

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 October 2022

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 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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 October 2022
JEL No. C11,C53,D22,D42,L11,L93

ABSTRACT

Firms often involve multiple departments for critical decisions that may result in coordination failures. Using data from a large U.S. airline, we document the presence of important pricing biases that differ significantly from dynamically optimal profit maximization. However, these biases can be rationalized as a “second-best” after accounting for department decision rights. We show that assuming prices are generated through profit maximization biases demand estimates and that second-best prices can persist, even under improvements to pricing algorithm inputs. Our results suggest caution in abstracting from organizational structure and drawing inferences from firms’ pricing decisions alone.

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

Large firms often involve multiple departments for critical decisions. This is true for pricing, which often involves specialized teams and complex optimization systems. Each department is responsible for a particular sub-decision. One department might manage procurement and inventory, another department specializes in demand predictions, and an additional department manages competitive response. For example, major airlines pioneered the use of pricing algorithms and operationalized these systems by delegating the decision rights for each of multiple pricing inputs to distinct departments. Moving beyond the simple model of firms as unitary decision-makers demonstrates the potential for coordination failures and the possibility that departments introduce biases with large welfare effects.

We study the decisions of multiple departments that control pricing at a large U.S. airline and quantify their impacts using rich internal data and the actual pricing algorithm.¹ We show that pricing differs significantly dynamically optimal profit maximization. Departments are miscoordinated and introduce multiple pricing biases. With demand estimates and all department input decisions, we show that any unilateral bias reduction for a single department reduces revenues. That is, each department's observed input is nearly optimal for the firm, holding other departments' inputs fixed. We call such outcomes "second-best optimal for the firm." They need not follow the global, first-best outcome, and we demonstrate how accounting for their presence affects how we interpret prices. We show that assuming prices are first-best and generated through a canonical model of profit maximization overstates consumer price sensitivity compared to both our demand estimates and the firm's internal demand estimates. We compare market outcomes to the canonical model. We discuss why large firms facing complex objectives must simplify the problems they face, giving rise to second-best prices, and highlight why departing from current practices is difficult.

¹The airline has elected to remain anonymous.

We begin with an overview of airline pricing practices, describing the transition from regulated prices to the organizational structure and implementation of algorithmic pricing that arose post-deregulation. The organization structure has remained unchanged since the 1980s. It features knowledge hierarchies in the spirit of Garicano (2000), with decisions being passed vertically to the next department as an input. The first department (network planning) decides where to fly and assigns initial capacities. We do not model these decisions. Next, 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 demand analysis and maintaining the pricing algorithm, which combines all inputs. The algorithm does not actually decide price, rather, it allocates the amount of inventory to sell at each discrete price level. That is, “pricing” involves prices and quantities set by separate departments.

The organizational structure at airlines has been durable and pervasive. Using job listings, we show that all major airlines, ultra low-cost carriers, and recently founded airlines have the same organizational structure and department responsibilities. We show that cruises (e.g., MSC Cruises, Carnival Corporation), hotels (e.g., InterContinental Hotels Group, Universal Orlando Resort), and car rentals (e.g., Hertz Corporation, Avis Budget Group) have also adopted the exact same organizational structure and department responsibilities.² Therefore, we believe our findings likely hold for the airline industry broadly and in other industries where firms sell perishable inventory.

Our insights are derived from comprehensive data provided to us through a research partnership with a large U.S. airline. In addition to observing daily prices and quantities, we also observe all department input decisions, the demand model, granular demand forecasts, internal pricing documentation, and hundreds of millions of consumer interactions or “clicks” on the airline’s website. The data cover over 300,000 flights and 450 domestic routes over a span of two years. In addition, we observe the

²Archived copies of all job postings, network profiles, and patents available upon request.

pricing algorithm’s exact design (code) and use it in our counterfactual analyses.

We reveal that the firms actual pricing practices differs significantly from canonical models of pricing. For example, we show that the RM department uses a pricing algorithm with several critical abstractions. There are no modeled cross-price elasticities of any kind, including to other products (cabins) within a flight or to other substitute flights. Contrary to common perception, the algorithm does not internalize competitive decisions at all. The heuristic is also not fully forward-looking. These simplifications are necessary because solving a standard profit maximization at scale is not feasible.

The heuristic requires inputs from multiple departments. We show that these inputs are subject to miscoordination. First, we show that there is a “glitch” in that the heuristic sometimes (11% of the sample) allocates inventory to fares set by the pricing department that do not exist. Second, the pricing department does not rely on demand analysis in constructing itineraries and commonly sets fares that are on the inelastic side of the RM departments demand models. In standard economic models, the optimal price would never be on the inelastic side of the demand curve, so this form of miscoordination would be inconsequential. However, as we describe in Online Appendix A, the pricing heuristic does not equate marginal revenue to marginal cost and its “default” choice is to allocate seats to the lowest fare. We show that if the heuristic allocated seats to the lowest fares set by the pricing department, demand would be inelastic according to the RM department’s demand models for 98% of the sample.

The RM department knows that low filed fares can result in inelastic demand according to its own analysis. Moreover, the heuristic is well-known for potentially understating opportunity costs (e.g., Wollmer, 1992; Cooper, Homem-de Mello, and Kleywegt, 2006). It has implemented a clever solution to counteract these forces within its decision rights—it biases its own inputs. We show that RM analysts input persistently upward biased forecasts that raise opportunity costs, and hence, prices.³ We

³For additional perspectives on miscalibrated firm expectations, see Massey and Thaler (2013); Akepanidtaworn, Di Mascio, Imas, and Schmidt (2019); Ma, Ropele, Sraer, and Thesmar (2020).

show that this workaround, or kludge (Ely, 2011), has a profound effect, reducing the potential for inelastic prices by up to 60% (from over 95% to 35%).⁴

We establish that the pricing biases we document affect all routes, regardless of market structure. Due to the additional complexity of modeling competitive interactions, our subsequent analysis concentrates on routes where our air carrier is the only airline operating nonstop. Nonetheless, our results may be informative for routes with competition. Routes with existing nonstop competition are priced using the same single-product heuristic and feature more frequent inelastic prices based on the RM department’s own analysis. This widespread presence of inelastic prices cannot be rationalized by standard models of profit maximization and likely suggests that all routes follow second-best outcomes as we describe below.

In order to quantify the impacts of alternative pricing inputs, we must estimate unbiased preferences for air travel. We use a recently proposed demand methodology (Hortaçsu, Natan, Parsley, Schweg, and Williams, 2022). We consider a model in which “leisure” and “business” travelers arrive according to time-varying Poisson distributions. Conditional on arrival, consumers solve a standard discrete choice problem. We provide new descriptive evidence using search data to motivate our assumptions.

We address the identification challenge of estimating preferences in models with aggregate demand uncertainty by leveraging consumer arrival and bookings data. Our approach allows us to avoid assuming that observed prices are optimal (e.g., Williams, 2022; Aryal, Murry, and Williams, 2022; D’Haultfœuille, Février, Wang, and Wilner, 2022; Pan and Wang, 2022; Cho, Lee, Rust, and Yu, 2018).⁵ We show that assuming observed prices are optimal, i.e., they are generated through a benchmark dynamic

Berman and Heller (2020) present a theory on why biases can persist, including underestimating price elasticities, in a model of competition.

⁴According to Ely (2011), “A kludge is a marginal adaptation that compensates for, but does not eliminate, fundamental design inefficiencies.”

⁵D’Haultfœuille, Février, Wang, and Wilner (2022) consider a partial identification approach in a dynamic pricing environment that can also accommodate firm learning.

pricing problem, results in biased demand estimates. To stress magnitudes, imposing supply-side optimality conditions results in elasticities near -4. Both our preferred demand estimates and internal demand estimates suggest near unit-elastic demand.

We then flexibly recover preferences consistent with the RM department's forecasts, i.e., we conduct demand estimation using forecasted demand rather than realized demand. We show that these recovered preferences understate preference heterogeneity across consumer types, days before departure, departure dates, and routes, compared to our demand estimates. This is due to not only the kludge, but also because the demand models from which the forecasts are derived have limited flexibility. The models assume a constant elasticity of substitution within consumer type, and each set of parameter estimates is applied to hundreds of routes.

With demand estimated, we conduct a series of counterfactuals where departments select alternative pricing inputs. We first establish that current department decisions are nearly optimal (within 1% of total revenue) for the firm, holding the other department's input fixed. This allows us to present a rational explanation for pricing biases, including pricing on the inelastic side of demand, that has been noted in other industries (e.g., Levitt, 2016; DellaVigna and Gentzkow, 2019; Dubé and Misra, 2021). We establish this result by considering deviations in decisions for each department in isolation and showing that such deviations are not profitable for the firm. In addition, we show that relaxing the restrictive assumptions placed on demand hardly change outcomes, exemplifying the notion of limited gains of unilateral change (Milgrom and Roberts, 1990, 1995).

We also quantify market outcomes of the hypothetical first-best using a standard dynamic programming problem (that also accounts for cross-price elasticities) and route-specific demand estimates. Although we argue that such a pricing system is not currently implementable, the counterfactual allows us to sign the welfare effects of moving beyond the second best. We show that second-best prices lead to higher

load factors and limit price targeting compared to the first-best. First-best prices result in an increase in dead-weight loss of 6%; the firm and leisure consumers benefit, and business consumers face higher prices.

Finally, we discuss explanations for the persistence of second-best prices. Departing from current practices is remarkably difficult, even though tracked key-performance indicators (load factors, revenue yield, etc.) do not suggest department incentives are misaligned.⁶ We discuss why current practices are reinforced due to computational barriers and environmental factors (Siggelkow, 2001), including the influential role of intermediaries (e.g., online travel agencies) that discipline firm pricing practices.

This 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 Sections 5-7. Analysis of counterfactuals appear in Section 8.

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.⁷ Airlines pioneered pricing technologies that are now commonly used in perishable goods markets. We briefly describe airline pricing practices. Vinod (2021) provides a detailed, historical account.⁸

Prior to deregulation of the airline industry, the networks and fares of airlines were federally controlled. In order to facilitate the 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

⁶See Atkin, Chaudhry, Chaudry, Khandelwal, and Verhoogen (2017); Sacarny (2018) for work on misalignment in incentives within an organization.

⁷See <https://www.iata.org/en/iata-repository/publications/economic-reports/the-united-states-value-of-aviation/>. Sep 28, 2022.

⁸See McGill and Van Ryzin (1999) and Talluri and Van Ryzin (2004) for additional details.

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 common in many 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.⁹ American Airlines continued to develop the algorithm as computing power increased (Vinod, 2021). This innovation 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. These two departments, pricing and revenue management (RM), are now observed at legacy carriers (e.g., Smith, Leimkuhler, and Darrow (1992) discusses American Airlines), low-cost carriers (e.g., Frontier Airlines), and start-up airlines (e.g., Azul Airlines in Brazil). In fact, job listings show that all major airlines are organized in this way.¹⁰

Today, the pricing and RM departments maintain the same vertical structure and responsibilities of the 1980s. The pricing department decides which itineraries to offer and assigns a discrete menu of fares and fare restrictions for every itinerary.¹¹ The RM department then allocates remaining seats to each of the prices in the menu. These two departments are ever-present in the travel industry. The same organizational structure and department responsibilities can be found in car rentals, hotels, cruises, trains, buses, and air cargo.¹² We discuss reasons why in Section 8.

⁹The 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.

¹⁰See Footnote 2.

¹¹Advance purchase (AP) fares are common at 7, 14, and 21 days before departure.

¹²As an example, at Universal Parks & Resorts, the pricing department “develop(s) price recommendations for both leisure and group rate programs” whereas the RM department “analyze(s) revenue management system output for forecast and optimization anomalies through the use of statistical and mathematical optimization models.” This is analogous to what occurs at our airline.

2.1 Organization Durability and Department Responsibilities

Why is this bifurcated organizational structure the predominant way firms have organized in industries selling perishable goods? One reason is that the level of responsibility differs. Pricing departments control the types of products (e.g., itineraries) offered.¹³ It designs itineraries and sets fare menus, but it does not control which fare is offered at any point in time. In addition, it gathers and interprets competitor prices and initiates/responds to industry changes, e.g., implementing a fuel surcharge. It is a fact at our airline that the pricing department assigns fares without the use of short- or long-run demand models. The latter is not estimated by the RM department at all.

The RM department develops strategies regarding segments rather than itineraries. This entails maximizing flight-level metrics taking fares/itineraries as given. RM departments use both utilization and revenue yield as key performance metrics.

A second, related reason that the departments are separated is due the difficulty of jointly optimizing fares and inventory controls (prices and quantities). The smallest flights in our sample involve over 1.5 million inventory allocations. Combining inventory allocation with the simplest pricing decision, i.e., deciding among two possibilities for each discrete fare class, would result in more than 10^{1000} potential choices.

A third reason that separate departments are maintained is due to external compatibility. Airlines still rely on 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. See Section 8 for details. At our airline, this requires using entirely separate IT systems for pricing and RM activities.

¹³The 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.

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. Reoptimization of remaining inventory is conducted 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 discount requirements. Fares have dozens of attributes, e.g., a footnote details if it is being used for a sale. We focus on nonstop/non-connecting fares. Adjustments to the menus are common—the probability that a fare menu changes day-over-day for a given route (at least one change to the menu) is 18%.

(5) *Forecasting data*: The RM department forecasts demand based on short-run demand estimates using conditional choice models. We use “demand model” to denote the baseline choice 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 search/booking classification algorithm. We observe the demand model, parameter estimates, daily 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 over 450 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

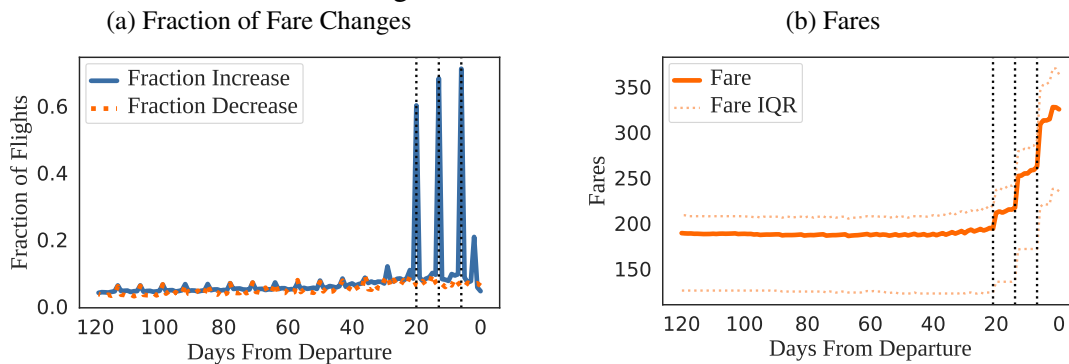
Data	Variable	Mean	Std. Dev.	Median	5th pctile	95th pctile
Fares						
	One-Way Fare (\$)	201.3	139.4	163.3	88.0	411.1
	Num. Fare Changes	9.3	4.2	9.0	3.0	17.0
	Fare Change Inc.	50.4	73.0	31.2	2.2	164.5
	Fare Change Dec.	-53.0	75.5	-32.2	-175.2	-4.3
Bookings						
	Booking Rate-OD	0.2	0.7	0.0	0.0	1.0
	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 cleaned 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 advance purchase (AP) discount opportunities (see Figure 1-a). Over 60% of all price adjustments occur outside these time windows. In Figure 1-(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 1: 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%. 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 and account for any bookings that occur earlier than 120 days before departure in our analysis.

3.3 Motivating Evidence

Airline Pricing Practices and the use of Algorithmic Pricing

Reviewing the algorithm code/revisions and supporting documentation, we find that airline pricing differs significantly from dynamically optimal profit maximization and recent work on algorithmic pricing (e.g. Calvano, Calzolari, Denicolo, and Pastorello, 2020; Brown and MacKay, 2021). The heuristic used in practice does not account for cross-price elasticities, including across cabins with 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. The heuristic does not solve or approximate a dynamic program. Ancillary revenue, including baggage fees, upgrade charges, etc., are not considered when pricing tickets.¹⁴ Fares are not personalized; loyalty metrics are not used in pricing and RM activities. We confirm that the algorithm has hardly changed over the past 25 years.

Studying department decisions is important because they may be reinforcing or offsetting. For example, suppose the pricing department wants to run a sale (or match a competitor) by offering a \$50 fare. This input is then passed to the RM department, through the demand forecast, and then into the pricing algorithm. If the algorithm 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. That is, the sale (or price matching) will not occur. We confirm that observing equal prices across flights within a day (or matching a competitor) is not a strategic choice of the algorithm. Observing identical fares across flights is made possible due to the pricing department's fare menu decisions, but it is not enforced by the algorithm. Rather, it is a consequence of independent actions by the algorithm.

¹⁴This is not uncommon. For example, Japan, Etihad, Philippine, Flydubai, Korean, Jeju, Frontier, Malaysia, All Nippon, Hawaiian, and Lufthansa all use an RM solution offered by PROS that does not allow for joint ancillary fare revenue maximization. Patent data suggest other RM solutions closely resemble our airline's systems.

New Facts on Airline 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 consumers observe a single price per flight because 91% of consumers purchase the lowest available fare. We find that consumers 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. 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 forecasts demand. These labels are mechanically linked to attributes of the ticket, e.g., the number of days before departure it was booked. They are not attached to traveler characteristics, e.g., passenger status, booking channel, etc. We observe the exact algorithm used by the RM department.

We find evidence that supports using a static discrete choice model in our setting. 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 12-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 search for an itinerary once (including shoppers referred to the airline increases the percentage). Interestingly, among the remaining 10% (see Figure 12-b), only 20% ever observe a lower fare for at least one flight in later searches (most search spells end within five days). We estimate that only 2% of shoppers who may be strategically waiting actually obtain a lower price. This estimate is lower than the IQR of between 5.2% to 19.2%

found by Li, Granados, and Netessine (2014) without 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 searched/purchased within the previous week. Only 8% would have benefited from delaying. We note, however, that imposing static demand can affect estimates of willingness to pay (e.g., Hendel and Nevo, 2006; Sweeting, 2012).

4 Pricing Biases and Miscoordination of Pricing Inputs

We provide examples of pricing biases and miscoordination. Additional empirical evidence can be found in Online Appendix C.

4.1 Heuristic Bias

The RM department uses a pricing heuristic for tractability reasons. The heuristic relies on well-known optimization techniques in operations. We detail the algorithm which our firm's is based on in Online Appendix A. We do not identify the exact algorithm for confidentiality reasons, however, we use the exact algorithm in our analysis.

As previously mentioned, the heuristic does not allow for cross-price elasticities, including own substitute and competitor options. This is true, regardless of market structure. In Online Appendix C.2, we show that prices are not even indirectly affected by demand of substitute flights. The heuristic is also partially dynamic in that it considers demand today versus the sum of future expected demand (the "second period"). It internalizes scarcity in this way and then assigns remaining inventory to each fare level. Lowest priced units are assumed to sell first. Although the heuristic shares some features of dynamic pricing problems, e.g., scarcity will drive prices upward, its objective is not explicitly to maximize revenues. Rather, it focuses on the booking rate. The heuristic can understate opportunity costs because internalizing that future prices

will adjust raises the value of a seat (e.g., Cooper, Homem-de Mello, and Kleywegt, 2006), and accounting for cross-price elasticities gives additional pricing power.

4.2 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 8.

4.3 Fares on the Inelastic Side of Demand Curves

A second form of miscoordination that we observe is that the pricing department often files fares that are too low according to the RM department’s internal demand models. Although this would be inconsequential if the pricing heuristic explicitly modeled revenues—it would not allocate seats on the inelastic side of the demand curve—in practice, it cannot prevent “inelastic fares” from being offered to consumers. When capacity is not constrained, it will default to the lowest fare, regardless of what it is. This is true even if the more expensive, profit maximizing fare is included in the set.

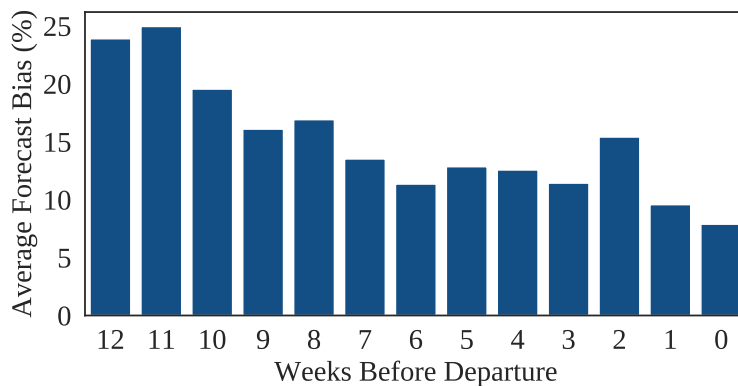
Using the RM department’s continuous and differentiable demand models, $Q(p)$,

we calculate the elasticity of demand, $e(p)$. We plug the lowest fare filed by the pricing department. We find that if the heuristic selected these fares to be offered to consumers, demand would be inelastic according to the RM department’s demand models in 98% of the sample. Importantly, we note that the lowest fares are not “sale” fares. Sale fares have special attributes and are only active for short periods of time.

4.4 Using Persistently Biased Forecasts

The RM department knows about the heuristic’s biases and that the presence of low menu fares can cause consumers to be offered inelastic prices. It has introduced a clever workaround, or kludge, to counteract these forces. In order to raise prices (via increasing opportunity costs) within their decision rights, all RM analysts input persistently upward biased forecasts. This is done by first taking the original demand models and then adjusting them several ways. Over six adjustments are used in practice; the most common one scales up/down multiple routes’ demand curves simultaneously. The end result are demand forecasts that are inputted into the heuristic.

Figure 2: 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.

In Figure 2, we plot the average forecast bias (the demand models after analyst adjustments) by week before departure. We calculate the forecast bias for a particular

week before departure as

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

where forecasted and realized demand account for the price 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. We discuss reasons for the observed decrease in the forecast bias in a footnote below.¹⁵ This pattern is observed across all routes, regardless of route performance. 79% of flights have overforecasted demand 120 days before departure. Interestingly, we find that routes with nonstop competitors feature slightly larger forecasting bias compared to single-carrier routes.

How does the forecast bias potentially affect the decisions made by the pricing algorithm? 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 95% of flights we previously estimated if the heuristic allocated inventory to the lowest fares. This is driven by both the kludge as well as realized demand shocks that cause the lowest fares to not be offered. We separate the effect of the kludge from demand shocks in a counterfactual (see Section 8). One might wonder why forecasts are not even more

¹⁵We 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 lower bias is necessary to raise prices.

biased to ensure that no inelastic prices are offered. We show in Section 8 that the optimal bias is substantially higher and likely impermissible given its impact.

We do not consider the remaining presence of inelastic prices to be a strategic decision along the lines of Goolsbee and Syverson (2008) and Sweeting, Roberts, and Gedge (2020) because we observe inelastic prices based on the RM department’s demand analysis broadly. Moreover, inelastic prices occur even more frequently in routes with existing, direct nonstop competition, including routes with existing legacy and/or low-cost carrier competition. This finding is difficult to rationalize with standard economic models. We believe our findings that prices are second-best optimal for the firm may hold for all routes served, not just the ones we study in our structural analysis.

5 Empirical Model of Air Travel Demand

In order to quantify how department input decisions affect market outcomes, we need to estimate a model of unbiased preferences. We utilize both the demand model and estimation approach of Hortaçsu, Natan, Parsley, Schwieg, and Williams (2022), which allows us to capture rich substitution patterns, including seasonality effects, day-of-week effects, etc. 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, $j = 0$. Our model covers all bookings, regardless of booking channel.

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 as 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 are short-lived and do not strategically choose flights based on remaining capacity, $C_{j,t,d,r}$. 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 unobservable that is potentially correlated with price, and $\varepsilon_{i,j,t,d,r}$ is a random component of utility and 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 $u_{i,j,t,d,r} \geq u_{i,j',t,d,r}, \forall j' \in J \cup \{0\}$.

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

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.4 for supporting evidence); (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 $q_{j,t,d,r} = \min\{\tilde{q}_{j,t,d,r}, C_{j,t,d,r}\}$.

6 Estimation

6.1 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 dates, 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 ($\lambda_{t,r}$), and we observe strong departure-date effects ($\lambda_{d,r}$).

Because we do not observe all searches for all bookings (e.g., a booking through a travel agency), we adjust our estimated arrival rates. Additional descriptive analyses in Online Appendix C.4 motivate adjusting arrival rates differently over time. 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.¹⁶ That is, we use the relative fraction of L (or B) sales and searches across channels to scale up L (B) arrivals. This logic follows the simpler case with a single consumer type: if searches account for 20% of bookings, and unobserved searches involve the same underlying demand distributions, we can scale up estimated arrival rates by 5×. We conduct robustness to this specification (see Online Appendix D.2) and obtain quantitatively similar demand estimates.

6.2 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 . We build upon 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.

¹⁶We 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.

6.3 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 the price coefficients are identified from exogenous variation introduced by instruments.

We use the carrier’s shadow price of capacity as reported by the pricing algorithm, advance purchase indicators, and total number of inbound or outbound bookings from a route’s hub airport as our demand instruments. 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 14 in Online Appendix C.3). 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.¹⁷

Our identification argument does not rely on optimal pricing. The pricing heuristic is fully determined by its opportunity costs. We use these opportunity costs as instruments in a flexible way using a polynomial expansion. Our instruments are relevant and highly correlated with price as the pricing equations approximate the heuristic’s decisions. Pseudo first-stage R^2 s average 0.72 across routes. For additional flexibility, we also allow for the variance in the pricing equation unobservable to vary over

¹⁷For 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.

time. This allows us to account for the fact that the RM department manages flights differently over time in an observable way.

7 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 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 ODs.

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.

Our demand model accurately matches aggregated arrivals at both the day before departure and departure date level. We find a slight downward bias in quantity demanded of between 6-8%. This is largely driven by infrequent large group bookings that are hard to predict. We accurately match average booking rates even though 88% of observations have zero sales. We obtain noisy estimates at very granular levels due to both the discrete nature of the data and the high number of zeros, e.g., we estimate pseudo- R^2 s of around 0.23 for predicting particular (j, t, d, r) demand.

We present a summary of our demand 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, on average, or less than the spread in prices often 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. This is not true for all routes, and it does not always reflect seasonal variation in demand. We estimate this spread to be \$52 on average.

Figure 3: 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 3-(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 rate is not only due to changing preferences, as found in previous airline studies. In panel (b), we plot the average own-price elasticities for the mean, median, and interquartile range over routes. 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

departure (which is also observed using the firm’s demand models, see Section 4).

We refer to our demand estimates as “Model E” for the remainder of our analysis.

7.1 Preferences Consistent with Demand Forecasts

Next, we recover the arrival and preference parameters consistent with the RM department’s demand forecasts to use in our counterfactual exercises. Although the forecasts in theory can be reverse engineered to be used in counterfactuals directly, doing so would require us to explicitly model the decisions of analysts’ adjustments over many dimensions over time. Instead, we use an intuitive idea—we conduct a demand estimation using forecasted demand curves instead of realized demand observations. We call these recovered demand estimates “Model F.” 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 Model E. 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 $j \in J_{d,r}$.¹⁸ 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,

$$\tilde{Q}_{j,t,d,r}^\ell(k) = \tilde{\lambda}_{t,d,r}^\ell s_{j,t,d,r}^\ell(k). \quad (1)$$

Plugging in a different price, k' , results in another matched equation for the same forecast j, t, d, r . Taking logs of the Equation 1 and subtracting the log of the outside good share, we use the inversion of Berry (1994) to obtain¹⁹

$$\log\left(\frac{\tilde{Q}_{j,t,d,r}^\ell}{\tilde{\lambda}_{t,d,r}^\ell}\right) - \log(s_{0,t,d,r}^\ell) = \log(s_{j,t,d,r}^\ell) - \log(s_{0,t,d,r}^\ell) = \tilde{\delta}_{j,t,d,r}^\ell. \quad (2)$$

This demand inversion allows us to use the same restrictions imposed in Model E, i.e., mean utility differs across consumer types is only on the price coefficient.²⁰ Defining the left-hand side of Equation 2 above as $\tilde{\delta}$, we obtain the linear estimating equation

¹⁸Instead, we could assume arrivals are $\tilde{\lambda}_{t,d,r}^\ell/J$, so that each flight receives $1/J$ of arrivals. This increases product shares and results in consumers estimated to be more price insensitive.

¹⁹This 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).

²⁰We 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

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

(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 Model E. We use the mean of the posterior for that observation taken from Model E. This adjustment does not greatly impact our findings. In total, we recalibrate all 20,000 preference parameters estimated in Model E.

7.2 Comparing Demand Models

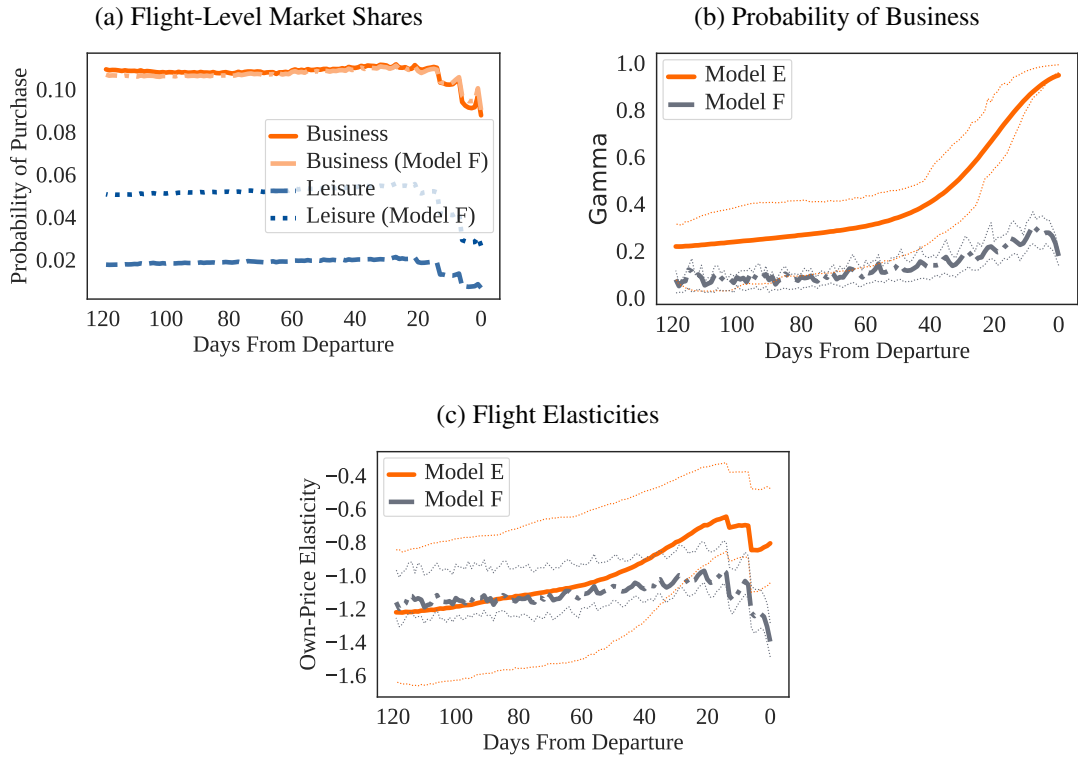
The predictions of Model E and Model F differ substantially, which we highlight in Figure 4. In panel (a), we plot product shares for both passenger types over time. Model F results in consumer types being "closer together" than under Model E, with leisure travelers being more price inelastic under Model F. In panel (b), we show that Model F also results in demand over time being "closer together," with a much smaller change in the types of consumers arriving over time. Finally, we plot own-price elasticities over time in panel (c). Model E produces elasticities that are generally increasing (to between -0.8 and -1.0) as more business consumers arrive over time. Model F yields more compressed and flatter elasticities until close to the departure date, where large price increases reduce estimated elasticities.

Four factors help us to rationalize the differences in predictions of Model E versus Model F. First, the kludge applied to the forecasts inflates early-arriving demand the most, resulting in higher leisure traveler willingness to pay. Second, Model F con-

Thus, we obtain the following inversion,

$$\log\left(\frac{\tilde{Q}_{t^*}^\ell}{\sum_{t \in t^*} \tilde{\lambda}_{t,t}^\ell}\right) - \log(s_0^\ell) = \log(s^\ell) - \log(s_0^\ell) = \tilde{\delta}^\ell. \quad (3)$$

Figure 4: Comparison of Demand Predictions



Note: (a) Comparison of product shares across consumer types, over time. (b) Estimates of γ_t versus those calculated using the passenger assignment algorithm. (c) Comparison of own-price elasticities over time. (b) and (c) contain the 25th and 75th percentiles. Results are reported averaging over all observations in the data.

strains variation in arriving consumer types over time because no consumer-level characteristics enter the passenger assignment algorithm. Model E more flexibly captures route-specific changes in the arrival processes over time. Third, Model E rationalizes greater heterogeneity in demand across time, departure dates, and routes because in practice, the demand model used to construct the forecasts has limited flexibility. One set of parameter estimates is applied to over 450 origin-destination pairs on average. Fourth, the demand forecasts are based on a model that assumes constant elasticity of substitution within consumer type. This results in an overstatement of bookings at low prices and an understatement of bookings at high prices. Thus, consumer type heterogeneity is understated.

7.3 Demand Estimates Assuming Optimal Pricing

Supply-side optimality conditions are often imposed when estimating empirical models of demand (e.g., Berry, Levinsohn, and Pakes, 1995). In the context of dynamic pricing, these conditions are used to separate demand shocks from underlying preferences. We have noted several biases and forms of miscoordination between departments that cause airline pricing to diverge from dynamically optimal profit maximization. What happens to our estimates of demand if we maintain the common assumption that prices are optimal? To answer this question, we re-estimate demand following Williams (2022) by assuming that the firm solves a canonical dynamic programming program (see also, Lazarev, 2013; Aryal, Murry, and Williams, 2022; Pan and Wang, 2022; Cho, Lee, Rust, and Yu, 2018). We include in estimation the conditional choice probabilities (CCP) that rationalize observed prices as optimal based on differences in choice-specific value functions. Though this estimation exercise is standard, we are able to highlight how prices which depart from the first best significantly alter inferences about demand.

We find that demand estimates imposing price optimality differ substantially from both Model E and Model F. For example, whereas Model E/F yield average demand elasticities close to -1, imposing optimality results in average demand elasticities close to -4. This is largely driven by the substantial fraction of observed low/inelastic prices. Imposing optimality through a standard dynamic programming problem requires demand to be elastic, which pushes price sensitivity coefficients to be more negative relative to both our and the firm's own analysis.

8 Empirical Analysis of Department Input Decisions

We quantify the welfare effects of alternative department input decisions through several counterfactuals. We focus on two questions: (1) Can current, biased inputs be

rationalized by accounting for department decision rights, and (2) How do current prices and market outcomes compare to a benchmark, dynamic pricing problem?

8.1 Counterfactual Implementation

For each counterfactual, we simulate flights based on the empirical distribution of observed remaining capacity 120 days before departure. For each vector of initial remaining capacities (departure date), we then draw preferences and arrival rates given our demand estimates (Model E). We simulate 10,000 flights for every departure date. Like our demand model, we do not endogenize connecting (or flow) bookings. Therefore, we handle connecting bookings through exogenous decreases in remaining capacity, based on Poisson rates estimated using connecting bookings.²¹ Consumers are assumed to arrive in a random order within a period. If demand exceeds remaining capacity, consumers are offered seats in the order they arrive.²²

8.2 Second-Best Optimal Outcome for the Firm

Our benchmark specification uses current department inputs: the RM department's forecasts (Model F), the observed fare choices by the pricing department, and the firm's pricing heuristic. We consider two deviations. First, we consider a hypothetical scenario in which there is feedback from the RM department to the pricing department—recall, currently decisions are passed vertically from the pricing department to the RM department, and the pricing department does not currently use demand models in their decision making. In this counterfactual, the pricing department uses the RM department's forecasts to adjust its fare menu decisions by eliminating fares that are on the inelastic side of the RM department's forecasts (Model F). That is, we consider

²¹Alternatively, we could subtract off observed connecting bookings from the initial capacity condition. However, this constrains initial capacity and results in higher prices than what we observe.

²²Using 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.

a simple deviation in which the pricing department addresses miscoordination. In the second deviation, we consider the deviation in which the RM department adjusts up or down its forecasts—similar to what we observe in practice—holding the pricing menus fixed. We define the exact procedure more precisely below.

We report counterfactual results in Table 3 where we aggregate over all flights and routes. Our baseline model—used to approximate present day airline pricing practices—is shown in the first row. We normalize the outcomes in this baseline to 100% for all welfare measures (leisure and business consumer surplus, quantity sold, revenues, and welfare).

Table 3: Counterfactual Estimates and a Comparison to Present Practices

Counterfactual	CS_L	CS_B	Q	Rev	W
1) Observed Inputs	100.0	100.0	100.0	100.0	100.0
2) Pricing Department Deviation	59.8	98.7	76.5	95.1	95.1
3) RM Department Deviation	<i>see Figure 5</i>				

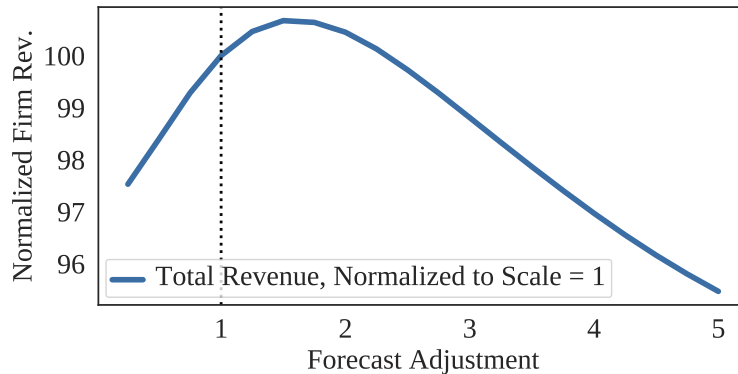
Note: In counterfactual (1), we approximate current pricing practices. Counterfactual (2) and (3) address a single organizational team bias, but leave others in place. Finally, in counterfactual (4), we consider a scenario in which RM and pricing department decisions are coordinated.

In row 2, we consider the pricing department deviation. Although this counterfactual addresses miscoordination between departments, we find that it actually reduces overall revenues by 5%. The reason is that although the deviation removes fares inconsistent with the RM department’s forecasts (Model F), it involves coordination to a biased input. Moreover, because the kludge used by the RM department raises demand expectations far from departure the most, coordination in this counterfactual ends up removing fares on the elastic side of Model E demand estimates. Higher prices result in a substantial drop in quantity, lowering leisure consumer surplus. Business consumer surplus drops slightly, and we estimate an increase in dead-weight loss of 5%. Because this is not a profitable deviation, miscoordination is preferred.

Next, we consider the deviation in which the RM department reassesses the bias

it has introduced into its forecasts, holding the pricing department decisions fixed. Because the RM department typically scales up or down the forecast for an entire route (or many routes) with a scaling parameter, we implement this counterfactual in a similar way. We have the RM department solve for inventory allocations using the pricing heuristic and its forecasts scaled uniformly through a scaling parameter χ . We consider $\chi \in \{0.25, 0.5, 0.75, \dots, 5.0\}$. For each χ , the RM department solves for inventory allocations using Model F, scaled by χ , as the input, but simulates market outcomes using unscaled, Model F demand. The optimal χ is the scaling factor that maximizes total expected revenues.

Figure 5: Optimal Forecast Bias, Holding Fare Menus Fixed



Note: Counterfactual revenues under alternative forecast bias. The forecast bias is a scalar multiple of Model F demand, with 1.0 being estimated demand using the firm's forecasts.

Figure 5 plots expected revenues as a function of the forecast bias, where we normalize results to the current forecast bias. We find that the current bias is nearly optimal for the firm. On the one hand, reducing the bias results in a decline in revenues, confirming that 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 the optimal forecast bias is an impermissible input. This is likely due to the fact that forecast bias is a key-performance indicator.

Manipulating the forecast bias has a modest impact on revenues because the heuristic is significantly impacted by other department inputs (as well as its own biases). In this deviation, we keep the fare menus unchanged and hold initial capacity fixed. Under alternative scaling factors, the opportunity costs derived from the heuristic are not so different as to warrant restricting availability of certain fare classes. As a result, the heuristic often selects the same fare as in the baseline scenario.

Our analysis of department input deviations establish that current inputs constitute a “second-best.” That is, although inputs are biased and miscoordinated, they are nearly optimal (within 1% of revenue) for the firm, holding other department decisions fixed. This suggests a rational and strategic reason for observing pricing biases and miscoordination. We next investigate a benchmark dynamic pricing problem and then discuss why departing from second-best outcomes is difficult.

The current demand framework used at the firm has important limitations: it considers products in isolation, it assumes a constant elasticity of substitution within consumer type, and heterogeneity in preferences across routes is not fully captured due to the number of routes using the same set of demand parameters. Would market outcomes differ substantially if the firm implemented a rich demand model that accounts for route-specific heterogeneity? We consider a second type of unilateral deviation that uses our new, external demand model: the RM department replaces Model F with Model E (Poisson Random Coefficients) demands, holding all remaining inputs fixed. We find that outcomes are again largely unchanged. Consumer surplus is within 0.2% of current practices; output declines 0.2%, and revenues rise less than 2%. Akin to the RM department forecast deviation, we find that a more accurate and flexible demand model does not greatly affect the value of remaining capacity because capacity is not sufficiently scarce and the heuristic ignores substitutes (we set substitute flights to observed, average prices). As a result, the heuristic allocates seats to the same low fares of the baseline scenario. This finding supports classic theories in organizational eco-

nomics of limited productivity gains under unilateral change when complementarities are important (Milgrom and Roberts, 1990, 1995).

8.3 First-Best Optimal Outcome for the Firm

We implement first-best prices in a hypothetical counterfactual where prices are determined by solving a dynamic programming (DP) problem using route-specific demand estimates (Model E). This counterfactual allows us to sign the potential distributional effects of removing pricing biases at a large firm by investigating how canonical models of pricing reallocate capacity across consumers. However, we emphasize that implementing the first-best scenario more broadly is not feasible due to the constraints that we will discuss below.

We retain discrete prices and use the distribution of fares currently offered to construct fare menus. 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). The highest price is the observed full-fare, refundable ticket taken from the data. This is our preferred menu for several reasons. First, the use of discrete prices allows for significantly faster computation than the use of continuous prices. Second, the pricing department has determined that this highest fare is viable. Consumers already face these fares in practice. This counterfactual allows us to quantify how using a DP would adjust the distribution of current fares offered. We adjust the lower end of the fares because our analysis of Model F suggests fares may sometimes be too high.²³

Formally, we implement first-best prices using a standard DP (j, d, r subscripts

²³Restricting solely to observed menus can force misallocation in that leisure consumers are not offered seats at lower price points even in situations where capacity is unlikely to be constrained.

suppressed), where the firm solves

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

such that unsold units are scrapped, and $C_0 \geq 0$ is given. In principle, we can accommodate an arbitrary number of substitutes; however, markets with greater than two flights contain more than 15 million states with a transition kernel size of one trillion. We implement the DP for routes with at most twice daily service.

In Table 4, we report counterfactual results comparing first-best prices to second-best prices, which are also normalized to observed practices (Row 1 of Table 3). We find that DP results in substantially different outcomes than the previous counterfactuals in which a single pricing input is adjusted. Here, the revenue improvement is over fourteen times the revenue effect of the RM department deviation. Implementing first-best prices results in a substantial change in the allocation of capacity over time. Two offsetting forces affect leisure consumers, highlighting the ambiguous welfare effects even within consumer type. On the one hand, DP results in higher capacity costs that act to raise the distribution of prices offered. On the other hand, Model E demand estimates suggest leisure consumers are more price sensitive than under Model F, which lowers the optimal price. The latter force dominates, and we estimate that leisure consumer surplus increases by 8%.

Table 4: Counterfactual Estimates of the First-Best Outcome for the Firm

Counterfactual	CS_L	CS_B	Q	Rev	W
Observed Inputs	100.0	100.0	100.0	100.0	100.0
4) First-Best Optimal Outcome	108.1	76.7	97.8	114.3	94.3

Note: Comparison in market outcomes of first-best prices to second-best prices. We also report results using Model E demand estimates in place of Model F.

In general, business consumers face significantly higher prices, reducing their wel-

fare by 23% and overall output by 2%. First-best prices increases dead-weight loss by a similar magnitude to the pricing department deviation counterfactual. These results are driven by the enhanced ability to segment markets. For the routes we study, moving from the current heuristic to a full DP does not significantly increase the importance of responding to demand shocks. These welfare estimates support descriptive work on airlines that argue price movements are largely driven by market segmentation motives instead of scarcity pricing (Puller, Sengupta, and Wiggins, 2015).

8.4 Barriers to Implementing the First Best

We have established that current inputs are second-best optimal for the firm and have characterized large differences in outcomes under a hypothetical first-best scenario. We briefly discuss constraints that prevent effective implementation of the first best.

The current algorithm is a necessary simplification because solving the dynamically optimal profit maximization problem is impractically slow for routes with greater than two flights a day. The current system—simplified to a single product problem—takes more than five hours to complete optimization every day. Moreover, our analysis abstracts from several key features of airline markets, including network considerations, overselling, and cancellations. Incorporating these features makes the optimal decision rule even harder to characterize, and is one of the key reasons current practices remain in place (Vinod, 2021).

Competition is another key factor in many airline markets, and current practices delegate these considerations entirely to the pricing department. In both our implementation of the first-best as well as under current practices, the pricing rule does not accommodate competitive interactions. Incorporating competition in dynamic environments is non-trivial. The current bifurcation of responsibilities enables the pricing department to react to competitor choices. However, as we have pointed out, the current organizational structure also allows these decisions to be offset by either decisions

of the RM department or the heuristic itself downstream, e.g., a sale fare does not receive positive inventory allocations.

Finally, we stress that important environmental factors exist which affect organizational structure and pricing. Global distribution systems that connect airline reservation systems to third parties, including travel agencies, require airlines to use a particular protocol. This necessitates the use of discrete prices and the design of several of the systems used at the firm. Abandoning GDSs might allow a firm greater flexibility in implementing dynamic pricing systems, but leaving the GDSs prevents a firm from being able to sell tickets via OTAs and conventional travel agents.²⁴ Indeed, currently indirect channels make up approximately one half of the tickets sold in our sample.

9 Conclusion

We study pricing at a large U.S. airline that has delegated the decision rights for pricing inputs to distinct departments. We show that although departments are miscoordinated and introduce multiple pricing biases, each department's pricing input is nearly optimal within its decision rights. Our results emphasize the importance of accounting for organizational details in drawing inferences about firm pricing behavior. We establish a rational explanation for observing pricing biases and miscoordination at a sophisticated and large U.S. firm. We show that assuming prices are derived from standard models of profit maximization can lead to incorrect inferences about demand fundamentals. Moreover, observing granular department decisions, bookings, and search data allow us to answer the fundamental question of how dynamic pricing systems

²⁴Lufthansa recently announced the introduction of "Continuous Pricing," which may suggest that it has implemented a pricing system similar to our first-best scenario. However, we have obtained documents from Lufthansa that show that Continuous Pricing not only requires the use of the current revenue management system, but that its capabilities cannot be integrated with the GDSs. In practice, Continuous Pricing allows for adjusting the filed fares by the pricing department by a small increment depending on the opportunity cost calculated by the heuristic. These adjusted fares are only available to customers who search on lufthansa.com. Copies of the documents are available on request.

affect welfare in practice. We compare observed, second-best optimal outcomes to several counterfactual scenarios. We find under a hypothetical first-best scenario that dynamic prices could exacerbate price targeting but note that current practices are reinforced due to computational and environmental factors.

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.
- Aryal, G., C. Murry, and J. W. Williams (2022): “Price Discrimination in International Airline Markets,” Working Paper.
- 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.
- 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.
- DellaVigna, S., and M. Gentzkow (2019): “Uniform pricing in us retail chains,” *The Quarterly Journal of Economics*, 134(4), 2011–2084.
- 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.
- Garicano, L. (2000): “Hierarchies and the Organization of Knowledge in Production,” *Journal of political economy*, 108(5), 874–904.
- Goolsbee, A., and C. Syverson (2008): “How do incumbents respond to the threat of entry? Evidence from the major airlines,” *The Quarterly Journal of Economics*, 123(4), 1611–1633.
- Hendel, I., and A. Nevo (2006): “Measuring the implications of sales and consumer inventory behavior,” *Econometrica*, 74(6), 1637–1673.
- Hortaçsu, A., O. Natan, H. Parsley, T. Schwieg, and K. Williams (2022): “Incorporating Search and Sales Information in Demand Estimation,” Working Paper.
- 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.
- Lazarev, J. (2013): “The Welfare Effects of Intertemporal Price Discrimination: an Empirical Analysis of Airline Pricing in US Monopoly Markets,” Working Paper.
- 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.
- 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*.
- Puller, S. L., A. Sengupta, and S. N. Wiggins (2015): “Does Scarcity Drive Intra-Route Price Dispersion in Airlines?,” *working paper*.
- 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.
- Smith, B. C., J. F. Leimkuhler, and R. M. Darrow (1992): “Yield management at American airlines,” *Interfaces*, 22(1), 8–31.
- Sweeting, A. (2012): “Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets,” *Journal of Political Economy*, 120(6), 1133 – 1172.
- Sweeting, A., J. W. Roberts, and C. Gedge (2020): “A Model of Dynamic Limit Pricing with an Application to the Airline Industry,” *Journal of Political Economy*, 128(3), 1148–1193.
- Talluri, K., and G. Van Ryzin (2004): *The Theory and Practice of Revenue Management*, vol. 68. Springer Science & Business Media.
- 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	The Presence of Pricing Frictions	14
C.4	Booking Trends Across Booking Channels	15
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

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 6. 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 6: 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 7: Nonstop, One-stop and Connecting Traffic

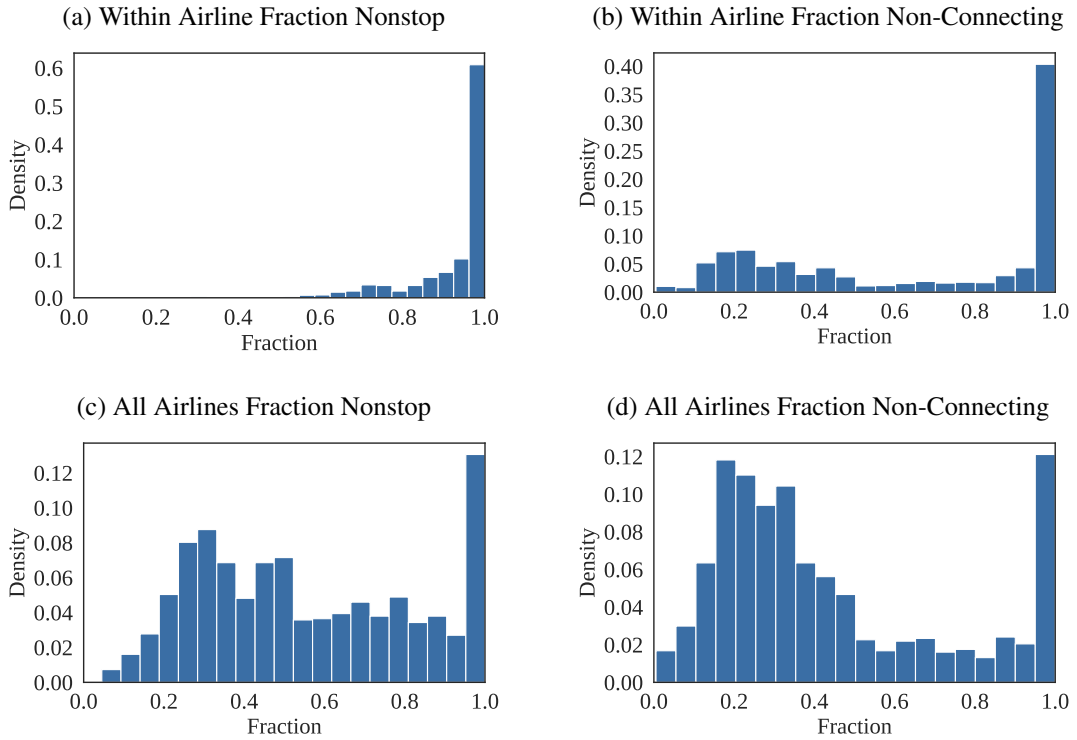


Note: We use the term nonstop to denote the sold black line, or passengers solely traveling between (Origin, Destination). Unless otherwise noted, we will use directional traffic, labeled $O \rightarrow D$. Non-directional traffic is specified as $O \leftrightarrow D$. The blue, dashed lines represent 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 7 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 8 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 8: Route Selection Using Bureau of Transportation Statistics Data

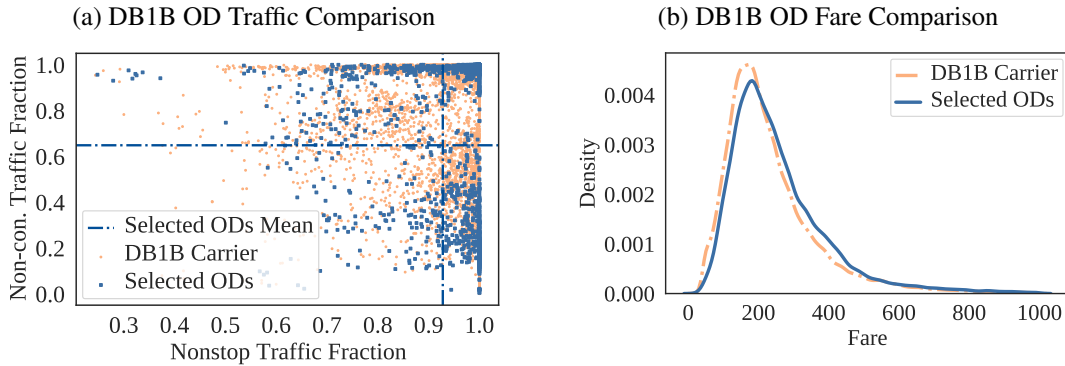


Note: Density plots over the fraction of nonstop traffic and the fraction of non-connecting traffic for the selected routes using DB1B data. "Within" means passengers flying on our air carrier. "Total" means all air carriers on a given origin-destination pair. Within nonstop and total nonstop coincide 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 9: 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 9-(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 9-(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

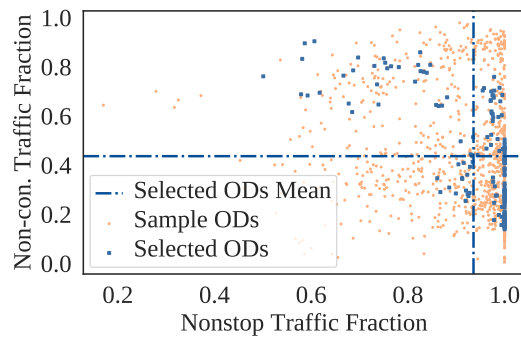
only by our air carrier. We report percentage differences between our estimation routes and the entire sample for key characteristics below in Table 5. Figure 10 shows a two-way plot of the fraction of nonstop and non-connecting traffic for the routes selected for estimation relative to the entire sample. Figure 11 recreates Figure 8, 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%

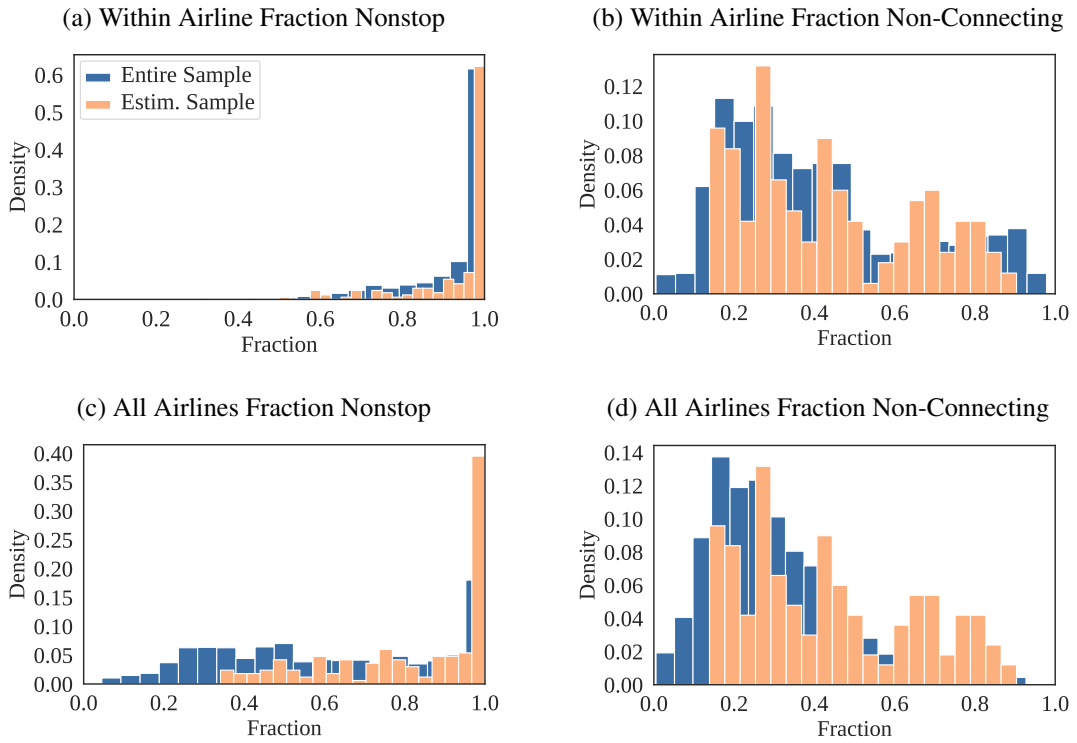
Note: Statistics calculated using the DB1B data for the years 2018-2019.

Figure 10: 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 11: Estimation Route Comparison

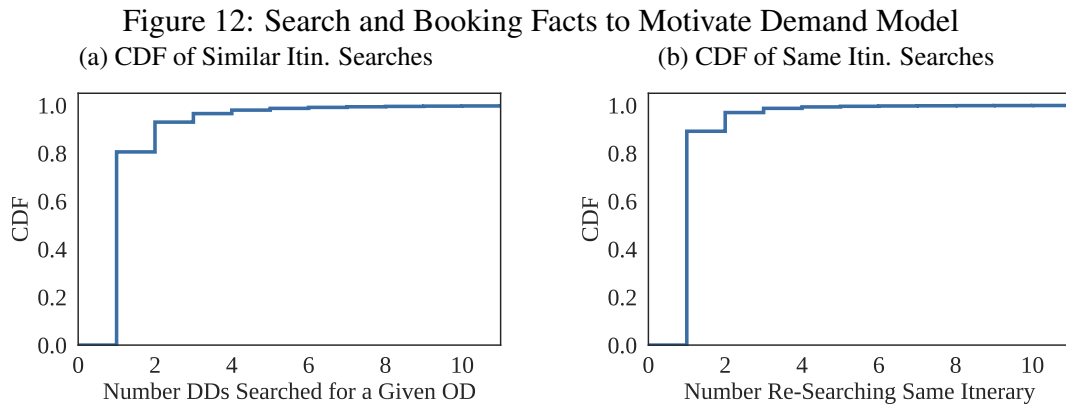


Note: Density plots over the fraction of nonstop traffic and the fraction of non-connecting traffic for the selected routes using DB1B data. "Within" means passengers flying on our air carrier. "Total" means all air carriers on a given origin-destination pair. Within nonstop and total nonstop coincide 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 12 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.



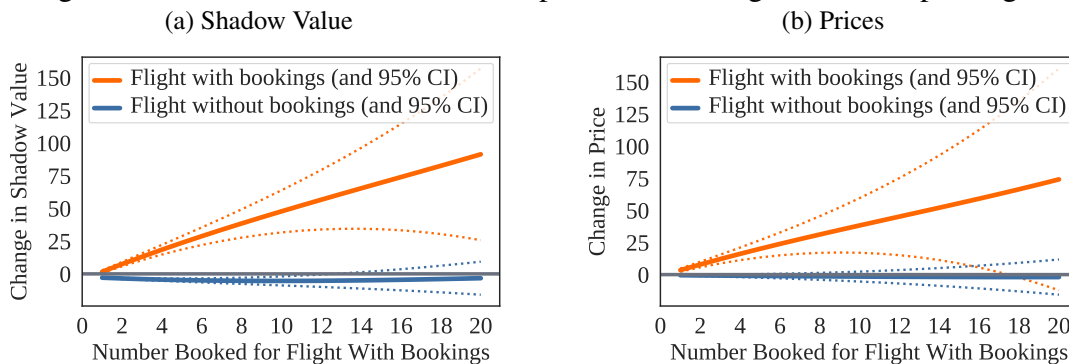
(a) CDF of the number of departure dates searched. (b) CDF of the number of days from departure searched for a given itinerary.

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 14); (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 13-(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 13-(a) confirms substitute shadow values are unaffected by bookings. Panel (b) shows that there is no price response.

Figure 13: Shadow Value and Price Response to Bookings with Multiple Flights



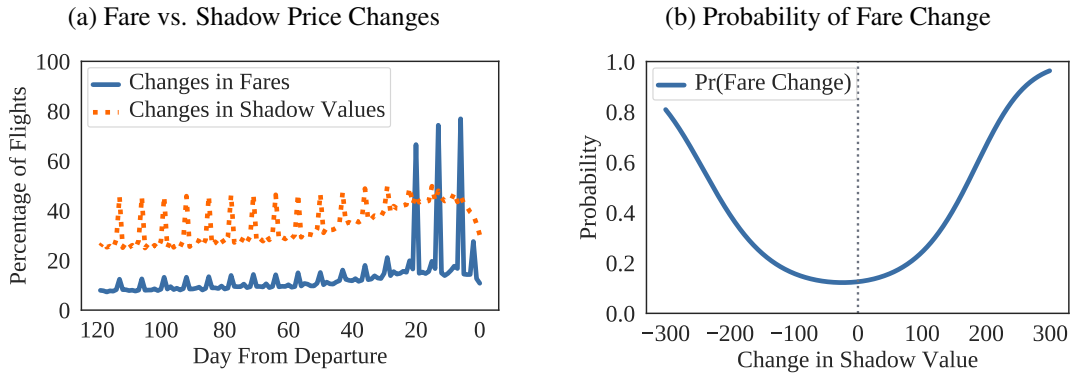
Note: (a) The orange line denotes the average change in shadow value for a flight with bookings. The blue line is the average change to shadow value when a sale occurs for the substitute product. (b) This panel depicts the same as panel a, but instead of changes in shadow value it depicts changes in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

C.3 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 14-(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.

Figure 14: Fare Adjustments in Response to Shadow Value Changes

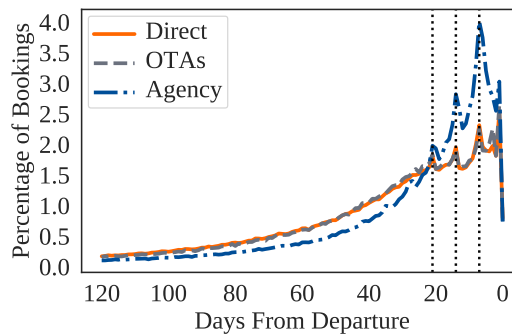


Note: (a) The fraction of flights that experience changes in the fare or the shadow value of capacity over time. (b) The probability of a fare change, conditional on the magnitude of the shadow value change.

Figure 14-(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 demand expectations. The process of reforecasting demand leads to a larger fraction of flights experiencing a change in the value of remaining capacity.

C.4 Booking Trends Across Booking Channels

Figure 15: 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 15 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:

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

$$A_{t,d}^B \sim \text{Poisson}(\lambda_{t,d}\tilde{\gamma}_t\zeta_t^B),$$

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

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

$$q_{j,t,d} = \min\{\tilde{q}_{j,t,d}, C_{j,t,d}\}.$$

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:

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

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

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

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 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 ν 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, ν , 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 $Y_k = (\nu, \xi)$ 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.