

NBER WORKING PAPER SERIES

ORGANIZATIONAL STRUCTURE AND PRICING:
EVIDENCE FROM A LARGE U.S. AIRLINE

Ali Hortaçsu
Olivia R. Natan
Hayden Parsley
Timothy Schwieg
Kevin R. Williams

Working Paper 29508
<http://www.nber.org/papers/w29508>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2021, Revised January 2023

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research. We thank the anonymous airline for giving us access to the data used in this study. Under the agreement with the authors, the airline had "the right to delete any trade secret, proprietary, or Confidential Information" supplied by the airline. We agreed to take comments in good faith regarding statements that would lead a reader to identify the airline and damage the airline's reputation. All authors have no material financial relationships with entities related to this research. We thank Stefano DellaVigna, Bob Gibbons, Ariel Pakes, Tobias Salz, Steve Tadelis, and the participants and discussants at several conferences and seminar presentations. We gratefully acknowledge financial support from the Becker Institute, the Tobin Center, the University of Chicago, the Yale School of Management, and NET Institute, www.NETinst.org. Jose Betancourt provided excellent research assistance.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

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

Organizational Structure and Pricing: Evidence from a Large U.S. Airline
Ali Hortaçsu, Olivia R. Natan, Hayden Parsley, Timothy Schwieg, and Kevin R. Williams
NBER Working Paper No. 29508
November 2021, Revised January 2023
JEL No. C11,C53,D22,D42,L11,L93

ABSTRACT

Firms facing complex objectives often decompose the problems they face, delegating different parts of the decision to distinct subordinates. Using comprehensive data and internal models from a large U.S. airline, we establish that airline pricing is inconsistent with canonical dynamic pricing models. However, we show that observed prices can be rationalized as an equilibrium of a game played by departments who each have decision rights for different inputs that are supplied to the observed pricing heuristic. Incorrectly assuming that the firm solves a standard profit maximization problem as a single entity understates overall welfare actually achieved but affects business and leisure consumers differently. Likewise, we show that assuming prices are set through standard profit maximization leads to incorrect inferences about consumer demand elasticities and thus welfare.

Ali Hortaçsu
Kenneth C. Griffin Department of Economics
University of Chicago
1126 East 59th Street
Chicago, IL 60637
and NBER
hortacsu@uchicago.edu

Olivia R. Natan
Haas School of Business
2200 Piedmont Avenue
Berkeley, CA 94720
olivia.natan@berkeley.edu

Hayden Parsley
University of Texas at Austin
parsley@utexas.edu

Timothy Schwieg
University of Chicago
5757 S University Ave
Chicago, IL 60637
timothyschwieg@uchicago.edu

Kevin R. Williams
Yale School of Management
165 Whitney Avenue
New Haven, CT 06520
and NBER
kevin.williams@yale.edu

1 Introduction

Firms facing complex decisions often need to simplify the problems they face. They may rely on heuristics and delegate decision rights to individual departments that are responsible for particular sub-decisions. This can be true for firms' pricing decisions, which frequently involve specialized teams and complex optimization systems. One department might manage procurement and inventory, another department specializes in demand predictions, and an additional department manages competitive response. Moving beyond the simple model of the firm as a unitary decision-maker which solves a standard economic model creates the potential for coordination failures across sub-units, and the use of heuristics may cause pricing decisions to differ significantly from those predicted by benchmark economic models.

Using granular data and internal models from a large U.S. airline, we empirically demonstrate the importance of accounting for organizational structure, department decisions, and the use of heuristics in modeling pricing decisions.¹ We find that airline pricing is subject to important pricing biases that impact all flight prices—most notably, and for tractability reasons, prices are set by a heuristic that differs substantially from traditional, dynamic profit maximization. We establish that observed prices are inconsistent with standard profit maximization by the firm. However, prices can be rationalized as an equilibrium of a game played by departments who each have decision rights for different inputs that are supplied to the observed pricing heuristic. Although departments have the same profit-maximization objective as the firm, providing biased inputs is a mutual best-response given the pricing heuristic utilized by the firm. In other words, the firm ends up making decisions that are boundedly rational, even though each department is acting rationally (Simon, 1955). This has important implications for our understanding of welfare. Incorrectly assuming that the firm solves a standard profit maximization problem as a single entity understates welfare actually achieved,

¹The airline has elected to remain anonymous.

but also results in significant differences in predictions of consumer surplus and revenues. Interestingly, business travelers benefit and leisure travelers are made worse off under observed practices compared to the benchmark dynamic pricing model.

We begin with an overview of airline pricing practices, describing the transition from regulated prices to the use of pricing heuristics post-deregulation. In the spirit of Simon (1962) and Radner (1993), the observed organizational structure features decentralized decision-making by departments, each responsible for particular sub-decisions.² The “network planning department” decides where to fly and assigns initial capacities. We do not model these decisions. Given the network, the “pricing department” designs itineraries and chooses a menu of discrete prices that consumers may face. Finally, the “revenue management (RM) department” is responsible for analyzing demand, monitoring flight-level performance, and making adjustments to demand predictions. The heuristic does not actually decide price, rather, it allocates inventory to each discrete price level. That is, “pricing” involves prices, quantities, and capacities set by separate departments.

Although the separation of pricing responsibilities across departments is known to exist (e.g., Vinod, 2021), we establish just how prevalent the observed organizational structure is. We collect job listings that show that all major airlines, ultra low-cost carriers, and recently founded airlines have the same organizational structure and department responsibilities. Moreover, we show that cruises, hotels, and car rentals have also adopted the same organizational structure and department responsibilities.³ Therefore, we believe our insights likely hold broadly in industries where firms face the complex problem of dynamically pricing perishable inventory.

²The organizational structure also possesses similarities to more recent theoretical work in organizational economics (e.g., Alonso, Dessein, and Matouschek, 2008; Rantakari, 2008; Dessein, Galeotti, and Santos, 2016).

³Examples of firms with the same organizational structure include: MSC Cruises, Carnival Corporation, InterContinental Hotels Group, Universal Orlando Resort, Hertz Corporation, and Avis Budget Group. Archived copies of all job postings, network profiles, and patents available upon request.

Our insights are derived from comprehensive data provided to us by a large U.S. airline. In addition to observing daily prices and quantities, we also observe all department decisions. We observe the demand models used, internal demand estimates, and the pricing heuristics’s exact design (code), which we use in counterfactuals. We also observe all consumer interactions (clickstream data) on the airline’s website. Our sample covers 300,000 flights and 470 domestic routes over a span of two years.

Using the data and internal models, we establish three facts. These facts show that airline pricing differs significantly from standard, dynamic profit maximization and highlight how department decisions can affect pricing. We use the facts to guide our modeling choices.

First, the airline uses a pricing heuristic that does not solve or approximate canonical dynamic pricing models for computational tractability. The heuristic abstracts from key market features—including cross-price elasticities—and is therefore inherently biased.⁴ The heuristic does not internalize substitution across cabins within a flight, across own substitute flights (within a day and across days), and across all competitor options. No competitor information enters the heuristic at all. Prices for each flight are optimized independently.⁵ Although an expansive literature studies firms through the lens of optimal dynamic decision making (Rust, 2019), the actual pricing heuristic does not internalize that it will revisit its decisions at all—it is not dynamically consistent. This is a striking fact given that airlines have been cited as using the most sophisticated pricing systems of all firms (McAfee and Te Velde, 2006).

Second, we show that department input decisions are also biased and subject to what we call department “miscoordination.” For example, the RM department uses single-product demand models that also do not account for any form of substitution.

⁴This can affect what inferences can be made about market conditions. For related theoretical work, see, e.g., Bohren (2016) and Heidhues, Kőszegi, and Strack (2018).

⁵We discuss why the heuristic used differs substantially from proposed models of algorithmic pricing (e.g. Asker, Fershtman, and Pakes, 2021; Calvano, Calzolari, Denicolo, and Pastorello, 2020; Brown and MacKay, 2021).

We show that departments do not internalize all of the decisions made by other departments. Moreover, we show that department decisions are incompatible: many prices in the pricing department’s fare menus are misaligned with the RM department’s demand forecasts—they are on the inelastic side of internal demand analysis. We show and discuss why this leads the heuristic to frequently offer consumers “inelastic” fares.

Third, we show that departments actively respond to organizational constraints and the limitations of the pricing heuristic by manipulating their own inputs, in similar spirit to manipulating information in the seminal behavioral theory of the firm of Cyert and March (1963). More precisely, we show that the RM department addresses what it views as prices that are suboptimally too low by distorting its own demand models, as a sort of workaround or kludge (Ely, 2011).⁶ We find that RM department analysts inflate internally estimated demand models to effectively raise prices that appear to reduce the percentage of flights that are priced on the inelastic side of demand.⁷

Analysis of internal models and data establish that pricing of all routes, regardless of market structure, depart from standard profit maximization, e.g., all flights rely on the same single-product heuristic. We then compare pricing biases across market structure and show that biases are even more pronounced in routes with nonstop competitors. Following the existing dynamic pricing literature, we focus on single-carrier routes in our structural analysis.⁸

We calculate payoffs of a team-theoretic game where departments supply inputs to the observed pricing heuristic. For a given set of inputs, we simulate counterfactual

⁶According to Ely (2011), “A kludge is a marginal adaptation that compensates for, but does not eliminate, fundamental design inefficiencies.” See Wollmer (1992) and Cooper, Homem-de Mello, and Kleywegt (2006) for analysis on why heuristics can lead to mis-pricing.

⁷For additional perspectives on miscalibrated firm expectations, see Massey and Thaler (2013); Akepanidaworn, Di Mascio, Imas, and Schmidt (2019); Ma, Ropele, Sraer, and Thesmar (2020). Berman and Heller (2020) present a theory on why biases can persist, including underestimating price elasticities, in a model of competition.

⁸With the exception of Hortaçsu, Öry, and Williams (2022), the literature on dynamic pricing in perishable goods markets has abstracted from oligopolistic competition where stage game payoffs depend on competitor scarcity.

outcomes based on unbiased and flexible demand estimates. We use a recently proposed demand methodology (Hortaçsu, Natan, Parsley, Schwieg, and Williams, 2022) that considers “leisure” and “business” travelers arriving according to time-varying Poisson distributions. Conditional on arrival, consumers solve a standard discrete choice problem. We provide new descriptive evidence to motivate our demand assumptions. We address the identification challenge of estimating preferences in models with aggregate demand uncertainty by leveraging arrivals and bookings data. We discuss why our instrumental variables are valid in the presence of pricing biases.

With demand estimated, we then consider counterfactual deviations from current department input decisions. In one counterfactual, we allow the pricing department to “coordinate” its fare menus to the RM department’s demand models by removing fares on the inelastic side of demand. In another counterfactual, we allow the RM department to adjust how it manipulates its demand models, holding the pricing department’s fare menus fixed. We show that these unilateral deviations are not profitable. Therefore, although observed pricing practices differ significantly from standard profit maximization, realized prices can be rationalized as an equilibrium of a common interests model where departments have delegated decision rights around the observed heuristic. The internal data and models imply suboptimal pricing, perhaps due to mistakes or behavioral frictions (e.g., Levitt, 2016; DellaVigna and Gentzkow, 2019; Dubé and Misra, 2021). However, within observed constraints, we establish that observed prices follow a model of bounded rationality.

We further use these counterfactuals to highlight the organizational theories of limited gains of unilateral change (Milgrom and Roberts, 1990, 1995). We show that the impact of large changes in a department’s inputs can result in small revenue effects because of how the pricing heuristic works. Moreover, we discuss why the observed equilibrium is not driven by a misalignment in incentives within the organization (Atkin, Chaudhry, Chaudry, Khandelwal, and Verhoogen, 2017; Sacarny, 2018).

While it is possible that many input combinations constitute mutual best responses, formally analyzing the relevant team-theoretic game is difficult because computing all potential department choices is computationally intractable. Instead, we compare observed prices to those predicted under dynamically optimal profit maximization. Although this counterfactual can only be implemented for some routes for computational reasons, it allows us to contrast welfare under observed practices with those expected under the optimal dynamic pricing solution. Implicitly, the counterfactual also provides a lower bound on the costs (re-organization, computational burden, etc.) and environmental factors (Siggelkow, 2001) that constrain pricing decisions.

We find that observed prices differ, sometimes substantially, from those predicted under the optimal dynamic pricing solution (“dynamic prices”). As a result, incorrectly assuming that the firm solves the canonical pricing model understates welfare actually achieved by 6%. Dynamic pricing overstates revenues by 14%. Moreover, the differences result in significant welfare effects for business and leisure travelers. Leisure travelers would actually prefer if dynamic pricing were possible because current pricing biases inflate early prices (8% higher surplus). However, dynamic prices exacerbate scarcity effects and price targeting, driving up fares for late-arriving, less-elastic business travelers. As a result, business travelers benefit from the observed organizational structure and use of heuristics (23% higher surplus).

In a complementary exercise, we also show that imposing the common assumption that firms are unboundedly rational when estimating demand—and therefore, abstracting from organizational structure, department decisions, and the use of heuristics—leads to incorrect inferences about consumer demand elasticities and thus welfare. We compare our demand estimates to those obtained by additionally imposing profit maximization as advocated by, e.g., Berry, Levinsohn, and Pakes (1995), and commonly used in dynamic pricing studies (e.g., Williams, 2022; Aryal, Murry, and Williams,

2022; Pan and Wang, 2022; Cho, Lee, Rust, and Yu, 2018).⁹ We find significant differences in estimating demand curves that again understate welfare for the routes studied.

Finally, we discuss the broad economic importance of our findings. First, we reveal that advanced pricing systems rely on consequential simplifications, e.g., single-product demand models and dynamically inconsistent optimization tools. Second, we establish that observed pricing practices are not well approximated by a benchmark model of dynamic pricing. We offer an empirical quantification on the differences between actual firm pricing practices and what Rust (2019) describes as the questionable, yet maintained, assumption that firms behave “as if” they have solved the dynamic problems they face. Third, we show that a model of bounded rationality (Simon, 1955) involving delegated decision-making and use of heuristics explains firm pricing decisions. Fourth, we show that there are significant welfare differences between observed prices and the optimal dynamic pricing solution. Our results show that understanding how firms structure the complex problems they face is crucial for the inferences we draw from firms’ pricing decisions and how we model firm behavior.

The paper is organized as follows. In Section 2, we discuss industry pricing practices. We discuss the data in Section 3 and pricing biases in Section 4. We present the demand model, estimation details, and parameter estimates in Section 5. We investigate department decisions in Section 6. We consider dynamic pricing in Section 7.

2 Industry Setting and Organizational Structure

We study the US airline industry, an industry that directly supports over 2.2 million jobs and contributes over \$700 billion to the US economy.¹⁰ We briefly describe airline pricing practices. McGill and Van Ryzin (1999) and Vinod (2021) provide detailed,

⁹D’Haultfœuille, Février, Wang, and Wilner (2022) consider a partial identification approach.

¹⁰See <https://www.iata.org/en/iata-repository/publications/economic-reports/the-united-states-value-of-aviation/>. Sep 28, 2022.

historical accounts.

Prior to deregulation of the airline industry, the networks and fares of airlines were federally controlled. In order to facilitate ticket purchases, airlines collectively started ATPCO, a corporation that gathers and disseminates fares to Global Distribution Systems (GDSs), or travel reservation systems that merge and process fare and ticket availability for travel agents. “Pricing” meant filing new binders of fares. Post-deregulation, competition intensified, which resulted in airlines lowering fares, particularly for consumers who shopped early. These were the first advance purchase (AP) fares that are now commonly observed in advance purchase markets.

In 1972, BOAC Airlines developed the first tractable inventory management system that controlled how many discounted fares to offer based on expected demand.¹¹ American Airlines developed a related optimization tool as computing power increased (Vinod, 2021). The adoption of these optimization tools resulted in a bifurcation of pricing responsibilities. One department continued to decide the set of itineraries and fares to offer. The newly formed department—revenue management—worked on algorithm development, created demand forecasts, and allocated inventory to the pricing menu. Using job listings, we show that all major airlines, including legacy carriers, low-cost carriers, and start-up airlines maintain separate pricing and RM departments.¹² In fact, we show that the same organizational structure and department responsibilities is used in car rentals, hotels, cruises, trains, and buses.

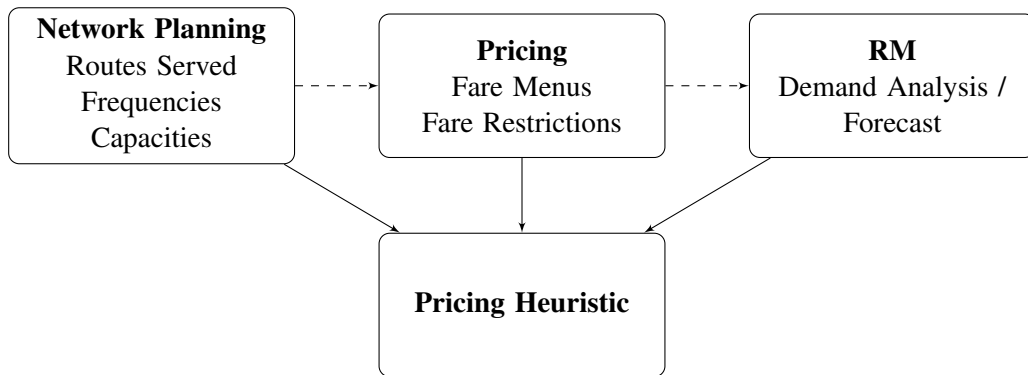
Figure 1 depicts the organizational structure we study. First, the network planning department decides routes served, flight frequencies, and capacities. We do not model these decisions. The core responsibility of the pricing department is to decide a discrete menu of fares and fare restrictions for every itinerary. The most common restriction is an advance purchase (AP) requirement, meaning that a fare needs to be

¹¹The BOAC employee, Ken Littlewood, developed what is now referred to as Littlewood’s rule. The central intuition of his model is that if future demand is expected to be strong, an airline should offer fewer seats at lower prices today.

¹²See Footnote 3.

purchased before a certain number of days before departure.¹³ The pricing department does not simply choose fares, it designs the types of itineraries and tickets that can be purchased.¹⁴ The department gathers and interprets competitor prices and initiates/responds to industry-level changes, e.g., implementing a fuel surcharge. The RM department is responsible for demand analysis and monitoring flight-level performance. The RM department estimates short-run demand models; the department does not maintain long-run demand estimates. As demand is realized, RM analysts make adjustments to their forecasts. The pricing heuristic combines all inputs—the demand analysis (flight-level forecasts) made by the RM department, the fare menu decisions of the pricing department, and the capacity constraint set by the network planning department. It allocates how many (remaining) seats can be sold at each price level.¹⁵ The key performance metrics used at the firm do not suggest that these departments have misaligned incentives.

Figure 1: Division Responsibilities at all Airlines



Note: Key departments, responsibilities, and decision-making process at all airlines. The dashed arrows show the flow of information that can be leveraged by departments. The solid arrows show the flow of decisions, i.e., decisions are based to the pricing heuristic.

¹³Advance purchase (AP) fares are common at 7, 14, and 21 days before departure.

¹⁴The pricing department decides fares that cover different classes of service, connecting options, blackout dates, etc. Each fare has dozens of characteristics that can be adjusted. The coarsest level of a fare is its fare class, e.g., discounted economy versus full-fare economy. We observe over 20 million fares for economy-class fares in our sample.

¹⁵Note, for expositional purposes, we refer to the heuristic as determining price.

Organizational inertia, coupled with significant computational constraints, are reasons why the observed bifurcation of pricing responsibilities is the predominant way firms have organized in industries selling perishable goods. Although the divisions in pricing responsibilities developed as optimization techniques improved, there has been little advancement in solving the general airline pricing problem at scale (the code edits we observe show that the airline’s heuristic has not changed substantially for decades). Accounting for substitution patterns in dynamic models is difficult and optimizing remaining inventory across the entire network is computationally intractable given the size of the airline’s operations.¹⁶ To emphasize how difficult the airline pricing problem is, take the organizational structure as given and abstract from all network considerations. The smallest flights in our sample involve over 1.5 million inventory allocations by the pricing heuristic. Combining inventory allocation with the simplest pricing decision, i.e., deciding among two possibilities for each discrete fare class, results in an objective with more than 10^{1000} potential choices. That is, the scope of the joint pricing and inventory decision is intractable for even the smallest flights.

Similar to organizational structure, external compatibility is another factor that constrains airline pricing. Airlines still rely on the same external systems for publishing prices (e.g., ATPCO) and managing bookings across booking channels (direct, via OTAs, etc.). These systems ensure access to the same fares across channels, but it necessitates the use of discrete fares and handling inventory in a unified manner. Hotels, car rentals, etc. also leverage these systems. These systems discipline department decisions—for example, the number of fares in the pricing department’s menu choices is limited because of external compatibility.

¹⁶In practice, the firm leverages a procedure built on top of the pricing heuristic that attempts to correct for aspects of network demand. We do not study this aspect of pricing because it requires data on the entire network and is very costly to compute. In general, we may understate the value of capacity in our analysis; however, we attempt to minimize this aspect of demand through our route selection criteria.

3 Data and Summary Analysis

3.1 Data Overview

We use comprehensive data provided to us by a large U.S. airline. To maintain anonymity, we exclude some details. We combine several data sources, which we refer to as: (1) bookings, (2) inventory, (3) search, (4) fares, and (5) forecasting data.

(1) *Bookings data*: We observe all tickets purchased, regardless of booking channel, e.g., the airline’s website, travel agency, etc. Key variables included are the fare paid, the number of passengers, the particular flights included in the itinerary, the booking channel, and the purchase date. We focus on nonstop, economy class tickets.

(2) *Inventory data*: The inventory data detail the number of seats the airline is willing to sell at each fare level, at each point in time. We observe the pricing algorithm code and algorithm output, including the opportunity cost of selling a seat. Prices are reoptimized, given remaining inventory, daily.

(3) *Search data*: We observe all internet activity on the airline’s website for two years. The search data contain hundreds of millions of data points. Tracked actions include, but are not limited to, search queries, bookings, referrals from other websites, and the sets of flights that appear on every page that the consumer visits.

(4) *Fare data*: The fare data contain the pricing department’s decisions. A fare denotes a price and ticket restrictions, including any advance purchase requirements. We focus on nonstop/non-connecting fares and observe all fare menu changes.

(5) *Forecasting data*: The RM department forecasts demand based on short-run demand estimates. We use “demand model” to denote the observed, baseline demand model and “demand forecasts” to denote the demand model’s predictions after analyst adjustments. Adjustments allow for reacting to changing market conditions, including the performance of recently departed flights. The demand forecasts are still demand curves, i.e., quantity demanded for a given price. The RM department maintains

separate models/forecasts for “business” and “leisure” travelers based on an observed classification algorithm. We observe the demand model, parameter estimates, analyst adjustments, and the forecasts themselves.

3.2 Summary Analysis

We do not study all routes served due to data size constraints. Instead, we select 470 routes. In Online Appendix B, we discuss route selection. On average, the routes we study have a higher fraction of nonstop traffic, fewer flights per day, and smaller total capacity compared to the airline’s overall domestic network. Nonetheless, our analyses cover a diverse set of routes in terms of competition, seasonality, frequencies, and traffic flows. The sample contains large “trunk routes” between major cities as well as routes from metropolitan areas to small cities. We focus on domestic routes.

Table 1: Summary Statistics

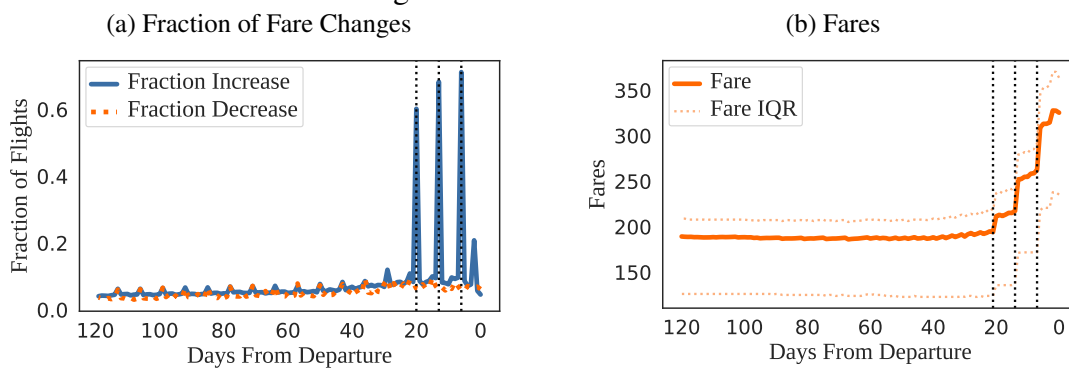
Data	Variable	Mean	Std. Dev.	Median	5th pctile	95th pctile
Fares						
	One-Way Fare (\$)	201.3	139.4	163.3	88.0	411.1
	Num. Fare Changes	9.3	4.2	9.0	3.0	17.0
	Fare Change Inc.	50.4	73.0	31.2	2.2	164.5
	Fare Change Dec.	-53.0	75.5	-32.2	-175.2	-4.3
Bookings						
	Booking Rate-OD	0.2	0.7	0.0	0.0	1.0
	Booking Rate-All	0.6	1.4	0.0	0.0	3.0
	Load Factor (%)	82.2	21.4	90.0	36.0	102.0
Searches						
	Search Rate	1.9	4.8	0.0	0.0	9.0

Summary statistics for the data sample. The booking rates are for non-award, direct travel on nonstop flights and for all traffic on nonstop flights (including passengers who connect onward), respectively. The number of passengers denotes the number of passengers per booking. Load factor includes all bookings, including award and connecting itineraries. The search rate is for origin-destination queries at the daily level.

Table 1 provides a basic summary of the nearly 300,000 flights in our sample. We focus on the last 120 days before departure due to the overwhelming sparsity of search and sales observations earlier in the booking horizon.

Average fares in our sample are \$201, with large dispersion across routes and over time. Typically, prices for a particular flight adjust nine times. Many fare adjustments occur at specified times, such as after expiration of AP opportunities (see Figure 2-a). Over 60% of all price adjustments occur outside these time windows. In Figure 2-(b), we plot average fares over time. Fares increase by over 70% in 120 days. More than 25% of routes see fares more than double. For a few routes, fares triple in 120 days.

Figure 2: Fares Time Series



Note: Fraction of fare changes and average fares by day before departure. Also included is the IQR across fares.

The booking rate (sales per flight-day) is low; the percentage of zero sales is 80%. Importantly, the highest booking rates occur when prices are the highest (within the last 7 days before departure). The average load factor at departure is 82.2%. Although 5% of flights eventually oversell, we abstract from this possibility because we do not observe denied boarding/no show information. We use the flight’s “authorized” capacity, which is the capacity the heuristic uses when it allocates remaining inventory.

3.3 Motivating Evidence on Demand

We provide new descriptive evidence to motivate many of our modeling assumptions. The bookings data suggest that unit demand is a reasonable assumption. The average passengers per booking is 1.3, and the median is 1. We assume that consumers observe a single price per flight because 91% of consumers purchase the lowest available

fare (because of how the pricing heuristic works, consumers can often purchase more expensive tickets that may or may not share the same attributes, e.g., refundability).¹⁷ Special fares, such as corporate or government discounts, are rare in the routes studied.

We adopt a two-type consumer model, corresponding to “leisure” and “business” travelers, because that is how the RM department models demand. The labels “leisure” and “business” are mechanically linked to attributes of the ticket, e.g., the number of days before departure it was booked and are not attached to traveler characteristics, e.g., passenger status or travel purpose.

We find evidence that supports using a static discrete choice model to model demand (the following assumptions are also made in the RM department’s demand models). We “daisychain” the clickstream data, linking searches across devices and cookies for hundreds of millions of clicks. We assume that consumers consider a single departure date because 82% of customers search a single departure date (see Figure 13-a in Online Appendix C.1). Among the remaining 18%, the average time lag between these searches is 45 days, suggesting different trips. We do not model consumers strategically waiting to purchase tickets because 90% of consumers complete their search activity in a single day (including shoppers referred to the airline increases the percentage). Interestingly, among the remaining 10% that search over time (see Figure 13-b), only 20% ever observe a lower fare for at least one flight in later searches (most search spells end within five days). Our estimates suggest that only 2% of shoppers who may be strategically waiting actually obtain a lower price. This estimate is lower than the IQR (5.2% to 19.2%) found by Li, Granados, and Netessine (2014) who do not have access to search data. We also do not model strategic timing of arrival because otherwise the data suggests consumer mistakes may be very common. Over 62% of consumers would have received a lower price if they purchased a week earlier. Only

¹⁷Consumers with the highest status are the most likely to pay more than the lowest price available. This finding complements Orhun, Guo, and Hagemann (2022), who show that loyal consumers tend to fly longer itineraries than necessary in order to obtain status at an airline whose loyalty program is mileage based.

8% would have benefited from delaying their purchasing decision.

4 Pricing Biases in Airline Markets

We provide several examples why airline pricing practices differs from dynamically optimal profit maximization. We use the term “pricing biases.” We provide additional empirical evidence in Online Appendix C, including an example of “miscoordination” in which departments do not necessarily internalize the decisions of other departments when determining inputs and that airline pricing is subject to pricing frictions (marginal cost changes that do not trigger price adjustments).

4.1 Heuristic Bias

Prices are determined by a heuristic for tractability reasons. We detail the heuristic which our firm’s is based on in Online Appendix A. We do not identify the exact heuristic for confidentiality reasons. However, we use the exact heuristic in our counterfactuals. The edits to the heuristic’s code confirms that it has hardly changed over the past 25 years.

There are several reasons why airline pricing differs significantly from dynamically optimal profit maximization. Examining the heuristic’s code, one of the most salient differences is that the heuristic is dynamically inconsistent—it does not solve a standard dynamic program at all. Rather, the heuristic simplifies the world to two periods, demand today and the sum of future expected demand (the “second period”). By aggregating all of future demand into a single period, the heuristic avoids solving a dynamic program. The heuristic does not internalize that it will revisit its allocation decisions daily. Given remaining capacity (and all input decisions), it works to ensure that enough seats are reserved to meet future demand—the second period.

The pricing heuristic is also inherently biased because it cannot account for cross-

price elasticities of any kind, including across cabins within a flight, other flight options, and competitors. All flights, regardless of market structure, etc., are priced using the same single-product heuristic. Competitor prices do not enter the algorithm at all. In Online Appendix C.2, we show that the pricing heuristic does not even indirectly react to demand realizations of substitute flights. This precludes direct responses to competitors.¹⁸ We discuss how only the pricing department can partially internalize competitor prices in the next subsection. Finally, the heuristic does not explicitly maximize revenues—rather, it focuses on the booking rate.

The biases we document affect all decisions the heuristic makes—pricing decisions do not follow standard economic models, regardless of market structure. The heuristic’s design will tend to cause it to understate opportunity costs.¹⁹ This is because accounting for (any) own-product substitution and internalizing that prices will be optimized over time will tend to increase opportunity costs. Addressing that the heuristic may yield suboptimally low prices is difficult because adjustments should be specific to every flight’s state variables (time and capacity remaining). We examine how RM analysts address perceived mispricing in Section 4.3.

4.1.1 How Department Decisions Interact with the Pricing Heuristic

Before continuing, we provide two illustrative examples on how department input decisions may interact with the pricing heuristic. The examples emphasize that inputs can be “offsetting” or “reinforcing”—that is, oppose one department’s desired affect on pricing decisions or compound one department’s goal, given the heuristic’s design.

Suppose the pricing department wants to run a sale (or match a competitor) by

¹⁸Note that fares are not personalized, and loyalty metrics are not used in pricing and RM activities. The heuristic does not consider ancillary revenue, including baggage fees, upgrade charges, etc., when it makes its decisions. This is not uncommon in the industry. For example, Japan, Etihad, Philippine, Flydubai, Korean, Jeju, Frontier, Malaysia, All Nippon, Hawaiian, and Lufthansa all use an RM optimization solution offered by PROS that does not allow for joint ancillary fare revenue maximization.

¹⁹Cooper, Homem-de Mello, and Kleywegt (2006) shows that this can even happen in single-product settings using a similar heuristic.

offering a \$50 fare. This fare decision is fed through the RM department's demand forecast within the pricing heuristic. If the heuristic decides that such a low fare would not save seats for future demands based on the forecasts, it will not allocate any seats to the sale fare. This means that the RM department's input decision and the heuristic will offset the pricing department's intentions. As a result, the sale (or price match) will not occur.

If the heuristic optimizes prices for every flight in isolation and does not consider competitor information, how is it possible that multiple flights offered by the same airline can have the same price? How is it possible that multiple airlines charge the same price simultaneously? Observing identical prices is made possible by department decisions, but it is not enforced by the heuristic. That is, if a competitor is offering a \$50 fare, the pricing department can "match" it by filing a \$50 fare, but it is up to the heuristic if that fare is selected given the other inputs. The same is true for multiple flights offered by the same airline—this again is made possible by the pricing department, because fare menus can vary by departure date but not at the flight-level. The heuristic decides inventory allocations for each flight in isolation without any information on substitute flights (own and competitor options).

As a second example, suppose the pricing department institutes a fuel surcharge, raising all fares, and the RM department simultaneously adjusts its demand predictions upward due to predicting a strong travel season. Both actions independently would (tend to) raise price, but when combined, the price increase would (tend to) be larger.²⁰

4.2 Fares on the Inelastic Side of Demand Curves

Next, we show that pricing inputs are subject to department "miscoordination." More specifically, we observe that the pricing department often files fares that are too low

²⁰We add the words "tend to" because the objective is nonstandard and both capacity and fare choices are discrete. It could be that certain adjustments yield the opposite effect, e.g., increasing the price of very high fares, leaving lower priced fares the same, which could cause the distribution of prices to fall.

according to the RM department’s internal demand models. Although this would be inconsequential if the pricing heuristic solved a standard profit maximization problem—it would not allocate seats on the inelastic side of the demand curve—in practice, it may not prevent “inelastic fares” from being offered to consumers. When capacity is not sufficiently constrained, it will default to the lowest fare on the menu, regardless of what it is. This is true even if the more expensive, profit maximizing fare is included in pricing department’s menus.

To quantify this form of miscoordination, we use the RM department’s continuous and differentiable demand models, $Q(p)$. We calculate the elasticity of demand, $e(p)$, and plug in the lowest fare filed by the pricing department. We find that if the heuristic allocated seats to these fares, consumer demand would be inelastic according to the RM department’s demand models in 98% of the sample.²¹


4.3 Using Persistently Biased Forecasts

The RM department knows about the heuristic’s biases and that the presence of low prices in the fare menus can cause consumers to be offered inelastic prices according to its own analysis. It has introduced a workaround, or kludge, to counteract these forces. In order to raise prices (via increasing opportunity costs) within its organizational decision rights, we find that all RM analysts manipulate their own demand models. They incorporate an upward bias that results in systematically overpredicting demand. This is most commonly done by scaling up/down multiple routes’ demand models simultaneously. The demand models, after RM analyst adjustments, are supplied to the pricing heuristic.

In Figure 3, we plot the average forecast bias by week before departure. We calcu-

²¹One might conjecture that low fares remain on the menus so they can be activated in case of a sale. However, the lowest fares are not “sale” fares. Sale fares are observable to us. They have special attributes and are only active for short periods of time.

Figure 3: Forecast Bias by Day Before Departure



Note: Forecast bias is calculated by comparing the sum of expected bookings for each flight (and price) to realized bookings, by week before departure.

late the forecast bias for a particular week before departure as

$$\text{Forecast Bias} := 100 \cdot \frac{\overbrace{\sum_{j,d,t} \mathbb{E}Q_{j,d,t}(p_{j,d,t})}^{\text{Forecasted Demand}} - \overbrace{\sum_{j,d,t} Q_{j,d,t}(p_{j,d,t})}^{\text{Realized Demand}}}{\underbrace{\sum_{j,d,t} Q_{j,d,t}(p_{j,d,t})}_{\text{Realized Demand}}},$$

where forecasted and realized demand account for the current and expected future prices offered, and we sum over all flights (j), departure dates (d), and days before departure (t), for a given week. The average forecast bias is 15% higher than actual realized demand. We find that the bias shrinks from nearly 25% of expected sales early on to 8% close to the departure date.²² This pattern is observed across all routes, regardless of route performance or market structure. 79% of flights have overforecasted demand 120 days before departure. We find that routes with nonstop competitors fea-

²²We may expect the optimal bias to decrease over time for two reasons. First, the heuristic tends to deflate opportunity costs the most well in advance of the departure date. As time to departure decreases, the effect of the heuristic bias decreases because there are fewer opportunities to reoptimize remaining inventory. That is, the understatement of opportunity costs decreases. Second, as remaining capacity decreases, each additional seat sold will tend to have a larger effect on opportunity costs, meaning smaller bias is necessary to raise prices.

ture slightly larger forecasting bias compared to single-carrier routes.

How effective are these manipulations to demand forecasts at addressing perceived mis-pricing? We again calculate the elasticity of demand using RM departments demand models, $Q(p)$, but instead plug in realized prices. We find that 38% of flights are actually priced on the inelastic side of demand, rather than the 98% of flights we previously estimated if the heuristic defaulted to offering the lowest fares. Interestingly, routes with existing nonstop competition feature more frequent inelastic demand based on internal demand models (yet, also feature larger forecast biases) than single-carrier routes. We study RM analyst decisions more formally in our counterfactual analyses.

5 Empirical Model of Air Travel Demand

While our descriptive evidence highlights actual airline pricing practices, it does not allow us to quantify the differences between observed pricing practices and the canonical single-firm dynamic pricing model. To do so, we need to estimate a model of unbiased preferences. We cannot use the internally estimated demand models because they are misspecified, e.g., they assume that demand is single-product. We utilize both the demand model and estimation approach of Hortaçsu, Natan, Parsley, Schweg, and Williams (2022). This allows us to capture rich substitution patterns, including seasonality effects, day-of-week effects, etc.

In the model, the definition of a market is an origin-destination (r), departure date (d), and day before departure (t) tuple. The booking horizon for each flight j leaving on date d is $t \in \{0, \dots, T\}$. The first period of sale is $t = T$, and the flight departs at $t = 0$. In each market t , arriving consumers choose flights from the choice set $J(r, t, d)$ that maximize their individual utilities, or select the outside option, $j = 0$. Note that our model covers all bookings, regardless of how the ticket was purchased.

5.1 Utility Specification

Arriving consumers are one of two types, corresponding to leisure (L) travelers and business (B) travelers. An individual consumer is denoted by i and her consumer type is denoted by $\ell \in \{B, L\}$. The probability that an arriving consumer is a business traveler is equal to $\gamma_{t,r}$. We incorporate two assumptions to greatly simplify the demand system. First, we assume that consumers do not choose flights based on remaining capacity, $C_{j,t,d,r}$. This allows us to avoid modeling infrequent events where a consumer may otherwise choose a less preferred option because there is a higher probability of securing a seat. Second, we incorporate random rationing if demand exceeds remaining capacity.

We assume that indirect utilities are linear in product characteristics and given by

$$u_{i,j,t,d,r} = \begin{cases} X_{j,t,d,r} \beta_r - p_{j,t,d,r} \alpha_{\ell(i),r} + \xi_{j,t,d,r} + \varepsilon_{i,j,t,d,r}, & j \in J(t, d, r) \\ \varepsilon_{i,0,t,d,r}, & j = 0 \end{cases},$$

where $X_{j,t,d,r}$ denote product characteristics other than price $p_{j,t,d,r}$, and preferences are denoted by $(\beta_r, \alpha_{\ell,r})_{\ell \in \{B, L\}}$. The term $\xi_{j,t,d,r}$ denotes an unobserved demand shock that is potentially correlated with price, and $\varepsilon_{i,j,t,d,r}$ is a random component of utility which is assumed to be distributed according to a type-1 extreme value distribution. All consumers solve a straightforward utility maximization problem: consumer i chooses flight j if and only if $u_{i,j,t,d,r} \geq u_{i,j',d,t,r}, \forall j' \in J \cup \{0\}$.

The distributional assumption on the idiosyncratic error term leads to analytical expressions for the individual choice probabilities (Berry, Carnall, and Spiller, 2006). The probability that consumer i wants to purchase a ticket on flight j is equal to

$$s_{j,t,d,r}^i = \frac{\exp\left(X_{j,t,d,r} \beta_r - p_{j,t,d,r} \alpha_{\ell(i),r} + \xi_{j,t,d,r}\right)}{1 + \sum_{k \in J(t,d,r)} \exp\left(X_{k,t,d,r} \beta_r - p_{k,t,d,r} \alpha_{\ell(i),r} + \xi_{k,t,d,r}\right)}.$$

Since consumers are one of two types, we define $s_{j,t,d,r}^L$ be the conditional choice probability for a leisure consumer (and $s_{j,t,d,r}^B$ for a business consumer). Integrating over consumer types, we obtain market shares, $s_{j,t,d,r} = \gamma_{t,r} s_{j,t,d,r}^B + (1 - \gamma_{t,r}) s_{j,t,d,r}^L$.

5.2 Arrival Processes and Integer-Valued Demand

Whereas the empirical literature commonly assumes a market size, we model consumer arrivals and estimate arrival rates based on search data. We assume that both consumer types arrive according to a time-varying Poisson distribution. We assume: (i) arrivals are distributed Poisson with rate $\lambda_{t,d,r}$, (ii) arrivals are independent of price (see Online Appendix C.5 for supporting evidence) and $\xi_{j,t,d,r}$; (iii) consumers have no knowledge of remaining capacity; (iv) consumers solve the above utility maximization problems. With these assumptions, conditional on prices and product characteristics, demand for flight j is equal to

$$\tilde{q}_{j,t,d,r} \sim \text{Poisson}(\lambda_{t,d,r} \cdot s_{j,t,d,r}).$$

Realized demand is equal to $q_{j,t,d,r} = \min\{\tilde{q}_{j,t,d,r}, C_{j,t,d,r}\}$.

5.3 Empirical Specification

We assume that consumer utility is given by

$$u_{i,j,t,d,r} = \beta_{0,r} - \alpha_{\ell(i),r} p_{j,t,d,r} + \text{FE}_r(\text{Time of Day } j) + \text{FE}_r(\text{Week}) + \text{FE}_r(\text{DoW}) + \xi_{j,t,d,r} + \varepsilon_{i,j,t,d,r},$$

where "FE" denotes fixed effects for the variable in parentheses. We parameterize the probability an arrival is of the business type as

$$\gamma_{t,r} = \frac{\exp(f_r(t))}{1 + \exp(f_r(t))},$$

where $f_r(t)$ is an orthogonal polynomial basis of degree five with respect to days from departure. This specification allows for non-monotonicities while producing values bounded between zero and one. We specify the arrival processes using a multiplicative relationship between day before departure and departure date fixed effects, i.e., $\lambda_{t,d,r} = \exp(\lambda_{t,r} + \lambda_{d,r})$, because consumer arrivals are observed at the (t, d, r) level of granularity. This specification captures that searches tend to increase over time or evolve discontinuously ($\lambda_{t,r}$), and we observe strong departure-date effects ($\lambda_{d,r}$).

Our method accounts for the fact that we do not observe all searches, but we do observe all bookings (e.g., a booking through a travel agency). Essentially, our estimation procedure scales up estimated search rates based on what fraction of bookings our search data account for.²³ Using properties of the Poisson distribution, we assign $A_{t,d,r}^L \sim \text{Poisson}(\lambda_{t,d,r}(1 - \tilde{\gamma}_{t,r})/\zeta_{t,r}^L)$, and $A_{t,d,r}^B \sim \text{Poisson}(\lambda_{t,d,r}\tilde{\gamma}_{t,r}/\zeta_{t,r}^B)$, where $\tilde{\gamma}_{t,r}$ is derived from the RM department's passenger assignment algorithm, and $\zeta_{t,r}^\ell$ is the fraction of bookings that occur through indirect channels.²⁴ That is, we use the relative fraction of L (and B) sales and searches across channels to scale up L (and B) arrivals. This logic follows the simpler case with a single consumer type: if channels with observed searches account for 20% of bookings, and unobserved searches involve the same underlying demand distributions, we can scale up estimated arrival rates by 5 \times . The reason for our more complex adjustment is that our search data under/over rep-

²³Additional descriptive analyses in Online Appendix C.5 motivate adjusting arrival rates differently over time.

²⁴We use time intervals early on due to data sparsity. Closer to the departure date, the intervals become length one. We smooth the calculated fractions using a fifth-order polynomial approximation.

resent certain consumers, e.g., some business consumers tend to book through travel agencies. Our method accommodates this data feature. We also conduct robustness to this specification (see Online Appendix D.2) and obtain quantitatively and qualitatively similar demand estimates.

5.4 Estimation Procedure

We use a hybrid-Gibbs sampler to estimate route-specific parameters. Our model allows us to rationalize the large number of zero-sale observations while maintaining a Bayesian IV correlation structure between price and the aggregate demand shock ξ . Our approach builds on the estimation procedure developed by Jiang, Manchanda, and Rossi (2009) by incorporating arrival processes, Poisson demand, and censored demand. Additional details can be found in Online Appendix D.1. A complete treatment can be found in Hortaçsu, Natan, Parsley, Schweg, and Williams (2022).

5.5 Identification and Instruments

Estimating models of aggregate demand uncertainty require separably identifying shocks to arrivals from shocks to preferences. We address this complication by using arrivals data. Conditional on market size, preference parameters are identified using the same variation commonly cited in the literature on estimating demand for differentiated products using market-level data. The flight-level characteristic parameters are identified from the variation of flights offered across markets, and we identify the price coefficient using instrumental variables.

We use the shadow price of capacity as reported by the pricing heuristic, advance purchase indicators, and total number of inbound or outbound bookings from a route's hub airport as our demand instruments. The shadow price is the firm's estimate of "marginal cost"—we instrument for price based on the heuristic's estimate of opportu-

nity costs. The advance purchase indicators account for the fact that prices may adjust even in situations where opportunity costs are not observed to change (see Figure 15 in Online Appendix C.4). The number of inbound/outbound bookings to a route’s hub airport is a congestion instrument that captures the change in opportunity costs driven by demand shocks in other markets.²⁵

Note that our identification argument does not rely on optimal pricing. The pricing heuristic is fully determined by its calculated opportunity costs, which we use as instruments in a flexible way (by also adding quadratic and cubic terms). Our instruments are relevant and highly correlated with price because our pricing equations essentially approximate the heuristic’s decisions. For example, suppose that the forecast for a flight is persistently biased upward, but a flight’s demand shocks tend to be low. Our approach rationalizes this scenario because large opportunity costs rationalize higher prices. Our estimates recover low ξ s which differ in a flexible way from the pricing equation unobservables (see Online Appendix D). In fact, these correlations may even be negative. We estimate pseudo first-stage R^2 s average 0.72 across routes.²⁶

5.6 Demand Estimates

We select a subset of routes for estimation where our air carrier is the only airline providing nonstop service. Our estimation sample includes routes with varying flight

²⁵For a route with origin O and destination D , where D is a hub, the total number of outbound bookings from the route’s hub airport is defined as the following; $\sum_{i=1}^K Q_{D,D'}$. Where $Q_{D,D'}$ is the total number of bookings in period t , across all flights, for all K routes where the origin is the original route’s destination. If the route’s origin is the hub, we calculate the total number of inward bound bookings, which equals $\sum_{i=1}^K Q_{O',O}$. Where $Q_{O',O}$ is the total bookings from all K routes where the original routes origin is the destination. For example, for a flight from A to B , where B potentially provides service elsewhere and is a hub, we use all traffic from B onward to other destinations C or D . We assume demand shocks are independent across markets, so shocks to $B \rightarrow C$ and $B \rightarrow D$ are unrelated to demand for $A \rightarrow B$. Thus, a positive shock to onward traffic, out of hub B , will raise the opportunity cost of serving $A \rightarrow B \rightarrow C$ or $A \rightarrow B \rightarrow D$. This propagates to price set on the $A \rightarrow B$ leg.

²⁶For additional flexibility, we also allow for the variance in the pricing equation unobservable to vary over time. This allows us to account for the fact that the RM department manages flights differently over time in an observable way.

frequencies, importance of seasonality, and percentage of nonstop and non-connecting traffic. Online Appendix B discusses the estimation sample in more detail. In total, we estimate nearly 20,000 demand parameters across 39 routes.

Table 2: Demand Estimates Summary across Markets

Parameter	Mean	Std. Dev.	Median	25th Pctile.	75th Pctile.
Day of Week Spread	32.53	19.61	28.19	17.55	39.81
Flight Time Spread	74.99	59.29	45.45	34.70	95.95
Week Spread	52.35	61.90	35.12	21.98	56.62
Intercept	-1.095	1.274	-0.777	-1.405	-0.509
α_B	0.286	0.167	0.277	0.165	0.376
α_L	1.764	0.736	1.834	1.169	2.199

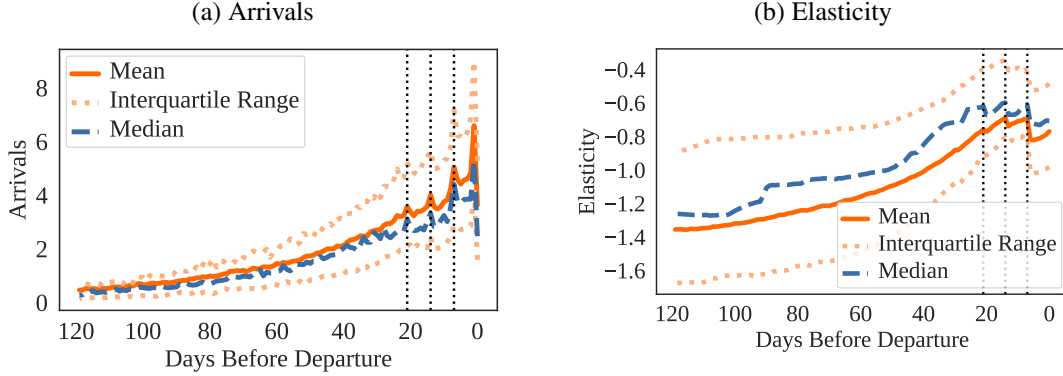
Note: Spread refers to the dollar amount a leisure consumer would pay to move from the least preferred time or day offered to the most preferred time or day of week. Arrival parameters refer to the variation in search across flight departure day of week.

We summarize our demand results given the number of parameters estimated. We find that our demand model accurately matches aggregated arrivals at both the day before departure and departure date level. We accurately match average booking rates even though 88% of observations have zero sales. However, we note a slight downward bias in quantity demanded which is driven by very infrequent large group bookings that are hard to predict. We obtain noisy estimates at very granular levels (predicting demand for a flight at a specific point in time) due to both the discrete nature of the data and the high number of zeros. For example, we estimate pseudo- R^2 s of around 0.23 for predicting particular (j, t, d, r) demand.

We summarize parameter estimates in Table 2. The first panel describes the spread in willingness to pay (in dollars) for a leisure consumer to switch between the most and least-preferred option (day of the week, time of the day, week of the year). We estimate this spread to be \$75. This tends to be less than the spread in flight prices (within a day) observed in the data. Time of day preferences tend to be stronger than day of week preferences (a spread of \$33). Consumers generally prefer morning and late afternoon departure times. We estimate that some weeks have systematically higher

demands than other weeks—for example, major holidays. However, this is not true for all routes, and it does not always reflect seasonal variation in demand.

Figure 4: Aggregate Arrivals and Elasticities



(a) Estimated arrival rates aggregated over all 39 routes. (b) Estimated Own Price Elasticity of demand aggregated over all 39 routes.

In Figure 4-(a), we plot arrival rates for the average route as well as the interquartile range across routes. Although levels of arrivals vary—the interquartile range spans more than a doubling of arrivals—overall, search increases as the departure date approaches. This is important because it means that the observed increase in booking rates is not entirely driven by late-arriving, price insensitive consumers. In panel (b), we plot the average own-price elasticities for the mean, median, and interquartile range over routes. We find that demand elasticities increase (toward zero) due to a significant shift in demand towards business customers. The decline in elasticities close to the departure date mostly reflect very significant price increases. We estimate the average overall elasticity to be 1.05. We frequently find inelastic demand close to the departure. This is also observed using the firm’s demand models (see Section 4).

We briefly note that these estimates are robust to alternative specifications. These include alternative fixed effects (day before departure) in demand, alternative instruments (fuel costs, remaining capacity), and different arrival scaling factors (see Online Appendix D.2). All of these specifications provided similar elasticity patterns to the

results reported here.

6 Analysis of Department Pricing Input Decisions

With demand estimated, we conduct counterfactual simulations where departments supply input decisions to the pricing heuristic. Prices are set by the observed pricing heuristic, and we simulate market outcomes using our demand estimates. We first provide additional details on our counterfactual implementation. We then analyze deviations from current department input decisions.

6.1 Counterfactual Implementation

For each counterfactual, we simulate flights based on the empirical distribution of remaining capacity 120 days before departure. We simulate 10,000 flights for every departure date. Like our demand model (and the RM department’s demand model), we do not endogenize connecting (or flow) bookings. We handle connecting bookings through exogenous decreases in remaining capacity based on Poisson rates estimated using connecting bookings data. After the heuristic determines price based on department input decisions, consumers arrive (in discrete time) according to our model estimates. Demand is then realized. If demand exceeds remaining capacity, consumers are offered seats in the order they arrive.²⁷

Our baseline counterfactual results use the observed department input decisions. Recall that the pricing department’s inputs are discrete fare menus, and the RM department’s observed decisions are demand forecasts. The forecasts are the original single-product demand models that include RM analyst adjustments. Using the demand forecasts directly in our counterfactuals is difficult to two reasons. First, the

²⁷Using the pricing heuristic, if the lowest-priced fare has a single seat and is sold immediately, the next arriving consumer within a period is offered the next least-expensive fare. This occurs rarely.

forecasting data is so large that the firm does not keep an entire record for every flight. The data are not sparse as we observe hundreds of thousands of observations for each route. However, the data do not provide a complete panel for every flight. Second, we cannot model all seven analyst adjustments to the demand models in our counterfactual analysis. However, because the most commonly applied adjustment is a uniform scaling factor that affects several routes and departure dates simultaneously, we take the lesser used adjustments (the other six) as given and endogenize the most commonly used adjustment used in counterfactuals.

The forecasting data are inputted into the heuristic in our counterfactuals using an auxiliary demand model. More precisely, we leverage the ideas of the empirical demand estimation literature by conducting demand estimation using (forecasted) demand curves, rather than observed demand quantities. Using this procedure, we “re-estimate” the same (20,000) parameters from the demand model presented in Section 5 using the forecasting demand. This allows us to fill in any missing values in the forecasts. We discuss differences between our demand estimates and the ones derived from the forecasting data in the next section. For additional details of this procedure, see Online Appendix E.

6.2 Establishing that Current Inputs are Mutual Best Responses

We compare market outcomes using current department input decisions to two sets of department deviations. In the first set of counterfactuals, we consider the pricing department solving the “coordination” problem. Recall, the pricing department commonly files fares on the inelastic side of the RM department’s demand models. In this deviation, the pricing department removes all fares inconsistent with standard profit maximization according to the RM department’s demand models. We adjust the menus so that no fares can be offered below the point where marginal revenue equals zero. Fares below that are inconsistent with profit maximization when capacity costs

are zero. We retain all fares higher than this threshold, which may be offered based on both demand expectations and realized demand.

We consider this pricing department deviation for every route-departure date-day before departure tuple and then aggregate over all simulations. We report counterfactual market outcomes in Table 3, where we normalize outcomes under current inputs to 100 (leisure and business consumer surplus, quantity sold, revenues, and welfare). At first, our results may seem counterintuitive: addressing the miscoordination problem by removing fares on the inelastic side of the RM department’s demand models lowers revenues by 5% and also reduces both leisure and business consumer surplus.

Coordinating the fare menus to the RM department’s input is not a profitable deviation because it entails coordination to biased inputs. The demand model used in practice is misspecified, e.g., it assumes single-product demand; in addition, the heuristic itself is inherently biased, e.g., it also considers every flight in isolation and is dynamically inconsistent. As a result, we find it is optimal for the pricing department to set choose fares that are inconsistent with basic economic theory (“inelastic fares”) because it improves revenues relative to removing those fares from the menus. Because this deviation is not profitable, we conclude that the pricing department’s inputs are best responses given its organizational decision rights.

Table 3: Counterfactual Estimates and a Comparison to Present Practices


Counterfactual	CS_L	CS_B	Q	Rev	W
Observed Inputs	100.0	100.0	100.0	100.0	100.0
Pricing Department Deviation	59.8	98.7	76.5	95.1	95.1

Note: In counterfactual (1), we approximate current pricing practices. Counterfactual (2) examines the pricing department input deviation.

In the second set of deviations, we empirically examine the manipulations the RM department applies to its demand models, holding the pricing menus fixed. As previously stated, the RM department typically scales up or down demand for many routes

simultaneously with a single scaling parameter. Therefore, we implement this counterfactual in a similar way: the RM department scales demand by χ , and this input is given to the pricing heuristic. We consider $\chi \in \{0.25, 0.5, 0.75, \dots, 5.0\}$. We select the χ that results in the highest total expected revenues for the routes studied.

Figure 5: Optimal Forecast Bias, Holding Fare Menus Fixed



Note: Counterfactual revenues under alternative RM department forecasts.

We summarize our analysis of RM department deviations in Figure 5, where we plot revenues as a function of the forecast bias (χ). We normalize results to the current forecast bias. These counterfactuals establish two important results. First, the current bias is *nearly* optimal for the firm. On the one hand, reducing the bias results in a decline in revenues, confirming that unilateral bias reduction is suboptimal. On the other hand, we find that the bias would have to be increased over 50% in order to obtain a less than 1% increase in expected revenues. Although we do observe such overstated forecasts in the data (forecasts overstated by over 50%), analysts tend to quickly revise these forecasts downward, suggesting that very biased inputs—including the optimal forecast bias calculated here—is an impermissible input. Based on our discussions with the firm, we postulate that large forecast biases are discouraged because forecast bias is a key-performance indicator (KPI) in addition to revenue yield.

Second, RM department deviation counterfactuals also emphasize the classic theories in organizational economics of limited productivity gains under unilateral change

when complementarities are important (Milgrom and Roberts, 1990, 1995). The results in Figure 5 can be viewed through another lens: significant adjustments to the current demand models when the heuristic and fare menus are fixed yields fairly incremental revenue changes. This occurs because department input decisions can be reinforcing or offsetting, as we discussed in Section 4.1.1.

Our department deviation counterfactuals establish that providing biased inputs is a mutual best-response given the pricing heuristic. Observed prices can be rationalized as an equilibrium of a team-theoretic game where departments have decision rights for different inputs that are supplied to the observed pricing heuristic. Departments are making decisions that are boundedly rational, e.g., the pricing department does not leverage demand models in its decision making; the RM department use single-product demand models are all routes. Nonetheless, each department is acting rationally given its decision rights and available information.

This naturally leads to several follow-up questions, including: are departments at a “saddle point” in the team-theoretic game, or, what input decisions are in the set of optimal equilibria? Unfortunately, answering these questions is exceedingly difficult. For example, in order to examine the pricing department’s best response to a change in the forecast bias, this would require optimizing over sets of menus (each containing hundreds of entries), where each fare chosen can affect revenues endogenously through the opportunity cost of capacity. In fact, every adjustment by any department can affect how the heuristic allocates capacity. This highlights the complexities of airline pricing. Because the complex objective the firm faces is an intractable problem, it has decomposed the problem it faces through delegated decision-making. However, department decisions are just as complex because pricing inputs are complementary—prices depend on all input decisions. Optimizing over all combinations of pricing inputs—menus and forecasts—is computationally infeasible and beyond the scope of this paper. Therefore, we instead ask a complementary question: how far are observed

outcomes from what is predicted by a model of unbounded rationality. We investigate the canonical dynamic pricing model that internalizes across-flight substitution, is dynamically consistent, and represents a single-entity determining price in the next section.

7 Analysis under a Benchmark Dynamic Pricing Model

Our final counterfactual implements a benchmark model of dynamic pricing (DP) assuming that the firm is a centralized entity using unbiased inputs. We abstract from observed department decisions, such as the use of single-product demand models and RM analyst adjustments that are used to counteract pricing biases within organizational decision rights. In this counterfactual, we use our estimated demand system to simulate consumer purchasing decisions and as the demand model from which dynamic prices are determined. The scenario mimics standard practices in the empirical literature of flexibly estimating demand and assuming standard profit maximization. Relative to observed pricing practices, DP internalizes that prices will be reoptimized over time as well as non-zero cross-price elasticities based on our demand estimates. The firm solves (j, d, r subscripts suppressed)

$$V_0(C_0) = \max_{\{p_t\}_{t=0}^T} \mathbb{E} \sum_{t=0}^T p_t \cdot \min\{\tilde{q}_t(p_t), C_t\}, \quad (1)$$

such that $C_{t+1} = C_t - \min\{\tilde{q}_t(p_t), C_t\}$, unsold units are scrapped, and $C_0 \geq 0$ is given.

We make two additional modeling choices. First, although this model is well-defined for an arbitrary number of flights, in practice, the model becomes intractable quickly—routes with greater than two flights contain more than 15 million states and a transition kernel size of one trillion. Therefore, we restrict this analysis to routes with at most twice daily service. Second, we retain discrete prices for computational

tractability, as the use of discrete prices allows for significantly faster computation than the use of continuous prices. We construct fare menus based on the distribution of observed fares. This implementation allows us to quantify how DP would adjust the distribution of consumer-realized prices rather than select prices outside the typical set faced by consumers.²⁸

Table 4: Counterfactual Estimates under a Benchmark Dynamic Pricing Model

Counterfactual	CS_L	CS_B	Q	Rev	W
Observed Inputs	100.0	100.0	100.0	100.0	100.0
Benchmark Dynamic Pricing Model	108.1	76.7	97.8	114.3	94.3

Note: Comparison of market outcomes under a benchmark dynamic pricing model to the current observed pricing practice. The DP results use our demand estimates for optimization. Both results compute surplus using our demand estimates.

In Table 4, we report counterfactual results comparing the benchmark DP model to observed practices, which we have normalized to 100 in an analogous way to Table 3. Our key finding is that incorrectly assuming that the firm solves the canonical pricing model with unbiased inputs understates welfare, overstates leisure consumer surplus, and understates business consumer surplus. The magnitudes are significant.


DP suggests significant changes in capacity allocation over time compared to observed practices. Current department input decisions result in overstating early demand, not only relative to internally estimated demand models, but also relative to our demand estimates. Therefore, DP results in lower prices early on because it suggests that early opportunity costs are actually lower than those calculated using observed pricing practices. As a result, leisure consumer surplus is overstated by 8% if we incorrectly assume the firm uses DP to set prices based on our estimated demand model.

²⁸We construct the fare menus as follows. We define the lowest price on the menu to be the fare such that demand is unit elastic on the day before departure with the most price sensitive demand (typically the earliest period). We use this fare because analysis of our estimated demand system and the observed fare menus show that fares are often too high. We specify the highest fare on the menu to be the most expensive economy-class fare observed in the data. Thus, the menus are truncated versions of observed fare menus.

We find the opposite welfare result for business consumers. DP overstates how current practices segment markets. This is shown in Figure 6-(a), which plots the distribution of fares offered under current practices with those predicted under DP for an example route. DP results in far greater dispersion in prices offered, with lower prices initially offered and higher prices closer to departure. In Figure 6-(a), there is a mass of DP prices between \$300 and \$400. Higher prices under DP close to the departure date reflect an additional economic force. The lower fares offered early on under DP exacerbate scarcity effects toward the departure date. DP overstates the magnitude of price changes over time and understates business consumer surplus compared to actual outcomes by 23%. Two additional findings are notable. First, DP overstates revenues by 14%. Second, in total, we estimate that DP understates the efficiency of the airline routes we study by 6%.

Fundamentally, this counterfactual emphasizes the importance of accounting for organizational structure, department decisions, and the use of heuristics in rationalizing firms' pricing decisions. Consider the RM department's single-product demand models after RM analysts adjustments (forecasts). While the previous counterfactual establishes that these input decisions are rational given the department's decision rights and constraints of the heuristic, these inputs also distort demand predictions relative to the truth. In Figure 6-(b), we plot average own-price elasticities at observed prices using our estimated demand system and the implied elasticities from the forecasting data. Both lines share some common patterns, such as aggregate price responsiveness decreasing over time and elasticities decreasing (more negative, as expected) after significant observed price increases. However, the elasticities also differ significantly in levels and variability, where observed input decisions lead to understating consumer heterogeneity due to misspecification (single-product demand) and functional form (one set of parameter estimates is applied to roughly 450 routes; see also Figure 17 in Online Appendix E). In addition, demand is often inelastic close to the departure

date based on our model estimates (also true for 38% of observations based on internal demand estimates; see Section 4). Critically, the observed heuristic does not approximate DP, which requires that demand is at least unit elastic. Optimally pricing under DP shifts the (orange) elasticity curve to at or below -1 , raising some prices—shifting the second grey peak in Figure 6-(a) to the right. Compared to the observed persistently biased forecast, DP also lowers some prices, moving the first grey peak in Figure 6-(a) to the left. We further emphasize that correctly modeling firms’ pricing decisions has important welfare consequences in the next subsection.



Note: (a) Simulated prices using current practices and DP for an example route. (b) Comparison of own-price elasticities over time based on estimated demand model and the forecasting data.

7.1 Imposing Dynamically Optimal Profit Maximization in Demand Estimation

Our results in the previous subsection show that incorrectly assuming that the firm optimizes according to the benchmark dynamic pricing problem yields incorrect predictions of market outcomes. We elaborate on this insight through a complementary exercise by quantifying how imposing a model of standard profit maximization in demand estimation affects estimates of willingness to pay.

Supply-side optimality conditions offer useful restrictions when estimating empiri-

cal models of demand (e.g., Berry, Levinsohn, and Pakes, 1995). In order to investigate how imposing the assumption that the firm solves a DP affects demand estimates, we follow the estimation strategy of Williams (2022) in the context of dynamic pricing. Similar ideas are used in, e.g., Aryal, Murry, and Williams (2022), Pan and Wang (2022), and Cho, Lee, Rust, and Yu (2018). We exploit restrictions imposed by Equation 1 in demand estimation, which is analogous to using the first-order condition of a static Bertrand-Nash game in differentiated product markets with competition. We estimate the Poisson-Random Coefficients model by incorporating conditional choice probability (CCP) restrictions that rationalize observed prices as optimal based on differences in choice-specific value functions (see, e.g., Williams, 2022).

We find that demand estimates imposing DP differ substantially from our estimated demand system because DP is inconsistent with observed pricing practices. Our demand estimates (and internal demand estimates) result in average elasticities that are close to -1, we find that imposing optimality with DP results in average demand elasticities close to -4. The large discrepancy is due to the difficulty in rationalizing observed prices as solutions to the benchmark model of dynamic pricing. For example, our descriptive evidence showed that when department input decisions are combined by the pricing heuristic, resulting prices are often on the inelastic side of internally estimated demand models. They are also on the inelastic side of our estimated demand system. In order to rationalize these prices as “optimal” in DP, demand must be at least unit elastic. Imposing optimality in demand estimation results in biasing the demand estimates toward being too elastic, resulting in, on average, understating consumer willingness to pay and thus consumer welfare. This analysis does not suggest that supply-side assumptions are not useful in demand estimation. Rather, supply-side assumptions that do not approximate how the firm actually structures the problem it faces can lead to biased demand estimates, and, consequently, misleading estimates of welfare.

8 Conclusion

This paper presents five main findings. First, we document the extent to which large firms decompose the problem of pricing to distinct departments. Using job listings, we show that all major airlines, cruises, car rentals, and hotels leverage an organizational structure in which departments have distinct pricing-related decision rights. Second, using data and internal models from a large U.S. airline, we demonstrate several ways in which pricing practices diverge from dynamically optimal profit maximization. For example, the firm relies on a heuristic for tractability reasons that does not solve or approximate a dynamic program. Flights are priced independently, even though flight substitution, including available nonstop options, are measurably important drivers of demand. Third, we show that although pricing is inconsistent with benchmark models of dynamic pricing, observed prices are consistent with a team-theoretic game in which departments determine pricing inputs that are supplied to the observed pricing heuristic. Fourth, accounting for organizational structure, department input decisions, and the use of heuristics is critical for accurately modeling firm pricing behavior. Our counterfactuals establish that we mismeasure welfare when we assume that the firm has is unboundedly rational. We understate overall welfare, overstate revenues, and mismeasure surplus across consumer types. Fifth, we highlight this result by showing that estimates of willingness to pay and market power may be significantly biased if we assume that observed prices are derived from the firm solving canonical models of dynamic pricing.

Studying firms solving complex problems is inherently difficult; however, understanding how firms actually structure the problems they face offers important insights into their decision-making processes and provides guidance on how to more accurately measure market outcomes.

References

- Akepanidaworn, K., R. Di Mascio, A. Imas, and L. Schmidt (2019): “Selling fast and buying slow: Heuristics and trading performance of institutional investors,” Working Paper.
- Alonso, R., W. Dessein, and N. Matouschek (2008): “When does coordination require centralization?,” *American Economic Review*, 98(1), 145–79.
- Aryal, G., C. Murry, and J. W. Williams (2022): “Price Discrimination in International Airline Markets,” Working Paper.
- Asker, J., C. Fershtman, and A. Pakes (2021): “Artificial Intelligence, Algorithm Design and Pricing,” .
- Atkin, D., A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen (2017): “Organizational barriers to technology adoption: Evidence from soccer-ball producers in Pakistan,” *The Quarterly Journal of Economics*, 132(3), 1101–1164.
- Belobaba, P. (1987): “Air Travel Demand and Airline Seat Inventory Management,” Ph.D. thesis, Massachusetts Institute of Technology.
- Berman, R., and Y. Heller (2020): “Naive analytics equilibrium,” *arXiv preprint arXiv:2010.15810*.
- Berry, S., M. Carnall, and P. Spiller (2006): “Airline hubs: costs, markups and the implications of customer heterogeneity,” *Advances in airline economics*, 1, 183–214.
- Berry, S. T. (1994): “Estimating Discrete-Choice Models of Product Differentiation,” *The RAND Journal of Economics*, pp. 242–262.
- Berry, S. T., J. Levinsohn, and A. Pakes (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, pp. 841–890.
- Bohren, J. A. (2016): “Informational herding with model misspecification,” *Journal of Economic Theory*, 163, 222–247.
- Brown, Z. Y., and A. MacKay (2021): “Competition in pricing algorithms,” Discussion paper, National Bureau of Economic Research.
- Calvano, E., G. Calzolari, V. Denicolo, and S. Pastorello (2020): “Artificial intelligence, algorithmic pricing, and collusion,” *American Economic Review*, 110(10), 3267–97.
- Cho, S., G. Lee, J. Rust, and M. Yu (2018): “Optimal Dynamic Hotel Pricing,” Working Paper.

- Cooper, W. L., T. Homem-de Mello, and A. J. Kleywegt (2006): “Models of the spiral-down effect in revenue management,” *Operations research*, 54(5), 968–987.
- Cyert, R. M., and J. G. March (1963): *A behavioral theory of the firm*, vol. 2. Englewood Cliffs, NJ.
- DellaVigna, S., and M. Gentzkow (2019): “Uniform pricing in us retail chains,” *The Quarterly Journal of Economics*, 134(4), 2011–2084.
- Dessein, W., A. Galeotti, and T. Santos (2016): “Rational inattention and organizational focus,” *American Economic Review*, 106(6), 1522–36.
- D’Haultfœuille, X., P. Février, A. Wang, and L. Wilner (2022): “Estimating the Gains (and Losses) of Revenue Management,” Working Paper.
- Dubé, J.-P., and S. Misra (2021): “Personalized Pricing and Consumer Welfare,” Working Paper.
- Ely, J. C. (2011): “Kludged,” *American Economic Journal: Microeconomics*, 3(3), 210–31.
- Heidhues, P., B. Kőszegi, and P. Strack (2018): “Unrealistic expectations and misguided learning,” *Econometrica*, 86(4), 1159–1214.
- Hortaçsu, A., O. Natan, H. Parsley, T. Schweg, and K. Williams (2022): “Incorporating Search and Sales Information in Demand Estimation,” Working Paper.
- Hortaçsu, A., A. Öry, and K. R. Williams (2022): “Dynamic Price Competition: Theory and Evidence from Airline Markets,” .
- Jiang, R., P. Manchanda, and P. E. Rossi (2009): “Bayesian Analysis of Random Coefficient Logit Models using Aggregate Data,” *Journal of Econometrics*, 149(2), 136–148.
- Levitt, S. D. (2016): “Bagels and donuts for sale: A case study in profit maximization,” *Research in Economics*, 70(4), 518–535.
- Li, J., N. Granados, and S. Netessine (2014): “Are consumers strategic? Structural estimation from the air-travel industry,” *Management Science*, 60(9), 2114–2137.
- Littlewood, K. (1972): “Forecasting and control of passenger bookings,” *Airline Group International Federation of Operational Research Societies Proceedings*, 1972, 12, 95–117.
- Ma, Y., T. Ropele, D. Sraer, and D. Thesmar (2020): “A quantitative analysis of distortions in managerial forecasts,” Working Paper.
- Massey, C., and R. H. Thaler (2013): “The loser’s curse: Decision making and market efficiency in the National Football League draft,” *Management Science*, 59(7), 1479–1495.

- McAfee, P., and V. Te Velde (2006): “Dynamic Pricing in the Airline Industry,” *Handbook on Economics and Information Systems*.
- McGill, J. I., and G. J. Van Ryzin (1999): “Revenue management: Research overview and prospects,” *Transportation science*, 33(2), 233–256.
- Milgrom, P., and J. Roberts (1990): “Rationalizability, learning, and equilibrium in games with strategic complementarities,” *Econometrica: Journal of the Econometric Society*, pp. 1255–1277.
- (1995): “Complementarities and fit strategy, structure, and organizational change in manufacturing,” *Journal of accounting and economics*, 19(2-3), 179–208.
- Orhun, A. Y., T. Guo, and A. Hagemann (2022): “Reaching for Gold: Frequent-Flyer Status Incentives and Moral Hazard,” *Marketing Science*, 41(3), 548–574.
- Pan, Q., and W. Wang (2022): “Costly price adjustment and automated pricing: The case of Airbnb,” *Available at SSRN 4077985*.
- Radner, R. (1993): “The organization of decentralized information processing,” *Econometrica: Journal of the Econometric Society*, pp. 1109–1146.
- Rantakari, H. (2008): “Governing adaptation,” *The Review of Economic Studies*, 75(4), 1257–1285.
- Rust, J. (2019): “Has dynamic programming improved decision making,” *Annual Review of Economics*, 11(1), 833–858.
- Sacarny, A. (2018): “Adoption and learning across hospitals: The case of a revenue-generating practice,” *Journal of health economics*, 60, 142–164.
- Siggelkow, N. (2001): “Change in the presence of fit: The rise, the fall, and the renaissance of Liz Claiborne,” *Academy of Management journal*, 44(4), 838–857.
- Simon, H. A. (1955): “A behavioral model of rational choice,” *The quarterly journal of economics*, pp. 99–118.
- Simon, H. A. (1962): “The Architecture of Complexity,” *Proceedings of the American Philosophical Society*, 106.
- Vinod, B. (2021): *Evolution of Yield Management in the Airline Industry*. Springer.
- Williams, K. R. (2022): “The Welfare Effects of Dynamic Pricing: Evidence from Airline Markets,” *Econometrica*.
- Wollmer, R. D. (1992): “An airline seat management model for a single leg route when lower fare classes book first,” *Operations research*, 40(1), 26–37.

Online Appendix

Organizational Structure and Pricing: Evidence from a Large U.S. Airline

by Hortaçsu, Natan, Parsley, Schwieg, and Williams

Table of Contents

A	Details on the Pricing Heuristic	2
A.1	Littlewood's Rule	2
A.2	EMSR-b Algorithm	3
B	Route Selection	7
B.1	Estimation Sample Comparison	10
C	Additional Descriptive Evidence	13
C.1	Search Patterns	13
C.2	Heuristic Bias	13
C.3	Allocating Inventory to Fares that do not Exist	14
C.4	The Presence of Pricing Frictions	15
C.5	Booking Trends Across Booking Channels	16
D	Additional Details on Demand Estimation	17
D.1	Demand Estimation Procedure	17
D.2	The Impact of the Scaling Factor on Demand Estimates	23
E	Model of Demand Forecasts in Counterfactuals	23

A Details on the Pricing Heuristic

The pricing algorithm used at the firm is based on the well-known heuristic called Expected Marginal Seat Revenue-b or EMSR-b (Belobaba, 1987). The heuristic relies on a concept called Littlewood's Rule (detailed below) and was developed in order to avoid solving highly complex dynamic pricing problems. The heuristic simplifies the firm's decision in each period by aggregating all future sales into a single future period. It requires a single-product demand model. We describe the heuristic below and show how to incorporate Poisson demand in EMSR-b. The heuristic provides an allocation over a given finite set of prices, instead of providing the optimal price itself given any flight's state. EMSR-b associates each price with a fare-class then chooses a maximal number of sales to be offered for each fare-class. This means more than a single price is offered in any given period, however, in practice, consumers almost always choose the cheapest available option. When one class is closed, the next higher priced class opens.

A.1 Littlewood's Rule

EMSR-b is a generalization of Littlewood's rule (Littlewood, 1972), which is a simple case where a firm prices two time periods and uses two fare classes. A firm with a fixed capacity of goods (seats) wants to maximize revenue across two periods, where leisure (more elastic) consumers arrive in the first period and business (less elastic) consumers arrive in the second period. The firm sets a cap on the number of seats b it is willing to sell in the first period to leisure passengers. This rule returns a maximum number of seats for leisure when the price to both leisure and business customers has already been decided; it does not determine optimal pricing.

The solution equates the price of a seat sold in the first period (to leisure travelers) to the opportunity cost of lowering capacity for sales in the second period (to business

travelers). Given prices p_L , p_B , capacity C , and the arrival CDF of business travelers F_B , Littlewood’s rule equates the fare ratio to the probability that business class sells out. The fare ratio is the marginal cost of selling the seat to leisure (the lower revenue p_L) which is set equal to the marginal benefit—the probability that the seat would not have sold if left for business customers only. Littlewood’s rule is given by

$$1 - F_B(C - b) = \frac{p_L}{p_B}.$$

This equation can then be solved for b , the maximum number of seats to sell to leisure customers in period one. This solution is exact if consumers arrive in two separate groups and there are only two time periods and two consumer types.

A.2 EMSR-b Algorithm

The EMSR-b algorithm (Belobaba, 1987) extends Littlewood’s rule to multiple fare levels or classes. For each fare class, all fare classes with higher fares are aggregated into a single fare-class called the “super-bucket.” Once this bucket is formed, Littlewood’s rule applies, and can be done for each fare class iteratively. Rather than just comparing leisure and business classes, the algorithm now weights the choice of selling a lower fare-class ticket against an average of all higher fare classes.

We apply the algorithm for K sorted fare-classes such that $p_1 > p_2 > \dots > p_K$. Each fare class has independent demand with a distribution F_k . Under our specification, the demand for each fare class is distributed Poisson with mean μ_k that is given by future arrivals times the share of the market exclusive to that bucket.

The super-bucket is a single-bucket placeholder for a weighted average of all higher fare-class buckets. Independent Poisson demand simplifies this calculation, as the sum of independent Poisson distributions is itself Poisson. The mean of the super-bucket is the sum of the mean of each higher fare-class bucket. The price of the super-bucket

is a weighted average of the price of each higher-fare class, using the means as the weight.

For each fare class, Littlewood's Rule is then applied with the fare-class taking the place of leisure travel, and the super-bucket in place of business travel. It is assumed that all future arrivals appear in a single day. The algorithm then describes a set of fare-class limits b_k that define the maximum number of sales for each class before closing that fare class. We denote the remaining capacity of the plane at any time by C . The algorithm uses the following pseudo-code:

```

for  $t > 2$  do
  |
  | for  $k \leftarrow K$  to 1 by  $-1$  do
  | |
  | | i) Compute un-allocated capacity  $C_{k,t} = C - \sum_{i=k}^K b_i$ ,
  | |
  | | ii) Construct the super-bucket
  | |
  | | 
$$\mu_{sb} = \sum_{i=1}^{k-1} \mu_i, \quad p_{sb} = \frac{1}{\mu_{sb}} \sum_{i=1}^{k-1} p_i \mu_i, \quad F_{sb} \sim \text{Poisson}(\mu_{sb}),$$

  | |
  | | iii) Apply Littlewood's Rule using the super-bucket distribution as the demand for
  | | business
  | |
  | | 
$$C_{k,t} - b_k = \min \left\{ F_{sb}^{-1} \left( 1 - \frac{p_k}{p_{sb}} \right), C_{k,t} \right\}.$$

  | |
  | end
end

```

In the case where $t = 1$, dynamics are no longer important, so there is no longer a need to trade off based on the opportunity cost. As a result, we limit the fare of the highest revenue class to all remaining capacity, and set limits of all other classes to zero.

A.2.1 Fare Class Demand

What remains is computing the mean μ_k for each fare class bucket. We detail the process in this section. Demand in each market is an independent Poisson with arrival rate $\exp(\lambda_t^t + \lambda_d^d)s_j(p)$. For readability, we suppress the subscript r —all parameters are route-specific. Note that this p is a vector of the prices of all flights in the market. We assume that the firm believes other flights will be priced at their historic average over the departure date and day before departure. This allows us to construct a residual demand function $s_j(p_j)$ that is a function of the price of the current flight only. We will treat this as the demand for the flight at a given bucket's price for the remainder of this section.

Each fare class has a set price p_k , at any time t , departure date d we will see $\exp(\lambda_t^t + \lambda_d^d)$ arrivals, of which $s(p_k)$ are willing to purchase a fare for bucket k . However, $s(p_{k-1})$ are willing to purchase a fare for bucket $k - 1$ as well, since they will buy at the higher price p_{k-1} . Only $\exp(\lambda_t^t + \lambda_d^d)[s_t(p_k) - s_t(p_{k-1})]$ are added by the existence of this fare class with price $p_k < p_{k-1}$. Note that this is a flow quantity—the amount of purchases in time t , but EMSR-B requires stock quantities: How many will purchase over the remaining lifetime of the sale?

What is the distribution of future purchases then? Each day t is an independent Poisson process split by the share function. Independent split Poisson processes are still Poisson, so we may compute the mean of purchases solely in a fare class by summing arrivals over future time t , and taking the difference in shares between price p_k and p_{k-1} . For time t and departure date d , the stock demand for fare-class k is given by

$$\sum_{i=1}^t \exp(\lambda_i^t + \lambda_d^d)[s_t(p_k) - s_t(p_{k-1})],$$

where $s_t(p_0) = 0$ for notational parsimony.

This demand distribution is only used to compute the super-bucket demand distri-

bution. Note that we only include future stock demand in the super bucket, and thus only sum arrivals until time $t - 1$. For fare-class k . The super bucket's stock demand is given by

$$\mu_{sb} = \left(\sum_{i=1}^{t-1} \exp(\lambda_i^t + \lambda_d^d) s_i(p_{k-1}) \right)$$

$$p_{sb} = \frac{1}{\mu} \sum_{j=1}^{k-1} p_j \sum_{i=1}^{t-1} \exp(\lambda_i^t + \lambda_d^d) [s_i(p_j) - s_i(p_{j-1})].$$

The updated pseudo-code for the EMSR-b algorithm is:

for $t > 2$ **do**

for $k \leftarrow K$ **to** 1 **by** -1 **do**

 i) Compute un-allocated capacity $C_{k,t} = C - \sum_{i=k}^K b_i(t)$,

 ii) Construct the super-bucket

$$\mu_{sb} = \left(\sum_{i=1}^{t-1} \exp(\lambda_i^t + \lambda_d^d) \right) s_i(p_{k-1}),$$

$$p_{sb} = \frac{1}{\mu_{sb}} \sum_{j=1}^{k-1} p_j \sum_{i=1}^{t-1} \exp(\lambda_i^t + \lambda_d^d) [s_i(p_j) - s_i(p_{j-1})],$$

$$F_{sb} \sim \text{Poisson}(\mu_{sb}),$$

 iii) Apply Littlewood's Rule using the super-bucket distribution as the demand for business.

$$C_{k,t} - b_k(t) = \min \left\{ F_{sb}^{-1} \left(1 - \frac{p_k}{p_{sb}} \right), C_{k,t} \right\}.$$

end


end

For $t = 1$ we continue to allocate the highest revenue fare class to the entire remaining capacity. Note that for this allocation rule, $b_k(t, d)$ is a function of time since

the arrivals are changing over time. This policy can be computed for each time t and remaining capacity c , for all departure dates d and arrival rates λ .

The algorithm determines the number of seats to assign to each bucket and in particular, the lowest bucket to receive a positive allocation. This bucket is referred to as the lowest available class (LAC). We plot the LAC for an example flight in Figure 7. On the vertical axis, we note the discrete set of fares set by the pricing department, with bucket one being the least expensive and bucket twelve being the most expensive. Little variation in color over days from departure for a given bucket shows that the bucket prices themselves are mostly fixed. However, in the bottom right of the graph, the white space shows that the pricing department has restricted the availability of the lowest fares close to the departure date. Given all pricing inputs, the white line marks the LAC.

Figure 7: Fare Bucket Availability and Lowest Available Fare




Note: Image plot of fare availability over time as well as the active lowest available fare. Bucket1 is the least expensive bucket; Bucket12 is the most expensive bucket. The color depicts the magnitude of prices—blue are lower fares, red are more expensive. White space denotes no fare availability. The white line depicts the lowest available fare.

B Route Selection

We use publicly available data to select markets to study. The DB1B data are provided by the Bureau of Transportation Statistics and contain a 10% sample of tickets sold.

The DB1B does not include the date purchased nor the date traveled and is reported at the quarterly level. Because the DB1B data contain information solely for domestic markets, we limit our analysis to domestic markets as well. Furthermore, we use the air carrier’s definition of markets to combine airports within some geographies.

Figure 8: Nonstop, One-stop and Connecting Traffic

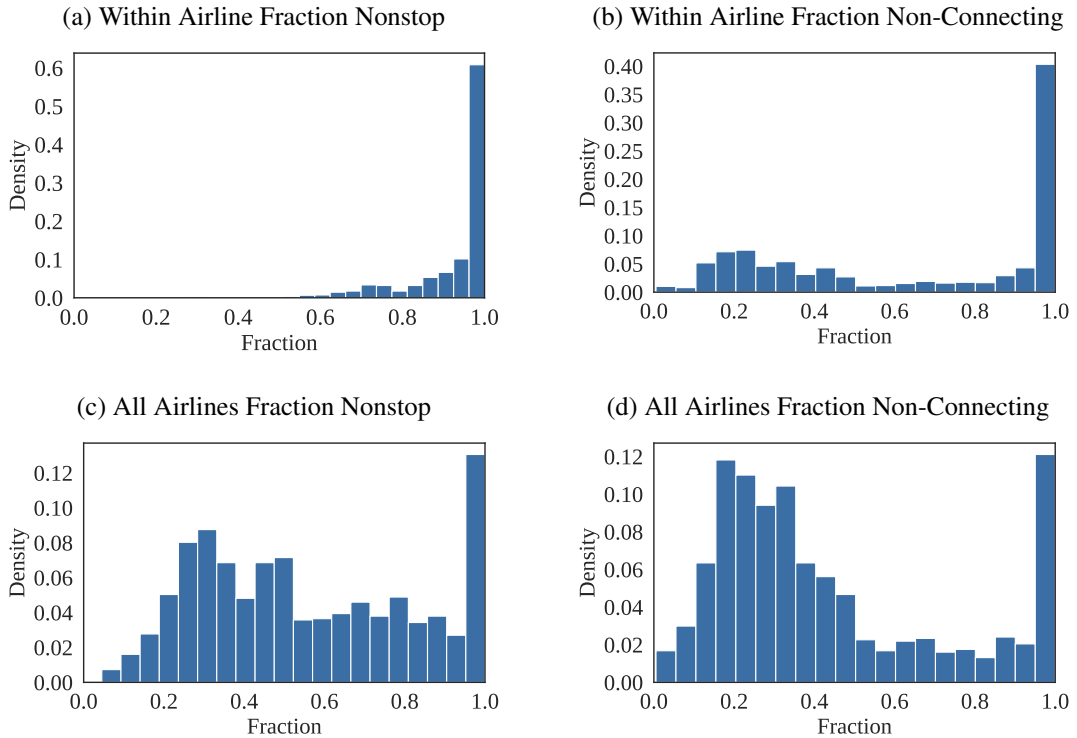


Note: We use the term nonstop to denote the sold black line, or passengers solely traveling between (Origin, Destination). Unless otherwise noted, we will use directional traffic, labeled $O \rightarrow D$. Non-directional traffic is specified as $O \leftrightarrow D$. The blue, dashed lines represent passengers flying on $O \leftrightarrow D$, but traveling to or from a different origin or destination. Finally, one-stop traffic are passengers flying on $O \leftrightarrow D$, but through a connecting airport.

We consider two measures of traffic flows when selecting markets: traffic flying nonstop and traffic that is non-connecting. Both of these metrics are informative for measuring the substitutability of other flight options (one-stop, for example) as well as the diversity of tickets sold for the flights studied (connecting traffic). Figure 8 provides a graphical depiction of traffic flows in airline networks that we use to construct the statistics. We consider directional traffic flows from a potential origin and destination pair that is served nonstop by our air carrier. The first metric we calculate is the fraction of traffic flying from $O \rightarrow D$ nonstop versus one or more stops. This compares the solid black line to the dashed orange line. Second, we calculate the fraction of traffic flying from $O \rightarrow D$ versus $O \rightarrow D \rightarrow C$. This compares the solid black line to the dashed blue line.

Figure 9 presents summary distributions of the two metrics for the markets (ODs) we select. In total, we select 407 ODs for departure dates between Q3:2018 and Q3:2019. The top row measures the fraction of nonstop and connecting traffic for

Figure 9: Route Selection Using Bureau of Transportation Statistics Data

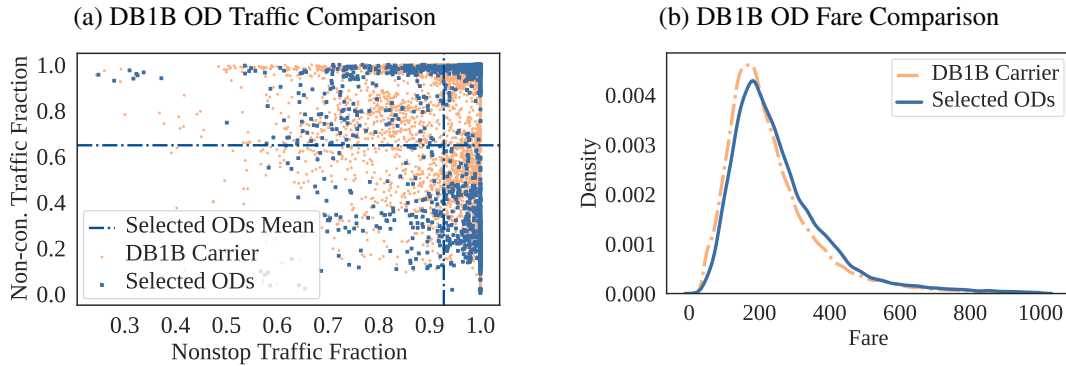


Note: Density plots over the fraction of nonstop traffic and the fraction of non-connecting traffic for the selected routes using DB1B data. "Within" means passengers flying on our air carrier. "Total" means all air carriers on a given origin-destination pair. Within nonstop and total nonstop coincide if our carrier is the only carrier flying nonstop.

tickets sold by our our carrier. The left plot shows that, conditional on the air carrier operating nonstop flights between OD, an overwhelming fraction of consumers purchase nonstop tickets instead of purchasing one-stop connecting flights. The right panel shows that fraction of consumers who are not connecting to other cities either before or after flying on segment OD. There is significant variation across markets, with the average being close to 50%.

The bottom panel repeats the statistics but replaces the denominator of the fractions with the sum of traffic flows across all air carriers in the DB1B. Both distributions shift to the left because of existence of competitor connecting flights and sometimes direct competitor flights. In nearly 75% of the markets we study, our air carrier is the only firm providing nonstop service. Our structural analysis will only consider single carrier

Figure 10: DB1B Comparison



Note: (a) A scatter plot of the fraction nonstop and fraction non-connecting for all origin-destination pairs served by our air carrier. The blue dots show selected markets; the orange dots show non-selected markets. (b) Kernel density plots of all fares in the DB1B data for our air carrier; the blue line shows the density for our selected markets.

markets.

In Figure 10-(a), we show a scatter plot of the fraction of nonstop traffic and the fraction of non-connecting traffic for all origin-destination pairs offers by our air carrier in the DB1B. The orange dots depict routes non-selected markets and the blue dots show the selected markets. We see some dispersion in selected markets, however this is primarily on non-connecting traffic. An overwhelming fraction of the selected markets have high nonstop traffic, although this is true in the sample broadly. Essentially, conditional on the air carrier providing nonstop service, most passengers choose nonstop itineraries. In Figure 10-(b) we show the distribution of purchased fares in the DB1B for our carrier along with our selected markets. The distribution of prices for the selected sample are slightly shifted to the right, which makes sense since we primarily select markets where the air carrier is the only airline providing nonstop service.

B.1 Estimation Sample Comparison

Our estimation sample contains 39 markets. Compared to the overall sample, these routes tend to be smaller in terms of total number of passengers, larger in terms of percentage of nonstop and non-connecting passengers, and nonstop service is provided


only by our air carrier. We report percentage differences between our estimation routes and the entire sample for key characteristics below in Table 5. Figure 11 shows a two-way plot of the fraction of nonstop and non-connecting traffic for the routes selected for estimation relative to the entire sample. Figure 12 recreates Figure 9, separating the estimation sample from the entire sample.

Table 5: Estimation Sample Comparison

Characteristic	Percentage Difference from Mean
Number of Nonstop Passengers	-38.8%
Total Number of Passengers	-33.4%
Number of Local Passengers	-37.7%
Fraction of Traffic Nonstop	1.02%
Fraction of Traffic Non-Connecting	5.91%

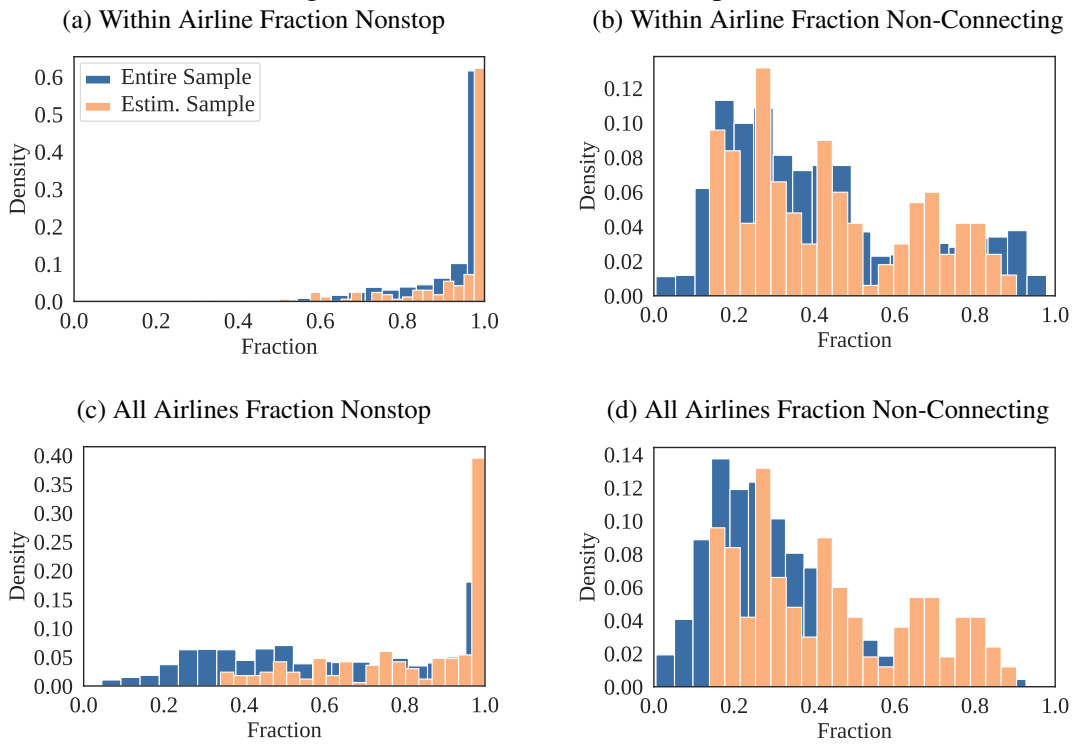
Note: Statistics calculated using the DB1B data for the years 2018-2019.

Figure 11: Route Estimation Selection using DB1B Data



Note: A scatter plot of the fraction nonstop and fraction non-connecting for all origin-destination pairs served by our air carrier. The blue dots show markets used for estimation; the orange dots show non-selected markets.

Figure 12: Estimation Route Comparison




Note: Density plots over the fraction of nonstop traffic and the fraction of non-connecting traffic for the selected routes using DB1B data. "Within" means passengers flying on our air carrier. "Total" means all air carriers on a given origin-destination pair. Within nonstop and total nonstop coincide if our carrier is the only carrier flying nonstop. Blue denotes the entire sample; orange denotes the estimation sample.

C Additional Descriptive Evidence

C.1 Search Patterns

In Figure 13 we plot CDFs on distributions of repeat shoppers. In panel (a) we consider if consumers search multiple departure dates. The plot shows that 80% of consumers search a single departure date. In panel (b) we consider if consumers shop for the same itinerary across days from departure—waiting to purchase. The plot shows that 90% of consumers single once. We do not consider consumers who were referred to the airline’s website, e.g., from meta search engines.



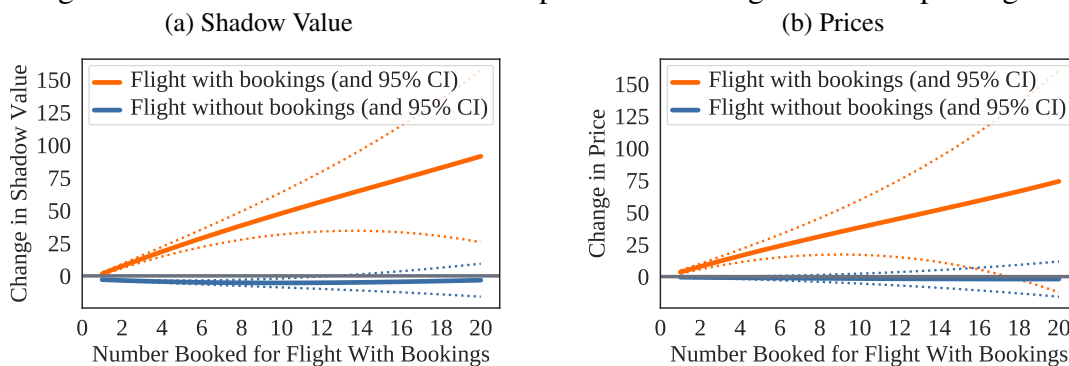
(a) CDF of the number of departure dates searched. (b) CDF of the number of days from departure searched for a given itinerary.

C.2 Heuristic Bias

We select observations that satisfy the following conditions: (i) the firm offers two flights a day; (ii) we include periods where demand is not being reforecasted (the observed spikes in Figure 15); (iii) the total daily booking rate is low (less than 0.5); and (iv) one flight receives bookings and the other flight does not. By considering markets where the total booking rate is low, we can apply theoretical results of continuous time (as well as a discrete time approximation) pricing models. In Figure 14-(a), we plot

the average change in shadow values (opportunity costs) for the flights that receive bookings and for the flights that do not receive bookings (the substitute option) using flexible regressions. In standard dynamic pricing models, every time a unit of capacity is sold, prices jump. This is also true in environments with multiple products—any sale causes all prices to increase. Figure 14-(a) confirms substitute shadow values are unaffected by bookings. Panel (b) shows that there is no price response.

Figure 14: Shadow Value and Price Response to Bookings with Multiple Flights



Note: (a) The orange line denotes the average change in shadow value for a flight with bookings. The blue line is the average change to shadow value when a sale occurs for the substitute product. (b) This panel depicts the same as panel a, but instead of changes in shadow value it depicts changes in price.

C.3 Allocating Inventory to Fares that do not Exist

We observe a form of miscoordination in which the RM department uses obsolete fare menus, meaning that the set of active fares in the market differs from what is inputted into the algorithm. Although this form of miscoordination can be seen as a “glitch,” its presence and prevalence suggests difficulty in processing algorithm inputs across completely separate IT systems.

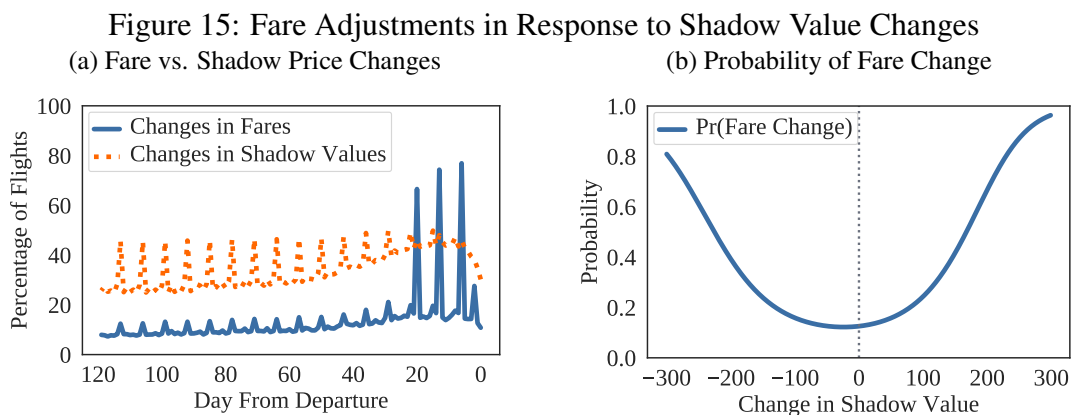
Examining the menu of active fares and the resulting inventory allocations, we observe inventory allocations to fares that do not exist. This affects 11.7% of observations. This is possible because fare validation—the process of ensuring a fare is active—occurs when consumers search to obtain tickets. 32.6% of routes feature

persistent “phantom allocations.” Although we do not isolate the effect of this miscoordination in our counterfactuals, we do quantify how alternative fares affect the heuristic’s decisions in Section 6.

C.4 The Presence of Pricing Frictions

The pricing heuristic requires a discrete set of fares as an input. This naturally gives rise to pricing frictions as fares change in discrete levels but the value of a seat can be any positive value. Sometimes the pricing frictions can be large in magnitude.

In Figure 15-(a), we plot the fraction of flights that experience changes in price or shadow value (as reported by the heuristic) over time. Opportunity costs change much more frequently than do prices. In panel (b), we run a flexible regression of the change in costs on an indicator function of a price adjustment occurring. As the figure shows, changes in opportunity costs exceeding \$100 lead to price adjustments with only a 20% probability.




Note: (a) The fraction of flights that experience changes in the fare or the shadow value of capacity over time. (b) The probability of a fare change, conditional on the magnitude of the shadow value change.

Figure 15-(a) shows noticeable spikes that occur at seven day intervals. This arises because the RM department has chosen to reforecast demand on a 7-day interval. Outside of these periods, remaining inventory is reoptimized without updating future de-

mand expectations. The process of reforecasting demand leads to a larger fraction of flights experiencing a change in the value of remaining capacity.

C.5 Booking Trends Across Booking Channels

Figure 16: Bookings Across Booking Channels



Percentage of bookings, across days from departure, for each channel. Direct refers to bookings that occur on the air carrier’s website, OTAs are bookings made on online travel agencies, and Agency are bookings made through travel agencies.

Figure 16 shows the the distribution of bookings within channel (direct, OTAs, and agency) over days before departure. The distribution of bookings for tickets purchased on OTAs, or online travel agencies, very closely follows the distribution of bookings via the direct channel. However, they do not coincide. The agency curve—which includes corporate travel bookings—is more concentrated closer to departure. There are small spikes in the booking rates across all channels when AP fares expire. Although this may suggest some consumers strategically time market participation, we also find support for the assumption that current time periods simply have higher demands. We partition the data sample into two groups, one group includes routes that do not have a 7-day AP requirement, and the other contains routes where the pricing department files 7-day AP fares. We find that both search and bookings bunch at the 7-day AP requirement, regardless of their existence. Booking rate returns to the pre-bunching levels (or, even higher levels) within one to two days after AP opportunities expire. In fact, the day with the highest booking rate corresponds to the day with the highest

prices—right before departure.

D Additional Details on Demand Estimation

D.1 Demand Estimation Procedure

We provide an overview on the implementation details of each stage the MCMC routine for demand parameter estimation. For readability we suppress the subscript r —all parameters are route-specific. Simultaneously drawing from the joint distribution of our large parameter space is infeasible, therefore, we use a Hybrid Gibbs sampling algorithm. The algorithm steps are shown below. At each step of the posterior sampler, we sequentially draw from the marginal posterior distribution groups of parameters, conditional on other parameter draws. Where conjugate prior distributions are unavailable, we use the Metropolis-Hastings algorithm, a rejection sampling method that draws from an approximating candidate distribution and keeps draws which have sufficiently high likelihood. Additional detail can be found in Hortaçsu, Natan, Parsley, Schwieg, and Williams (2022).

- 1: **for** $c = 1$ to C **do**
- 2: Update arrivals λ (Metropolis-Hastings)
- 3: Update shares $s(\cdot)$ (Metropolis-Hastings)
- 4: Update price coefficients α (Metropolis-Hastings)
- 5: Update consumer distribution γ (Metropolis-Hastings)
- 6: Update linear parameters β (Gibbs)
- 7: Update pricing equation η (Gibbs)
- 8: Update price endogeneity parameters Σ (Gibbs)
- 9: **end for**

Algorithm 1: Hybrid Gibbs Sampler

Sampling Arrival Parameters

We start the sampling procedure by drawing from the posterior distribution of arrival parameters, $\lambda_{t,d}$. The posterior is derived by defining the joint likelihood of arrivals for each consumer type and quantities sold, conditional on product shares. Recall that arriving consumers have likelihood based on their type:

$$\begin{aligned} A_{t,d}^L &\sim \text{Poisson}(\lambda_{t,d}(1 - \tilde{\gamma}_t)\zeta_t^L), \\ A_{t,d}^B &\sim \text{Poisson}(\lambda_{t,d}\tilde{\gamma}_t\zeta_t^B), \end{aligned}$$

where $\tilde{\gamma}_t$ is the probability a consumer is of the business type as derived from the passenger assignment algorithm, and ζ_t^ℓ is the fraction of bookings that do not occur on the direct channel for each consumer type (leisure and business). The purchase likelihood is a function of shares and arrivals and is equal to

$$\begin{aligned} \tilde{q}_{j,t,d} &\sim \text{Poisson}(\lambda_{t,d} \cdot s_{j,t,d}), \\ q_{j,t,d} &= \min\{\tilde{q}_{j,t,d}, C_{j,t,d}\}. \end{aligned}$$

This directly accounts for censored demand due to finite capacity. Since arrivals are restricted to be non-negative, we restrict the set of fixed effects by transforming the multiplicative fixed effects to be of the form $\lambda_{t,d} = \exp(W_{t,d}\tau)$. We select a log-Gamma prior for τ . We sample from the posterior distribution by taking a Metropolis-Hastings draw from a normal candidate distribution.

Sampling Shares and Utility Parameters

Updating shares. We treat product shares as unobserved, since the market size may be very small and lead to irreducible measurement error. We use data augmentation to treat shares as a latent parameter that we estimate. Conditional on all other parameters

$(\lambda, \alpha, \gamma, \beta, \eta, \Sigma)$, product shares are an invertible function of the demand shock, ξ . If we conditioned additionally on ξ , shares would be a deterministic function of data and other parameter draws. Instead, we leverage the stochastic nature of ξ , which we explicitly parameterize. The distribution of unobserved ξ is the source of variation for constructing a conditional likelihood for shares:

$$\left. \begin{aligned} \xi_{j,t,d} &= f^{-1}(s_{j,t,d} | \beta, \alpha, \gamma, X) \\ v_{j,t,d} &= p_{j,t,d} - Z'_{j,t,d} \eta \end{aligned} \right\} \left| \kappa = k \sim \mathcal{N}_{\text{iid}}(0, \Sigma_k) \right.$$

such that $\Sigma_k = \begin{pmatrix} \sigma_{k,11}^2 & \rho_k \\ \rho_k & \sigma_{k,22}^2 \end{pmatrix}$.

Here, κ is a mapping from days to departure t to an interval (block) of time. That is, the pricing error and the demand shock have a block-specific joint normal distribution. Conditional on the pricing shock v , the distribution of ξ , $f_{\xi_{j,t,d}}(\cdot)$, is

$$\xi \left| v, \kappa = k \sim \mathcal{N} \left(\frac{\rho_k v}{\sigma_{k,11}^2}, \sigma_{k,22}^2 - \frac{\rho_k^2}{\sigma_{k,11}^2} \right).$$

The density of shares is then given by the transformation $f_{s_{j,t,d}}(x) = f_{\xi_{j,t,d}}(f^{-1}(x)) \cdot \left| \mathcal{J}_{\xi_{j,t,d} \rightarrow s_{j,t,d}} \right|^{-1}$, where $\mathcal{J}_{\xi_{j,t,d} \rightarrow s_{j,t,d}}$ is the Jacobian matrix of model shares with respect to ξ . To produce the full joint conditional likelihood of shares, we also include the mass function for sales, which are a product of shares and arrivals:

$$\prod_t \prod_d \prod_{j=1}^{J(t,d)} \left[\phi \left(\frac{f^{-1}(s_{j,t,d}) - \frac{\rho_k v}{\sigma_{k,11}^2}}{\sqrt{\sigma_{k,22}^2 - \frac{\rho_k^2}{\sigma_{k,11}^2}}} \right) \frac{(\lambda_{t,d} s_{j,t,d})^{q_{j,t,d}} \exp(-\lambda_{t,d} s_{j,t,d})}{q_{j,t,d}!} \right] \cdot \left| \mathcal{J}_{\xi \rightarrow s} \right|^{-1},$$

where $\phi(\cdot)$ is the standard normal density function. We draw from the posterior based on a uniform prior distribution and normal candidate Metropolis-Hastings draws.

Updating price coefficients, α_B, α_L . We construct the conditional likelihood (and thus the conditional posterior distribution) for $\alpha = (\alpha_B, \alpha_L)$ in a similar manner to the product shares. For any candidate value of price sensitivity, we recover a residual ξ , invert the demand system, and recover a likelihood. Conditional on λ , shares, η , β , and Σ , we compute the distribution of ξ and determine the likelihood of a particular draw of α , given by

$$\prod_t \prod_d \prod_{j=1}^{J(t,d)} \left[\phi \left(\frac{f^{-1}(s_{j,t,d}) - \frac{\rho_k v}{\sigma_{k,11}^2}}{\sqrt{\sigma_{k,22}^2 - \frac{\rho_k^2}{\sigma_{k,11}^2}}} \right) \right] \cdot |\mathcal{J}_{\xi \rightarrow s}|^{-1},$$

where $\phi(\cdot)$ is the standard Normal density function. We impose a log-Normal prior on α , and impose $\alpha_B < \alpha_L$ to avoid label-switching. To draw from the conditional posterior, we take a Metropolis-Hasting step using a normal candidate distribution.

Updating the distribution of consumer types, γ . We allow for the mix of consumer types to change over the booking horizon t . We define γ from a sieve estimator of the booking horizon t , and we sample the sieve coefficients, ψ , according to

$$\gamma_t = \text{Logit}(G(t)' \psi),$$

where $G(t)$ is a vector of Bernstein polynomials. The logistic functional form ensures that the image of γ is in the interval $(0, 1)$. The inversion procedure used to construct the likelihood is similar to α and shares. It yields a likelihood for sieve coefficients ψ of the form

$$\prod_t \prod_d \prod_{j=1}^{J(t,d)} \left[\phi \left(\frac{f^{-1}(s_{j,t,d}) - \frac{\rho_k v}{\sigma_{k,11}^2}}{\sqrt{\sigma_{k,22}^2 - \frac{\rho_k^2}{\sigma_{k,11}^2}}} \right) \right] \cdot |\mathcal{J}_{\xi \rightarrow s}|^{-1}.$$

We use a uniform prior on ψ , and we sample from the posterior with a Metropolis-Hastings step using a normal candidate draw.

Updating remaining preferences, β . To sample the remaining preferences that are common across consumer types, we impose a normal prior on β , with mean $\bar{\beta}_0$ and variance V_0 . We adjust for price endogeneity to conduct a standard Bayesian regression. Define $\delta_{j,t,d} = X_{j,t,d}\beta + \xi_{j,t,d}$, which is evaluated at the ξ computed in the prior step. We normalize each component of δ by subtracting the expected value of ξ and dividing by its standard deviation. The normalized equations have unit variance and are thus conjugate to the normal prior. Let $\sigma_{k,2|1} = \sqrt{\sigma_{k,22}^2 - \frac{\rho_k^2}{\sigma_{k,11}^2}}$ be the variance of ξ conditional on v and Σ . We center and scale δ :

$$\frac{\delta_{j,t,d} - \frac{\rho_{\kappa_t}}{\sigma_{\kappa_t,11}^2} v}{\sigma_{\kappa_t,2|1}} = \frac{1}{\sigma_{\kappa_t,2|1}} X_{j,t,d} \bar{\beta} + U_{j,t,d}^\beta,$$

where $U^\beta \sim \mathcal{N}(0, 1)$. Then, the posterior distribution of β is $\mathcal{N}(\beta_N, V_N)$, where

$$\begin{aligned} \beta_N &= (\hat{X}'\hat{X} + V_0^{-1})^{-1} (V_0^{-1}\beta_0 + \hat{X}'\hat{\delta}), \\ V_N &= (V_0^{-1} + \hat{X}'\hat{X})^{-1}, \\ \hat{X}_{j,t,d} &= \frac{X_{j,t,d}}{\sigma_{\kappa_t,2|1}}, \\ \hat{\delta} &= \frac{\delta_{j,t,d} - \frac{\rho_{\kappa_t}}{\sigma_{\kappa_t,11}^2} v}{\sigma_{\kappa_t,2|1}}. \end{aligned}$$

Given this normalization, we can draw directly from the conditional posterior distribution of β using a Gibbs step.

Sampling Price-Endogeneity Parameters

Updating pricing equation, η . We use a linear pricing equation of the form

$$p_{j,t,d} = Z_{j,t,d}\eta + v_{j,t,d}.$$

Conditional on shares, λ , γ , α , and β , ξ is known. Therefore, we use the conditional distribution of v given ξ to perform another Bayesian linear regression in a similar manner to β . We impose a Normal prior and normalize prices. Define $\sigma_{\kappa_t,1|2} = \sqrt{\sigma_{\kappa_t,11}^2 - \frac{\rho_{\kappa_t}^2}{\sigma_{\kappa_t,22}^2}}$. It follows that

$$\frac{p_{j,t,d} - \frac{\rho_{\kappa_t}}{\sigma_{\kappa_t,22}} \xi_{j,t,d}}{\sigma_{\kappa_t,1|2}} = \frac{1}{\sigma_{\kappa_t,1|2}} X_{j,t,d} \bar{\eta} + U_{j,t,d}^\eta,$$

where $U^\eta \sim \mathcal{N}(0, 1)$. Just as we did for β , we can draw from the posterior of η from a linear regression with unit variance. This step allows us to directly sample from the posterior of η rather than using a Metropolis-Hastings step.

Updating the price endogeneity parameters, Σ . We flexibly model the joint distribution of ξ and v by allowing for a route-specific, time-varying correlation structure. We divide the booking horizon into four equally sized 30-day periods, and each block is indexed k . We restrict the price endogeneity parameters Σ , which determine the joint distribution of ξ, v , to be identical within these blocks. Within each block, the pricing and demand residual follow the same joint distribution. We draw the variance of this normal distribution with a typical Inverse-Wishart parameterization. Our prior for Σ_k is $IW(\nu, V)$ where k refers to the block. Define the vector $Y_k = (v, \xi)$ to be the collection of residual pairs conditional on block k , and $Y_k \sim \mathcal{N}(0, \Sigma_k)$. The posterior

for the covariance matrix Σ_k is then

$$\Sigma_k \sim IW(\nu + n_k, V + Y_k' Y_k).$$

Block k has n_k observations. This Gibbs step is repeated for each block k , and we sample directly from the conditional posteriors of Σ .

D.2 The Impact of the Scaling Factor on Demand Estimates

We consider alternative specifications on our scaling factor ζ in order to understand how changes in imputed market size affect our demand estimates. Our biggest concern is that our scaling factor may understate the presence of price-sensitive consumers who primarily shop with online travel agencies. For each route, we adjust our leisure scaling factor by multiplying the original scaling factor by 1.5, 2, 3, 5 and 10. We find that between 1.5 to 3 times the original scaling factor, our demand estimates are largely unchanged. For larger scaling factors—between 5 and 10—we find that demand becomes less price sensitive far from departure and more price sensitive close to departure. The parameters most affected by this scaling are the parameters governing the probability of business, γ . As we scale up the leisure arrival process, our estimated probability of business falls. The change in consumer types over time is reduced, however, we still estimate average elasticities to be similar to the baseline model.

E Model of Demand Forecasts in Counterfactuals

We use an auxiliary demand model to represent the RM department’s actual demand forecasts in counterfactuals. We conduct demand estimation using the forecasting data. We proceed in two steps.

First, we recalibrate our estimated compound Poisson distributions using the RM

department's passenger assignment algorithm, assuming the forecasts are generated using the same model and total intensity of consumer arrivals. We recalibrate the composition of arriving customers ($\gamma_{t,r}$) as

$$\gamma_{t,r}^{\text{forecast}} = \frac{\sum \text{Arrivals}_{t,r}^B}{\sum \text{Arrivals}_{t,r}^B + \sum \text{Arrivals}_{t,r}^L},$$

where $\text{Arrivals}_{t,r}^B$ is the total number of arrivals classified as business for route r using the passenger classification algorithm (L is similarly defined). With these estimates, the adjusted arrival processes are $\lambda_{t,r} \lambda_{d,r} \gamma_{t,r}^{\text{forecast}}$ for business passengers and $\lambda_{t,r} \lambda_{d,r} (1 - \gamma_{t,r}^{\text{forecast}})$ for leisure traffic. We label these Poisson rates $\tilde{\lambda}_{t,d,r}^B$ and $\tilde{\lambda}_{t,d,r}^L$.

Second, we recover preferences consistent with the RM department's forecasts. Recall that the forecasts are predictions of sales quantities at the flight, departure date, passenger type, day before departure, and price level—they are demand curves at discrete prices. We impose the same demand specification as in our demand model. Because the firm uses a single-product demand model, we abstract from cross-price elasticities for this analysis and assume arrival rates are equal to $\tilde{\lambda}_{t,d,r}^\ell$ for each flight $j \in J_{d,r}$. Instead, we could assume arrivals are $\tilde{\lambda}_{t,d,r}^\ell / J$, so that each flight receives $1/J$ of arrivals. However, we find that this increases product shares and results in consumers estimated to be more price insensitive. Using the unconstrained forecast, $\tilde{Q}_{j,t,d,r}^\ell(k)$, which is simply the prediction of unit sales at a price of k for consumer type ℓ if capacity were not constrained, we match these curves to its corresponding model counterpart,

$$\tilde{Q}_{j,t,d,r}^\ell(k) = \tilde{\lambda}_{t,d,r}^\ell s_{j,t,d,r}^\ell(k). \quad (2)$$

Plugging in a different price, k' , results in another matched equation for the same forecast j, t, d, r . Taking logs of the Equation 2 and subtracting the log of the outside

good share, we use the inversion of Berry (1994) to obtain

$$\log\left(\frac{\tilde{Q}_{j,t,d,r}^\ell}{\tilde{\lambda}_{t,d,r}^\ell}\right) - \log(s_{0,t,d,r}^\ell) = \log(s_{j,t,d,r}^\ell) - \log(s_{0,t,d,r}^\ell) = \tilde{\delta}_{j,t,d,r}^\ell. \quad (3)$$

This inversion is only possible because the forecasting data are at the consumer-type level. Otherwise, we would have to use the contraction mapping in Berry, Levinsohn, and Pakes (1995) and Berry, Carnall, and Spiller (2006).

This demand inversion allows us to use the same restrictions imposed in our demand model, i.e., mean utility differs across consumer types is only on the price coefficient. However, we must also confront a data limitation in that our forecasting data is not necessarily at the t level, but rather, at a grouping of ts the firm uses for decision making. The number of days in a grouping varies. We address this data feature in the following way. Note that our demand model does not have t -specific parameters—preferences do not vary by day before departure. Therefore, if \tilde{Q}^ℓ is the forecast for consumer type i for multiple periods, the model analogue to this is

$$\tilde{Q}_{\cdot,t^*}^\ell = \sum_{t \in t^*} \tilde{\lambda}_{\cdot,t}^\ell s_{\cdot,t}^\ell(\cdot) = \left(\sum_{t \in t^*} \tilde{\lambda}_{\cdot,t}^\ell \right) s_{\cdot,t^*}^\ell(\cdot).$$

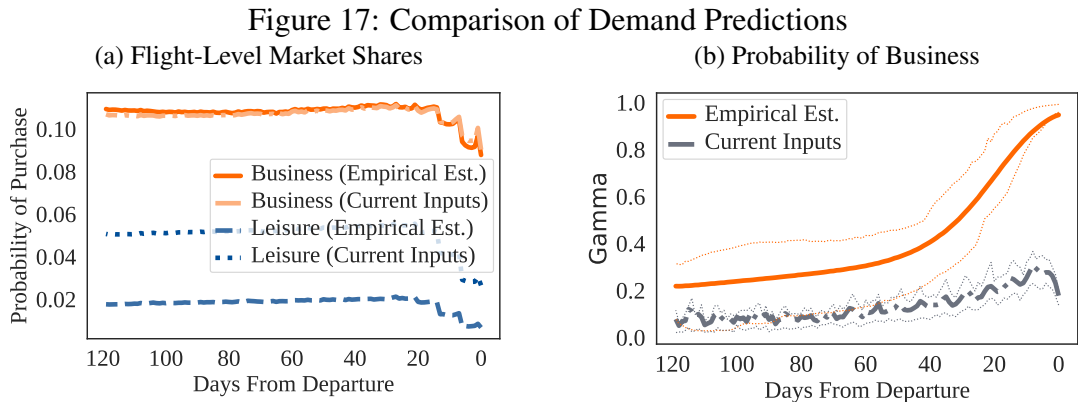
We can simply sum over the relevant time indices for arrival rates because the time-index does not enter within-consumer type shares, and the forecasting data assumes a constant price within a grouping of time. This is important because we can then define consumer-type product shares as

$$\frac{\tilde{Q}_{\cdot,t^*}^\ell}{\sum_{t \in t^*} \tilde{\lambda}_{\cdot,t}^\ell} = s_{\cdot,t^*}^\ell(\cdot).$$

Thus, we obtain the following inversion,

$$\log\left(\frac{\tilde{Q}_{\cdot,t^*}^\ell}{\sum_{t \in t^*} \tilde{\lambda}_{\cdot,t}^\ell}\right) - \log(s_0^\ell) = \log(s_{\cdot,t^*}^\ell) - \log(s_0^\ell) = \tilde{\delta}^\ell. \quad (4)$$

Defining the left-hand side of Equation 3 above as $\tilde{\delta}$, we obtain the linear estimating equation (suppressing subscripts) $\tilde{\delta} = X\tilde{\beta} - \tilde{\alpha}p + u$, where $\tilde{\beta}, \tilde{\alpha}^B, \tilde{\alpha}^L$ are preferences to be estimated. One caveat to this approach is that we estimate a " ξ " that also differs across consumer types through u . We set these residuals equal to zero and include our estimated ξ to be consistent with our demand model. We use the mean of the posterior for that observation taken from our demand estimates. This adjustment does not greatly impact our findings. In total, we recalibrate all 20,000 preference parameters estimated in our original demand specification. We compare the demand systems in the figure below.



Note: (a) Comparison of product shares across consumer types, over time. (b) Estimates of γ_t versus those calculated using the passenger assignment algorithm. Results are reported averaging over all observations in the data.