secure a seat. Consequently, business consumer surplus is 13.3 percent lower under dynamic pricing than under uniform pricing. Aggregating over consumer types, I estimate that dynamic pricing results in 6.3 percent lower total consumer surplus compared to uniform pricing.

Dynamic pricing leads to substantially higher revenues than those under uniform pricing (7.6 percent higher). This primarily comes from a transfer of business consumer surplus to the firm. Allocative efficiency is also improved. Output expands as the firm offers lower fares to early-arriving customers, and for peak flights, capacity is allocated to those with the highest willingness to pay. Because

the two counterfactuals are nested, I calculate

$$\frac{\text{Revenue under IPD} - \text{Revenue under Uniform Pricing}}{\text{Revenue under Dynamic Pricing} - \text{Revenue under Uniform Pricing}} \quad (7.1)$$


to measure the importance of intertemporal price discrimination versus the importance of responding to demand shocks in explaining the revenue increase. Using this decomposition, I find that 65.7 percent of the revenue gains associated with dynamic pricing over uniform pricing come from intertemporal price discrimination. The remaining 34.3 percent comes from the ability to respond to demand shocks. The increase in revenues is greater than the aggregate consumer surplus decline.

I estimate the overall welfare effect of dynamic pricing is a one percent increase in surplus compared to a world without price discrimination, with a stark reallocation of capacity across consumers. Dynamic pricing leads to a 7.2 percent increase in tickets purchased by leisure consumers and a 5.8 percent decrease in the number of tickets purchased by business travelers.

Dynamic airline pricing increases welfare in aggregate, but not for each market individually. Figure 5 graphically shows the welfare effects of dynamic pricing for each market separately; each dot denotes the total welfare of dynamic pricing over the welfare of uniform pricing on the vertical axis. On the horizontal axis, I plot the calculation in Equation 7.1. I find that dynamic pricing lowers welfare in seven of the markets studied and increases welfare in fifteen of the markets studied, thus highlighting that the welfare predictions of dynamic pricing are in general ambiguous.

As Figure 5 highlights, the directionality of the overall welfare effect depends on which sources of price adjustments drive revenues. Welfare declines under dynamic pricing when price changes are mainly in response to changes in willingness to pay and not in response to demand shocks. On the other hand, dynamic pricing improves overall welfare relative to uniform pricing only if demand is subject to many unpredictable shocks.

Figure 5: Welfare Effects of Dynamic Pricing due to IPD



Note: Each dot represents counterfactual results for a single market. The vertical axis is welfare under dynamic pricing over the welfare under uniform pricing. The horizontal axis computes the percentage of revenue gains from uniform pricing to dynamic pricing attributed to intertemporal price discrimination. Figure 13 in Online Appendix A presents an alternative figure with market labels and reports the frequency of sell outs for each market.

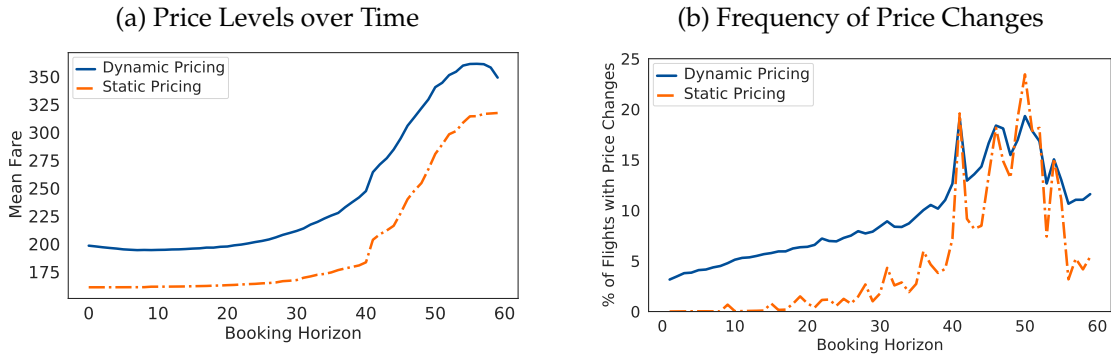
7.2 The Sources of Price Adjustments in Airline Markets

In this section, I examine the importance of the two reasons for price adjustments with dynamic pricing: changes in the willingness to pay and changes in the opportunity cost of a seat. I consider a model of static pricing where the cost of capacity is (close to) zero.³⁶ This implies that prices only change in response to changes in demand. The revenue maximization problem in this counterfactual is simply the baseline pricing model with the discount factor set equal to zero; that is, $\max_{p_t} \mathbb{E}_{y_t} [p_t \min \{Q_t(p_t, y_t), c_t\}]$.

The static pricing model results in substantially lower fares because there is no opportunity cost of capacity. Prices still rise significantly over time due to changes in willingness to pay, conditional on the firm having seats remaining. Figure 6-(a) shows that both models produce qualitatively similar patterns. Fares nearly double in sixty days and the slopes are similar. There is a level shift because the

³⁶Because the model is not in continuous time, more than a single seat may be demanded in a single period. This would raise prices (opportunity costs are positive); however, the estimated arrival rates are sufficiently low that this does not significantly impact the results. Results are very similar to a model where the firm does not take into account its capacity constraint when pricing.

Figure 6: Decomposition of the Source of Price Adjustments



Note: (a) Average fares over time under dynamic pricing and static pricing. (b) The percentage of flights that experience price changes from t to $t + 1$ under dynamic pricing and static pricing.

opportunity cost of capacity is close to zero under static pricing, lowering fares. This establishes that the primary source for increasing prices in airline markets is intertemporal price discrimination. Although average fares are nearly 30 percent lower under static pricing, output increases by only 3.5 percent and the number of seats sold to business travelers decreases by 41.9 percent (leisure sales increase by 17.8 percent).

In Figure 6-(b), I plot the percentage of flights that experience price changes over time. Comparing the two lines provides insights on the sources for price adjustments. Under static pricing, the first significant price hike occurs twenty-one days before departure, or when there is a significant change in the composition of arriving consumers according to the model estimates. These price adjustments occur regardless of whether the firm internalizes scarcity. Under dynamic pricing, there are substantially more price adjustments at all other times. This occurs well in advance of departure date (well over five times more price adjustments), even though preferences are not changing. There are also over three times as many price adjustments close to the perishability date. Both of these findings are consistent with the raw data and show that the early price adjustments are primarily in response to demand shock realizations that allow the firm to better reoptimize

remaining inventory for future (and increasingly price insensitive) arrivals.

8 Conclusion

This paper investigates two major determinants of airfare fluctuations, demand shocks and intertemporal variation in willingness to pay. The main contribution of this paper is to jointly study these pricing forces to quantify their welfare implications. I do so by examining US monopoly markets using novel flight-level data. I show that dynamic airline pricing expands output, increases revenues, and lowers total consumer surplus relative to uniform pricing. Leisure consumers benefit from dynamic pricing. Although business consumers are ensured seats, they are also targeted with high prices. In aggregate, I find welfare is higher under dynamic pricing than under uniform pricing. The results at the route level highlight that the welfare effects of dynamic pricing are ambiguous. In markets where price adjustments are primarily in response to changes in willingness to pay, the intertemporal price discrimination force dominates, and welfare decreases under dynamic pricing.

References

- AGUIRRE, I., S. COWAN, AND J. VICKERS (2010): "Monopoly price discrimination and demand curvature," *American Economic Review*, 100(4), 1601–15.
- AGUIRREGABIRIA, V. (1999): "The dynamics of markups and inventories in retailing firms," *The Review of Economic Studies*, 66(2), 275–308.
- 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," *working paper*.
- 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 (2021): "Intertemporal Price Discrimination in Sequential Quantity-Price Games," *working paper*.
- DANA, J. D. (1998): "Advance-Purchase Discounts and Price Discrimination in Competitive Markets," *Journal of Political Economy*, 106(2), 395–422.
- 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.
- GERSHKOV, A., B. MOLDOVANU, AND P. STRACK (2018): "Revenue-maximizing mechanisms with strategic customers and unknown, markovian demand," *Management Science*, 64(5), 2031–2046.
- GOWRISANKARAN, G., AND M. RYSMAN (2012): "Dynamics of consumer demand for new durable goods," *Journal of political Economy*, 120(6), 1173–1219.
- 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.
- 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,” *working paper*.
- 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 (2020): “Dynamic oligopoly pricing: Evidence from the airline industry,” *International Journal of Industrial Organization*, 71, 102639.
- SOBEL, J. (1991): “Durable goods monopoly with entry of new consumers,” *Econometrica*, 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.

9 Appendix

Table 4: Parameter Estimates

Variable		AUSBOS	BILSEA	BOIPDX	BOSJAX	BOSMCI	BOSSAN	BZNPDX	CHSSEA
<u>Logit Demand</u>									
DoW Prefs	β^0	6.769	3.642	4.529	6.404	7.536	9.365	3.412	4.911
		0.197***	0.300***	0.131***	0.170***	0.046***	0.349***	0.314***	0.193***
	β^1	6.275	3.246	3.993	5.165	7.231	8.597	2.579	2.374
		0.166***	0.285***	0.122***	0.173***	0.050***	0.286***	0.276***	0.318***
	β^2	6.126	3.110	4.088	5.343	7.432	8.609	2.803	3.371
		0.172***	0.286***	0.129***	0.169***	0.051***	0.317***	0.290***	0.237***
	β^3	6.414	3.291	4.780	6.662	7.400	9.679	3.696	4.120
		0.200***	0.299***	0.139***	0.164***	0.052***	0.359***	0.300***	0.230***
	β^4	6.031	3.365	4.922	6.515	7.550	9.270	3.777	5.692
		0.245***	0.308***	0.140***	0.193***	0.052***	0.339***	0.311***	0.191***
	β^5	6.169	3.388	3.737	6.090	8.397	9.634	2.470	-10.958
		0.181***	0.302***	0.120***	0.219***	0.050***	0.346***	0.291***	0.001***
	β^6	6.367	3.874	4.928	7.033	7.906	10.141	4.219	6.996
		0.255***	0.324***	0.143***	0.228***	0.048***	0.382***	0.294***	0.214***
Leis. Price Sens.	α_L	-4.087	-1.647	-3.268	-3.285	-3.994	-2.409	-2.056	-1.661
		0.099***	0.094***	0.078***	0.067***	0.022***	0.103***	0.125***	0.054***
Bus. Price Sens.	α_B	-1.721	-0.722	-1.498	-2.157	-1.636	-1.707	-0.693	-0.968
		0.044***	0.055***	0.044***	0.040***	0.011***	0.061***	0.062***	0.040***
Pr(Bus.) Cons.	γ_0	-1.561	-4.194	-6.661	-3.369	-0.452	-0.705	-18.985	-5.982
		0.152***	0.234***	0.152***	0.110***	0.004***	0.197***	2.263***	0.463***
Pr(Bus.) Linear	γ_1	-0.020	0.005	0.119	0.103	-0.026	-0.018	0.614	0.061
		0.003***	0.007	0.005***	0.007***	7.0E-04***	0.008**	0.093***	0.021***
Pr(Bus.) Quad.	γ_2	0.001	0.002	3.5E-04	-3.7E-04	5.2E-04	0.002	-0.004	0.001
		6.6E-05***	8.7E-05***	5.4E-05***	9.2E-05***	1.5E-05***	1.3E-04***	8.7E-04***	2.8E-04***
<u>Poisson Rates</u>									
> 21 Days	μ_1	3.531	1.451	0.804	1.338	1.344	1.178	1.193	2.094
		0.463***	0.070***	0.010***	0.037***	0.008**	0.034***	0.062***	0.068***
14 to 21 days	μ_2	2.948	1.661	1.349	1.028	1.484	0.939	0.831	1.580
		0.292***	0.108***	0.023***	0.033***	0.012***	0.036***	0.056***	0.067***
7 to 14 days	μ_3	2.279	1.418	1.237	0.745	1.287	0.770	0.533	1.140
		0.169***	0.103***	0.023***	0.025***	0.008***	0.038***	0.047***	0.053***
< 7 days	μ_4	0.809	1.798	1.218	0.442	0.986	0.406	0.851	1.303
		0.070***	0.143***	0.025***	0.019***	0.010***	0.029***	0.096***	0.056***
DoW Effect	μ^1	1.089	0.977	0.901	1.347	1.141	0.995	1.062	2.559
		0.033***	0.076***	0.030***	0.044***	0.008***	0.048***	0.068***	0.530***
	μ^2	1.039	1.144	1.024	1.252	1.190	1.173	1.036	1.602
		0.032***	0.060***	0.014***	0.035***	0.007***	0.044***	0.055***	0.101***
	μ^3	1.181	1.207	0.997	1.032	1.209	1.124	0.822	1.205
		0.035***	0.063***	0.012***	0.031***	0.007***	0.044***	0.038***	0.052***
	μ^4	1.425	1.174	0.905	1.051	1.128	1.232	0.871	0.956
		0.043***	0.066***	0.013***	0.045***	0.006***	0.043***	0.050***	0.043***
	μ^5	1.024	0.936	0.801	0.981	0.451	0.824	0.891	8.493
		0.029***	0.067***	0.029***	0.064***	0.003***	0.036***	0.045***	0.001***
	μ^6	1.366	0.932	0.826	0.973	0.941	1.160	0.695	0.691
		0.050***	0.048***	0.011***	0.050***	0.006***	0.060***	0.041***	0.048***
<u>Firm Shock</u>									
	σ	0.670	0.370	0.118	0.499	0.311	0.862	0.182	1.068
		0.030***	0.039***	0.005***	0.014***	0.001***	0.025***	0.012***	0.035***
LogLike		-47,557	-79,749	-266,052	-101,409	-60,597	-61,387	-53,669	-35,926
Number of Flights		238	423	1,626	478	304	317	333	173
Number of Dep. Dates		119	106	106	120	108	120	106	87
Number of Obs.		13,645	25,180	96,774	26,673	22,652	18,294	19,848	10,198

Note: Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Prices are scaled to \$100.

Table 5: Parameter Estimates

Variable		CMHSEA	FATPDX	GEGPDX	GTFSEA	HLNSEA	ICTSEA	MSOPDX
<u>Logit Demand</u>								
DoW Prefs	β^0	3.715	3.270	21.913	12.182	6.107	1.386	5.627
		0.343***	0.255***	0.604***	1.497***	0.327***	0.234***	0.446***
	β^1	3.319	3.017	21.325	12.105	5.667	1.768	5.043
		0.302***	0.216***	0.601***	1.491***	0.312***	0.215***	0.412***
	β^2	3.406	2.761	21.605	12.132	5.517	2.057	4.971
		0.305***	0.204***	0.603***	1.500***	0.304***	0.187***	0.409***
	β^3	3.849	3.032	22.216	12.175	6.077	2.463	5.433
	0.362***	0.239***	0.601***	1.496***	0.331***	0.209***	0.418***	
	β^4	3.780	3.212	22.152	12.247	6.350	1.845	5.166
		0.349***	0.253***	0.605***	1.502***	0.359***	0.249***	0.410***
	β^5	3.295	3.186	21.178	12.111	5.416	2.407	4.838
		0.325***	0.238***	0.601***	1.491***	0.297***	0.295***	0.404***
	β^6	3.951	3.323	22.225	12.261	6.593	0.837	5.430
		0.389***	0.261***	0.604***	1.498***	0.362***	0.316***	0.445***
Leis. Price Sens.	α_L	-1.288	-2.215	-17.718	-6.131	-3.321	-1.146	-3.266
		0.079***	0.097***	0.488***	0.729***	0.165***	0.043***	0.227***
Bus. Price Sens.	α_B	-0.552	-0.837	-8.825	-2.755	-1.691	-0.240	-1.538
		0.047***	0.049***	0.235***	0.348***	0.087***	0.021***	0.151***
Pr(Bus.) Cons.	γ_0	-6.175	-4.325	-2.734	-2.276	-12.088	-2.553	-1.742
		0.378***	0.309***	0.089***	0.438***	2.080***	0.128***	0.162***
Pr(Bus.) Linear	γ_1	0.115	0.022	-0.015	-0.360	-4.022	0.024	-0.029
		0.015***	0.011**	0.007**	0.181**	2.406*	0.005***	0.006***
Pr(Bus.) Quad.	γ_2	2.2E-04	0.001	0.002	0.010	0.107	7.2E-05	0.002
		1.4E-04	1.2E-04***	1.4E-04***	0.005**	0.061*	4.8E-05	1.8E-04***
<u>Poisson Rates</u>								
> 21 Days	μ_1	2.712	1.624	1.019	0.833	0.930	3.128	0.930
		0.143***	0.065***	0.015***	0.029***	0.020***	0.320***	0.025***
14 to 21 days	μ_2	3.278	2.606	0.614	0.588	0.492	3.448	0.690
		0.186***	0.165***	0.021***	0.064***	0.031***	0.372***	0.039***
7 to 14 days	μ_3	1.831	2.513	0.482	0.572	0.376	3.821	0.535
		0.189***	0.142***	0.021***	0.049***	0.021***	0.410***	0.037***
< 7 days	μ_4	2.501	2.861	0.277	0.946	0.234	6.489	0.366
		0.290**	0.158***	0.013**	0.089***	0.014***	0.727***	0.060***
DoW Effect	μ^1	0.990	0.777	0.895	0.869	0.882	0.678	0.942
		0.053***	0.039***	0.020***	0.059***	0.041***	0.071***	0.047***
	μ^2	0.925	0.838	0.876	1.120	0.946	0.643	1.050
		0.055***	0.038***	0.021***	0.038***	0.047***	0.052***	0.034***
	μ^3	0.979	0.928	0.890	1.168	0.928	0.528	0.987
		0.051***	0.024***	0.025***	0.037***	0.044***	0.040***	0.039***
	μ^4	1.006	0.970	0.917	1.218	1.009	0.765	1.140
	0.056***	0.024***	0.023***	0.047***	0.050***	0.096***	0.042***	
	μ^5	0.928	0.820	0.769	0.915	0.943	0.519	0.929
		0.086***	0.021***	0.018***	0.052***	0.047***	0.131***	0.050***
	μ^6	1.054	1.114	0.825	0.982	0.883	1.624	1.202
		0.061***	0.024***	0.025***	0.082***	0.035***	0.188***	0.048***
<u>Firm Shock</u>								
	σ	0.678	0.228	0.260	0.194	0.197	0.264	0.086
		0.073***	0.020***	0.004***	0.018***	0.012***	0.016***	0.010***
LogLike		-40,301	-78,041	-304,409	-57,210	-29,791	-39,109	-26,969
Number of Flights		210	424	1,718	412	212	212	208
Number of Dep. Dates		106	106	106	106	106	106	106
Number of Obs.		12,398	25,248	102,292	24,532	12,644	12,606	12,390

Note: Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Prices are scaled to \$100.

Table 6: Parameter Estimates

Variable		OKCSEA	PDXRNO	PDXSBA	PDXSTS	SBASEA	SEASTS	SEASUN
<u>Logit Demand</u>								
DoW Prefs	β^0	-0.834	6.991	3.173	4.479	3.698	4.591	9.845
		0.501*	0.203***	0.370***	0.183***	0.436***	0.333***	1.071***
	β^1	-0.590	6.045	2.463	2.917	2.799	3.099	9.279
		0.429	0.204***	0.280***	0.235***	0.356***	0.354***	1.086***
	β^2	-0.567	6.051	2.484	2.416	2.930	3.206	9.498
		0.441	0.206***	0.276***	0.226***	0.364***	0.352***	1.101***
	β^3	-0.403	6.888	2.942	4.313	3.638	4.421	8.822
		0.463	0.195***	0.352***	0.184***	0.437***	0.326***	0.993***
	β^4	-0.579	7.394	3.265	4.504	3.864	4.665	9.967
		0.549	0.216***	0.378***	0.183***	0.455***	0.347***	1.064***
	β^5	-0.449	6.256	2.647	3.359	3.458	2.772	8.767
		0.458	0.195***	0.314***	0.228***	0.410***	0.411***	1.007***
	β^6	-0.697	9.448	3.857	4.555	4.429	5.608	10.033
		0.479	0.277***	0.444***	0.193***	0.553***	0.388***	0.995***
Leis. Price Sens.	α_L	-0.715	-3.658	-1.603	-2.333	-1.850	-1.873	-7.118
		0.037***	0.102***	0.098***	0.077***	0.147***	0.096***	0.760***
Bus. Price Sens.	α_B	-0.029	-2.037	-0.741	-0.860	-0.747	-0.979	-2.551
		0.029	0.065***	0.066***	0.050***	0.073***	0.067***	0.254***
Pr(Bus.) Cons.	γ_0	-3.454	-3.739	-3.370	-6.307	-7.502	-7.091	-1.334
		0.219***	0.441***	0.401***	0.473***	0.368***	0.833***	0.078***
Pr(Bus.) Linear	γ_1	0.044	-0.114	0.008	0.058	0.160	0.020	-0.025
		0.004***	0.019***	0.011	0.015***	0.017***	0.019	0.005***
Pr(Bus.) Quad.	γ_2	-3.0E-04	0.004	0.001	0.001	-1.6E-04	0.002	6.0E-04
		4.7E-05***	2.2E-04***	1.7E-04***	1.9E-04***	1.8E-04	3.9E-04***	1.1E-04***
<u>Poisson Rates</u>								
> 21 Days	μ_1	11.419	1.137	1.790	1.603	2.418	1.234	0.991
		8.256	0.026***	0.133***	0.044***	0.126***	0.050***	0.032***
14 to 21 days	μ_2	13.810	1.151	1.862	1.662	3.172	1.353	0.832
		9.989	0.030***	0.173***	0.073***	0.245***	0.069***	0.047***
7 to 14 days	μ_3	13.188	0.951	1.795	1.015	2.521	0.995	0.821
		9.446	0.032***	0.179***	0.056***	0.217***	0.061***	0.050***
< 7 days	μ_4	18.045	1.106	2.305	1.248	2.895	1.036	0.895
		12.784	0.052***	0.241***	0.083***	0.346***	0.068***	0.157***
DoW Effect	μ^1	0.763	0.879	0.944	0.881	0.905	1.353	1.100
		0.159***	0.031***	0.052***	0.086***	0.054***	0.094***	0.047***
	μ^2	0.730	0.907	0.918	1.075	0.972	1.348	1.009
		0.156***	0.032***	0.061***	0.144***	0.062***	0.086***	0.047***
	μ^3	0.687	0.918	0.994	0.852	1.039	1.051	1.149
		0.163***	0.029***	0.042***	0.035***	0.036***	0.052***	0.046***
	μ^4	0.860	0.940	0.953	0.917	0.976	1.101	1.011
		0.207***	0.038***	0.036***	0.040***	0.038***	0.053***	0.044***
	μ^5	0.640	0.835	0.868	0.730	0.867	1.437	0.847
		0.147***	0.028***	0.046***	0.061***	0.042***	0.154***	0.063***
	μ^6	0.932	0.645	0.867	1.137	0.884	0.999	0.722
		0.205***	0.024***	0.053***	0.040***	0.104***	0.055***	0.132***
<u>Firm Shock</u>								
	σ	0.262	0.430	0.309	0.395	0.464	0.623	0.194
		0.013***	0.015***	0.036***	0.024***	0.065***	0.052***	0.015***
LogLike		-36,660	-143,894	-38,303	-66,833	-43,454	-58,125	-26,587
Number of Flights		212	721	212	384	212	306	168
Number of Dep. Dates		106	106	106	106	106	106	85
Number of Obs.		12,405	42,989	12,623	22,833	12,582	18,182	9,994

Note: Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Prices are scaled to \$100.

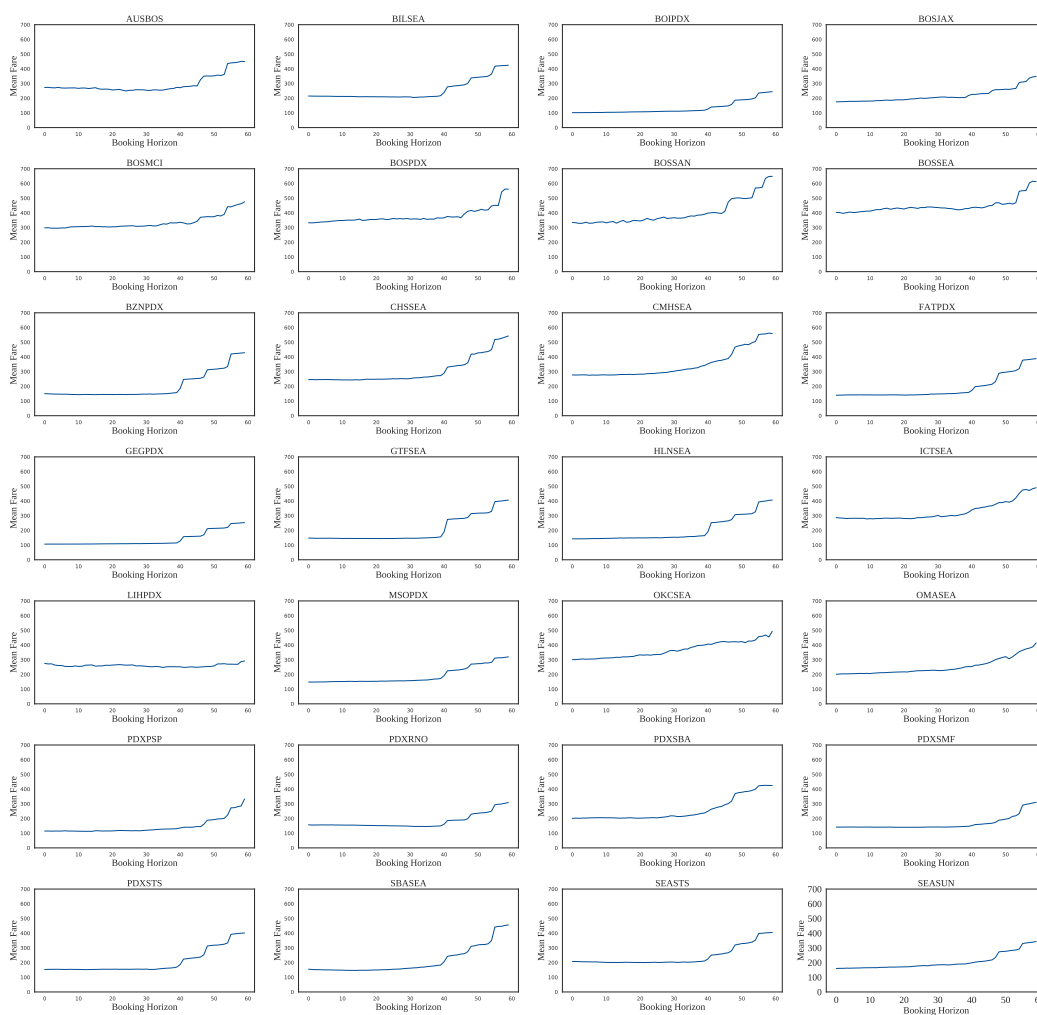
Online Appendix

The Welfare Effects of Dynamic Pricing: Evidence from Airline Markets

by Kevin R. Williams


A Additional Figures

Figure 7: Average Fares over Time by Market



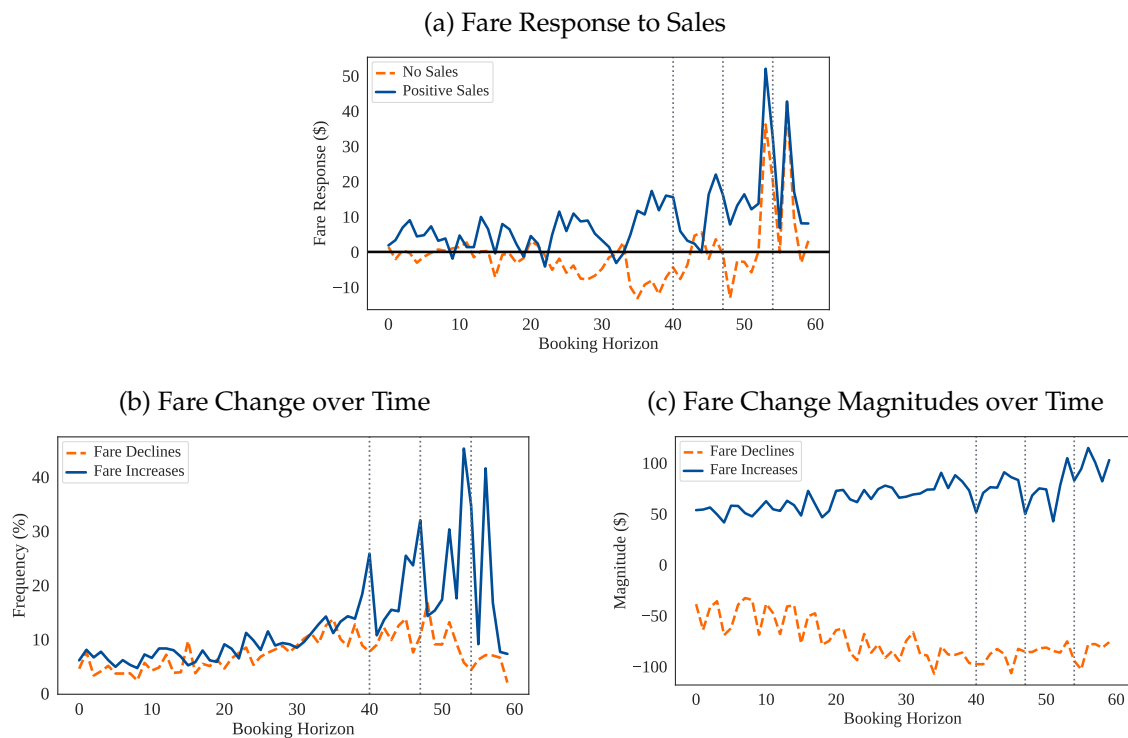
Note: Average fares over time for each market separately. This analysis combines origin-destination and destination-origin fares. Both axes are common across all plots.

Figure 8: Percentage Change in Fares over Time by Market



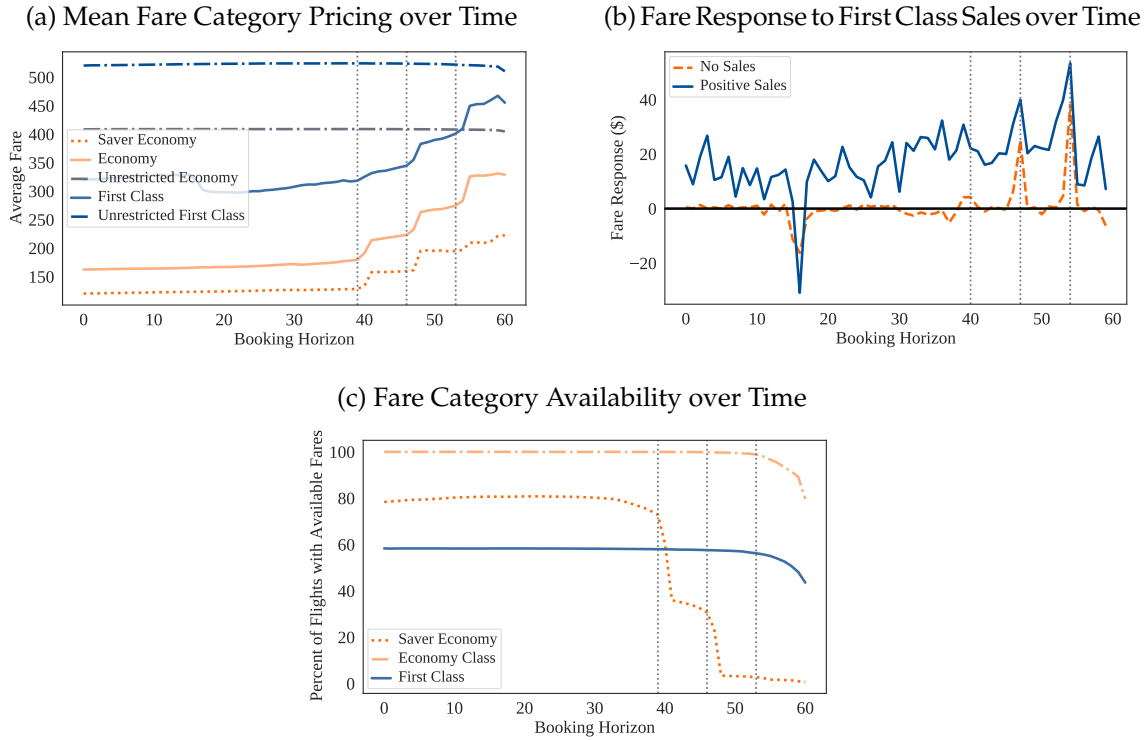
Note: Percent change in average fares over time for each market separately. This analysis combines origin-destination and destination-origin fares. Both axes are common across all plots.

Figure 9: Fare Dynamics in Competitive Markets



Note: Recreation of Figure 1-(b) through Figure 1-(d) for markets with non-stop competition. (a) Fare response to own bookings (no bookings) over time. (b) Frequency of fare increases and decreases over time. (c) Magnitude of fare increases and decreases over time.

Figure 10: Fare Category Pricing Dynamics




Note: (a) Mean prices of different fare categories over time. Full-fare (refundable) tickets for both economy and first class are flat over time. Average prices for saver-economy, economy, and first-class tickets rise over time. The gap in prices between saver-economy and economy prices grows as the departure date approaches.

(b) Recreation of Figure 1-(b) for first class. Compared to economy class, the presence of APDs is diminished in first class, and fare increases are more pronounced throughout the booking horizon.


(c) Percentage of flights that offer observed fare categories over time. First-class denominator is the number of flights with first class, not the number of flights in the sample. Close to departure, economy fare availability abruptly drops suggesting that the spike in load factor shown in Figure 1-(a) captures last-minute bookings. Economy fares rise and saver economy availability declines.

Figure 11: Pricing Effects on Other Itineraries

(a) Connecting Fare Response to Direct Sales



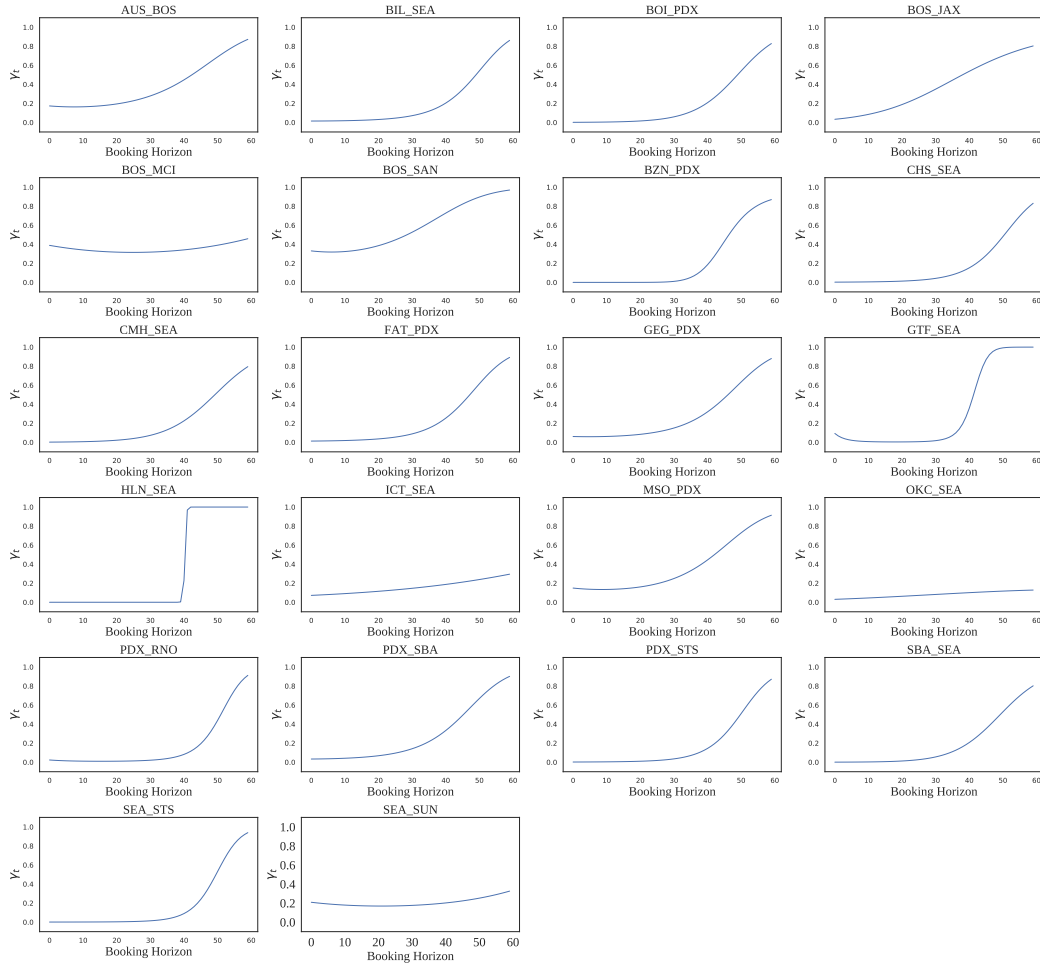
(b) Multiple Nonstop Prices and Load Factors



Note: (a) Recreation of Figure 1-(b), but with connecting fares instead of direct fares. The connecting fare is the average fare among connecting flight options for the same carrier, departure date, and booking date. Evidence suggests that connecting fares are unaffected by nonstop bookings.

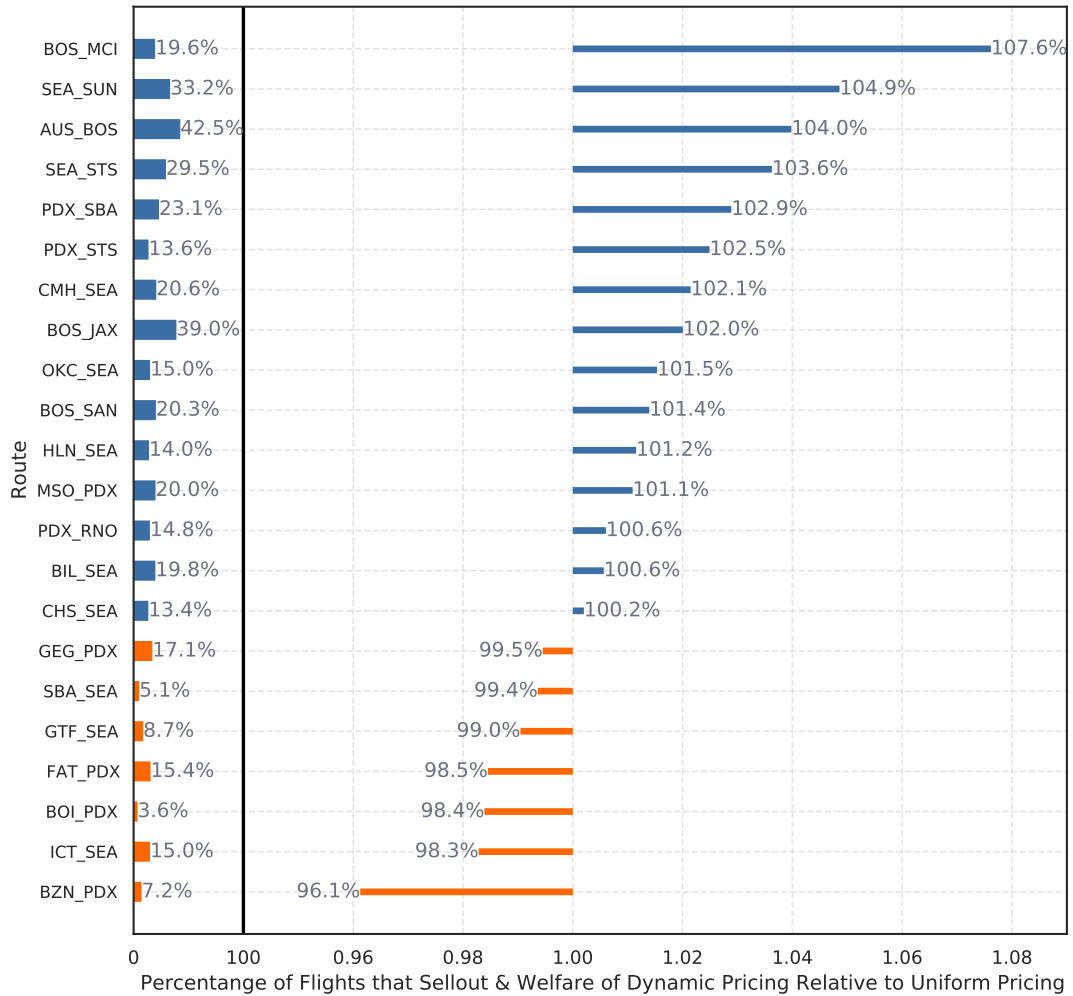
(b) Fourth order polynomial fit of a regression of the percent difference in fares on the percent difference in load factor when a carrier operates two nonstop flights a day. When flights have the same load factor, average difference in fares is 0.6 percent. The line is upward sloping, meaning that the flight with the higher load factor is, on average, more expensive.

Figure 12: Fitted Values of γ_t over Time for Each Market



Note: Probability that an arriving consumer is of the business type over time, by market.

Figure 13: Welfare Effects of Dynamic Pricing for Each Market



Note: (Left) The percentage of flights that sellout on or before the departure date. (Right) Welfare under dynamic pricing over welfare under uniform pricing. Numbers above 100% (top, blue) indicate welfare is higher under dynamic pricing than uniform pricing.

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).³⁷ Next, I focus on ODs with a nonstop option; this reduces the number of potential markets studied from 73,000 to 9,800. Over 40 percent of these markets have a single carrier providing nonstop service and this subset makes up a total of 14 percent of OD traffic in the United States. I then implement the following cleaning criteria: (1) total quarterly traffic, including connecting traffic with up to four stops, exceeds 600 passengers;³⁸ (2) a single carrier operates nonstop on the OD leg. This reduces the number of potential markets by over half, to roughly 3,900.

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

³⁷This 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).

³⁸This 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.

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

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

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

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

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 76 percent and 57 percent as compared with 83 percent and 61 percent for competitive markets (medians of 82 percent, 56 percent, 88 percent, 62 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 ($\rho = -0.33$), each is correlated with distance. The correlation between percentage non-connecting and distance is 0.24; ODs that are closer together have higher connecting traffic. The correlation between percentage nonstop and distance is -0.52 ; ODs that are closer together have a higher percentage of nonstop traffic.

Markets with high nonstop percentages and low connecting percentages are ideal because changes in seat maps are likely to be attributed to the correct itinerary, and hence, fare. One important caveat to this approach is that markets with a high nonstop percentage are also closer together, which implies there may be alternative modes of transportation, e.g., a train, that is relevant for airline demand. For

example, in 2019, there exist 556 ODs with nonstop and non-connecting fractions above the 95 percent threshold. Of those ODs, 523 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 meet threshold criteria (for nonstop and non-connecting traffic) are Alaska Airlines and JetBlue Airways.

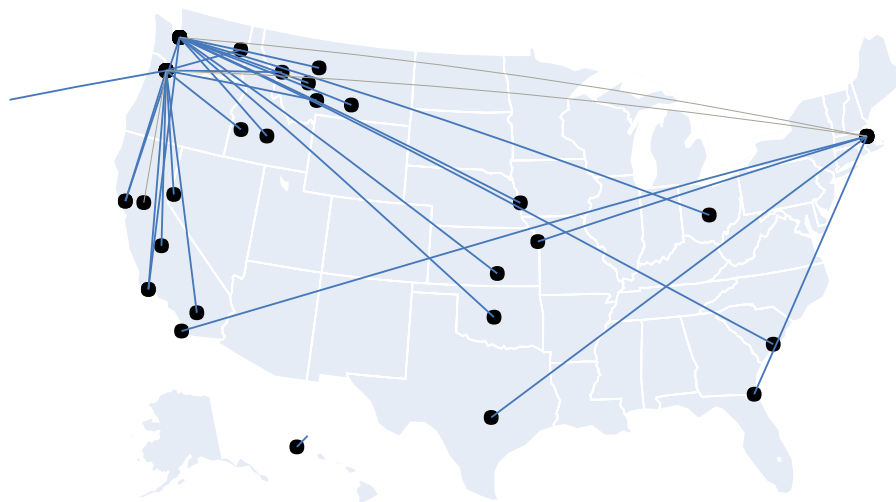
I select fifty ODs and concentrate data collection on two carriers, JetBlue Airways and Alaska Airlines, such that both seat map and airfare data could be collected. The other carriers included in the data are Delta Air Lines and Frontier Airlines. In addition, for a comparison in the descriptive analysis, I collect data on six duopoly markets.³⁹ Figure 14 maps the markets and Table 7 provides a dictionary for the airport codes. The data were collected in two phases: The data on markets operated by Delta and 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 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 because of the fee charged to select seats. These 556 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

³⁹The 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 exited the market.

Figure 14: Markets of Study

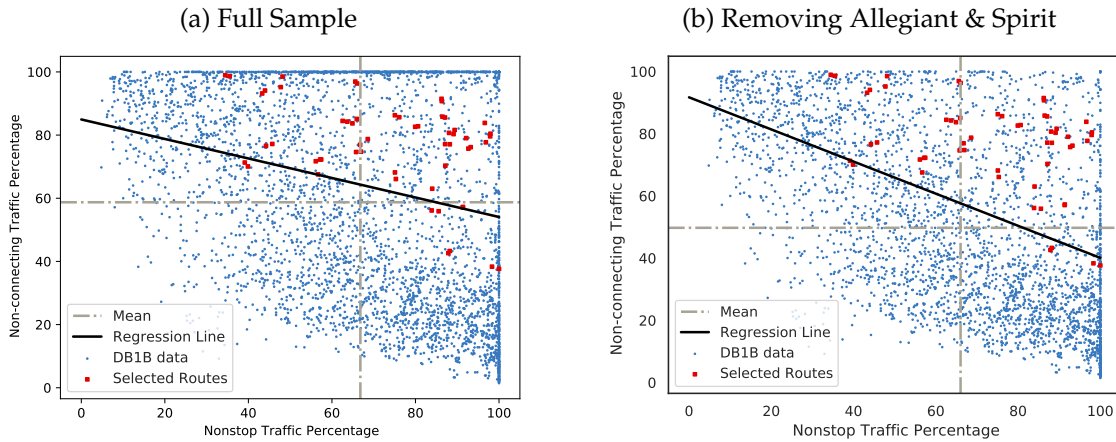


Note: Map of the markets selected for study. All of the markets either start or end at Seattle, WI; Portland, OR; and Boston, MA.

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 eight (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.

Table 8 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.

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



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 7: 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 8: 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

Seat maps may not accurately represent flight loads 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. Ideally, the severity of measurement error can be measured by matching changes in seat maps with bookings; however, this is impossible with publicly available data.

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. I make two adjustments. First, because I do not observe first class cabins in the 2012 sample, I assume first class goes out at 100% full and subtract off this passenger number using the size of the first class cabin as recorded from the plane types in the T100. Second, because the number of observed flights

can differ, e.g., due to cancellations, flight number changes cause data collection to end, or flights are not tracked for 60 days, I reconcile any differences in the number of departures by adding or subtracting the average observed flight load times the count difference. Figure 16 provides a scatter plot that compares the two statistics. Most points closely follow the 45-degree line, and I find seat maps overstate recorded enplanements, with the median difference being three percent. 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 allowed 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 2012. 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 zero to five percent across days, or at

most four seats. This suggests seat maps are useful for recovering bookings as the departure date approaches.

D 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.

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 AP fare expires. Table 9 reports regression results under three fixed effects specifications. I find insignificant bunching at the fourteen-day AP expiration. I find negative bunching at the three-day and twenty-one-day AP expiration, meaning sales are lower prior to the price increases. Finally, I find a positive and significant coefficient for the seven-day AP requirement; 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. For example, Li, Granados, and Netessine (2014) estimate that between 5 and 20 percent of consumers dynamically substitute across days.

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.

Table 9: Consumer Bunching Regressions

	(1)	(2)	(3)
APD3	-0.504** (0.144)	-0.502** (0.142)	-0.502** (0.142)
APD7	0.200* (0.0725)	0.202* (0.0697)	0.201* (0.0697)
APD14	-0.0717 (0.0459)	-0.0720 (0.0414)	-0.0719 (0.0413)
APD21	-0.131* (0.0412)	-0.131* (0.0397)	-0.130* (0.0397)
$m(t)$	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	738,625	738,625	738,625
R^2	0.609	0.618	0.748

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

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.

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

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

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

where EU^{wait} is the expected value of waiting one more period. This expected

utility can be written as

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

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

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

For this marginal consumer, the expected gains from waiting are

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

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

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

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

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

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

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

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

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

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

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

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

Moving the expectation operator, the last line above equals

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

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

For consumers who would purchase today, the gains from waiting are equal to $\bar{\phi}$, 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, $\bar{\phi}/\alpha = \Pr(p_{t+1} < p) \mathbb{E}[(p - p_{t+1}) \mid p_{t+1} < p]$ defines a transaction


cost such that waiting is never optimal. For these costs to be calculated, the information set of consumers needs to be defined. I assume consumers form expectations given current prices and time, but they do not forecast the changes in number of seats remaining across time. This seems reasonable given that remaining capacity is not reported to consumers. With these assumptions, I find the median and mean transaction costs to be \$5.85 and \$5.75, respectively. These costs are based on the most extreme case—the consumer who is indifferent between purchasing today or delaying the decision.

E Additional Counterfactuals

E.1 Initial Capacity and Approaching Static Pricing

I compute optimal dynamic prices and simulate outcomes for a wide range of initial capacity values in order to investigate how large initial capacity has to be in order for static pricing to be a reasonable approximation of the environment.

Figure 18: Initial Capacity Counterfactual



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 0.5%. The right panel shows average prices over time for those three scenarios (average less ten, average, and the limiting case).

Figure 18 demonstrates the counterfactual. 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 right vertical axis plots total expected revenues by initial capacity. The (grey) vertical line depicts the average observed initial capacity. The (light orange) square denotes revenues with six fewer seats than the average (a row of a plane); the (orange) triangle denotes the minimum initial capacity such that approximate a static pricing model (revenues are within 0.5 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 then compare initial observed capacities to the calculated thresholds. I find that 31.9 percent of the observed flights can be approximated by static pricing.

E.2 Frequent Price Adjustments

I explore the use dynamic pricing, with the restriction that prices must be maintained for k days. I conduct six counterfactuals, corresponding to $k = 2, 3, 6, 10, 20, 30$.

Figure 19 plots the revenue loss compared to revenues under daily re-optimization. The results show that the ability to update prices just once reduces the revenue loss compared uniform pricing by more than half (30-day adjustments). An additional price adjustment yields another 1.3 percent gain. Re-optimization in with time intervals less than one week result in similar revenues, meaning several demand shocks can be observed before re-optimization is required.

Figure 19: 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.