

Consequences of dynamic pricing in competitive airline markets*

Nan Chen[†]

Przemyslaw Jeziorski[‡]

January 27, 2023

Abstract

Dynamic pricing is a common tool to maximize sales in markets for perishable goods with limited capacity and stochastic demand. We look at the relationship between dynamic pricing and oligopolistic competition in the airline industry. For this purpose, we estimate a dynamic oligopoly model of dynamic price competition in a market with carrier exit using flight-level data. We discover that dynamic pricing results in a Pareto improvement, increasing firm profits, and consumer welfare. We break down the impact into two categories: price discrimination and pricing on residual capacity (revenue management). We find that price discrimination softens competition, whereas revenue management intensifies it.

JEL: D22, D25, D43, D61, L11, L13, L93.

Keywords: airline industry; dynamic pricing; dynamic oligopoly; perishable good; price discrimination; revenue management; demand uncertainty; capacity constraint.

*We thank you Miguel Villas-Boas, Ganesh Iyer, Minjung Park, Kei Kawai, Severin Borenstein, Jean-Pierre Dubé, Pedro Gardete, Günter Hirsch, Ben Handel, Wesley Hartmann, Yuichiro Kamada, Philipp Strack, and seminar participants at Berkeley Haas, Chicago Booth, Imperial College, NUS, and Stanford GSB. This paper was part of the Nan Chen's PhD dissertation titled "Essays on Quantitative Marketing and Economics" defended in 2018, available at <https://escholarship.org/uc/item/2586h4px>.

[†]nanchentrueman@gmail.com

[‡]UC Berkeley, przemekj@berkeley.edu

1 Introduction

Dynamic pricing is a common technique for optimizing sales in markets with limited capacity and fixed utilization deadlines. This form of pricing is applied in various industries, including hospitality, entertainment, and fashion; nevertheless, the airline industry is considered a pioneer and leader in this area. Airlines offer a fixed number of seats for the specific departure date. With the help of sophisticated algorithms, the price of a seat fluctuates as the departure date approaches and as a function of the plane’s remaining capacity. In extreme circumstances, the identical airline seat advertised on the day of departure may be several times more expensive than if quoted in advance. The resulting price dispersion raises questions over the degree of market dominance and the capability of rent extraction facilitated by sophisticated dynamic pricing (see [Borenstein and Rose, 1994](#)).

In this paper, we study the effect of dynamic pricing on consumer welfare and firms’ profits in the presence of oligopolistic competition. Specifically, we investigate the incentive for firms to engage in dynamic pricing and the impact of sophisticated pricing strategies on the distribution of rents between customers and competing firms. We propose a dynamic oligopoly model in which companies set prices over time based on the capacity utilization of themselves and their competitors. We employ daily-level data on prices and capacity utilization to quantify demand uncertainty and inter-temporal preference heterogeneity. Using the estimates, we recompute pricing equilibria with and without dynamic pricing. We subsequently compare consumer welfare and firms’ profits in both cases. We also highlight two mechanisms that determine the split of surplus: price discrimination and revenue management.

The ability to alter the price over time provides at least two substantial benefits. First, since the demand elasticity of consumers changes as the departure date approaches, time-based dynamic pricing enables price discrimination. For instance, if inelastic business consumers arrive late, the firm may wish to raise prices as the departure date approaches. Second, given the stochastic nature of the demand for airplane seats, the airline may wish to alter the price based on previous demand realizations and remaining capacity – a practice called revenue management. For example, if historical sales were exceptionally high and resulted in minimal remaining capacity, an airline is driven to increase the price. Using the structural model, we decompose the effects of dynamic pricing into price discrimination and

revenue management by examining an intermediate Nash Equilibrium in which time-based price discrimination is permitted, but revenue management is not.

To identify the demand elasticity, we conduct an event study of the carrier exit in the Seattle-Tuscon route. Seattle-Tuscon is a monopoly route that is seasonally operated by two carriers. We concentrate on the brief time window surrounding the exit event to minimize price endogeneity. In addition, we apply difference-in-differences identification by incorporating a San Diego-Boston duopoly control market with nearly identical pricing pre-trend.

We discover that leisure clients who arrive early are more price sensitive than business customers who arrive late. Furthermore, we demonstrate that corporate customers are more brand loyal. This heterogeneity rationalizes increasing fares as the departure date approaches, which is ubiquitous in most airline routes. Subsequently, we use the preference estimates to conduct a counterfactual that eliminates the ability to price dynamically. Specifically, we compute a Nash Equilibrium in which all firms are constrained to set uniform prices as the departure date approaches. We find that dynamic pricing leads to Pareto improvement raising consumer welfare and firms' profits. Specifically, consumer welfare grows by 3%, profits increase by 8%, and total welfare soars by 5%. In contrast, [Williams \(2022\)](#) concludes that dynamic pricing reduces consumer welfare in a monopoly market. Since that work considers different airline markets, we also provide an apples-to-apples analysis of the interaction between market structure and dynamic pricing. For this purpose, we demonstrate the effect of dynamic pricing after imposing a merger to a monopoly, keeping other parameters the same. We find that dynamic pricing leads to a 14% decrease in consumer welfare, an 8% increase in profits, and a nearly 5% decrease in total surplus. The comparison stresses the importance of competition when allowing dynamic pricing.

To disentangle the effects of price discrimination and revenue management, we explore an intermediate case in which revenue management is prohibited, but price discrimination is permitted. For this purpose, we compute a Nash Equilibrium in which the companies must commit to the pricing schedule as the departure date approaches. As a result, airlines are barred from adjusting rates based on past sales and residual capacity but can still engage in third-degree price discrimination. Perhaps not surprisingly, we observe lower prices in the early elastic market and higher prices in the late inelastic market. Notably, the prices under price discrimination are, on average, higher than the uniform price and uniformly higher than the prices under simultaneous price discrimination and revenue management. Also, the

quantity supplied in the intermediate case is lower than under uniform and fully dynamic pricing. Consequently, we find that price discrimination results in a 10% loss in customer welfare (compared to a 3% rise when revenue management is included) and an 8% increase in profits compared to uniform pricing. Putting all the results together, we conclude that price discrimination softens competition, whereas revenue management intensifies it.

The mechanisms behind our results rely on shifting the fixed capacity over time as the departure date approaches and the varying efficiency of capacity utilization. Price discrimination causes an initial decrease in prices and an increase in sales in the early market, which leads to extra scarcity in the late market. For example, for almost all periods, the probability of monopolization is greater under price discrimination than under uniform and dynamic pricing, which leads to upward pricing pressure and deadweight loss in the late market. As a result, profits improve, but the welfare of late-arriving business consumers declines by more than 11%. Revenue management leads to more efficient capacity utilization because firms can throttle prices in response to past sales and lose a portion of their commitment power. Specifically, our numbers indicate that monopolization under full dynamic pricing is unlikely, which removes some scarcity-related upward pricing pressure in the late market. As a result, business travelers break even after the introduction of both price discrimination and revenue management – they pay higher prices but are more likely to obtain a seat. At the same time, the welfare of leisure travelers doubles.

This research is related to several streams of literature. It adds to the work on airline pricing, competition, and air travel demand. The airline industry is a textbook example of dynamic pricing (McAfee and Te Velde, 2006). Though airlines’ pricing techniques have attracted advocates from many industries, their implications are not well understood (Borenstein and Rose, 2014). Additionally, the industry’s competitiveness has received attention from both the regulators and academics (Borenstein, 1992).

In an attempt to start filling the gap, we (i) implement a research design and collect high-frequency data and (ii) construct and estimate a large dynamic game. Previous work (e.g., Berry, 1990, Berry et al., 2006, and Berry and Jia, 2010) estimates air travel demand and supply using static Nash equilibrium models and aggregate data.

Other related work includes structural papers by Lazarev (2013), Williams (2022), and Aryal et al. (2022). These papers focus on the monopoly markets, which are more computationally tractable. As a result, these models may allow for some features we had to

assume away. Lazarev (2013) studies the implications of intertemporal price discrimination and allowing ticket resale. Notably, they explicitly model connecting traffic and forward-looking consumers. Williams (2022) studies the impact of removing dynamic pricing on welfare. They include more markets in their analysis, but as a result, they must make stronger assumptions regarding price endogeneity. Aryal et al. (2022) allow second-degree price discrimination across the economy and business fares and endogenize seat release policy.

In a complementary working paper, Hortaçsu et al. (2022) study dynamic pricing in duopoly airline markets. Our paper focuses on empirical analysis and provides a more complete treatment of unobserved heterogeneity and price endogeneity. Hortaçsu et al. (2022) focus on a theoretical analysis of a novel continuous-time pricing model, while our model is an adaptation of the empirical framework by Ericson and Pakes (1995) (1995). These approaches lead to different conclusions: we find dynamic pricing increases consumer welfare, while Hortaçsu et al. (2022) show consumer surplus falls. In addition to welfare analysis, our approach allows us to study the interaction between market structure and the effects of dynamic pricing through a merger counterfactual.

The methodology in this work relates to the structural modeling of dynamic oligopoly (Ericson and Pakes, 1995). In particular, it studies a non-stationary and stochastic setting of perishable good pricing. There is an adjacent line of research on pricing in a durable good oligopoly. This literature is recent but fruitful (e.g., Goettler and Gordon, 2011, Chen et al., 2013, and Hendel and Nevo, 2013). Some approaches focus on the dynamics in consumer demand (e.g., Hendel and Nevo, 2006; Gowrisankaran and Rysman, 2012).

Much of the existing work on perishable good pricing investigates the role of demand uncertainty (Gallego and Van Ryzin, 1994, Sweeting, 2012, and Sweeting, 2015). Building on those insights, we model demand with both stochastic variability and temporal heterogeneity. The demand system builds on a Poisson arrival process and a discrete choice framework. The two processes translate into a single, simple, and tractable sales process.

Our analysis builds on the theoretical work on dynamic pricing and revenue management. For example, Deneckere and Peck (2012) study a dynamic pricing model in which a continuum of firms produce one unit of output. In addition, a growing body of work in operations research uses game theory to model competitive revenue management (see Lin and Sibdari, 2009, Levin et al., 2009, and Gallego and Hu, 2014). A more recent paper by Dana Jr and Williams (2022) considers intertemporal competition with deterministic

demand and endogenous capacity. Our model considers exogenous capacity but allows for stochastic demand. In comparison to the theoretical work, we take an empirical approach and structurally uncover demand and supply primitives. The estimates allow us to assess the impact of dynamic pricing using counterfactual exercises.

The paper is organized as follows. Section 2 presents the empirical setting and data collection procedure. Section 3 describes our data. Section 4 sets up the model. Section 5 discusses the empirical strategies and identifications. Section 6 shows the estimation results. Section 7 performs the counterfactual analysis. Finally, Section 8 concludes.

Next section contains a brief overview of the airline industry.

2 Airline industry

The airline industry is an important contributor to economic development. In 2015, the industry generated \$767 billion revenue and transported 3.3 billion passengers according to the International Air Transport Association (IATA). Historically, air travel in the U.S. was viewed as a public good.

The industry operated under governmental subsidies and regulations (Borenstein and Rose, 2014). To avoid “destructive competition,” the Civil Aeronautics Board (CAB) controlled airfares using a nonlinear distance-based formula (called the Standard Industry Fare Level or SIFL). CAB disallowed discounts and promotions on the grounds they disadvantaged competitors or were unduly discriminatory across passengers.

In 1978, the Airline Deregulation Act abruptly removed restrictions on fares and entry. It transformed the strictly regulated airline industry to an innovative, dynamic, and complex one (Borenstein, 1992). Among others, the market liberation spurred unimpeded competition and strategic behavior among airlines, triggering airlines’ development and implementation of dynamic pricing systems.

Airlines are credited as the pioneers of revenue management. Revenue management was originally called yield management, although now the latter term is outdated. The meaning of revenue management can be very broad in many cases, however, the fundamental concept has not changed. Profits are optimized contingent on product availability, in other words, on the actual number of seats currently booked on a specific flight. Thanks to modern information technology, the practice of revenue management is common in today’s airline

industry. This has been documented by previous research (Escobari, 2012).

Figure 1 plots the average price path for JetBlue Flight 19 from March 01, 2016, to June 01, 2016. In line with the well-known industry feature, airfares are on average more expensive closer to departure date (McAfee and Te Velde, 2006). The increasing price path is viewed as evidence of (intertemporal) price discrimination. Generally, leisure consumers with higher price elasticities tend to arrive early. Business travelers who are less price sensitive tend to arrive late. Airlines are thus able to screen consumers based on the times of their arrivals.

There has been much discussion on whether the industry is subject to excessive competition or excessive market power. Airlines claim that the industry is “ultra-competitive.” This is supported by more than 100 bankruptcy filings in the U.S. airline industry since 1978. Every major U.S. interstate airline existent at the time of deregulation in 1978 has since filed a bankruptcy request.

The industry’s profitability may fluctuate due to cyclical demand, sticky fixed costs, and repeated disruptive business innovations. Factors like entry/exit and short-run profitability are not enough to understand the whole picture. One aspect of airlines’ competition is with respect to fares. The internet has facilitated great transparency in airfares, and consumers can search them at presumably small cost. If flight tickets are relatively undifferentiated across airlines, one would expect airlines’ price competition to be fierce. If airlines are sufficiently differentiated, monopoly powers may arise.

In an effort to mitigate competition, airlines introduce a loyalty inducing marketing device called the frequent flyer program (FFP). FFPs reduce travelers’ cross-price elasticities by encouraging them to buy tickets from a single airline. As a result, it increases brand loyalty and switching cost (Borenstein, 1992).

3 Data

We collect high-frequency panel data of posted prices and seat maps from airlines’ websites. We utilize a scraping script to search for all nonstop flights departing within 100 days for a specified set of routes. We track the rates and remaining seat count for each flight. We focus on economy class and disregard seat and fare distinctions. In particular, we utilize the lowest available economy fare to approximate pricing and the number of remaining economy seats to approximate capacity. The scraping script runs daily from March 17, 2016, to April

15, 2016. For approximately 5% of the data we scrapped the prices twice per day. There is less than 14% chance that the price changes in the middle of the day. If the consumers arrive uniformly using daily pricing data may affect less than 7% consumers.¹

A *seat map* is a public interface connected to airlines’ real-time availability data. Airlines maintain real-time inventory data (CRSs) in their computerized reservation systems. Most airlines have outsourced their computer reservation systems (CRSs) to global distribution systems (GDS) providers such as Amadeus, Sabre, Galileo, and Worldspan. GDS businesses strive for accurate seat maps. According to the website of Amadeus, “Amadeus Interactive Seat Map allows airlines and travel agents to deliver superior service to passengers by displaying real-time and accurate seat information... Enhance customer service by providing them with seat choices based on accurate, real-time availability.”² Despite this, it is possible that the seat map data will produce a noisy measure of residual capacity. For instance, airlines are prone to reserve better seats for elite passengers; consequently, some seats that appear to be occupied may be accessible.³ Conversely, certain tickets may receive a delayed seat assignment, which may cause us to overestimate the remaining capacity.⁴ Seat maps may also conflate new bookings and cancellations. Specifically, if two seats were sold and one seat was canceled, our seat maps would yield a sale of a single seat. We rarely detect negative sales, which we reset to zero.

Notwithstanding its shortcomings, according to our knowledge, seat map data is the best publicly available disaggregated data on airline sales. As a matter of fact, airlines use web scraping to obtain competitive intelligence on schedules and fares. The pricing and revenue management team examines data from competitors’ websites and employs algorithms to determine if fares should be adjusted.⁵ Consequently, airlines’ and our measurement errors are likely to be comparable when evaluating the capacity of competitors.

An alternative data set used in the airline literature is the “Airline Origin and Destination Survey” (DB1B) collected by the U.S. Department of Transportation (DOT). The data

¹Internet-collected high-frequency price and sales data have previously been used for academic research, see McAfee and Te Velde (2006), Lazarev (2013), Escobari (2012), Clark and Vincent (2012), and Williams (2022).

²See www.amadeus.com for an overview of the Amadeus interactive seat map.

³JetBlue tends to reserve a constant number of seats as a function of the number of days till departure. The average number of reserved seats is 10.6 one day before departure and 11.5 fifty days in advance.

⁴Measurement error in capacity due to seat blocking and late assignment may cause us to underestimate the importance of revenue management. This is because supply-side pricing equations may suffer from attenuation bias, and because seat-blocking is a substitute to revenue management.

⁵See *Airline Management Magazine*, 2007, Issue 2, p. 18.

reports a random sample of 10% of all domestic airline tickets at the quarter-route level. This data set is unfortunately highly aggregated. For example, it does not report the date a ticket is priced or purchased or the date at which the flight departs. Therefore, DB1B would not be enough to investigate dynamic pricing.

3.1 Route selection

We follow a systematic route selection procedure similar to Williams (2022). To accomplish this, we utilize two data sources. First, we analyze the DB1B data collected for the first quarter of 2014 to obtain summary statistics for the majority of domestic airline routes in the United States. Second, we manually collect additional information from Google Flights. In particular, we searched for 17,392 randomly generated pairs of domestic airports and recorded the prices of the non-stop tickets. This search covers over 60% of all possible combinations among 237 major U.S. airports. Using DB1B and Google Flights data, we determine which routes match the following criteria.

Free seat assignment. We consider routes where consumers can select seats for free at the time of purchase. This selection allows us to minimize the measurement error in capacity caused by seats that were ticketed but not assigned. For example, we exclude Southwest, that does not allow any advance seat assignment and Spirit, that charges consumers for advance seat reservation. We keep the following airlines American, Delta, JetBlue, Alaska, United, Hawaii, and Virginia. We also exclude routes that offer “basic economy” fares, since these seats are typically not assigned in advance.

Pricing of round-trip tickets. We only consider routes where round-trip tickets are priced as two one-way tickets. Such pricing allows us to approximate round-trip tickets with two independent one-way tickets. Conveniently, this restriction applies to the majority of Alaska and JetBlue itineraries.

High non-connecting traffic. We select routes with at least 80% non-connecting traffic to avoid modeling network pricing.

Minimal number of daily flights. We are interested in the interaction between competition and dynamic pricing, but we aim to avoid computing equilibria with many active firms. As a result, we focus only on duopoly markets. In addition, to avoid the complexity of multi-product pricing, we select routes where each airline flies a single flight daily, with some exceptions during peak-demand departure dates.

The selection rules lead us to a final sample of five routes: Seattle-Tucson (SEA-TUS), San Diego-Boston (SAN-BOS), New York-Sarasota (NYC-SRQ), New York-Aguadilla (NYC-BQN), and New York-San Antonio (NYC-SAT). These routes are serviced by JetBlue, Delta, Alaskan, and United airlines, which are among the top six domestic carriers in the United States. We monitor all nonstop flights on the five routes March 17, 2016, to April 15, 2016. During this period, the dataset includes 57,077 observations of daily flight fares and inventory for 2,561 flights.

We aim to identify price elasticity using the exogenous variation in choice sets. For this purpose, we consider two routes: a control route in which there are no changes to the choice set and a treated route in which the choice set changes from duopoly to monopoly. We begin by identifying a treatment route. From December 2015 to March 2016, Alaska and Delta offer daily direct flights on the Seattle-Tucson route, making it a duopoly.

Alaska entered the Seattle-Tucson route in 2000 by offering year-round daily nonstop flights. Delta added a weekly nonstop flight from Seattle to Tucson International Airport in December 2014 as part of a significant service expansion from Seattle Airport. In the same year, Delta added Seattle as a hub, which created a “turf war” with Alaska, whose headquarters and central hub are in Seattle. Delta announced the Seattle-Tucson connection as seasonal, scheduled to operate between December 2015 and April 2016. Before March 31, consumers seeking nonstop connections could pick between Alaska and Delta; beginning in April, they could only fly Alaska. We refer to this transition from duopoly to monopoly as “exit,” although we acknowledge that the connection was seasonal, which was common knowledge. Our data includes departure dates preceding and following the exit.⁶

To adjust for any seasonality, we employ a control route, San Diego-Boston, whose average pricing trend closely parallels Seattle-Tucson before the exit. By combining the exit (treated) route with the control route, we can recover the price elasticity of demand using difference-in-differences style identification. JetBlue and Alaska are the relevant carriers on the San Diego-Boston route. JetBlue began service on this route in 2007, followed by Alaska in 2013.

As explained later, our identification assumption allows for common seasonality with the control route. Therefore, we do not permit demand discontinuities on March 31 that are exclusive to the Seattle-Tucson connection. To exclude confounding factors, we analyzed

⁶Conveniently, because Delta’s “exit” was anticipated and occurred more than once, Alaska did not have to learn how to price after the exit. Thus, learning and predatory pricing are not likely phenomena in our case.

supply at airports in Tucson and Seattle around March 31. During our sampling period, there were 23 daily nonstop flights to or from Tucson airport. Four of the 23 flights were seasonal, including Delta’s Seattle-Tucson connection, which ended in March 2015. The remaining 3 seasonal flights (Alaska’s Portland-Tucson, Delta’s Minneapolis-Tucson service, and Southwest’s Oakland-Tucson service) operated from November to June. Regarding the Seattle airport, Delta had 60 nonstop flights, 12 of which were seasonal. Two of the 12 seasonal flights ended in March, including the focal Seattle-Tucson connection. Four of the 12 flights concluded in April. The last six flights ended in August.

Seattle is the largest airport for Alaska, accounting for 18% of all of Alaska’s traffic in the first quarter of 2016. San Diego is the sixth largest airport for Alaska and accounts for 3% of Alaska’s traffic. Boston is the 26th largest airport for Alaska and accounts for 1% of Alaska’s traffic, while Tucson accounts for 0.4% of Alaska’s traffic. Boston is the largest airport for JetBlue, accounting for 13% of all of JetBlue’s traffic in the first quarter of 2016. On the other hand, San Diego only accounts for 0.8% of JetBlue’s traffic.

To avoid modeling network pricing, we keep the price of one-stop segments constant as we vary the price of direct routes. This assumption is plausible since the particular routes we study are relatively small and are unlikely to affect the prices of one-stop segments that contain much larger routes. As a result of this assumption, one-stop traffic enters an outside option in our choice model.

Using American Travel Survey,⁷ we compute the average share of self-reported business travelers from or to a given city pair (i.e., San Diego-Boston and Seattle-Tucson) as a measure for the share of business types (see [Borenstein, 2010](#)). The share of business types is 0.40 for San Diego-Boston and 0.30 for Seattle-Tucson. The average share of business types across all cities is 0.36, with a standard deviation of 0.10. Thus, the city pairs studied in the current paper are not outliers regarding the share of business travel.

3.2 Data description

Table 1 describes the research design with the Seattle-Tucson treated route and the Boston-San Diego control route. In particular, a market defined as a departure date for a particular city-pair and direction. We restrict the analysis to departure dates within a

⁷The survey was carried out by the Bureau of Transportation Statistics of the U.S. Department of Transportation in 1995.

30-day window centered on the exit event, reducing price endogeneity due to unobserved route-level demand shocks. In other words, we only use flights that depart from March 17, 2016, to April 15, 2016. We consider a pricing window of 49 days before departure, which results in a triple panel. Specifically, three untreated blocks comprise

$$\underbrace{15}_{\text{departure dates}} \quad \underbrace{4}_{\text{carrier-directions}} \quad \underbrace{49}_{\text{pricing dates}} = \underbrace{2,940}_{\text{observations}}$$

As described above, the untreated duopoly block has 2,940 observations. The treated block has 1,470 observations since it contains only one operating carrier. Overall, we obtain 10,290 observations on daily flight-level prices and sales. Table 2 reports the summary statistics of the data. On average, a flight sells 1.11 seats each day. For more than half of the days, a flight sells no tickets. As the departure date approaches, the price climbs on average by \$7.5 per day. An average flight sells 47.39 seats seven weeks before departure.

The flight-level Gini coefficient in the bottom row of Table 2 captures the intertemporal price dispersion for a given flight. The mean of the flight-level Gini coefficient equals 0.21. In other words, an expected absolute difference between two randomly selected prices for the same flight is 42%. The large dispersion of airfares is consistent with prior research.

Figure 2 shows the trajectory of pricing and available seats as the departure date approaches for a representative flight departing on March 25, 2016. The flight had more than 100 remaining seats 90 days before departure. On the departure day, seventeen tickets remained unsold. Prices range between \$150 and \$600. The shaded regions highlight pricing patterns suggestive of revenue management. The capacity sells out quickly in the 8-6 weeks to the departure area, so the algorithm is sharply increasing the price. Subsequently, sales flatten in the 6-4 region, resulting in a price decrease. Lastly, prices in the 2-0 region increase dramatically, which is consistent with price discrimination.

We use two sources of price variation when estimating price elasticity. As previously noted, we employ a 2 × 2 difference-in-differences panel data. The 2 × 2 design allows the exit decision to correlate with the aggregate demand trend shared with the control route. We also limit the dataset to a small time frame surrounding the exit event. The restriction allows us to leverage on locally exogenous jump in supply reminiscent of regression discontinuity. This design mitigates the concern about the endogeneity of the timing of the exit event.

Figure 3 depicts Alaska’s average prices for the 2 × 2 blocks. Each dot represents the

average price for flights departing on a given date. We average the prices across pricing dates and flight directions. Note that Alaska raises its price significantly in the treated route but not in the control route. We utilize the impact of this price increase on sales to estimate the slope of the demand curve.

Further, we zoom in on time trends around the exit event. Figure 4 displays kernel estimates of industry pricing paths. We observe that the control and treatment routes exhibit comparable time trends before the exit. Despite the treated route being significantly cheaper, the price differential remains steady. Similar to Figure 3, we notice a price increase on the treated route, whereas there is no discernible price increase on the control route.

Next, we examine the dynamic pricing as the departure date approaches. Figure 5 depicts the dynamic pricing paths for each route before and after the exit. The exit does not impact the control route's dynamic price path. In the treated route, prices increase for almost every pricing date. Exit does not affect extreme fares far from the departure date and those close to the departure date, which may be an artifact of the fixed pricing buckets (letter fares). Price rises significantly for clients who arrive 2 to 5 weeks before departure, indicating that firms compete most intensely for leisure customers. Our structural analysis will explain this outcome through the interaction of revenue management and competition.

Our research design allows us to test for substitution between adjacent departure dates and verify our market definition. In particular, we utilize the fact that the prices increase after the exit in the treated route. Therefore, if there is a meaningful substitution between the adjacent dates, we should see an increase in demand for dates before the exit in the treated market compared to the control market. To test this assertion, we compare sales (demeaned by a route average) for flights departing N days before exit between treated and control markets. We obtain p-values 0.21, 0.51, and 0.22 for N equal to 1, 2, and 3. We also obtain similar p-values when comparing transaction prices. Despite a significant post-exit price hike, we conclude that the demand for departure dates before the exit stays the same. This result suggests negligible between-day substitution.

In the next section, we present the structural model.

4 Model

In this section, we set up a dynamic duopoly model in which two firms compete to sell limited capacity while facing stochastic demand.

Fix a *market*, defined as a combination of a route and departure date. Consider two airlines $j \in \{1, 2\}$ endowed with an initial capacity of homogeneous seats denoted by c_j^0 . The seats can be sold in T selling periods representing days as the departure date approaches. In each period, the available capacity c_j^t is determined by the number of seats unsold in the previous period. The final capacity c_j^T occurs after the flight is closed, and it cannot be monetized.

The remainder of this section describes the consumer demand for airline seats, single-period profit function, value function, and equilibrium concept. We discuss the underlying assumptions and motivating empirical evidence when developing the model.

4.1 Consumer demand

Each period t brings I^t potential consumers with unit demand. Consumers may purchase airline seats or choose the outside option. Consumers are short-lived and myopic; they leave the market after choosing the outside option.⁸⁹

The number of potential consumers I^t is distributed as Poisson with an arrival rate λ^t . Since λ^t depends on t , the average number of arriving consumers depends on the number of days to the departure. The airlines know the distribution and arrival rates but do not observe the exact number of potential consumers I^t . This assumption is motivated by the nature of price setting in the airline industry. Before setting the price for a certain period, companies will likely know the average number of customers arriving. However, it is unlikely that they will know the actual number of consumers shopping for flights.

⁸There are several empirical reasons why presuming myopic customers may not be restrictive in our application. Primarily, the equilibrium price paths are increasing; hence, consumers have no reason to wait for prices to fall. Also, since high-valuation consumers arrive late, waiting and pooling with higher types may not be optimal. Nonetheless, some consumers may be delaying finalizing their trip arrangements. Those consumers may strategically wait, trading off lower uncertainty of their travel plans for higher prices. We leave this case open for further investigation.

⁹Another implication of assuming that consumers are short-lived is not allowing price-matching if the fare for the same departure date decreases in the future. Price-matching policies varied across carriers, ex., Alaska allowed while JetBlue did not allow it during our sample. Nevertheless, for our routes, the prices typically increase over time. In the rare instances when the price decreases, we would not allow consumers to use price matching.

Each consumer who arrives t days before departure can purchase a seat if capacity is available or choose an outside option. The outside option contains all other alternatives to flying on a particular date. The alternatives include other means of transportation, other travel dates, and alternative destinations. We abstract from modeling joint price setting for multiple departure dates and routes.

Consumers have random utility for flights, which follows [Berry et al. \(1995\)](#), henceforth BLP. Similarly to BLP, each consumer i observes the current prices denoted by p_j^t , idiosyncratic utility shock vector, represented by ϵ_i^t , and aggregate utility shock, denoted by η_j^t . Consumer's utility from choosing the flight j is given by

$$u_{ij}^t = \epsilon_i^t + \eta_j^t + \text{FE}_{\text{treatment}} + \text{FE}_{\text{weekend}} + \beta_j^t + \gamma_{ij}^t = \epsilon_i^t + \eta_j^t \quad (1)$$

We normalize the outside option by setting $u_{i0}^t = 0$. Parameter β_j^t is an airline-specific random coefficient, while γ_{ij}^t is a random price coefficient. Denote the joint distribution of ϵ_i^t and η_j^t , as $F(\cdot; t)$. The dependence of F on t allows the own- and cross-elasticity to depend on the number of days to the departure. For instance, since business customers may arrive late, our specification allows consumers to become inelastic and brand-loyal as the departure date approaches.

The aggregate utility shock η_j^t varies across departure dates, which captures varying preferences towards departure dates. The utility function also contains a fixed effect $\text{FE}_{\text{treatment}}$, which allows for a common demand shock in the period after Delta leaves the treated route. This parameter does not vary by route, which embodies the common trend assumption in a difference-in-differences regression. We also explicitly allow for a common shock to the demand for weekend flights represented by $\text{FE}_{\text{weekend}}$.

Further, we assume to make a standard assumption that ϵ_i^t is identically and independently distributed as a Type-1 extreme value. Thus, the probability that each arriving consumer purchases a seat from the company j is given by

$$s_j^t = \int \frac{\exp(\epsilon_i^t)}{1 + \sum_k \exp(\epsilon_{ik}^t)} dF(\cdot; t)$$

Since customers arrive as Poisson distribution, the quantity sold q_j^t for every firm j , conditional on s_j^t , follows a Poisson distribution with arrival rate s_j^t . Moreover, the Poisson

distributions are independent across firms, conditional on \mathcal{I}^t .¹⁰ These features are valuable for reducing the computation and making the state transition numerically tractable.

The next remainder of this section describes the dynamic pricing game.

4.2 Stage payoffs

Given prices $(p_j^t; p_{-j}^t)$ firm j obtains the following static payoff

$$\int_{q_j} q_j (p_j^t - mc_j^t) dF(q_j; p_j^t; p_{-j}^t; c^t; \mathcal{I}^t);$$

where q_j has a Poisson distribution with an arrival rate λ_j^t that is censored at c_j^t . Non-trivial probability mass on c_j^t represents the event of selling out the flight. Marginal cost is given by $mc_j^t = \pi c_j + \epsilon_j^t$; where ϵ_j^t is an idiosyncratic cost shock observable to the firms but not to the econometrician.

4.3 Timing

The dynamic pricing game has the following timing:

1. Capacity vector c^t , demand shocks \mathcal{I}^t and marginal cost shocks ϵ^t are revealed to all firms.
2. Firms simultaneously choose prices $(p_j^t; p_{-j}^t)$.
3. \mathcal{I}^t consumers arrive. Each consumer makes a choice determining quantity vector $(q_j^t; q_{-j}^t)$.
4. Firms receive stage payoffs. Next period capacity is determined according to

$$c_j^{t+1} = c_j^t - q_j^t;$$

Next, we define a strategy space and value function.

¹⁰The proof is straightforward and can be found in the Appendix ??.

4.4 Strategies and value function

The industry state is given by a triple $(c^t; t; !^t)$. We consider non-stationary Markov pricing strategies g_j . Specifically, the prices are contingent on all current period payoff-relevant state variables. That is, the number of days to departure t , and $(c^t; t; !^t)$. Formally,

$$g_j : \underbrace{T}_{\text{time}} \quad \underbrace{C^J}_{\text{capacities}} \quad \underbrace{\Xi^J}_{\text{demand shocks}} \quad \underbrace{\Omega^J}_{\text{cost shocks}} \quad ! \quad \underbrace{R^+}_{\text{current price}} :$$

For each strategy profile $(g_j; g_{-j})$ define a non-stationary value function

$$V_j^t(c^t; t; !^t; g_j; g_{-j}) = E \sum_{s=t}^T q_j^s (p_j^s - mc_j^s)$$

The expectation is according to a controlled Markov process over prices, quantities, and marginal costs. The value prescribes the cash-flow for firm j with capacity c_j^t , and competitors capacity c_{-j}^t . We impose a terminal condition $V_j^T = 0$.

4.5 Equilibrium

We consider Markov Perfect Equilibria $(g_j; g_{-j})$, such that

$$V_j^t(c^t; t; !^t; g_j; g_{-j}) = V_j^t(c^t; t; !^t; g_j; g_{-j}); \quad \partial g_j; c^t; t; !^t; j; t \quad T:$$

To solve this game, we use a standard transformation leveraging the independence of states t and $!^t$ over time. With abuse of notation, define an ex-ante value function as $V_j^t(c) = E_{!,t} V_j^t(c; !^t)$.

We omit subscript t to simplify the notation and denote $t+1$ value function by V^0 . In the MPE, ex-ante value functions must satisfy the following Bellman equation

$$V_j(c) = E_{!,t} \max_{p_j} \int_q [q_j (p_j - mc_j) + V_j^0(c^0)] dF(q; p_j; p_{-j}; c_{-j}) \quad (2)$$

Some discussion of the modeling choices is warranted. The model contains random consumer arrival q , demand shocks Ξ , and supply shocks $!$. Since we observe the integer number of seats sold, not market shares, the model has to generate distribution over finite quantities. Poisson distribution is the simplest way to obtain computationally tractable stochastic

integer quantities.

Demand shocks ϵ_{jt} capture the possibility of price endogeneity. Because the firms are likely to have more information about the demand than the econometrician, it is likely that the departure dates with higher demand, such as holidays, would receive higher prices. This possibility is captured by the dependence of prices on the vector ϵ_{jt} , which varies across routes, departure dates and pricing dates. Note that, without ϵ_{jt} , all cross-sectional price variation, after controlling for observables would be loaded on the supply shocks because firms do not observe the realization of the Poisson arrival process before setting prices. Thus, assuming away ϵ_{jt} would rule out demand-driven price endogeneity and may lead to underestimating price elasticities.¹¹

Conversely, η_{jt} represents cost shocks that firms observe before setting prices. We include these shocks in the model to allow for supply-driven price endogeneity. We estimate supply and demand as a system of equations, tracing the changes in the pricing equilibrium. Incorporating ϵ_{jt} and η_{jt} in the model allows for simultaneous shifts in demand and supply curves across observed pricing equilibria. Without η_{jt} the unexplained price variation would be driven solely by the demand shocks, leading to overestimating price elasticities.

The following section contains the details of the estimation procedure.

5 Estimation

This section contains the details of the estimation and identification of the model. We start by describing the parametric specification. In general, we allow the parameters to be different for each route but we fix them across departure dates. The variation in demand across departure dates is absorbed by two utility fixed effects, $FE_{\text{treatment}}$ and FE_{weekend} , as well as by demand shocks ϵ_{jt} .

¹¹Alternatively, we could allow firms to observe the realizations of the consumer arrival process to generate price endogeneity. We decided against it because it is unlikely that firms observe the exact number of potential customers. Also, current airline pricing algorithms do not take into account information on web or call center traffic. Another possibility would be to replace ϵ_{jt} with shock to the arrival rate λ_{jt} , that are observed to the firm but not to the econometrician. However, given available data, λ_{jt} is hard to distinguish from the level of utility; thus, shocks to the utility and shocks to the arrival rate generate numerically close purchase patterns. As a result, we can accommodate only one of the shocks. In such case, and we decided to follow the established literature, such as BLP, and include shocks in the utility function.

5.1 Parametric specification

We use a third-order polynomial to approximate the consumer Poisson arrival rate:

$$(t; \text{arrival}) = \sum_{n=0}^3 \binom{n}{\text{arrival}} \left(\frac{t}{7}\right)^n :$$

We consider a finite number of latent consumer types, following [Hendel and Nevo \(2013\)](#) and [Berry and Jia \(2010\)](#). The latent types aim at capturing changing elasticity of consumer demand as the departure date approaches and brand preference heterogeneity. To capture these two features we introduce 4 latent segments. The segments follow a 2-by-2 specification, allowing both vertical and horizontal differentiation. To capture the former, we allow for two latent vertical types that differ by price sensitivity. We introduce a high type that represents business travelers with price-insensitive utility. We also introduce a low type that represents leisure traveler with price-sensitive utility. The probability that a consumer who arrives at time t is low type is approximated by a third-degree polynomial

$$\Pr_L(t; \text{vertical}) = \frac{1}{1 + \exp\left[\sum_{n=0}^3 \binom{n}{\text{vertical}} \left(\frac{t}{7}\right)^n\right]} :$$

The probability that is a high-type consumer is

$$\Pr_H(t; \text{vertical}) = 1 - \Pr_L(t; \text{vertical}) :$$

The logit transformation is chosen to bound probabilities in $[0; 1]$.

We allow for two horizontal types with potentially different brand preferences.¹² Horizontal types differ by their utility intercepts. Consider a vertical type H . The conditional probability that this consumer prefers firm 1 is given by $\frac{H,1}{\text{horizontal}}$. Conversely, the consumer prefers firm 2 with probability $\frac{H,2}{\text{horizontal}} = 1 - \frac{H,1}{\text{horizontal}}$. Similarly, consumer of low vertical type prefers firm 1 with probability $\frac{L,1}{\text{horizontal}}$, and firm 2 with probability $\frac{L,2}{\text{horizontal}} = 1 - \frac{L,1}{\text{horizontal}}$.

According to the conditional probability specification, the distribution of 4 latent types is given by the product of the time-varying probability distribution of the vertical types and

¹²We cannot identify brand preferences from preferences on other persistent product attributes. We do not distinguish these in the discussion.

the corresponding conditional probability distribution of horizontal types. In particular, at time t , high-type consumers leaning towards firm j arrive with probability

$$\Pr_H(t; \text{vertical}) = \frac{H_j}{\text{horizontal}}.$$

Similarly, low-type consumers leaning towards firm j arrive with a probability

$$\Pr_L(t; \text{vertical}) = \frac{L_j}{\text{horizontal}}.$$

We estimate two price coefficients β^H and β^L for high and low types, respectively. We also estimate 4 intercept vectors $\alpha^{H,1}$, $\alpha^{H,2}$, $\alpha^{L,1}$, and $\alpha^{L,2}$ that vary across both vertical and horizontal types. Each of the intercept vectors is a 2×1 vector. We normalize one of the intercepts in $\alpha^{L,2}$ to 0. This specification allows brand loyalty to differ across vertical types and firms. The such specification aims to capture differences in the frequent flyer programs across firms. We need to estimate 4 parameters of the vertical type arrival polynomial, 2 conditional probabilities for horizontal types, 2 price coefficients, and 7 intercepts.

We observe discrete realizations of quantities, not market shares; thus, BLP market share inversion is not applicable. Instead, we postulate parametric distributions of demand shocks and supply shocks ϵ . In particular, we assume shocks are distributed as normal random variables with mean 0 and variances σ^2 and τ^2 . The such specification allows us to integrate over ϵ instead of inverting the demand equations.

As explained in the following subsection, we estimate supply and demand jointly to be able to identify a flexible time-varying distribution of customer heterogeneity $\Pr_L(t; \text{vertical})$. This object is crucial for our conclusions regarding price discrimination and revenue management. We considered modeling shocks non-parametrically and adapting an inversion method that allows for a measurement error in market shares. In our case, such an approach is only practically possible if demand is not estimated jointly with the dynamic game. Consequently, to sustain flexible time-varying distribution of customer heterogeneity, we decided to focus on estimating the dynamic game and assume parametric distribution of the shocks.

5.2 Generalized method of moments

Demand and supply are estimated jointly using a Nested-fixed-point Simulated Generalized Method of Moments (GMM). There are 3 error terms in our model: consumer arrival, demand shocks and supply shocks. All these error terms are “structural,” that is they have an economic interpretation. In particular, we do not rely on “measurement error,” or “trembling hand” error terms.

Our data is a triple panel consisting of routes, departure dates, and pricing dates (number of dates to the departure). Let r denote a combination of route and departure date. As in the previous section, we denote pricing date as t . Following Rust (1987), the estimation routine consists of two loops. In the inner loop, we solve an MPE of the dynamic game fixing demand and supply parameters, θ . Using the equilibrium strategies, we simulate industry paths

$$c^{r;t}(\cdot); p^{r;t}(\cdot); q^{r;t}(\cdot); \theta; r; t \quad T$$

Let $Z^{r;t}$ be the set of instruments. We consider the following 5 groups of moments (all variables have $r; t$ superscripts, which we omit to simplify the notation)

1. Price and sales (4 moments per instrument):

$$E[p_j z] \quad \text{and} \quad E[q_j z]; \quad \theta_j \geq 1; 2$$

2. Price and sales for weekend departure dates (4 moments per instrument):

$$E[p_j \mathbb{I}_{\text{weekend}} z] \quad \text{and} \quad E[q_j \mathbb{I}_{\text{weekend}} z]; \quad \theta_j \geq 1; 2$$

3. Price times sales (4 moments per instrument):

$$E[p_j q_j z] \quad \text{and} \quad E[p_j q_{-j} z]; \quad \theta_j \geq 1; 2$$

4. Price times capacity (4 moments per instrument):

$$E[p_j c_j z] \quad \text{and} \quad E[p_j c_{-j} z]; \quad \theta_j \geq 1; 2$$

5. Sales times capacity (4 moments per instrument):

$$E[q_j c_j z] \quad \text{and} \quad E[q_j c_{-j} z]; \quad \forall j \geq 1; 2$$

Summing the above, we consider 20 groups of moments. The instruments follow a diff-in-diffs specification and consist of 4 dummy variables indicating

1. Route with exit (treated), departure date before exit
2. Route with exit (treated), departure date after exit
3. Route without exit (not treated), departure date before exit
4. Route without exit (not treated), departure date after exit

We interact the above dummies with 7 dummies indicating the number of weeks to the departure. As a result, we consider 28 instruments that indicate a combination treatment, departure date, and the number of weeks to the departure. All in all, we obtain $20 \times 28 = 560$ moments to estimate 45 parameters. Moments with zero empirical variance are removed resulting in 469 moments.¹³

Denote the vector of empirical moments as $\hat{\mu}$, and a vector of simulated moments as $\mu(\theta)$. Consider a weighting matrix W . The outer loop minimizes the following objective function

$$f(\theta) = \left[\hat{\mu}(\theta) - \mu(\theta) \right]' W \left[\hat{\mu}(\theta) - \mu(\theta) \right]:$$

We conduct two-step GMM. In the first step, we estimate a model with an identity weighting matrix. Using these estimates we compute a consistent estimator of an optimal weighting matrix denoted by \hat{W} . In the second step, we re-estimate the model using \hat{W} . Standard errors are obtained using a sandwich formula for an optimal weighting matrix.

5.3 Numerical considerations

The computational burden of the estimation depends on the computation speed of the MPE for a given θ . For each θ the equilibrium needs to be computed for every route,

¹³Note that for the route with exit and departure date after exit, the number of baseline moments is only 7 instead of 20. This condition generates $1 \times 7 \times 7 = 49$ moments whereas other 3 conditions generate $3 \times 20 \times 7 = 420$ moments.

separately for weekend/weekday departure dates, and for the departure dates that were treated. Because of this it is essential to employ a set of numerical techniques that accelerate the MPE computation.

To solve the Bellman equation (2) we employ a finite horizon version of Pakes et al. (1994) algorithm. In particular, we solve the game by backward induction. We start by setting $V^T = 0$. Then we solve a 1-period pricing game by computing V^{T-1} using the Bellman equation. We continue until we reach $t = 0$.

There are 3 main difficulties that complicate solving the Bellman equation. The equation contains a series of nested routines. We discuss these routines starting from the most inner computational layer. First, for every c and pair (i, j) , we need to solve a pricing game. We solve the pricing games using best response dynamics.

The second computational layer is the integration with respect to shocks (i, j) . The distribution of the shocks is a bi-variate normal. To perform the integration we use Gauss-Hermite quadrature with 3^4 grid points. As a result, to obtain an ex-ante value function for a fixed c , we need to solve 81 pricing games.

We have to perform backward induction by iterating on the Bellman equation. This step is relatively difficult because it has to be performed for each c and inner maximization and integration layers and computationally taxing. To lower the computational burden, instead of performing an exact value iteration, we perform parametric value iteration as in Judd (1998). In particular, we choose a 14×14 grid (subset) of feasible capacities and solve the equation (2) only on this grid, instead of for every $c < c^0$. This step is parallelized using a multi-processor server. We obtain V outside of the grid as necessary using cubic spline interpolation.

5.4 Identification

As mentioned earlier, the demand and supply is estimated jointly as a system of equations. In such case, the identification of demand elasticity requires exclusion restriction; that is, covariates that enter supply equations but are excluded from the demand equation. In our case, the number of players enters the supply equilibrium conditions, but is assumed to be uncorrelated with idiosyncratic demand shocks before and after the exit event on March 31st. Consequently, the price variation induced by the altered supply-side equations would be uncorrelated with $\ln p$ and would identify the price coefficients. Matching price

and sales moments interacted with treatment and control would result in an estimate of the elasticity. The intuition follows Figure 3, and matching corresponding difference-in-differences in quantity sold for the treated and control routes.

Following the BLP, the marginal cost is identified using the supply-side optimality conditions induced by the supply-side equation. That is, the estimation would yield estimates of marginal costs that rationalize prices observed in the data. In our case, the supply side equations cannot be easily inverted. Instead, we match unconditional price and quantity moments.

As mentioned earlier the identification of the consumer arrival rate from the demand intercept requires a normalization and we normalize one of the β 's to 0. The composition of consumer types as the departure date approaches is identified from heterogeneous impact of the exit event on the price and quantity paths. Previously discussed Figure 5 demonstrated heterogeneous impact of the exit event for early and later arriving customers. The price gap between treated and untreated market is a result of a complicated interplay of revenue management and price discrimination. We rely on our structural model to identify which pattern of consumer types replicates the curvature of the pricing gap between treated and control markets. Additionally, we rely on the panel structure of our data. Namely, the estimation assumes that β is IID across departure dates for a given group. This allows us to use remaining capacity as a price shifter by imposing $\rho_j c_j$ and $q_j c_j$ moments. This variation aids identification of changing customer composition on the route level.¹⁴

6 Results

In this section, we discuss the estimation results from the structural model. Table 3 reports the estimates from the structural model. In particular, Panel A of Table 3 shows the parameters of consumers' preference (β 's) in the exit and control routes. Panel B summarizes the parameters of consumers' arrival rates ($\beta_{arrival}$) and the parameters of the distributions of consumer types ($\beta_{horizontal}$) for the two routes. Panel C shows the estimates of demand shifters that vary by departure date, including "treatment" fixed effect (applicable to all

¹⁴Because we interact $\rho_j c_j$ and $q_j c_j$ with the two diff-in-diffs dummies, identification of a model with serially correlated β 's is theoretically possible. Nevertheless, because estimation of the model with persistent unobserved heterogeneity is numerically challenging, we leave it for further research. We speculate that persistent unobserved heterogeneity may be important in other markets with higher proportion of forward-looking consumers.

dates following the exit in both markets), weekend fixed effect and variance of the demand shocks . Panel D contains the estimates of the marginal cost.

The estimates are precise and significant at the 1% level, suggesting that the data parametrically identify the model. The estimates are intuitive overall. For instance, the price coefficient of consumers who arrive early is approximately three times that of those who arrive late. This discrepancy suggests that buyers who arrive later are “high-types” who are less price elastic.

Arrival rate parameters, $\lambda_{\text{arrival}}^{(n)}$, and joint distribution of types prescribed by $\lambda_{\text{vertical}}^{(n)}$ and $\lambda_{\text{horizontal}}$ determine the type-specific Poisson arrival rates, which are easier to interpret than raw parameters. Figure 6 visualizes the Poisson arrival rates implied by the estimated parameters and disaggregated by the consumer type. Notably, the absolute value of the arrival rate on the Y-axis is the consequence of a specific normalization of the low-type intercept in the utility function; hence, it should not be interpreted in isolation. Nevertheless, the relative arrival rates are identified.

The upper portion of Figure 6 depicts arrival rates for the exit route, whereas the lower portion depicts the rate for the control route. The consumer arrival pattern is comparable for both routes, which is to be expected given that both routes have the same time trend prior to exit. In both routes, the share of the high-type travelers increases as the departure date approaches. In the exit route, the proportion of high-type travelers is 7% seven weeks before departure and gradually rises to nearly 80% one week before departure. On the final day of ticket sales, around 87% of arriving passengers are high-type passengers.

Alaska has a relatively larger high-type segment in the exit route than in the control route. The difference is perhaps due to Alaska’s longer presence in the exit route than the control route. Longer presence facilitates building brand equity, which may result in higher brand loyalty for Alaska in the exit route.

Figure 6 also summarizes the relative arrival rates of consumers leaning towards a particular firm. In both routes, Alaska has more loyal consumers than JetBlue and Delta. Specifically, in the treated route, 88% of high-type and 84% of low-type consumers prefer Alaska to Delta. In the control market, 74% of high-type consumers and 62% of low-type consumers prefer Alaska to JetBlue.

Figure 7 shows own price elasticity and industry elasticity for the treated route. To highlight the variation in customer preferences across firms and over time, we evaluated

all elasticities at the constant average industry price. Consistently with the estimates of arrival rates, consumers become less price elastic as the departure date approaches. The estimated values are in line with prior research. [Gillen et al. \(2003\)](#) suggests that airline travel demands elasticity ranging from 0.181 – 2.01 with a median of 1.3 for a sample of 85 city pairs. Previous literature also points out that demand for long-haul routes is less elastic. The median own demand elasticity of a long-haul domestic leisure consumer is 1.228. Similarly, our estimates range from 0.4 – 2.2.

We also find that early on, Alaska consumers are more price inelastic than JetBlue consumers, but as the departure date approaches, the elasticities of Alaska and JetBlue converge. The minimal gap in elasticities in the late market suggests that even under the presence of disruptive low-cost carriers, such as JetBlue, the airlines retain market power as the departure date approaches.

Per-consumer “peanut costs” range between \$24 and \$32. These numbers include the monetary costs of serving an extra passenger and an opportunity cost of selling the seat as part of a multi-leg reservation (recall that our model only considers direct routes). Notably, the level of the marginal costs is consistent with the estimates provided by [Link et al. \(2009\)](#).

Next, we explore the model fit and focus our discussion on predicting the pricing variation before and after the exit, and the curvature of the price and quantity paths as the departure date approaches. [Figure 8](#) shows trajectories in the treated market before and following the intervention. The model replicates the level of prices before and after the exit and the curvature of the price path. Specifically, we successfully anticipate that competitive price trajectories will be concave, whereas monopoly price paths will be almost linear. Similarly, [Figure 9](#) compares the observed and model-predicted curvature of sales and prices for all observations as a function of the number of days till departure. As the departure date nears, we anticipate that sales will fall until week -2. We also precisely predicted a substantial increase in ticket sales in the last week.

7 Counterfactuals

In this section, we conduct counterfactual exercises utilizing the model estimates. The objective is to isolate the influence of dynamic pricing on welfare and firm profitability. Moreover, we aim to differentiate the effects of price discrimination and revenue management

and emphasize their relationship with market competition. We concentrate on the duopoly market of San Diego-Boston. To minimize computation, we set the demand errors and supply errors ϵ to zero.

We begin by generating a counterfactual in which dynamic pricing is ruled out. For this purpose, we compute a Bertrand-Nash equilibrium in which companies commit to a uniform price as the departure date approaches (henceforth, UPNE). In other words, we compute a price pair $p^{UPNE} = (p_1^{UPNE}; p_2^{UPNE})$, such that

$$p_j^{UPNE} \geq \operatorname{argmax}_{p_j} E \sum_{t=0}^{T-1} q_j^t (p_j^{UPNE} - mc_j^t) :$$

Note that firms choose prices to optimize the value function at $t = 0$ and that the entire quantity distribution over time is defined by the uniform pricing pair p^{UPNE} .

Denote the UPNE value function as $V_j^{UPNE}(c)$. The ranking of $V_j^{DP}(c)$ and $V_j^{UPNE}(c)$ is an empirical question, as any of the businesses may theoretically be worse off in equilibrium if they are permitted to use dynamic pricing. Uniform pricing is a specific example of dynamic pricing, hence in a monopoly market, $V_j^{DP}(c)$ is always bigger than $V_j^{UPNE}(c)$. Notably, the ranking of consumer surplus between uniform pricing and dynamic pricing is ambiguous regardless of the market structure.

UPNE has two implications. First, the firms cannot price discriminate between early-arriving elastic and late-arriving inelastic consumers. Therefore, when we permit dynamic pricing, which enables 3rd-degree price discrimination, the inelastic segment pays higher prices than UPNE, while the elastic segment pays lower prices. Figure 10 compares the pricing routes under uniform pricing (UPNE) and dynamic pricing (DPMPE), both of which were produced from the model. We observe that the uniform price is substantially higher for early-arriving customers and lower for late-arriving customers.¹⁵

Because it restricts revenue management, UPNE may lead to sub-optimal capacity allocation. Figure 11 compares the odds of monopolization in which at least one company sells out. Under UPNE, flights sell out earlier than under DPMPE. In our context, selling out early is inefficient because customers who arrive later have a higher willingness to pay.

The difference between rows (1) and (3) of Table 4 demonstrates the effect of dynamic

¹⁵Due to market share weighting, the UPNE pricing path increases as the departure date approaches. Alaska charges higher prices, and its market share increases.

pricing on the industry’s fundamentals. We find that under DPMPE, consumer welfare increases by 3%, industry profits increase by 8%, and total welfare increases by 6%, from \$69,783 to \$73,641. Permitting firms to implement dynamic pricing is a Pareto improvement since it increases customer welfare and profits. Nevertheless, firms benefit relatively more than consumers. Furthermore, we discover that dynamic pricing boosts industry output and decreases average prices. Notably, Alaska (a higher-priced airline) increases quantity while cutting prices, but JetBlue (a lower-price airline) raises both quantity and rates. This disparity suggests JetBlue relies more on revenue management while Alaska relies more on price discrimination. Alaska profits more from DPMPE than JetBlue, indicating that price discrimination may be more profitable in the equilibrium than revenue management. Later in this section, we examine this hypothesis by formally decomposing dynamic pricing into price discrimination and revenue management.

Columns (1) and (3) of Table 5 contain the impact of dynamic pricing on specific customer segments. We observe less than a 2% consumer welfare drop for inelastic consumers and more than a 100% increase in the welfare of the elastic segment. Thus, dynamic pricing leads to the reallocation of surplus from late-arriving to early-arriving customers.

Under DPMPE, The overall industry output increases by more than 15%. The concurrent increase in output and consumer welfare is reminiscent of Holmes (1989), who shows that increasing industry output is a sufficient condition for welfare increase under competition with 3rd-degree price discrimination. In our case, the change in output is a composition of price discrimination and better capacity utilization due to revenue management.

The rise in consumer surplus under DPMPE highlights the relationship between competition and dynamic pricing. According to Williams (2022), in monopoly markets, dynamic pricing produces less consumer surplus than uniform pricing. We discover the opposite. Further, to maintain internal validity and provide apples-to-apples comparison with monopolized market, we consider an acquisition of Alaska by JetBlue. Specifically, we recompute the equilibria for the monopoly firm with a joint capacity of both airlines and assign JetBlue’s preferences to all customers.¹⁶ Columns (1) and (3) in Table 6 contain the comparison of uniform pricing and Dynamic pricing by a monopolist. We replicate the results of Williams

¹⁶We also considered an acquisition of JetBlue by Alaska and obtained similar results. Nevertheless, the Alaska monopoly case is somewhat less interesting from the pricing perspective because the capacity is not binding in all pricing counterfactuals. Lack of meaningful capacity constraint occurs because Alaska monopolist sells predominantly to high types.

(2022) by showing that dynamic pricing in a monopoly airline market leads to greater profits but lower consumer surplus. We also discover that dynamic pricing by a single firm lowers total surplus.

The next step is decomposing dynamic pricing into price discrimination and revenue management. In particular, we segment the demand based on arrival time to obtain the intermediate case of 3rd-degree price discrimination without revenue management. We allow the firms to set different prices depending on the number of days to the departure, but not based on capacity. We compute a Nash equilibrium in which firms choose a fixed price path $p_j^{PDNE} = (p_j^{PDNE;0}, \dots, p_j^{PDNE;T-1})$ to maximize their value function at $t = 0$. When doing so, they keep the competitor's price path p_j^{PDNE} fixed. Formally,

$$p_j^{PDNE} \in \operatorname{argmax}_{p_j^0, \dots, p_j^{T-1}} E_{p^{PDNE}} \sum_{t=0}^{T-1} q_j^t (p_j^{PDNE;t} - mc_j^t) :$$

Compared to the DPMPE, we do not allow the prices to respond to the currently available capacities, c^t . Denote the intermediate equilibrium payoff by $V_j^{PDNE}(c)$. Similarly to the case of uniform and dynamic pricing, there is no ranking of $V^{PDNE}(c)$, $V^{UPNE}(c)$, and $V^{DPMPE}(c)$ for the oligopolistic competition. For the monopoly, $V^{PDNE}(c)$ is ranked in between of $V^{UPNE}(c)$ and $V^{DPMPE}(c)$. In other words, the monopolist can never do worse by price discriminating and can never do worse by adding revenue management on top of price discrimination.

Columns (1) and (2) in Table 4 contrast uniform pricing with price discrimination. Under price discrimination, consumers pay, on average, higher prices, and the total quantity supplied drops. Due to the capacity constraint, the theoretical results of (Holmes, 1989) cannot be directly applied to airlines; nonetheless, we find that reduced industry output under price discrimination co-occurs with a decline in consumer welfare. We also observe deadweight loss manifesting as a 6% decrease in total welfare. Table 5 displays the impact of price discrimination on welfare for each consumer segment. Comparing Columns (1) and (2) demonstrates the redistribution of welfare from late-arriving business consumers to early-arrival leisure travelers. Figure 12 compares the price trajectories under uniform pricing (UPNE) and price discrimination (PDNE) and confirms that prices in the early market decrease while prices in the late market sharply increase.

The impact of price discrimination is qualitatively similar to the monopoly case. For

example, Williams (2022) documents a drop in consumer surplus and an increase in profits after introducing price discrimination. Similarly, in Table 6, we observe a more significant drop in consumer surplus and more considerable incremental deadweight loss stemming from price discrimination in the monopoly compared to the oligopoly market. Even if competition mitigates the negative effects of price discrimination, it is insufficient to prevent deadweight loss.

The deadweight loss under PDNE is related to the interplay of capacity constraint and the degree of competition in the early and late markets. Under PDNE must commit to prices regardless of the realization of stochastic demand. Because of commitment, firms are incentivized to hedge and set higher prices early to save seats for late-arriving high-value customers.¹⁷ Such hedging generates upward pricing pressure, which, compared to price discrimination without capacity constraint, should reduce competition in an early market. However, despite the incentives to hedge, early consumers pay lower prices in PDNE compared to UPNE. According to Figure 11, lower prices in the early market result in a higher probability of a monopoly in the later market under PDNE than under UPNE.¹⁸ In other words, price discrimination softens competition for late-arriving customers via capacity pre-commitment. As a result, as shown in Table 5, late-arriving consumers' welfare decreases by more than 10%, which is the primary driver of deadweight loss under PDNE.

Next, we describe the impact of introducing revenue management on top of price discrimination. For this purpose, we compare PDNE with DPMPE. Figure 10 demonstrates that PDNE and DPMPE deliver differential pricing in the early and late markets; however, DPMPE leads to a nearly uniform price decrease. Moreover, according to Figure 13, DPMPE results in uniformly larger industry sales. Revenue Management decreases early market hedging since the firms can secure the capacity for high-value customers by increasing the price only when the demand is high. As a result, revenue management creates downward pricing pressure in the early market. Moreover, since the capacity utilization under DPMPE is more efficient, more residual capacity exists in the late market, creating further downward pricing

¹⁷For computational reasons, our model does not allow the firms to limit the number of seats on the market, a practice known as seat release policy (see Aryal et al. (2022)). We show that firms set higher prices early to save capacity for later periods. A seat release policy may be an alternative. Thus, assuming a seat release policy away may mute some of our results regarding price discrimination. Since price discrimination benefits the firms and hurts consumers, we speculate that dynamic pricing may lead to lower firms profits and greater consumer surplus.

¹⁸Notably, PDNE delivers a lower probability of monopolization in the last week when monetizing the inelastic demand.

pressure. Indeed, Figure 11 shows that revenue management is highly effective at preventing monopolization of the late market, beating not only PDNE but also UPNE.

Comparing Columns (2) and (3) in Table 4 confirms that revenue management leads to simultaneous increase in sales and decrease prices. Notably, after allowing revenue management, the industry profits decrease. Lower prices and industry profits contrast with the monopoly case. For example, Williams (2022) does not find a meaningful impact of revenue management on average price and documents an increase in the industry profits. Similarly, when comparing Columns (2) and (3) in Table 6, we find the negligible impact of Revenue Management on the monopoly market. The difference between monopoly and duopoly markets suggests that downward pricing pressure under revenue management interacts with market structure leading to fiercer competition. Comparing columns (2) and (3) of Table 5, we find that both elastic and inelastic customers benefit from revenue management. In other words, loosening the capacity constraint under revenue management leads the firms to compete away a portion of the rents from price discrimination and reduce deadweight loss. Interestingly, because of the reduced waste from unsold seats, firms can obtain profits comparable to those under price discrimination and simultaneously deliver a larger surplus to their customers.

We have examined the limits of our results when we alter the tightness of the capacity constraint. If we remove 25% of capacity we obtain 6% increase in the total welfare compared to 5% increase with the original capacity. Conversely, if we increase the capacity by 25% we obtain only 3.9% increase in total welfare. Thus, the tighter the capacity constraint the larger the relative impact of dynamic pricing on market efficiency.

Overall, revenue management (and dynamic pricing in general) requires a certain degree of competition to substantially increase market efficiency. The clues about the mechanism are contained in Lazarev (2019), who demonstrates that commitment may generate market power in competitive airline markets. Because revenue management removes some commitment power, rents are transferred to consumers. Dynamic pricing, as a combination of price discrimination and revenue management, becomes a win-win for consumers and firms if the market is competitive.

8 Conclusion

Sophisticated pricing strategies are becoming more accessible to businesses as a result of advances in information technology, enabling for prices to be adjusted in near real-time. As a result, price competition becomes increasingly dynamic. This study conducts an event study examining the consequence of dynamic pricing in an oligopolistic airline market. It delivers two main findings. First, inter-temporal price discrimination softens competition in the late market. It shifts rents from business to leisure travelers and allows airlines to extract more consumer surplus via capacity pre-commitment, reminiscent of [Kreps and Scheinkman \(1983\)](#). Second, the ability to smooth demand fluctuations (revenue management) boosts the efficiency of capacity utilization and undermines capacity pre-commitment. Removal of commitment leads to transferring some benefits from sophisticated pricing from airlines to consumers. Overall dynamic pricing is a Pareto improvement under a sufficient degree of competition.

The work on perishable good dynamic oligopoly is relatively scant compared to the durable good literature.¹⁹ The empirical work in this area faces the scarcity of disaggregated data, identification challenges caused by price endogeneity, and the complexity of solving dynamic oligopoly models. Under these hurdles, this work prioritizes internal over external validity by considering a slice of the airline industry, where robust identification of consumer preferences and their heterogeneity is feasible. With the increased availability of high-frequency data, applying our techniques to a more extensive set of markets is a promising area for further research.

Finally, we recognize a growing theoretical literature on selling perishable goods to forward-looking consumers. While this feature is unlikely to be pivotal in our setting, it may be a first-order issue in other airline markets. For instance, consumers may be more strategic on tourist routes with regular promotions and last-minute discounts. Considering forward-looking consumers is a possible direction for extending this research if the data on consumer waiting is available.

¹⁹In their theory work titled *Dynamic pricing of perishable assets under competition*, [Gallego and Hu \(2014\)](#) comment: "... The literature on competitive RM [revenue management] is scant. This is partly due to the challenges imposed by this complex game [...] not to mention even more thorny problems of time-varying demand structures."

References

- Aryal, G., C. Murry, and J. W. Williams (2022). Price discrimination in international airline markets. *arXiv preprint arXiv:2102.05751*.
- Berry, S., M. Carnall, and P. T. Spiller (2006). Airline hubs: costs, markups and the implications of customer heterogeneity. *Competition policy and antitrust*.
- Berry, S. and P. Jia (2010). Tracing the woes: An empirical analysis of the airline industry. *American Economic Journal: Microeconomics* 2(3), 1–43.
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica*, 841–890.
- Berry, S. T. (1990). Airport presence as product differentiation. *The American Economic Review* 80(2), 394–399.
- Borenstein, S. (1992). The evolution of us airline competition. *The Journal of Economic Perspectives* 6(2), 45–73.
- Borenstein, S. (2010). An index of inter-city business travel for use in domestic airline competition analysis. *NBER working paper*.
- Borenstein, S. and N. L. Rose (1994). Competition and price dispersion in the us airline industry. *Journal of Political Economy* 102(4), 653–683.
- Borenstein, S. and N. L. Rose (2014). How airline markets work... or do they? regulatory reform in the airline industry. In *Economic regulation and its reform: What have we learned?*, pp. 63–135. University of Chicago Press.
- Chen, J., S. Esteban, and M. Shum (2013). When do secondary markets harm firms? *American Economic Review* 103(7), 2911–34.
- Clark, R. and N. Vincent (2012). Capacity-contingent pricing and competition in the airline industry. *Journal of Air Transport Management* 24, 7–11.
- Dana Jr, J. D. and K. R. Williams (2022). Intertemporal price discrimination in sequential quantity-price games. *Marketing Science*.
- Deneckere, R. and J. Peck (2012). Dynamic competition with random demand and costless search: A theory of price posting. *Econometrica* 80(3), 1185–1247.
- Ericson, R. and A. Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of economic studies* 62(1), 53–82.
- Escobari, D. (2012). Dynamic pricing, advance sales and aggregate demand learning in airlines. *The Journal of Industrial Economics* 60(4), 697–724.
- Gallego, G. and M. Hu (2014). Dynamic pricing of perishable assets under competition. *Management Science* 60(5), 1241–1259.
- Gallego, G. and G. Van Ryzin (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science* 40(8), 999–1020.
- Gillen, D. W., W. G. Morrison, and C. Stewart (2003). Air travel demand elasticities: Concepts, issues and measurement. *Department of Finance, Government of Canada*.
- Goettler, R. L. and B. R. Gordon (2011). Does amd spur intel to innovate more? *Journal of Political Economy* 119(6), 1141–1200.
- 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.
- Hendel, I. and A. Nevo (2013). Intertemporal price discrimination in storable goods markets. *The American Economic Review* 103(7), 2722–2751.
- Holmes, T. J. (1989). The effects of third-degree price discrimination in oligopoly. *The American Economic Review* 79(1), 244–250.
- Hortaçsu, A., A. Öry, and K. R. Williams (2022). Dynamic price competition: Theory and evidence from airline markets.
- Judd, K. L. (1998). *Numerical methods in economics*. MIT press.
- Kreps, D. M. and J. A. Scheinkman (1983). Quantity precommitment and bertrand competition yield cournot outcomes. *The Bell Journal of Economics*, 326–337.
- Lazarev, J. (2013). The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in us monopoly markets. Technical report, Working paper.
- Lazarev, J. (2019). Getting more from less: understanding airline pricing. Technical report, Working paper.
- Levin, Y., J. McGill, and M. Nediak (2009). Dynamic pricing in the presence of strategic consumers and oligopolistic competition. *Management Science* 55(1), 32–46.
- Lin, K. Y. and S. Y. Sibdari (2009). Dynamic price competition with discrete customer choices. *European Journal of Operational Research* 197(3), 969–980.
- Link, H., W. Götze, and V. Himanen (2009). Estimating the marginal costs of airport operation using multivariate time series models with correlated error terms. *Journal of Air Transport Management* 15(1), 41–46.
- McAfee, R. P. and V. Te Velde (2006). Dynamic pricing in the airline industry. *forthcoming in Handbook on Economics and Information Systems, Ed: TJ Hendershott, Elsevier*.
- Pakes, A., P. McGuire, et al. (1994). Computing markov-perfect nash equilibria: Numerical implications of a dynamic differentiated product model. *RAND Journal of Economics* 25(4), 555–589.
- Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 999–1033.
- 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.
- Sweeting, A. (2015). A model of non-stationary dynamic price competition with an application to platform design.
- Williams, K. R. (2022). The welfare effects of dynamic pricing: Evidence from airline markets. *Econometrica* 90(2), 831–858.

A Figures & Tables

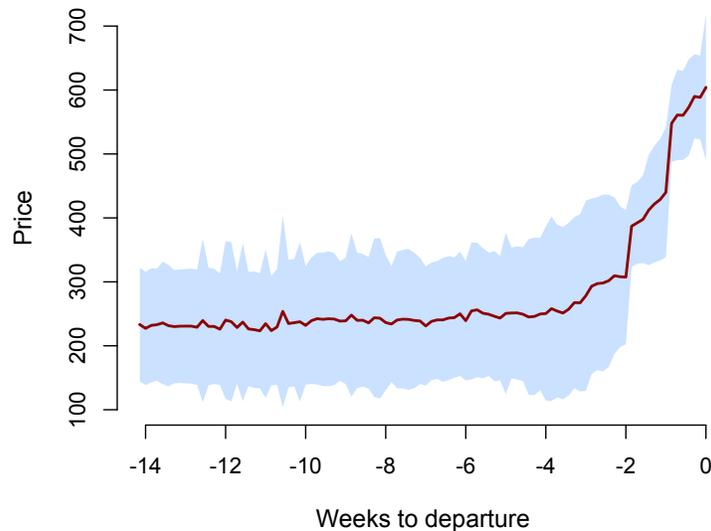


Figure 1: Example price path. This graph shows the average price path for JetBlue Flight 19 from March 01, 2016, to June 01, 2016. The red line shows the average price conditional on the number of days to departure. The light-blue area shows one standard deviation of prices.

	03/17 — 03/31	04/01 — 04/15
Exit route Seattle-Tucson	Delta Alaska N=2940	Alaska N=1470
Control route Boston-San Diego	Alaska JetBlue N=2940	Alaska JetBlue N=2940

Note: the number of observations in each cell is calculated by $N = \frac{\# \text{ rms} \# \text{ pricing-dates}}{\# \text{ departure-dates}}$.

Table 1: Research design. The table describes the difference in differences research design. The Seattle-Tucson route experienced the exit of Delta on 03/31/2016.

Statistic	Mean	Median	St. Dev.
Observations (N=10,290)			
Price (\$)	310.51	268.00	154.63
Number of seats sold daily	1.11	0	1.87
Price change across days (\$)	+7.49	0	62.03
Products (N=210)			
Capacity	64.27	64.00	23.28
Average total sales	47.39	46.00	21.12
Load factor	0.85	0.87	0.08
Gini coefficient	0.21	0.22	0.08

Table 2: Data summary. The upper part of the table contains descriptive statistics across the pricing events (routes, departure dates, and days to the departure). The lower part of the table contains statistics across routes and departure dates.

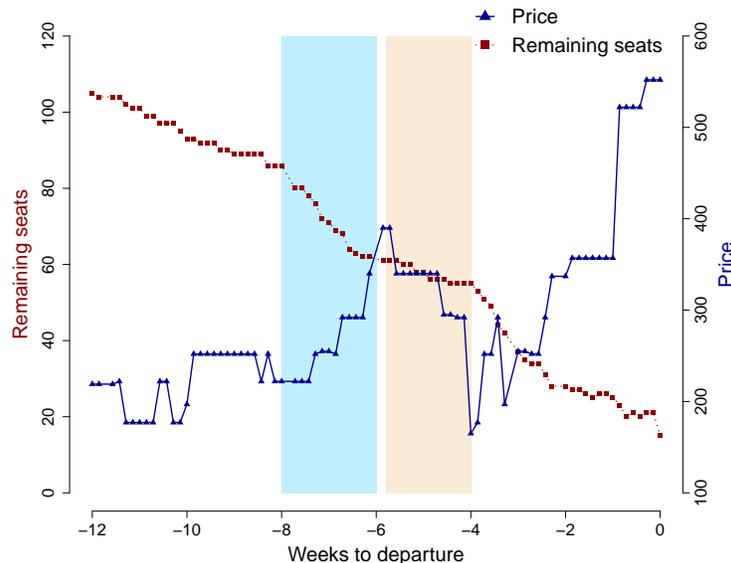


Figure 2: Price and remaining capacity. This graph shows the history of prices and remaining seats for a particular flight (JetBlue Flight 19), which departed on March 25, 2016. The red line with squares shows the path of the remaining seats. The blue line with triangles shows the path of prices. The shaded regions highlight some suggestive evidence of revenue management. In the light-blue area, seats were sold out quickly (the red line with squares dropped quickly), then the price (the blue line with triangles) increased. In the light-orange area, seats sold out slowly (the red line with squares was flat), then the price (the blue line with triangles) dropped.

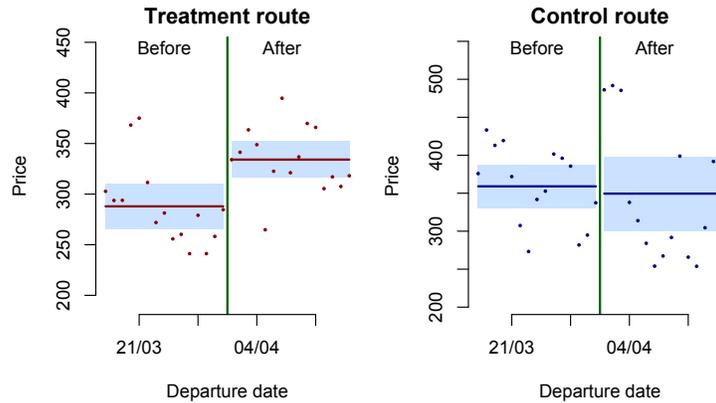


Figure 3: Alaska's price response to Delta's exit. This graph shows Alaska's average prices under the 2-by-2 treatment conditions (Table 1). Each dot is an average price for flights departing on a given departure date in a given route. It is averaged over $N = \#directions \#pricing-dates$ observations of prices. The light blue areas are 95% confidence intervals. Alaska raises its price significantly only in the route where its competitor exits.

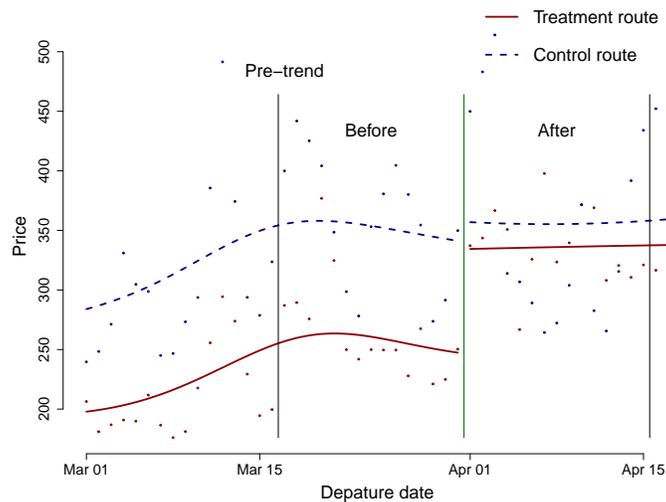


Figure 4: Parallel trend before the exit event. This graph shows price trends for the two routes before the exit. Each dot is an average price for flights departing on a given departure date in a given route. It is averaged over $N = \#firms \#directions \#pricing-dates$ observations of prices. Each line is smoothed using Gaussian kernels. Before March 31, 2016, the price gap was stable across the two routes. The gap changed after Delta's exit on April 1, 2016.

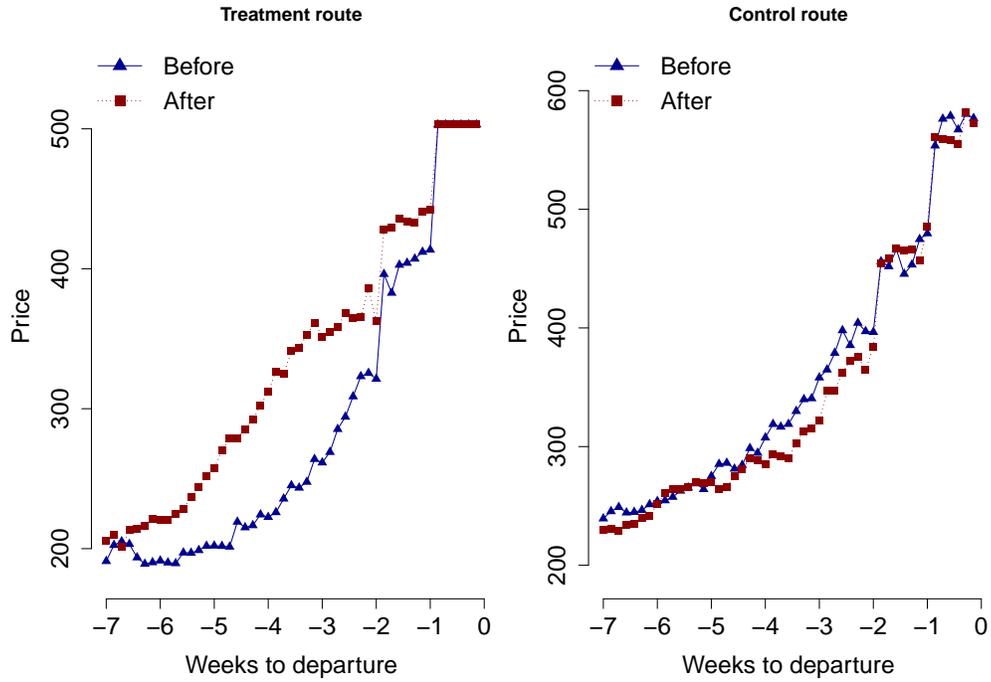


Figure 5: Heterogeneous effects on Alaska’s average price path. The graph shows Alaska’s average price path under the 2-by-2 treatment conditions. Each dot is an average price for flights departing under a given condition for a given number of days to departure. It is averaged over $N = \#directions \times \#departure-dates$ observations of prices. It illustrates heterogeneous effects across different numbers of days to departure.

Panel A				
	Exit route		Control route	
	High type	Low type	High type	Low type
A-1: Price coefficients β^H and β^L	-0.556 (0.007)	-1.528 (0.019)	-0.383 (0.001)	-1.535 (0.009)
A-2a: Firm-1-leaning consumers preference $\beta^{H,1}$ and $\beta^{L,1}$				
Firm 1	2.255 (0.042)	1.326 (0.072)	0.990 (0.010)	0.734 (0.017)
Firm 2	-0.288 (0.085)	-0.202 (0.042)	0.347 (0.009)	0.654 (0.009)
A-2b: Firm-2-leaning consumers preference $\beta^{H,2}$ and $\beta^{L,2}$				
Firm 1	-0.192 (0.065)		0.616 (0.026)	
Firm 2	4.077 (0.049)	2.497 (0.060)	3.213 (0.020)	3.170 (0.027)
A-3: Probability of leaning towards firm 1 $\beta^{H,1}$ horizontal and $\beta^{L,1}$ horizontal	0.889 (0.005)	0.830 (0.014)	0.740 (0.004)	0.623 (0.008)
Panel B				
	Exit route		Control route	
	$\binom{n}{n}$ arrival	$\binom{n}{n}$ vertical	$\binom{n}{n}$ arrival	$\binom{n}{n}$ vertical
n=0	6.775 (0.130)	-2.665 (0.021)	3.803 (0.013)	-1.245 (0.009)
n=1	0.520 (0.017)	0.472 (0.004)	1.283 (0.009)	0.369 (0.004)
n=2	-0.394 (0.004)	-0.147 (0.002)	-0.571 (0.004)	-0.155 (0.002)
n=3	0.052 (7e-4)	0.024 (4e-4)	0.063 (6e-4)	0.024 (2e-4)
Panel C				
	FE _{treatment}	FE _{weekend}		
Other demand parameters	0.059 (0.006)	-0.074 (0.006)	0.410 (0.010)	
Panel D				
	Alaska	Delta	JetBlue	σ
Cost parameters	32.100 (0.075)	24.130 (0.150)	30.937 (3.960)	11.514 (0.039)

Table 3: Estimates of the structural model. The table presents the estimates of the structural parameters. Panel A contains the estimates of consumer preferences. Panel B contains the coefficients of the polynomials gathering the Poisson arrival of consumers and the proportion of low-type customers as the arrival date approaches. Panel C contains the estimates of the fixed effects applied to departure dates after the exit event on 03/31/2016 and the weekends. It also contains the estimate of the variance of unobserved demand shifter σ – distributed normally with mean zero. Finally, panel D contains the estimates of marginal costs.

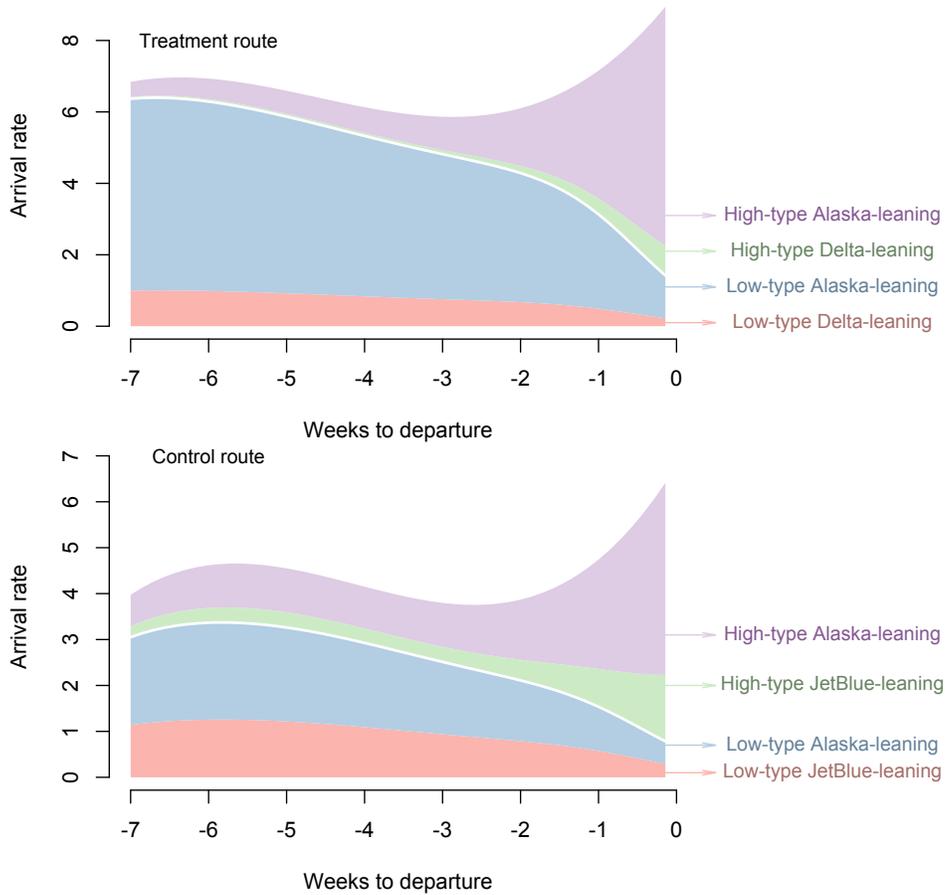


Figure 6: Estimated arrival by consumer segments. The graph plots Poisson arrival rates for different segments of consumers in the two routes. The upper graph shows the arrival distribution in the exit route. The lower graph shows the arrival distribution in the control route.

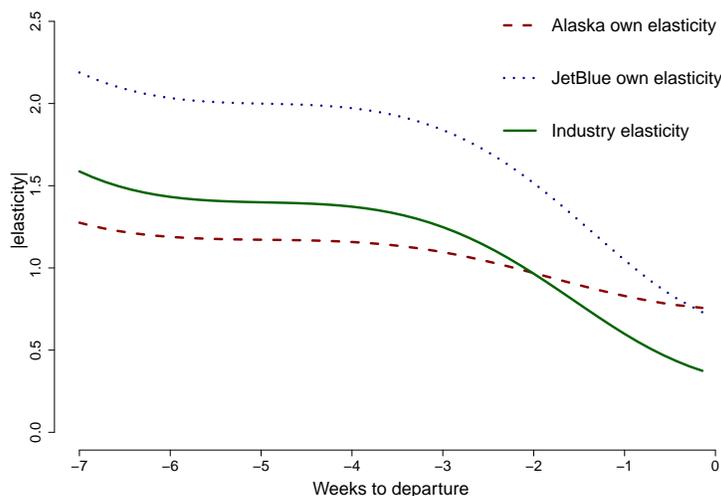


Figure 7: Demand elasticities. This graph shows own price elasticity and industry elasticity (in absolute values) evaluated at the average price. Own demand elasticity is the percentage change of own sales after a one percent increase in own price. Industry demand elasticity is the percentage change in industry sales when both firms increase their prices by one percent.

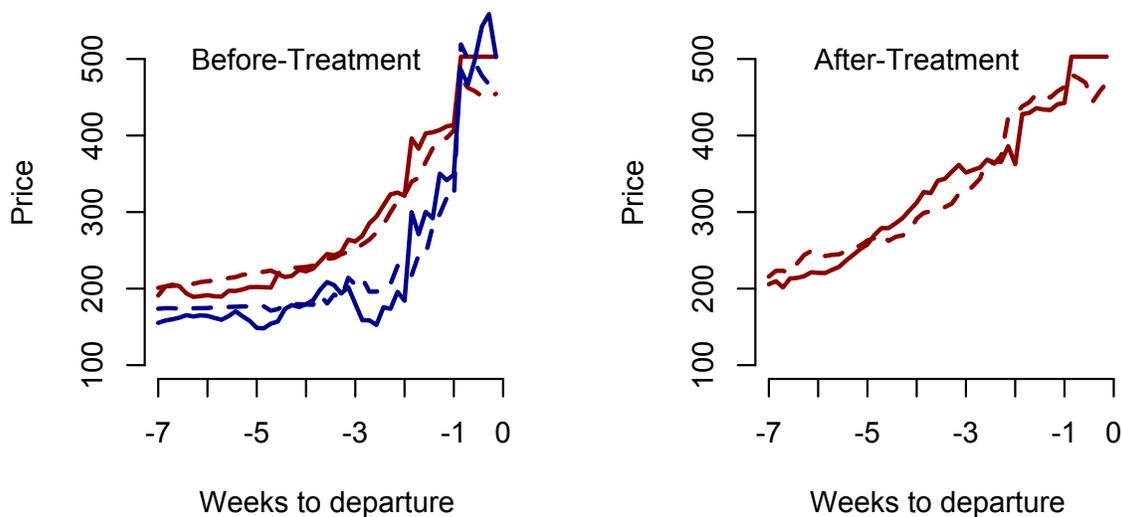


Figure 8: Model fit in the treated route. This graph compares the model-predicted prices with the actual prices conditional on the number of days to departure. The dashed lines are model-predicted prices, and the solid lines are the observed prices. The red curve is Alaska, and the blue curve is Delta.

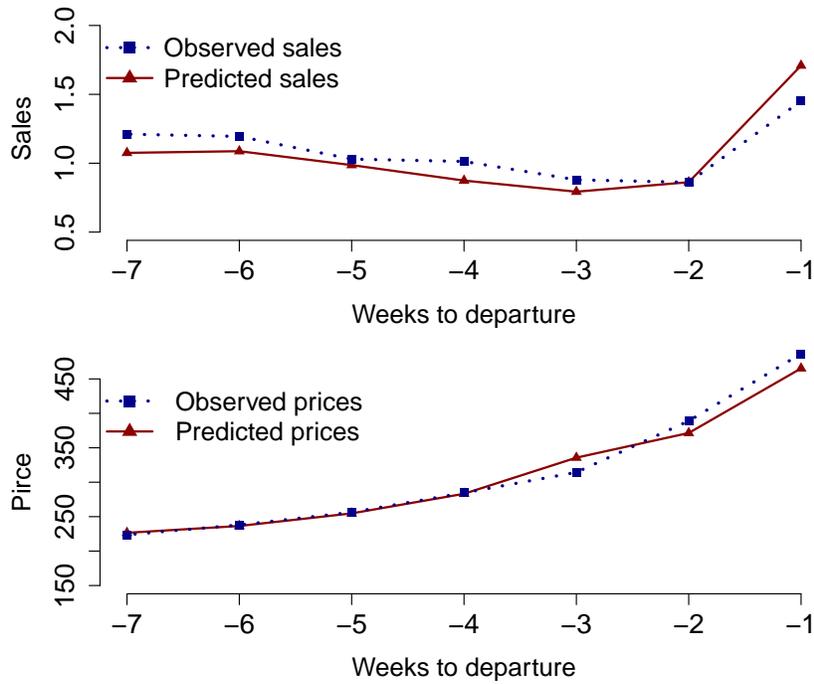


Figure 9: Model fit for price and sales. This graph shows the fit of price and sales averaged across routes, companies, and departure dates.

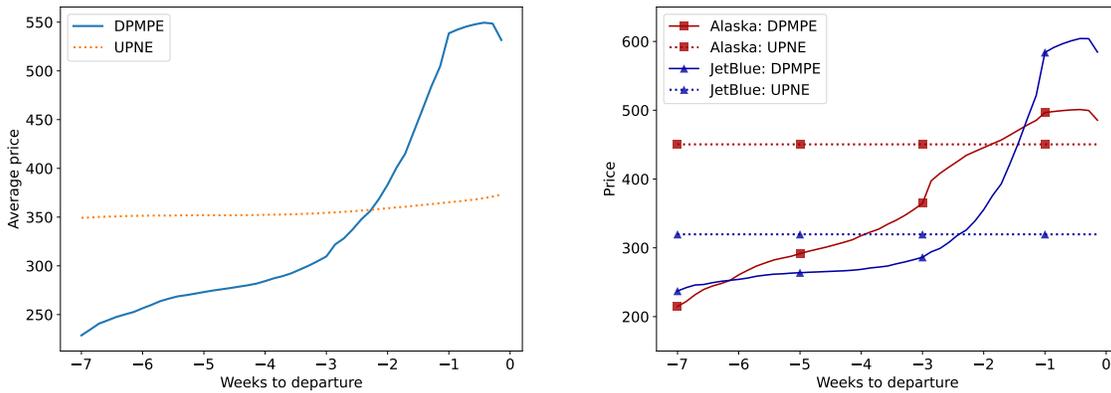


Figure 10: Impact of dynamic pricing on price paths. The graph depicts uniform pricing Nash-equilibrium (UPNE) and dynamic pricing (DPMPE) pricing paths. The prices are averaged across departure dates in the San Francisco-Boston route, assuming a duopoly market structure.

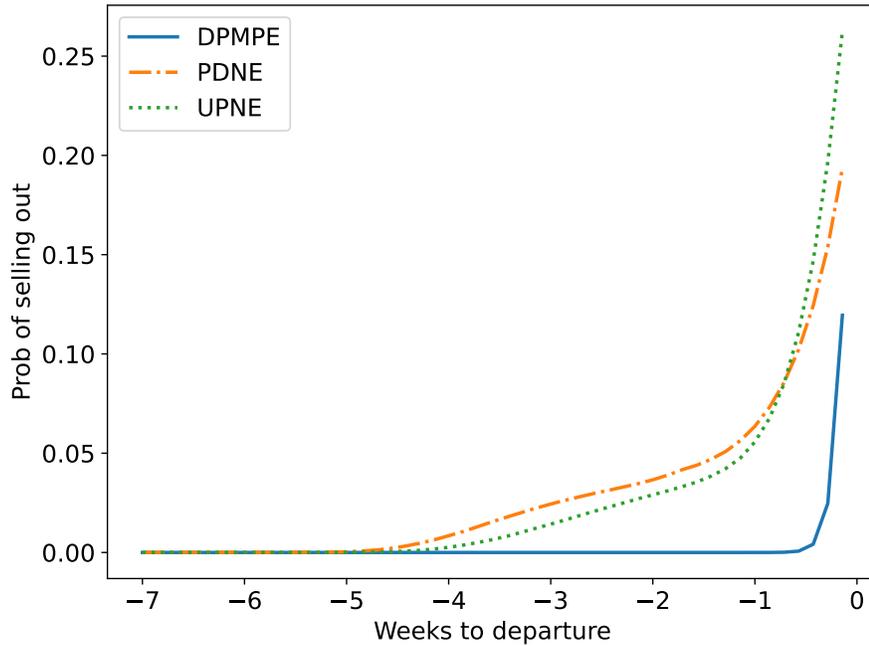


Figure 11: Likelihood of monopolization. The graph contains a cumulative likelihood of monopolization as the departure date approaches. Monopolization is defined as “at least one company selling out all capacity.” We compare monopolization for uniform pricing Nash-equilibrium (UPNE), price discrimination Nash-equilibrium (PDNE), and dynamic pricing Markov-perfect equilibrium (DPMPE). The prices are averaged across departure dates in the San Francisco-Boston route, assuming a duopoly market structure.

		(1)	(2)	(3)
		Uniform pricing	Price discrimination	Dynamic pricing
Alaska	Profits	\$11,503	\$13,071	\$13,031
	Sales	27.5	30.3	40.6
	Average price	\$450.4	\$463.5	\$352.8
JetBlue	Profits	\$18,695	\$19,850	\$19,726
	Sales	64.6	60.3	67.0
	Average price	\$319.7	\$360.4	\$325.5
Industry	Profits	\$30,198	\$32,921	\$32,758
	Sales	92.1	90.6	107.6
	Average price	\$358.7	\$394.9	\$335.8
Consumer surplus		\$39,585	\$35,726	\$40,883
Total welfare		\$69,783	\$68,648	\$73,641

Table 4: Impact of dynamic pricing. The table describes the effect of three pricing regimes on profits and welfare. The numbers are obtained in the San Francisco-Boston route, assuming a duopoly market structure.

		(1)	(2)	(3)
		Uniform pricing	Price discrimination	Dynamic pricing
High-type	Alaska-leaning	\$21,548	\$19,465	\$22,066
	JetBlue-leaning	\$16,213	\$14,124	\$15,076
	Both	\$37,761	\$33,590	37,143
Low-type	Alaska-leaning	\$159	\$205	\$787
	JetBlue-leaning	\$1,665	\$1,931	\$2,954
	Both	\$1,823	\$2,136	\$3,741
Consumer surplus		\$39,585	\$35,726	\$40,883

Table 5: Impact of dynamic pricing on vertical consumer segments. The table disaggregates the impact of pricing on consumer welfare by consumer segment. High types are inelastic late-arriving business travelers. Low types are elastic early-arriving leisure travelers. The numbers are averaged across departure dates for the treated route, assuming a duopoly market structure.

		(1)	(2)	(3)
		Uniform pricing	Price discrimination	Dynamic pricing
JetBlue	Profits	\$23,688	\$25,576	\$25,667
	Sales	82.5	76.9	79.3
	Average price	\$318.0	\$363.5	\$354.6
Consumer surplus	High-type	\$30,208	\$24,301	\$24,872
	Low-type	\$1,873	\$2,432	\$2,620
	Total	\$32,081	\$26,733	\$27,492
Total welfare		\$55,769	\$52,309	\$53,159

Table 6: Impact of dynamic pricing in the monopoly case. The table describes the effect of three pricing regimes on profits and welfare. The numbers are obtained in the San Francisco-Boston route, assuming a monopoly market structure. We obtain the monopoly market structure by conducting an acquisition of Alaska by JetBlue. In particular, we assume that the monopolist has a joint capacity of Alaska and JetBlue and that consumers have JetBlue parameters of preferences.

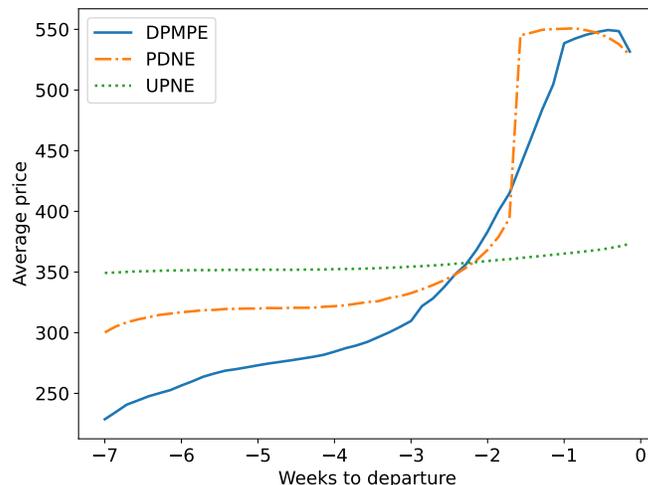


Figure 12: Comparison of price paths under various pricing regimes. We compare price paths as the departure date approaches for uniform pricing Nash-equilibrium (UPNE), price discrimination Nash-equilibrium (PDNE) and dynamic pricing Markov-perfect equilibrium (DPMPE). The prices are averaged across departure dates in the San Francisco-Boston route, assuming a duopoly market structure.

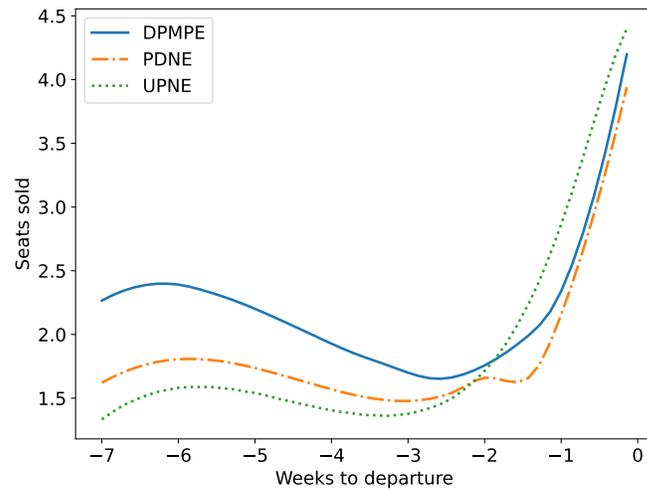


Figure 13: Comparison of seat sales under various pricing regimes. We compare daily seat sales as the departure date approaches for uniform pricing Nash-equilibrium (UPNE), price discrimination Nash-equilibrium (PDNE) and dynamic pricing Markov-perfect equilibrium (DPMPE). The sales are averaged across departure dates in the San Francisco-Boston route, assuming a duopoly market structure.