

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

By

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

November 2021

Revised March 2022

COWLES FOUNDATION DISCUSSION PAPER NO. 2312R



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY

Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

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

Ali Hortaçsu, *University of Chicago and NBER*
Olivia R. Natan, *University of California, Berkeley*
Hayden Parsley, *University of Texas, Austin*
Timothy Schwieg, *University of Chicago, Booth*
Kevin R. Williams, *Yale School of Management and NBER**

March 2022

Abstract

Although typically modeled as a centralized firm decision, pricing often involves multiple organizational teams that have decision rights over specific pricing inputs. We study team input decisions using comprehensive data from a large U.S. airline. We document that pricing at a sophisticated firm is subject to miscoordination across teams, uses persistently biased forecasts, and does not account for cross-price elasticities. With structural demand estimates derived from sales and search data, we find that addressing one team's biases in isolation has little impact on market outcomes. We show that teams do not optimally account for biases introduced by other teams. We estimate that corrected and coordinated inputs would lead to a significant reallocation of capacity. Leisure consumers would benefit from lower fares, and business customers would face significantly higher fares. Dead-weight loss would increase in the markets studied. Finally, we discuss likely mechanisms for the observed pricing biases.

JEL Classification: C11, C53, D22, D42, L10, L93

Keywords: Pricing, Organizational Structure, Revenue Management, Pricing Frictions, Behavioral IO

*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 also thank the seminar participants at Yale University, University of Virginia, Federal Reserve Bank of Minneapolis, University of Chicago, University of California-Berkeley, the NBER 2021 Summer Institute, the Virtual Quantitative Marketing Seminar, the NBER Organizations Meeting, Harvard University, the 2022 ASSA meetings, and Princeton University for comments. 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. Emails: hortacsu@gmail.com, olivia.natan@berkeley.edu, timothy-schwieg@uchicago.edu, parsley@utexas.edu, kevin.williams@yale.edu

1 Introduction

In economics, the pricing decision of the firm is typically modeled as the firm choosing a single price (or quantity) to maximize a known objective. In reality, large firms are not unitary decision makers. Rather, they are usually comprised of teams, each responsible for a particular sub-decision. Prices are often set through complex optimization systems designed around the organizational structure of the firm. One team might manage procurement and inventory, another team specializes in demand predictions, and an additional team monitors competitive response. Each team delivers its specific input to the entity determining price. For example, major airlines pioneered the use of pricing optimization algorithms and operationalized these systems by delegating the decision rights for each of multiple pricing inputs to distinct teams. Other industries, notably hotels, cruises, car rentals, entertainment venues, and retail, have adopted features of the airline pricing model. Given the prevalence and durability of this type of organizational structure, we may expect that economic forces lead to pricing outcomes that are “as if” each firm centrally chooses prices, as commonly modeled in economics.

In this paper we study organizational team inputs and pricing decisions by leveraging a data partnership with a large international air carrier based in the United States.¹ The granularity of the data allows us to understand how the decisions of individual teams affect prices without needing to assume that prices are optimally set. We document that the pricing at a sophisticated firm—one that employs advanced optimization techniques and relies heavily on automation—is subject to miscoordination and multiple biases. To quantify how biases affect the allocation of scarce capacity, we estimate a model of airline demand using sales and search information and contrast our predictions with the firm’s (biased) forecasts. In counterfactuals, we show that if one team addresses its own “frictions” in isolation, outcomes are largely unchanged. Moreover, we show that teams do not optimally account for biases introduced by other teams, that is, inputs are not second-best optimal for the firm. Finally, we estimate a significant reallocation of surplus would result if teams corrected bi-

¹The airline has elected to remain anonymous.

ases and coordinated on their pricing inputs. Leisure consumers would benefit from lower prices in advance, and business customers would face significantly higher fares.

We begin by discussing airline pricing practices and the decision rights of the organizational teams involved in pricing in Section 2. The organization structure is vertical; decisions are passed along to the next team as an input. The first team (the network planning department) decides where to fly and assigns initial capacities. We do not model these decisions. The next team (the pricing department) takes the network as given and is responsible for choosing a menu of discrete possible prices consumers may face as well as ticket restrictions, e.g., advance purchase discounts. A final team (the revenue management department) is responsible for the demand forecast and for the choice of the pricing algorithm which combines all inputs. The algorithm selected does not actually decide price, rather, it allocates the amount of inventory to sell at each discrete price level. That is, “pricing” involves both prices (as decided by the pricing department) and quantities (as decided by the revenue management department). We discuss the information used by teams when deciding their inputs. Using job listings, we show that all major airlines have similar organizational structures and team responsibilities. Therefore, we believe our discussion and subsequent empirical findings likely hold for the airline industry broadly and perhaps for other industries that have adopted similar pricing technologies and organizational structures.

We introduce the data in Section 3. In addition to observing prices and quantities, we also observe granular demand forecasts, the demand model, outputs of the pricing and allocation algorithms, the optimization code itself, and clickstream data that detail all consumer interactions on the airline’s website. The core data cover hundreds of thousands of flights spanning hundreds of domestic origin-destination pairs.

In Section 4, we provide evidence of pricing biases and miscoordination. Within teams, we show that the revenue management (RM) department maintains persistently biased forecasts over two years of data. The RM department has also chosen a pricing algorithm that is inherently biased. The selected algorithm does not consider cross-price elasticities, mean-

ing flights are priced independently. Across teams, we show teams' input decisions lead to pricing frictions, where changes in marginal (opportunity) costs do not lead to price adjustments. Although this naturally arises due to the use of coarse pricing, or "fare buckets", in the industry, the fact that opportunity costs sometimes adjust by hundreds of dollars without triggering a price adjustment suggests a mismatch between the fares chosen by the pricing department and demand fundamentals.

We present two forms of miscoordination. We show there is a "glitch" in that the pricing algorithm sometimes allocates inventory to fares that do not exist. This occurs because the pricing team has not informed the RM team that these fares have been removed from the menu. This affects 11% of observations. More importantly, we show that the fare choices of the pricing department are often "too low." When the fare choices are combined with the forecast, the RM department observes that some fares are assigned on the inelastic side of the RM department's demand estimates. Because of the way inventory is allocated by the pricing algorithm, it is possible that suboptimally low fares receive positive allocations. We find that prices offered to consumers are on the inelastic side of the RM department's forecasts in one third of the data sample. We observe no instances in which the pricing department removes these low fares from the system.

We document that pricing biases affect all routes, regardless of market structure, and are even more pronounced in competitive markets (larger forecast bias and more frequent "inelastic prices" based on the RM department's forecasts of residual demand). Due to the additional complexity of modeling competitive interaction, our subsequent analysis concentrates on routes where our carrier is the only airline providing nonstop service.

In the second stage of our analysis, we examine the pricing inputs using a structural approach. We estimate a model of consumer demand using a recently proposed demand methodology (Hortaçsu, Natan, Parsley, Schweg, and Williams, 2021). In Section 5, we consider a model in which "leisure" and "business" travelers arrive according to independent and time-varying Poisson distributions in discrete time. Consumers know their preferences and solve discrete choice maximization problems. Each short-lived consumer

chooses among the available flight options or an outside option.

We estimate the model using consumer search and bookings data. Aggregate search counts calculated from the clickstream data inform the overall arrival process, and we identify the price coefficient using instrumental variables (see Section 6). The estimates, presented in Section 7, reveal meaningful variation in demand, with a general increase in search for travel as the departure date approaches and substantial changes in the overall price sensitivity of consumers over time.

We then estimate “firm beliefs” about demand by matching the RM department’s forecasts to our demand model in Section 8. This allows us to flexibly recover the preference parameters consistent with the RM department’s demand predictions. We find that firm beliefs do not sufficiently distinguish between business and leisure willingness to pay and features too little change in the mix of arriving customers toward the departure date. Firm beliefs understate the heterogeneity in preferences within and across routes.

In Section 9, we perform counterfactual analyses using a pricing algorithm that closely follows the heuristic that the firm uses in practice. We show that if the RM department uses an unbiased forecast, with fare menus unchanged, or, if the pricing department aligned its pricing menus to the RM department’s biased forecasts, the pricing algorithm selects very similar inventory allocations. Consequently, outcomes are largely unchanged. We also consider outcomes as if the firm centrally chooses prices. With an unbiased forecast and fare menus tailored to this forecast, we estimate that prices charged to early-arriving leisure passengers would fall, and late-arriving business travelers would face higher prices. We find that revenue increases upward of 18% for some markets, and dead-weight loss may rise by over 10%.

Our results establish that teams are not best responding to the (biased) inputs of other teams—that is, outcomes are not near second-best optimal for the firm. We show that if the upstream team, the pricing department, were given unbiased demand estimates, the department could alter its fare menu choices—holding the RM department inputs fixed—and achieve 95% of the gap to the first-best outcome for the firm. In a second counterfactual,

we show that holding the pricing department’s fare inputs fixed, the RM department could improve overall revenue to 33% of the gap to the first-best outcome by manipulating its current demand forecasts. This requires substantially more biased forecasts to counteract that the pricing algorithm commonly allocates capacity to suboptimally lower fares inputted by the pricing department.

We close by discussing potential explanations for our findings. Although we cannot establish a causal link for the biases we explore, we argue that performance metrics, including a focus on load factors, limited transfer of information across teams, and limited experimentation are likely mechanisms. We rule out some explanations. For example, we argue that teams are not intentionally distorting pricing inputs so that the pricing algorithm considers long-run demand because of the algorithm choice, forecasting model, and forecast bias across types of tickets sold. Another hypothesis is that the firm manipulates inputs due to long-term competitive reasons; however, we argue the presence of (even more pronounced) biases in competitive markets suggests managers are not attempting to deter entry with limit pricing.

1.1 Related Literature

Our findings that prices are neither first nor second-best optimal for a large U.S. firm support classic theories in organizational economics which posit that coordination on complementary tasks may be difficult in practice (Milgrom and Roberts, 1990, 1995; Siggelkow, 2001). Although the adoption of information technology (IT) can increase productivity when complementary organizational and management practices are implemented alongside these investments (Bresnahan, Brynjolfsson, and Hitt, 2002; Bloom, Sadun, and Van Reenen, 2012), firms may not adopt technologies that increase productivity or revenues (Atkin, Chaudhry, Chaudry, Khandelwal, and Verhoogen, 2017; Sacarny, 2018).² In our context, the airline has adopted advanced IT, however, pricing biases prevent the carrier from achieving outcomes as if prices were centrally determined, absent frictions.

²Brynjolfsson and Milgrom (2012) provide an overview of this and related work.

Our analysis of biases introduced within teams complements studies on miscalibrated firm expectations (Massey and Thaler, 2013; Akepanidaworn, Di Mascio, Imas, and Schmidt, 2019; Ma, Ropele, Sraer, and Thesmar, 2020), and our analysis of biases introduced across teams allows us to explain deviations from optimal behavior shown in other settings (Blake, Nosko, and Tadelis, 2015; Levitt, 2016; Dubé and Misra, 2021). Our emphasis on miscoordination of pricing inputs highlights the problem of “implementation” as described in Gibbons and Henderson (2012). That is, prices are known to frequently be set on the inelastic side of (the RM department’s own estimates) of demand, removing these too low fares is necessary to optimize on price, but managers have not implemented changes to correct this bias. We compare current outcomes to the scenario devoid of pricing biases but we note that this outcome may not be achievable if the cost of communication is high. Dessein and Santos (2006) examine a related team-theoretic model theoretically.

The data allow us to directly measure pricing frictions, as observed marginal cost changes may not lead to price adjustments. These price rigidities have been argued to occur in other industries, including DellaVigna and Gentzkow (2019) and Hitsch, Hortaçsu, and Lin (2021) in retailing, Huang (2021) in peer-to-peer markets, Ellison, Snyder, and Zhang (2018) in online retailing.

Finally, we contribute to a large empirical literature on the airline industry, e.g., McAfee and Te Velde (2006); Li, Granados, and Netessine (2014); Puller, Sengupta, and Wiggins (2015); Aryal, Murry, and Williams (2021); Williams (2021), without imposing that prices are optimally set. We quantify the welfare effects of dynamic pricing in airline markets using heuristics used in the industry (Littlewood, 1972; Belobaba, 1987, 1989; Brumelle, McGill, Oum, Sawaki, and Tretheway, 1990; Belobaba, 1992; Wollmer, 1992).

2 Organizational Structure and Division Responsibilities

We study the US airline industry, an industry that directly supports over two million jobs and contributes over \$700 billion to the US economy.³ In 2019 alone, 811 million passengers flew within the United States.⁴ In addition to being an important industry in its own right, airlines have influenced the development of pricing technologies that are now used in other sectors—for example, in hospitality, retailing, and entertainment and sports events. Although the sophistication of these technologies has improved, many of the original yield management ideas described in McGill and Van Ryzin (1999) and Talluri and Van Ryzin (2004) remain in place today.

Fares at our air carrier depend on the actions of managers in several distinct departments. The organizational structure is vertical as decisions become increasingly granular, taking all previous departments' decisions as given. These decentralized decisions involve little coordination, as evidenced by their internal documentation, discussions with managers, and the data we describe below.

First, the network planning department decides the network, flight frequencies, and capacity choices. These decisions depend on operational constraints, traffic patterns, slot restrictions, and government approvals. We take the capacity constraints as given and do not endogenize network planning decisions. Having observed the network, the pricing department then sets a discrete menu of fares and fare restrictions for all possible itineraries.⁵ Finally, the revenue management (RM) department maintains the demand forecast and optimizes remaining inventory over the fares set by the pricing department. The diagram in Figure 1 depicts organizational boundaries involved with pricing decisions. In Online Appendix A, we use job listings to show that all major airlines are organized in this way.


The pricing team has two core responsibilities. The first is to facilitate the transmission

³See <https://www.iata.org/en/iata-repository/publications/economic-reports/the-united-states-value-of-aviation/>. July 1, 2021.

⁴See <https://www.bts.gov/newsroom/2019-traffic-data-us-airlines-and-foreign-airlines-us-flights-final-full-year>. July 1, 2021.

⁵Each filed fare contains an origin, destination, filing date, class of service, routing requirements, and other ticket restrictions. A common fare restriction decided by the pricing department is an advance purchase discount, which specifies an expiration date for a discounted fare to be purchased by. These discounts are commonly observed 7, 14, and 21 days before departure.

Figure 1: Division Responsibilities at all Airlines



Note: Key departments, responsibilities, and decision-making process at all airlines.

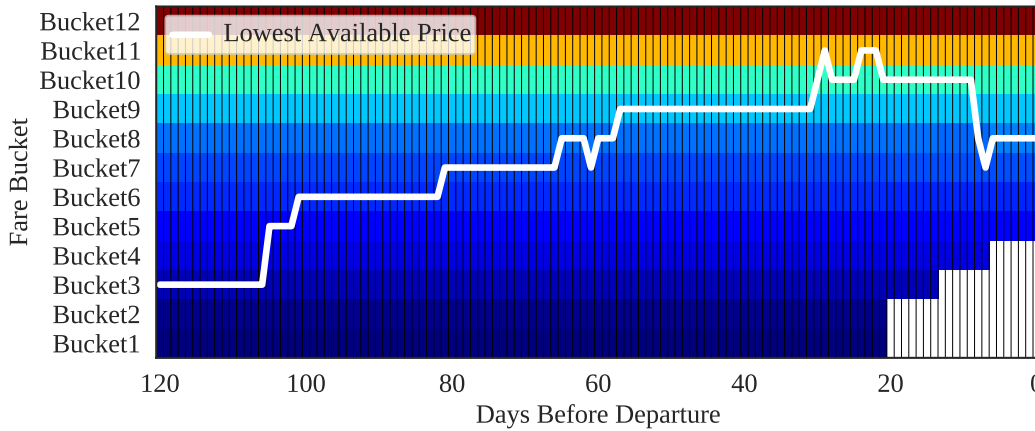
of fares through distribution channels so that consumers can retrieve and purchase tickets online. It is our understanding that this is a labor intensive process and involves real menu costs, as carriers are charged a small fee for each fare adjustment. The second core responsibility (and objective) is to assess the airline’s current fare positions relative to market conditions—including own connecting fares and competitor fares—and file fare updates. These decisions are made without the use of demand forecasts. Instead, the pricing team uses information on own historical fares and fares filed by other airlines for similar markets.

The RM department is responsible for allocating remaining inventory over time. It has decision rights over demand estimates, flight forecasts, and the pricing algorithm. The department does not have control over the fare and initial capacity decisions. The RM department aims to maximize revenues and closely monitors flight-level load factors and revenue yield as performance statistics (see Section 9.4 for additional details).

The forecasting model used by RM predicts flight-level bookings at granular intervals, e.g., quantity demanded by day before departure and price. The pricing algorithm chosen by the RM department is designed to maximize short-run, flight-level revenues. We observe exactly how the firm allocates remaining inventory (and hence, prices) over time. In our empirical analysis, we select a pricing algorithm that closely approximates the actual optimization tool. The heuristic we select is Expected Marginal Seat Revenue (EMSR-b). We provide additional details of the EMSR-b in Online Appendix B and outline the important features of the algorithm here. EMSR-b simplifies dynamic pricing problems with finite capacity and a deadline by assuming that all future demand will arrive tomorrow. The key trade-off, therefore, is to offer seats today versus reserve them for tomorrow. Given all of

the pricing inputs, it calculates the opportunity cost of a seat and then assigns the number of seats it is willing to sell at each price level. Lowest priced units are assumed to sell first. If expected future demand is high, it will restrict inventory at lower prices today. This raises the distribution of fares offered. When capacity is relatively unconstrained, the algorithm will allocate more capacity to lower fares.

Figure 2: 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.

To emphasize how the inputs come together in pricing, we show fare menu choices and the resulting pricing decision for an example flight in Figure 2. 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 lowest available price (LAP), or the lowest price with allocated inventory, by the pricing algorithm.

The pricing algorithm is influenced by all pricing inputs (initial capacity, prices, and the forecast) and it need not provide a good approximation to solving the dynamic program. For example, one could imagine removing the pricing department by simply inputting fares

in dollar increments, e.g., \$150, \$151, and so on. Because of the way the algorithm determines inventory allocations, granular fares can actually introduce distortions and yield worse outcomes relative to the use of coarse pricing. We show that fare menus greatly affect the pricing algorithm performance in Section 4 and Section 9.

The algorithm code and accompanying documentation detail what is included and excluded from optimization. All flights, regardless of market structure, flight frequencies, etc., are priced using the same algorithm. All forms of auxiliary revenue, including baggage fees and upgrade charges, are not considered when prices are determined.

3 Data and Summary Analysis

We use data provided by a large international air carrier based in the United States. To maintain anonymity, we exclude some data details. We do not study all routes served by the airline due to data size constraints; instead, we select over 400 routes. In Online Appendix C, 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 air carrier’s overall network. Nonetheless, our descriptive 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 major metropolitan areas to small cities. We focus solely on domestic flights.

3.1 Data Tables

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*: The bookings data contain details for each purchased ticket, regardless of booking channel, e.g., the airline’s website, travel agency, etc. Key variables included in these data are the fare paid, the number of passengers involved, the particular flights included in the itinerary, the booking channel, and the purchase date. Our analysis

concentrates on nonstop, economy class tickets.

(2) *Inventory data*: The inventory data contain the decisions made by the pricing algorithm. Inventory allocation is conducted daily. The data include the number of seats the airline is willing to sell for each fare class in economy and remaining capacity. We also observe output from the pricing algorithm, including the opportunity cost of a seat.

(3) *Search data*: We observe all consumer interactions on the airline’s website for two years. The clickstream data include any interactions or impressions that a consumer has on the website including, but not limited to, search queries, bookings, referrals from other websites, and the sets of flights that appear on every page the consumer visits. We link consumers across search sessions using a combination of web cookie(s) and login information tracked by the air carrier.

(4) *Filed fares data*: The filed fares data contain the decisions made by the airline’s pricing department. A filed fare contains the price, fare class, and all ticket restrictions, including any advance purchase discount requirements.

(5) *Forecasting data*: The RM department forecasts current and future demand at granular levels. The current-period forecast correspond to an economist’s definition of demand, i.e., quantity demanded for various prices, and future demand forecasts predict future bookings for given a particular forecasting period, i.e., quantity demand t days before departure as predicted $s > t$ days before departure. The firm maintains separate forecasts for “business” and “leisure” travelers. An algorithm classifies all search and bookings to these classifications. We also observe all managerial adjustments to the forecasts.

3.2 Data Summary

Table 1 provides a basic summary of the nearly 300,000 flights / 400 routes 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 flight fares in our sample are \$201, with large dispersion across routes and over time. Typically, prices for a particular flight adjust nine times and double in 120 days. Many

Table 1: Summary Statistics

Data Series	Variable	Mean	Std. Dev.	Median	5th pctl	95th pctl
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. Fares are for nonstop flights only. The booking rates are for non-award, direct travel on nonstop flights and for all traffic on nonstop flights, respectively. The number of passengers denotes the number of passengers per booking. The ending load factor includes all bookings, including award and connecting itineraries. The search rate is for origin-destination queries at the daily level.

price adjustments occur at specified times, such as after expiration of advance purchase (AP) discount opportunities. However, over 60% of price adjustments occur before the first set of AP fares expires.

In our sample, the average load factor is 82.2%. Although overselling is possible, we abstract from this possibility because we do not observe denied boarding/no show information. Our notion of capacity will be actual plane capacity plus the number of seats the airline is willing to sell over capacity (if any)—the observed “authorized” capacity.

3.3 Empirical Facts that Motivate Demand Assumptions

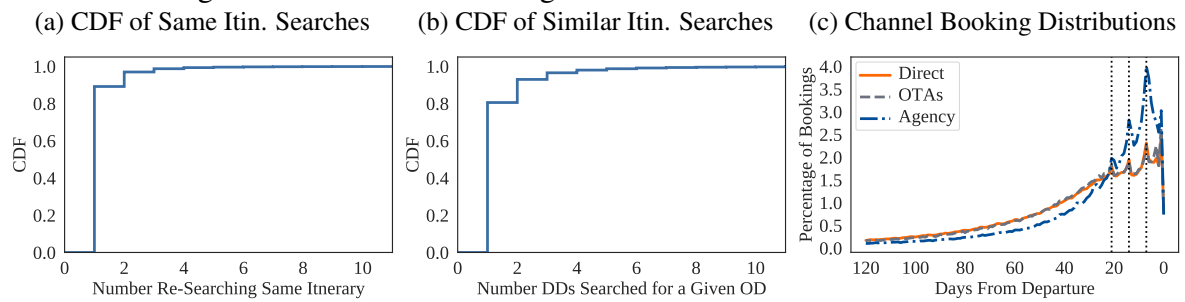
We summarize search and purchase patterns to motivate some of our demand assumptions.

The bookings data suggest that unit demand is a reasonable assumption. The average passengers per booking is 1.3. In addition, the bookings data confirm that overwhelmingly, consumers purchase the lowest available fare even though several fares may be offered at any point in time. We find that 91% of consumers purchase the lowest available fare. We verify that special fares, such as corporate or government discounts, are very rare in the routes studied. These fares are much more common in international markets.

Bookings and searches are sparse, which motivates using a model that accounts for low daily demand. We find that 60-80% of observations have zero observed searches. The fraction of zero sales is even higher (80% zeros). Zeros are not just present because we focus on nonstop demand. The fraction of zero sales for any itinerary involving a particular flight ranges between 40 and 80%.

We adopt a two-type consumer model, corresponding to “leisure” and “business” travelers, because that is how the firm forecasts demand.

Figure 3: Search and Booking Facts to Motivate Demand Model



(a) Empirical CDF of the number of days from departure searchers appear for a given itinerary. (b) Empirical CDF of the number of departure dates a given searcher looks for. (c) 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.

We use the clickstream data to explore consumer search patterns. We “daisychain” the data, meaning we link searches across devices and cookies whenever possible. We investigate the tendency for consumers to return to search for tickets, conditional on not being referred to the airline from other websites. In Figure 3-(a), we plot the CDF of number of times that consumers search for the same itinerary across days. 90% of consumers search for an itinerary (OD-DD pair) once. Panel (b) shows the CDF for the number of different departure dates (for the same OD) that consumers search over. 82% of customers search a single departure date. The average time lag between searches for different departure dates is 45 days, which likely suggests different purchasing opportunities (different trips).

Figure 3-(c) 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 di-

rect channel. However, they do not coincide. The agency curve—which includes corporate travel bookings—is more concentrated closer to departure. This motivates adjusting for unobserved searches differently over time (see Section 6.1). Note that Figure 3-(c) shows some bunching in bookings immediately before advance purchase opportunities expire. Although this may suggest consumers strategically time their purchasing decisions—they are forward looking—we find evidence that supports certain time periods simply have higher demands. Using the search data, we split the sample into two groups, one that includes routes that never have 7-day AP requirements, and one that includes these requirements. We find that search activity (and purchases) bunch at the 7-day AP requirement, regardless of their existence. Because arrivals increase regardless of price changes, we maintain the commonly used assumption that the market size is not endogenous to price. To account for this bunching in our model, we flexibly estimate arrivals as a function of days from departure and the departure date.

4 Pricing Biases

In this section, we provide evidence of pricing biases and miscoordination among teams managing pricing inputs.

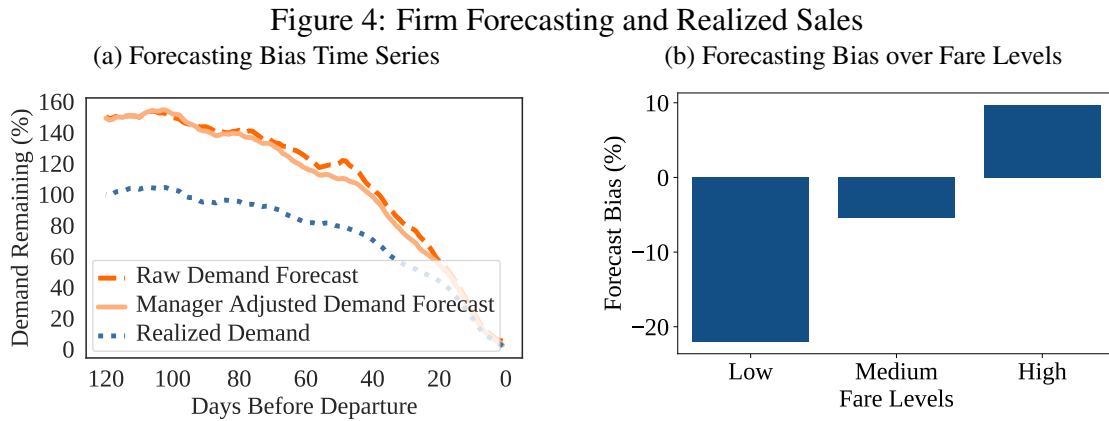
4.1 Pricing Biases within Organizational Teams

Using Persistently Biased Forecasts

We begin by discussing pricing biases introduced by a single team.

The demand forecasts maintained by the RM department are persistently biased in two years of data. In Figure 4, we plot the RM department’s forecasts of demand remaining along with realized future sales remaining. The lines decrease over time because fewer sales remain. On average, the forecasts are biased upwards from the true distribution of bookings for nearly the entire booking horizon. For the median observed forecast, the forecast is 10% higher than actual future demand, which is equivalent to predicting an extra

2.5 seats will be sold. Although the average forecast is biased upward, suggesting prices may be too high, the forecasts are also misaligned with observed demand at different prices (panel b). Low-fare transactions are underforecasted by 20%, and high-fare transactions are overforecasted by 10%.⁶ This suggests the forecasting model does not accurately reflect underlying demand.



Note: Forecasts and future sales are normalized by the aircraft's total coach capacity. Plots show 7-day moving average to smooth across strong day of week effects in the forecast sample. (a) Raw and managerial adjusted forecast of future demand compared to realized sales. (b) Average forecast bias for observations at differing fare levels.

We also observe all managerial adjustments to the forecasts, which are also plotted in Figure 4-(a). We find that manager adjustments reduced overall forecast bias, but the improvement is small in magnitude relative to the total absolute bias. Manager adjustments tend to deflate the forecasts for all flights for a given route (or routes) rather than react to individual flight realizations. These adjustments do not reduce forecast bias across the types of tickets (prices) sold.

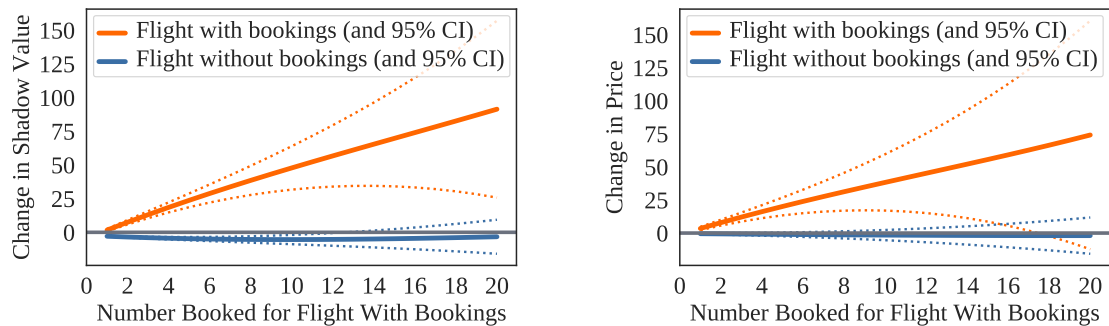
Forecasts are biased for all routes in the sample, and we do not find evidence that the firm is allocating more resources to reduce forecast bias in higher revenue generating routes. To the contrary, we find that routes with nonstop competitors feature slightly larger forecasting bias compared to single carrier routes. The frequency of managerial adjustments is similar across routes.

⁶Low fares refer to the bottom 3 fare classes; high fares refer to the top 3 classes.

Not Accounting for Cross-Price Elasticities

Reviewing the algorithm code, we confirm that the RM department has selected a pricing heuristic that does not internalize substitution to own products—for example, across cabins or to other flights—as well as substitution to other products, including all competitor flights. Therefore, another bias introduced within a team is that the RM department has selected a biased pricing algorithm that abstracts from cross-price elasticities. We use output from the pricing algorithm to demonstrate how this bias affects pricing decisions.

Figure 5: Shadow Value and Price Response to Bookings with Multiple Flights
(a) Shadow Value (b) Prices



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.

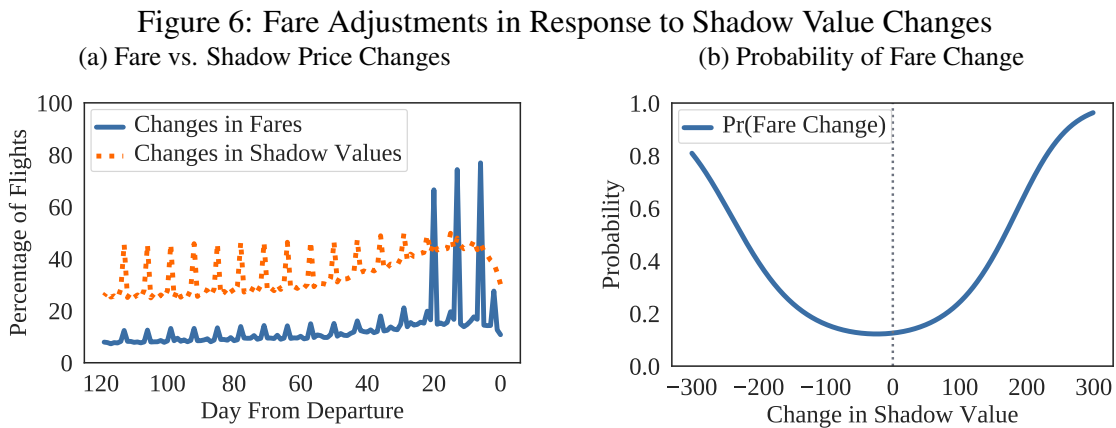
We select observations that satisfy the following conditions: (i) the firm offers two flights a day; (ii) we include periods where demand is not being reforecasted (the observed spikes in Figure 6); (iii) the total daily booking rate is low (less than 0.5); and (iv) one flight receives bookings and the 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 5-(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 5-(a) confirms substitute shadow values are unaffected by bookings. Panel (b) shows that there

is no price response.⁷

4.2 Pricing Biases across Organizational Teams

Pricing Frictions

Figure 6 demonstrates a pricing bias introduced across organizational teams. In panel (a), we plot the fraction of flights that experience changes in price or shadow value (as reported by the pricing algorithm) over time. The figure establishes that airline pricing is subject to pricing frictions in that costs change at a much higher frequency than do prices. This occurs because the pricing department assigns a discrete set of potential fares (fare buckets) for each route. This naturally creates a pricing friction because marginal costs may change by \$1 but the next fare may be \$20 more expensive. Our analysis suggests this friction is significantly more important. 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 marginal costs exceeding \$100 only lead to price adjustments with 20% probability. This suggests a mismatch between the fares set by the pricing department and underlying demand fundamentals for the routes in our sample.



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.

⁷We quantify the impacts of this bias in Online Appendix E.

Figure 6-(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.⁸ The process of reforecasting demand leads to a larger fraction of flights experiencing a change in the value of remaining capacity.

4.3 Miscoordination in Pricing Inputs

Allocating Inventory to Fares that do not Exist

One form of miscoordination present in the firm’s pricing decisions is that the pricing department does not necessarily ensure that the RM department has the correct set of active fares inputted into the pricing algorithm. Although this form of miscoordination can be seen as a “glitch,” its presence and prevalence suggests barriers in sharing information across teams.

Examining the menu of active fares (set by the pricing department) and the resulting inventory allocations (managed by the RM team), we observe inventory allocations to fares that do not exist. This affects 11.7% of observations. The most consequential form of this type of miscoordination is when the pricing algorithm allocates inventory to a fare lower than what is actually possible. For example, suppose there are two fares, at \$50 and \$100. The pricing algorithm may allocate seats to the \$50 fare, but if that fare does not actually exist, the realized price will be \$100. Consequently, the realized fare will be higher than what the pricing algorithm expects.⁹ We find that 71.2% of these “phantom allocations” occur to fares that are lower than the realized lowest available fare.

Smaller routes and routes with seasonal service are most affected by phantom allocations. 32.6% of routes feature persistent phantom allocations to at least one fare class throughout the entire data sample. We observe no actions that address this form of misco-

⁸We explore the idea of reforecasting/learning in Online Appendix E.

⁹Moreover, the pricing system never observes this allocation error because fare validation occurs after inventory allocation.

ordination.

Pricing on the Inelastic Side of the Demand Curve

A second form of miscoordination that we observe in the data is that the pricing department often files fares that are too low according to the RM department’s forecast. Although having some fares (within the menu) on the inelastic side of demand may seem innocuous, the chosen pricing algorithm cannot prevent allocating inventory to “too low” fares. This leads to mispricing—known to the RM department—yet we observe no instances where this miscoordination is corrected.

We provide a simple example to illustrate mispricing before presenting our evidence.

Example: Suppose the airline has 15 seats to sell over two days. Demand in the first period is equal to $Q_1(p_1) = 10 - 10p_1$, and demand in the second period is $Q_2(p_2) = 10 - p_2$. If the firm maximized revenues, using standard methods, we see that the capacity constraint would not bind, and optimal prices are equal to $(p_1, p_2) = (0.5, 5)$.

This outcome can also be achieved with separate pricing and RM departments and a pricing algorithm: The pricing department assigns prices to be $\{0.5, 5\}$ and $\{5\}$, and the RM department “forecasts” demand to be the functions above. Importantly, commonly used inventory management algorithms, including EMSR-b, only restrict inventory today (first period) to ensure that future demand can be satisfied tomorrow (second period). Because only 5 seats are needed tomorrow, it will allocate all seats to the price of 0.5 today. Tomorrow, the price will be 5—which can be viewed as an advance-purchase discount..

Instead, suppose the pricing department does not coordinate with the RM group and set prices equal to $\{.2, .5, 5\}$ and $\{5\}$. Because all first-period prices leave sufficient capacity available for the second period, EMSR-b will allocate all seats at the lowest price in the first period, which is now 0.2. This implies that the suboptimal price of 0.2 will be chosen even though the optimal price, 0.5, is included in the fare menu.

This simple example demonstrates a feature of the data that we observe often. Using the

RM department’s current-period demand forecasts, we calculate the elasticity of demand, $[(Q_1 - Q_2)/Q_2]/[(P_1 - P_2)/P_2]$, using the observed price as the base price along with the next higher fare.¹⁰ We find that 34% of flight observations are priced on the inelastic side of the demand forecasts estimated by the RM department. This is an example of the problem of implementation (Gibbons and Henderson, 2012), as the problem is identified, the solution is known, yet managers have not corrected this form of miscoordination.

We observe the (hashed) identities of the revenue management and pricing analysts involved in managing pricing inputs for each market, which allows us to explore if this form of miscoordination is limited to certain markets and/or analyst teams. We estimate regressions of the form $\mathbb{I}(\text{elasticity} > -1)_{r,i} = X_{r,i}\beta + u_{r,i}$, where X contains team identifiers as well as route characteristics. Although we find statistically significant differences across analyst teams, even the best teams are associated with setting “inelastic prices” 30% of the time. Higher traffic routes tend to have a larger percentage of inelastic prices. We conclude that miscoordination is widespread and not isolated to particular analyst teams and/or routes.

5 Empirical Model of Air Travel Demand

In order to quantify the welfare effects of pricing biases, we need to estimate a model of air travel demand. We utilize both the demand model and estimation approach of Hortaçsu, Natan, Parsley, Schwieg, and Williams (2021). We consider the demand for nonstop flights. 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 $t = 0$, and the flight departs at T . In each of the sequential markets t , arriving consumers choose flights from the choice set $J(r, t, d)$ that maximize their individual utilities, or select the outside option, $j = 0$.

¹⁰We also compute the arc elasticity, $[(Q_1 - Q_2)/(Q_1 + Q_2)/2]/[(P_1 - P_2)/(P_1 + P_2)/2]$. Using the arc elasticity, we find 52% of observations are priced on the inelastic side of the demand curve.

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 γ_t . We incorporate two assumptions to greatly simplify the demand system. First, we assume that consumers are not forward looking and do not strategically choose flights based on remaining capacity, $C_{j,t,d}$. This avoids the complication that consumers may choose a less preferred option in order to increase the chances of securing a seat. Second, we assume that when demand exceeds remaining capacity for a particular flight, random rationing ensures the capacity constraint is not violated.

We assume that the indirect utilities are linear in product characteristics and given by (suppressing the r subscript; all parameters are route specific)

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

where $X_{j,t,d}$ denote product characteristics other than price $p_{j,t,d}$. Consumer preferences over product characteristics and price are denoted by $(\beta, \alpha_{\ell})_{\ell \in \{B, L\}}$. For notational parsimony, we commonly refer to the collection $\{\alpha_B, \alpha_L\}$ as α . The term $\xi_{j,t,d}$ denotes an unobservable that is potentially correlated with price, and $\varepsilon_{i,j,t,d}$ is an unobserved 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,

$$u_{i,j,t,d} \geq u_{i,j',d,t}, \forall j' \in J \cup \{0\}.$$

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

is equal to

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

Since consumers are one of two types, we define $s_{j,t,d}^L$ be the conditional choice probability for a leisure consumer (and $s_{j,t,d}^B$ for a business consumer). Integrating over consumer types, we have

$$s_{j,t,d} = \gamma_t s_{j,t,d}^B + (1 - \gamma_t) s_{j,t,d}^L.$$

5.2 Arrival Processes and Integer-Valued Demand

We assume both consumer types arrive according to time-varying Poisson distributions. By explicitly modeling consumer arrivals, we can rationalize low or even zero sale observations. Specifically, we assume: (i) arrivals are distributed Poisson with rate $\lambda_{t,d}$, (ii) arrivals are independent of price (as argued in Section 3.3); (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

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

With the random rationing assumption, demand may be censored, i.e., $q_{j,t,d} = \min\{\tilde{q}_{j,t,d}, C_{j,t,d}\}$.

6 Estimation

6.1 Empirical Specification

Because consumer arrivals are observed at the t, d level, we cannot estimate the arrival process at the same granularity. Instead, we estimate the arrival process assuming a multiplicative relationship between day before departure and departure dates using the following specification,

$$\lambda_{t,d} = \exp(\lambda_t + \lambda_d).$$

We pursue this parameterization because searches tend to increase over time (λ_t) but there are also strong departure-date effects (λ_d). These parameters are route-specific.

In an ideal world, we observe all searches and estimate arrival rates using the sum of all leisure and business searches, i.e., $A_{t,d}^L + A_{t,d}^B$. However, we do not observe all searches—for example, a consumer that searches and purchases through a travel agency will result in an observed purchase without an observed search. Figure 3-(a) motivates adjusting for unobserved searches differently over time. We use the distributions of bookings and searches by passenger type as determined by the passenger-type classifier. Using properties of the Poisson distribution, we assign

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

where $\tilde{\gamma}_t$ is the firm's beliefs over the probability of business (see Section 8 for more details), and ζ_t^ℓ is the fraction of bookings that do not occur on the direct channel for each consumer type.¹¹ That is, we use the relative fraction of L (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 total bookings, and we assume unobserved searches follow the same underlying demand distributions, we can scale up estimated arrival rates by 5×. As we are concerned about the accuracy of this assignment algorithm, we conduct robustness to this specification in Online Appendix D.2.¹²

We assume consumer utility is given by

$$u_{i,j,t,d} = \beta_0 - \alpha_{\ell(i)} p_{j,t,d} + \text{FE}(\text{Time of Day } j) + \text{FE}(\text{Week}) + \text{FE}(\text{DoW}) + \xi_{j,t,d} + \varepsilon_{i,j,t,d},$$

where "FE" denotes fixed effects for the variable in parentheses. The flexibility in the utility

¹¹We use time intervals early on because of sparsity in searches and sales. The largest time window is composed of 14 days. Closer to the departure date, the intervals become length one. We smooth the calculated fractions using a fifth order polynomial approximation.

¹²An earlier version of this paper did not use the classifier and simply scaled up arrivals by the percentage of non-direct bookings. We obtained quantitatively similar demand estimates.

and arrival process specifications allows for rich substitution patterns, including seasonality effects, day-of-week effects, etc.

We parameterize the probability an arrival is of the business type as

$$\gamma_t = \frac{\exp(f(t))}{1 + \exp(f(t))},$$

where $f(t)$ 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.¹³

6.2 Estimation Procedure

We use a hybrid-Gibbs sampler to estimate route-specific parameters. With Poisson arrivals, we can rationalize zero-sale observations while maintaining a Bayesian IV correlation structure between price and ξ . Our approach builds upon the estimation procedure developed by Jiang, Manchanda, and Rossi (2009) by incorporating search, Poisson demand, and censored demand. Additional details on the estimation procedure can be found in Online Appendix D.1. A complete treatment can be found in Hortaçsu, Natan, Parsley, Schwiag, and Williams (2021).

6.3 Identification and Instruments

One difficulty in estimating a model with aggregate demand uncertainty is separably identifying shocks to arrivals from shocks to preferences. We address this complication by using search data. Conditional on market size, preference parameters are identified using

¹³We allow for the distribution of the random effects in demand, ξ and the residual of the pricing equation v (see Online Appendix for details), to vary by days before departure. We partition the distributions of shocks into four blocks of time. Within each block, each flight at the j, t, d level receives a demand shock drawn from a joint normal distribution, with the distributions themselves varying across blocks. This flexibility in the distribution of demand shocks allows us to capture observed varying managerial intervention in pricing over time. If we estimate the model with demand shocks exogenous to price, we estimate demand to be slightly more inelastic compared to this specification that allows for price endogeneity.

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.¹⁴ Online Appendix D outlines the Bayesian-IV procedure. The advance purchase indicators account for the fact that prices typically adjust in situations where the opportunity cost is not observed to change (see Figure 6). The total number of inbound or outbound bookings to a route’s hub airport captures the change in opportunity cost for flights that are driven by demand shocks in other markets. 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. Note that these instruments are relevant even if the teams that set prices are not optimally responding to information and each others’ actions. We are assuming only that prices tend to react to marginal cost and exogenous capacity changes.

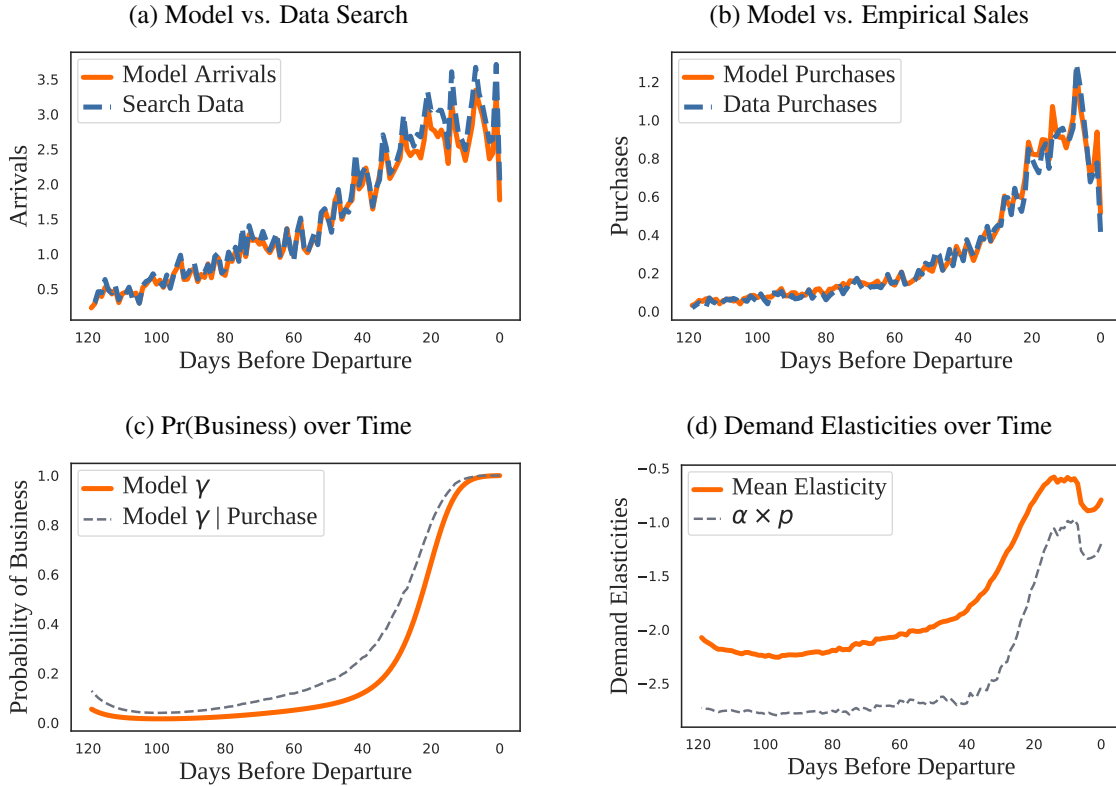
7 Demand Estimates

We select a subset of routes for estimation (39 ODs) where our air carrier is the only airline providing nonstop service. Our estimation sample includes routes with varying market characteristics, including flight frequency, importance of seasonality, and percentage of

¹⁴For 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 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.

nonstop and non-connecting traffic. Online Appendix C discusses the estimation sample in more detail.

Figure 7: Model Estimates for Example Route



Note: The horizontal axis of all plots denotes the negative time index, e.g. zero corresponds to the last day before departure. (a) Normalized model fit of searches with data searches. (b) Model fit of product shares with empirical shares. (c) Fitted values of γ_t over time, along with the probability a consumer is business conditional on purchase. (d) Mean product elasticities over time.

We first present results for an example route that demonstrates demand patterns we then confirm hold more broadly. For our example route, 88% of observations have zero product sales. It is not unusual to have so many zeros. The number of nonstop flights varies over the calendar year; typically, one or two flights are offered. In Figure 7-(a), we show that our arrival rates closely match the scaled up arrival data. Note that arrival rates are increasing toward the deadline. Panel (b) shows model and data booking rates over time. Model bookings closely follow the data and show a common pattern that purchases tend to increase as prices rise. This suggests demand becomes more inelastic, which we confirm in the bottom panels. Panel (c) reports our estimates of the probability of a business-type

consumer. We find a significant change in the composition of arriving consumers over time. This pattern is consistent with the airline demand estimates in Williams (2021). Recall that consumer types describe preferences, but not necessarily the reason for travel. In panel (d), we plot average own-flight price elasticities. Demand elasticities start at -2.1 and increase past -1.0 closer to the departure date.

Table 2: Demand Estimates Summary across Markets


Parameter	Mean	Std. Dev.	Median	25th Pctile.	75th Pctile.
Monday Arrivals	3.653	2.882	2.645	1.484	5.432
Tuesday Arrivals	3.001	2.260	2.030	1.352	4.827
Wednesday Arrivals	3.274	2.433	2.075	1.472	5.127
Thursday Arrivals	3.785	2.760	2.650	1.685	5.685
Friday Arrivals	4.395	3.432	2.995	2.007	6.119
Saturday Arrivals	3.085	2.412	2.175	1.285	4.490
Sunday Arrivals	4.286	3.393	3.426	1.764	6.466
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 the week.

In Table 2, we report variation in demand estimates across routes. The top panel shows average arrival rates for different days of the week. The interquartile ranges across routes confirm that average arrivals tend to be low. Friday and Sunday tend to be the busiest travel days for the routes in our estimation sample. The next 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). Time of day preferences tend to be stronger than day of week preferences. 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.

In Figure 8-(a), we plot arrival rates for the mean route as well as the interquartile range over routes. Although levels of arrivals vary—the interquartile range spans more than a doubling of arrivals—overall, search increases as the departure date approaches. In addition, demand tends to become significantly more inelastic over time, even though prices tend to rise. This is shown in panel (b), which shows average own-price elasticities for the mean, median, and interquartile range over routes. The drop off in elasticities close to the departure date mostly reflect very significant price increases after crossing advance purchase discount opportunities.

Figure 8: 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.

Within our estimation sample, we estimate a mean elasticity of -1.05 . We find that 56% of routes have inelastic demand at some point before departure. 82% of routes feature at least 10% of markets (departure date, days before departure) with inelastic demand. Just above half of the routes have inelastic demand on average. Inelastic demand tends to occur close to the departure date. We find that 85% of routes have inelastic demand in the final ten days before departure. We find no correlation between elasticity and number of searches for the route. In fact, although many of the inelastic routes tend to be routes from a large city to smaller regional cities, we find that the smallest and largest routes by search volume have elastic demand.

In Online Appendix D.2, we discuss demand estimates under alternative scaling param-

eters. Our findings are quantitatively similar under these alternative specifications.

8 Firm Beliefs about Demand

With our demand estimates in hand, we now ask, What does the firm believe demand looks like? To answer this question, we recover the preference parameters that best match the RM department’s demand forecasts.

We proceed in two steps. First, we recover “firm beliefs” on the arrival processes. We assume the RM department uses the same model of consumer arrivals and that the total intensity of demand is the same as our estimates, i.e., $\lambda_{t,d} = \lambda_t \lambda_d$. However, we allow the composition of arriving customers to vary over the booking horizon. For every route, we calibrate γ_t as

$$\gamma_t^{\text{beliefs}} = \frac{\sum \text{Arrivals}_t^B}{\sum \text{Arrivals}_t^B + \sum \text{Arrivals}_t^L},$$

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

Second, we recover preferences using the forecasts. These data are the predictions of sales quantities at the flight, departure date, passenger type, price, and day before departure. Whereas our previous analysis used the aggregate forecast (see Section 4.3), here we use the forecasts at the consumer-type level, $\ell \in \{L, B\}$.

We assume the RM department also uses a Poisson demand model, with the same specification as ours. Because the algorithm selected by the RM department considers single-product demand, we consider a single-product setting when recovering beliefs. We consider the unconstrained, cumulative forecast,

$$\tilde{Q}_{j,k,t,d}^\ell := \sum_{k' \geq k} E Q_{j,k',t,d}^\ell.$$

which is the forecast of (uncensored) unit sales at a price of k for consumer type ℓ . This forecast coincides to the definition of demand in economics, i.e., quantity demanded for a given price level. We assume the forecasting model assigns $\tilde{\lambda}_{t,d}^\ell$ as the arrival process for each flight $j \in J_d$.¹⁵ Our assumptions allow us to match the forecasting data to its corresponding model counterpart,

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

If we take logs of the equation above and subtract the log of the outside good share, we can use the inversion of Berry (1994) to obtain the following estimation equation

$$\log\left(\frac{\tilde{Q}_{j,k,t,d}^\ell}{\tilde{\lambda}_{t,d}^\ell}\right) - \log(s_{0,t,d}^\ell) = \log(s_{j,k,t,d}^\ell) - \log(s_{0,t,d}^\ell) = \tilde{\delta}_{j,k,t,d}^\ell. \quad (1)$$

This is only possible because the forecasting data is at the consumer-type level, which allows us to avoid using the contraction mapping in Berry, Levinsohn, and Pakes (1995) and Berry, Carnall, and Spiller (2006). Moreover, the inversion allows us to impose similar restrictions imposed in our model, i.e., the only difference in mean utility across consumer types is on the price coefficient.¹⁶

Defining the left-hand side of Equation 1 above as $\tilde{\delta}$, we obtain a linear estimating

¹⁵Instead, we could assume arrivals are $\tilde{\lambda}_{t,d}^\ell/J$, so that each flight receives $1/J$ of arrivals. This increases product shares and results in consumers estimated to be more price insensitive.

¹⁶We must also confront a data limitation in that our forecasting data is not necessarily at the t level, but rather, at a grouping of t s 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

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

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

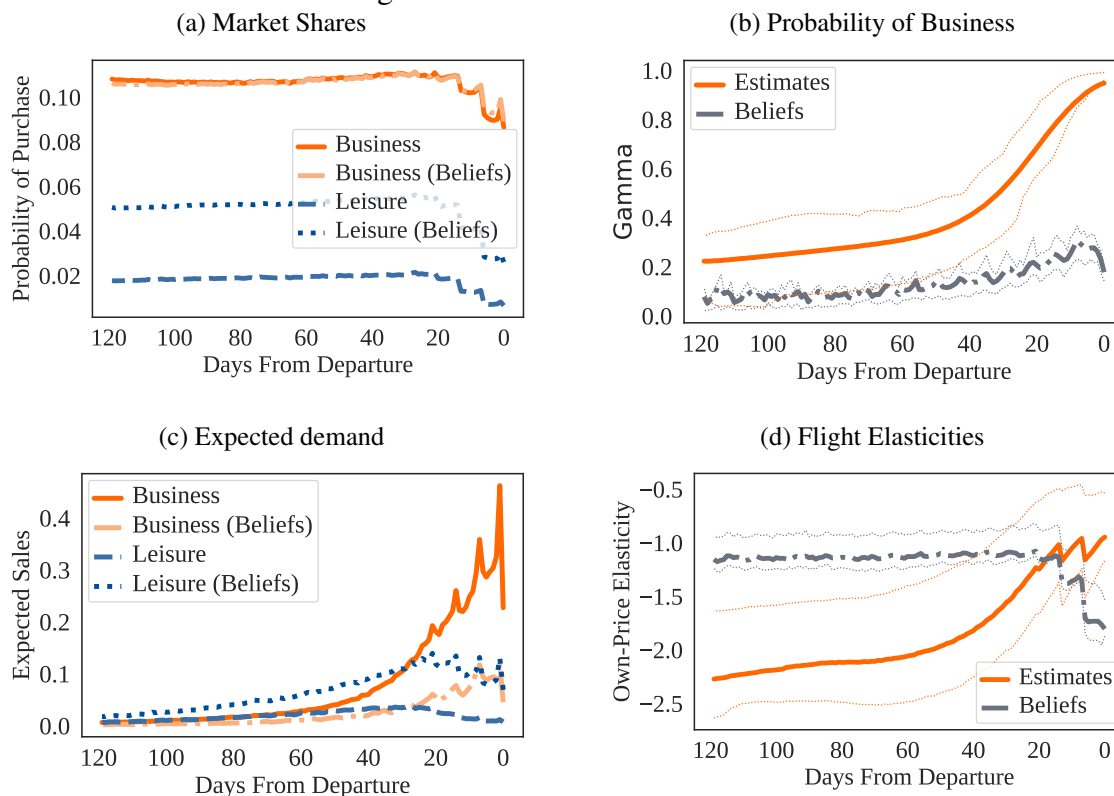
$$\frac{\tilde{Q}_{\cdot,t^*}^\ell}{\sum_{t \in t^*} \tilde{\lambda}_{\cdot,t}^\ell} = s^\ell(\cdot).$$

equation of the form

$$\tilde{\delta} = X\tilde{\beta} - \tilde{\alpha}p + \xi + u,$$

where $\tilde{\beta}, \tilde{\alpha}^B, \tilde{\alpha}^L$ are preferences to be estimated. We include our estimated ξ in the model, which is the mean of the posterior for that observation taken from our estimates. Thus, this approach also estimates a " ξ " that also differs across consumer types through u . We set these residuals equal to zero after recovering firm beliefs. These assumptions do not greatly impact our findings.

Figure 9: Firm Beliefs on Demand



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) Forecasted demand across consumer types, over time. (d) Comparison of own-price elasticities over time. (b) and (d) contain the 25th and 75th percentiles. Results are reported averaging over all observations in the data.

Thus, we obtain the following inversion,

$$\log\left(\frac{\tilde{Q}_{\cdot, t^*}^\ell}{\sum_{t \in t^*} \tilde{\lambda}_{\cdot, t}^\ell}\right) - \log(s_0^\ell) = \log(s^\ell) - \log(s_0^\ell) = \tilde{\delta}^\ell. \quad (2)$$

In Figure 9 we show that our demand estimates (Model E) are quite different than those recovered from the forecasting data (Model B). In panel (a) we plot product shares for both passenger types over time. Model B results in consumer types being “closer together” than under Model E, with leisure travelers being more price inelastic than under our estimates. In panel (b), we plot the probability that an arriving customer is a business traveler. Model E places more mass on business travelers and produces larger changes in the types of consumers arriving over the booking horizon. Model B suggests a significant drop in business consumer arrivals very close to departure. Panel (c) highlights that Model B inflates early arriving demand and understates business passenger demand close to departure. Finally, in panel (d), we plot own-price elasticities over time. Model E produces elasticities that are increasing (toward zero) as γ_t increases, whereas Model B results in mostly constant elasticities that then drop close to departure. This is due to both the probability of business declining close to departure and prices increasing substantially.

Overall, the two models are quite different. Model B yields more compressed demand elasticities where aggregate demand is more inelastic well in advance of the departure date than compared to Model E. This is driven by the upward bias in the forecasting data along with reduced variance in the forecasts across flights (relative to bookings). Model E suggests that there is more heterogeneity in demand across both flights and routes, with a more pronounced change in the overall price sensitivity of arriving consumers over time.

8.1 How compatible are the pricing and RM departments’ decisions?

Before turning to counterfactuals, we revisit the notion of miscoordination in that the fare choices of the pricing department are misaligned to the demand predictions by the RM department. We consider a simple scenario: Suppose capacity were sufficiently large so that capacity costs are zero. In this scenario, the dynamic pricing problem collapses to a static pricing problem. What would be the revenue maximizing price? The optimal price—according to the RM department’s forecasts—sets marginal revenue equal to zero, or $MR^{\text{Model B}}(p) = 0$. This price identifies the lowest price that the firm should ever charge

under Model B demand because scarcity would only increase the price.

We solve for the revenue maximizing price with zero capacity costs and compare these prices to the observed fares. We find that only 49.6% of observations involve fare menus where the minimum menu price exceeds the price that solves $MR^{\text{Model B}}(p) = 0$. Although we find that 85.1% of filed fares are higher than the lowest price that should ever be charged (higher fare classes are more expensive), 29.8% of observed prices are below the optimal price if capacity were unconstrained. These results match our descriptive evidence that fares are often too low, particularly close to departure when the booking rate increases as fares increase.

9 Empirical Analysis of Pricing Inputs

We quantify the impacts of the pricing biases we have documented on welfare through several counterfactuals. We explore three questions: (1) How does correcting biases individually affect welfare; (2) What are first-best optimal outcomes for the firm, (3) Are teams best responding to the (biased) inputs of other teams?

Our baseline model approximates current practices. We use the demand estimates based on the RM department’s forecasts (Model B) and the observed fare choices by the pricing department to compute baseline fares, allocations, revenues, and consumer surplus. We then consider two counterfactuals that investigate biases within teams and miscoordination as discussed in Section 4. First, we address forecasting bias by replacing Model B demand estimates with the Model E demand estimates. We leave the pricing menus unaltered. Second, to address miscoordination between teams, we tailor the pricing menus set by the pricing department to Model B demand estimates. To implement this counterfactual, we eliminate the ability for the pricing algorithm to choose any fare below $MR^{\text{Model B}} = 0$.¹⁷ This addresses the miscoordination described in Section 4 and Section 8.1 that the

¹⁷We have also implemented this counterfactual by setting the first fare equal to the price that solves $MR = 0$ using the Model B demand estimates. We then increase fares by scaling prices from $1\times$ to $2.5\times$ the minimum price spanning the number of observed buckets for each route. The results are similar.

pricing team often files fares that are on the inelastic side of the RM department’s forecasts (according to Model B). We leave the forecast set to Model B in this counterfactual.

These counterfactuals do not necessarily coincide with “second-best” outcomes for the firm, or where teams best respond to the biased inputs/processes of the other team. We explore the second best in Section 9.3. We also simulate a “first-best” outcome, where pricing is centralized under a single team that uses Model E demand estimates with fare menus tailored to the Model E demand forecast. These menus are derived by first solving $MR^{\text{Model E}} = 0$. This defines the lowest price on the menu. We then increase fares by scaling prices from $1\times$ to $2.5\times$ the minimum price spanning the number of observed buckets for each route.

One bias not addressed in these counterfactuals is that the pricing algorithm itself is also biased. We perform additional counterfactuals that allow for cross-price elasticities in a model of dynamic pricing in Online Appendix E.

9.1 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, we then draw preferences and arrival rates given our demand estimates (Model E). We simulate 10,000 flights for each combination of initial capacity and the drawn preference parameters. 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.¹⁸ 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. Therefore, if the lowest-priced fare has a single seat and is sold immediately, the next arriving consumer within a period is offered

¹⁸Alternatively, 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 in the data.

the next least-expensive fare.¹⁹

9.2 Welfare Comparison

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. Here, we use the demand model based on the firm’s forecasts (Model B), and the observed fares filed by the pricing team. 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 Fares, Model B Forecast	100.0	100.0	100.0	100.0	100.0
2) Observed Fares, Model E Forecast	99.9	99.8	99.7	101.9	100.6
3) Altered Fares to Model B Forecast	69.4	102.7	85.4	93.0	97.3
4) Altered Fares to Model E Forecasts	121.0	64.1	92.3	118.6	86.6

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 of Table 3, we investigate the impact of correcting the bias in the RM department’s demand forecasts, while preserving the coordination problem with the pricing team. To do this, we simulate outcomes keeping the filed fares at their observed levels, but replacing Model B demand estimates with Model E demand estimates. We find that total welfare under this counterfactual is within 0.6% of the calculated welfare for the benchmark case in Row 1. This occurs because the pricing algorithm generally expects that future demand

¹⁹Note that this differs from our demand model where all consumers are assumed to pay the same price within a period. However, because arrival rates are low, consumers very rarely pay different prices in our simulations. This is consistent with the data as well. We remove seven markets from our analysis that are estimated to have inelastic demand throughout time. These markets feature very low arrival rates and a very high percentage of zeros (over 95%). Our results are robust to including these routes, though the average revenue gains are over 5% higher with the inelastic routes included.

can be accommodated with remaining capacity under both Model E and Model B demands. Therefore, the opportunity cost of capacity estimated within the pricing algorithm is sufficiently low such that the algorithm typically allocates units to the lowest filed fare. Because the fare menus are the same in both counterfactuals, the lowest filed fare is made available in both cases. As a result, fixing the forecast in isolation does not affect market outcomes.

Row 3 addresses the miscoordination problem between the fare menu choices and Model B forecasts. Recall that in Section 8.1, we found that prices are commonly set on the inelastic side of the RM department's demand estimates (Model B). We find that removing fares on the inelastic side of the Model B forecasts actually reduces overall revenues by 7%. Although this finding may appear counterintuitive, recall that in this scenario fares are aligned to a biased forecast. Our results show that coordinating fares could actually provide worse outcomes for the firm (as well as consumers). We find that for 44% of the routes in the sample, revenues increase in this counterfactual. For over 56% of routes, revenues are within 1% of the baseline scenario. For the remainder of routes, and in particular, for a few routes that cater heavily to leisure customers, we find that current prices are actually too high (in terms of comparing $MR = 0$ under Model B and Model E). Removing the lowest fares on the menu exacerbates the problem of pricing too high, particularly well in advance of the departure date. In this counterfactual, we estimate that leisure consumer surplus would decline by 30.6%, and business consumer surplus would only increase by only 2.7%. Coordination without addressing all biases reduces overall welfare by 2.7%.

In Row 4, we consider a "first-best" scenario, where the firm uses the unbiased forecast (Model E) in conjunction with price menus coordinated to that unbiased forecast. This counterfactual represents the "as if" scenario where pricing is centrally decided at the firm. We estimate a significant reallocation of capacity compared to the benchmark case. The reason is that passenger types are further apart in terms of preferences according to Model E estimates. Because capacity is often not constrained, in this counterfactual fares early on tend to fall. Fares close to departure are much higher because business customers have higher willingness to pay according to Model E (see Figure 9). This leads to lower transacted

prices among leisure consumers, increasing leisure consumer surplus. Business consumer surplus falls sharply. Revenues would increase by 18% due to increased price targeting. Overall dead-weight loss would also rise in the markets studied.

9.3 Second Best Outcomes for the Firm

Our results establish that market outcomes are largely unchanged if a team corrects its own bias in isolation. We now ask: How close are current inputs to the second-best optimal inputs for the firm? We consider two scenarios. The first investigates how the upstream team—the pricing department—can correct for downstream forecast bias introduced by the RM department. In this scenario, we assume that the pricing department knows the Model E demand estimates and adjusts their fare menus (using the same process as in Section 9.1). These menu adjustments are made knowing that the RM department will use the biased Model B forecasts along with the pricing heuristic.

We also investigate the second-best outcome for the firm in which the RM department adjusts their forecast, holding upstream fare decisions by the pricing department fixed. Because the RM department typically scales up or down the forecast for an entire route with a scaling parameter, we implement this counterfactual in a similar way. We have the RM department solve for inventory allocations using ESMR-b and the Model B forecast scaled up uniformly by a scaling parameter χ . We consider $\chi \in \{1.0, 1.25, \dots, 13.75\}$. For each χ , the RM department solves for inventory allocations using the manipulated forecast and then simulates flights according to the non-manipulated, Model B forecast. The optimal χ is the scaling factor that maximizes expected revenues. For example, if the menus set by the pricing department tend to feature fares that are too low for Model B demand, the RM department can scale up/boost the forecast to inflate opportunity costs. This will raise the distribution of fares offered. We use the Model B forecast in this counterfactual because it uses information currently within the firm, and thus, is perhaps more realistic in terms of being implementable.

Results for these two counterfactuals appear in Table 4, which are also normalized to

Table 4: Counterfactual Estimates of Second-Best Outcomes for the Firm

Counterfactual	CS_L	CS_B	Q	Rev	W
5) Pricing Department uses Model E forecast	122.5	64.0	92.7	117.6	86.4
6) RM Department manipulates Model B forecast	77.7	94.0	89.7	106.2	97.2

Note: In Row 5 we consider a counterfactual in which the pricing department adjusts their fare menus according to Model E forecasts, and the RM department uses the Model B forecast when determining inventory allocations. In Row 6 we consider the counterfactual in which the RM department adjusts the Model B forecast using a single scaling parameter (up/down), and the pricing department keeps the fare menus constant.

observed practices (Row 1 of Table 3). We find that the second-best outcomes for the firm are not close to current outcomes. For example, if the pricing department best responded to the forecasting bias introduced by the RM department, we estimate that revenues would close 95% of the gap between Row 1 and Row 4 of Table 3. The fact that this second-best outcome is so close to the first-best outcome highlights the importance of the pricing menu decisions of the upstream team (the pricing department). We find that properly accounting for the average change in willingness to pay over time in the pricing menu decisions is significantly more important than having a pricing algorithm that also reoptimizes in response to demand shock realizations. As a result, market segmentation is the driving force for price adjustments rather than responding to scarcity.

In contrast, we find that if the RM department best responded to the misalignment between the pricing department’s fare menus and Model B demand (see Row 6 of Table 4), the forecast should be significantly *more* biased in order to raise the distribution of fares offered. The average optimal χ is 7.5, with a standard deviation of 5.4. The optimal adjustment to the current demand model goes against observed practices in that analysts tend to deflate the forecast (see Figure 4). The second-best outcome is to inflate the forecast more. Averaging across routes, we find that significantly more biased forecasts would close 33% of the gap between Row 1 and Row 4. This number is substantially lower than the other second-best outcome (95%) because inflating opportunity costs across the entire booking horizon is less effective at targeting specific consumer groups with higher prices. Note that leisure consumer surplus declines in this scenario, relative to Row 5, and business consumer surplus is relatively high compared to Row 5.

9.4 Discussion

In all counterfactuals except one, the firm would be better off if teams adjusted their pricing inputs. This raises a natural question: Why have economic forces not led to different outcomes in multiple years of data? Although we cannot establish a causal link to the pricing biases we explore, the data, code, and supporting documentation suggest some likely mechanisms.

One potential mechanism is that firms are concerned about non-revenue metrics which would tend to lead to low prices. Our estimates of the first-best (and second-best) optimal outcome would lead to a reduction in tickets sold. This conflicts with metrics airlines commonly emphasize in public reports. We collect and process all airline earnings calls between 2013 and 2019. In over half the calls, *all* major airlines emphasize load factor, or the percentage of seats occupied. Of course, maximizing load factor is not the same as maximizing revenues.

Load factors are also emphasized internally along with metrics on forecasting bias. Our counterfactual that investigates the second-best outcome if the RM department manipulated its forecast to account for the coordination problem is feasible (e.g., it does not require a new forecasting model), but firm norms may prevent its implementation. Not only does meeting load factor targets put downward pressure on prices—in the opposite direction of the second-best results—we found that the optimal scaling factor χ averages 7.5. It is unlikely that managers would accept forecasts that are biased upwards of 750%, when they are trying to reduce error.

In addition to performance metrics, another likely mechanism for the persistent pricing biases is the lack of information transmission across teams. Our descriptive evidence suggests miscoordination is never addressed in a two year data sample even though both an issue and fix are known to one team within the firm. It could be that the costs of improving coordination and aligning input decisions are especially high. For example, we observe that only 1.6% of prices are adjusted by the pricing department. Perhaps realigning team information sets and decision processes makes observed outcomes “optimal” in the presence of

substantial organizational costs.

We also note that the clickstream data suggest limited experimentation at the firm. Without experimental variation, managers may be unable to measure the substantial differences in willingness to pay across consumer types as our demand estimates suggest. We observe an identifier for consumers subject to experiments as well as a code to denote unique experiments. The number of experiments run, and the number of consumers subject to experiments, is very low.

We argue that some mechanisms are less plausible. Although it may be that the firm has long-run demand considerations in mind when determining prices (supporting lower prices), we note that the RM forecasts and the algorithm objective focus on short-run demand and revenues, respectively. Moreover, sometimes the forecasts of neighboring departure dates are adjusted in response to demand shocks—again highlighting short-run considerations (see Online Appendix E for an example). It is unclear why long-run demand estimates would yield persistently biased upward forecasts that in general understate (overstate) the number of low (high) priced tickets sold. Related to long-run demand, another reason to offer lower fares is to reward customers for loyalty. However, the data establish that more expensive tickets tend to be purchased by more loyal customers. If a large fraction of “business” customers do not pay for their tickets out of their own wallet, current mileage programs allow customers to reach status faster by purchasing more expensive tickets.

Another mechanism we argue is less likely to explain our findings is that the firm recognizes additional revenues are possible through increased market segmentation, but the firm is strategically choosing to offer lower prices as a response to the threat of a entry (Goolsbee and Syverson, 2008; Sweeting, Roberts, and Gedge, 2020). We view this as implausible because the biases we document in Section 4 affect all markets and are even more pronounced in markets with competition. It is unclear why a firm would also price on the inelastic side of demand when already facing direct competitors.

10 Conclusion

In this paper, we study pricing at a large U.S. airline where distinct organizational teams manage specific pricing inputs. We provide evidence of pricing biases and miscoordination. To quantify the welfare effects of these pricing biases, we estimate a structural demand model and conduct counterfactual experiments using a pricing heuristic that closely approximates what the firm uses in practice. We show that addressing pricing biases within teams individually does not substantially change market outcomes. Although biases are not reinforcing, they are not perfectly offsetting. Observed outcomes differ from optimal second-best outcomes for the firm. Finally, we show that if teams correct and coordinate on algorithm inputs, we find a significant reallocation of capacity across consumer types. Leisure consumers would benefit from lower fares, but business travelers would face significantly higher fares. Revenues would increase, but so would dead-weight loss.

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 (2021): “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.
- Belobaba, P. P. (1989): “OR practice-application of a probabilistic decision model to airline seat inventory control,” *Operations Research*, 37(2), 183–197.
- (1992): “Optimal vs. heuristic methods for nested seat allocation,” in *Presentation at ORSA/TIMS Joint National Meeting*.
- 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.
- Blake, T., C. Nosko, and S. Tadelis (2015): “Consumer heterogeneity and paid search effectiveness: A large-scale field experiment,” *Econometrica*, 83(1), 155–174.
- Bloom, N., R. Sadun, and J. Van Reenen (2012): “The organization of firms across countries,” *The quarterly journal of economics*, 127(4), 1663–1705.
- Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002): “Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence,” *The quarterly journal of economics*, 117(1), 339–376.
- Brumelle, S. L., J. I. McGill, T. H. Oum, K. Sawaki, and M. Tretheway (1990): “Allocation of airline seats between stochastically dependent demands,” *Transportation Science*, 24(3), 183–192.
- Brynjolfsson, E., and P. Milgrom (2012): “1. Complementarity in Organizations,” in *The handbook of organizational economics*, pp. 11–55. Princeton University Press.
- DellaVigna, S., and M. Gentzkow (2019): “Uniform pricing in us retail chains,” *The Quarterly Journal of Economics*, 134(4), 2011–2084.
- Dessein, W., and T. Santos (2006): “Adaptive organizations,” *Journal of political Economy*, 114(5), 956–995.
- Dubé, J.-P., and S. Misra (2021): “Personalized Pricing and Consumer Welfare,” Working Paper.
- Ellison, S. F., C. Snyder, and H. Zhang (2018): “Costs of managerial attention and activity as a source of sticky prices: Structural estimates from an online market,” Discussion paper, National Bureau of Economic Research, Working Paper.
- Gibbons, R., and R. Henderson (2012): *17. What Do Managers Do?* Princeton University Press.
- 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.
- Hitsch, G. J., A. Hortaçsu, and X. Lin (2021): “Prices and Promotions in US Retail Markets: Evidence from Big Data,” *Quantitative Marketing & Economics*.

- Hortaçsu, A., O. Natan, H. Parsley, T. Schwieg, and K. Williams (2021): “Incorporating Search and Sales Information in Demand Estimation,” Working Paper.
- Huang, Y. (2021): “Seller-Pricing Frictions and Platform Remedies: The Case of Airbnb,” 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.
- Levitt, S. D. (2016): “Bagels and donuts for sale: A case study in profit maximization,” *Research in Economics*, 70(4), 518–535.
- Li, J., N. Granados, and S. Netessine (2014): “Are consumers strategic? Structural estimation from the air-travel industry,” *Management Science*, 60(9), 2114–2137.
- Littlewood, K. (1972): “Forecasting and control of passenger bookings,” *Airline Group International Federation of Operational Research Societies Proceedings, 1972*, 12, 95–117.
- Ma, Y., T. Ropele, D. Sraer, and D. Thesmar (2020): “A quantitative analysis of distortions in managerial forecasts,” Working Paper.
- Massey, C., and R. H. Thaler (2013): “The loser’s curse: Decision making and market efficiency in the National Football League draft,” *Management Science*, 59(7), 1479–1495.
- McAfee, P., and V. Te Velde (2006): “Dynamic Pricing in the Airline Industry,” *Handbook on Economics and Information Systems*.
- McGill, J. I., and G. J. Van Ryzin (1999): “Revenue management: Research overview and prospects,” *Transportation science*, 33(2), 233–256.
- Milgrom, P., and J. Roberts (1990): “Rationalizability, learning, and equilibrium in games with strategic complementarities,” *Econometrica: Journal of the Econometric Society*, pp. 1255–1277.
- (1995): “Complementarities and fit strategy, structure, and organizational change in manufacturing,” *Journal of accounting and economics*, 19(2-3), 179–208.
- 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., 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.
- Williams, K. R. (2021): “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, Schweg, and Williams

Table of Contents

A All US Airlines have the same Organizational Structure	2
B Details on the Pricing Heuristic, EMSR-b	2
B.1 Littlewood’s Rule	3
B.2 EMSR-b Algorithm	4
C Route Selection	7
C.1 Estimation Sample Comparison	10
D Additional Details on Demand Estimation	11
D.1 Demand Estimation Procedure	11
D.2 The Impact of the Scaling Factor on Demand Estimates	17
E Additional Counterfactuals	17

A All US Airlines have the same Organizational Structure

Our description of airline pricing is not unique to the airline we study—all airlines have the same organizational structure and use similar pricing techniques. We show this by collecting job postings information for all the major carriers in the U.S.²⁰ We confirm that Alaska, American, Delta, JetBlue, Southwest, and United have a network planning, pricing, and revenue management department. As an example, JetBlue Airlines job postings show that the firm has three teams related to pricing: Future Schedules, Revenue Management-Pricing, and Revenue Management-Inventory. Job details delineate team responsibilities. The Revenue Management department at JetBlue has two separate teams, Pricing and Inventory. The Pricing team has ownership over fares by “monitoring industry pricing changes filed through a clearinghouse throughout the day, and determining and executing JetBlues response.”²¹ The Inventory team uses “inventory controls to determine the optimal fare to sell at any given moment in time to maximize each flights revenue.”²² American Airlines managers describe how inventory controls are implemented in Smith, Leimkuhler, and Darrow (1992)—they outline EMSR-b. Because all carriers have the same organizational structure and use similar algorithms, we believe our analysis characterizes the entire industry, rather than the perspective from a single firm.

B Details on the Pricing Heuristic, EMSR-b

We approximate the solution to a dynamic pricing (DP) problem using a well-known heuristic in operations research, Expected Marginal Seat Revenue-b or EMSR-b (Belobaba, 1987). The heuristic 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 before the deadline into a single future period. It also simplifies the demand system to be for only a single product, so competitive effects cannot be considered. We

²⁰Screenshots of the job postings are available on request.

²¹See <https://careers.jetblue.com/job/Long-Island-City-Analyst-Revenue-Management-NY-11101/737962800/>. July 1, 2021.

²²See <https://careers.jetblue.com/job/Long-Island-City-Analyst-Revenue-Management-NY-11101/737962800/>. July 1, 2021.

describe this process below and show how to incorporate Poisson demand in EMSR-b. It is important to note that EMSR-b provides an allocation over a given finite set of prices, instead of providing the optimal price itself given any state of the world. EMSR-b associates each price with a fare-class then chooses a maximal number of sales that can be made to each fare-class. This means that consumers may face different prices within a single pricing period when one class is closed and a higher priced class opens.

B.1 Littlewood’s Rule

EMSR-b is a generalization of Littlewood’s rule, which is a simple case where a firm prices two time periods 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.

B.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.

B.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 $\exp(\lambda_t^t + \lambda_d^d) s_j(p)$. 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 distribution. 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

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

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

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

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

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

 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 .

C 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 10: Nonstop, One-stop and Connecting Traffic

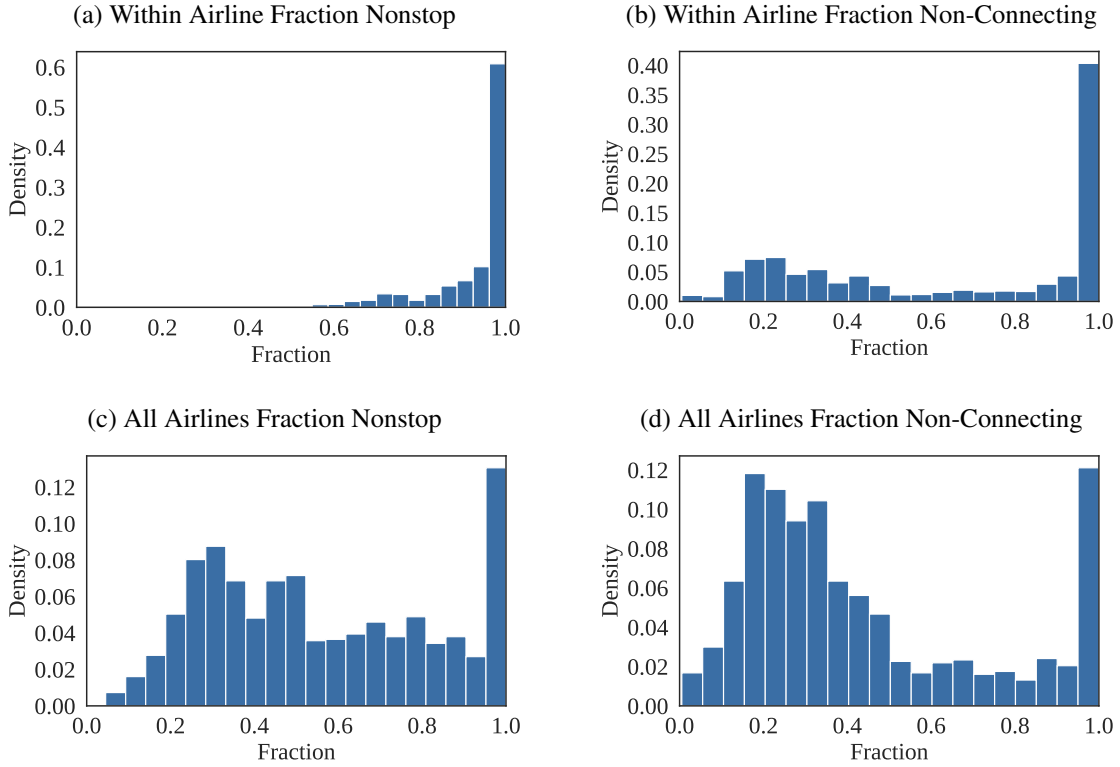


Note: We use the term nonstop to denote the solid black line, or passengers solely traveling between (Origin, Destination). Unless otherwise noted, we will use directional traffic, labeled $O \rightarrow D$. Non-directional traffic is specified as $O \leftrightarrow D$. The blue, dashed lines represent 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 10 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 11 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 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

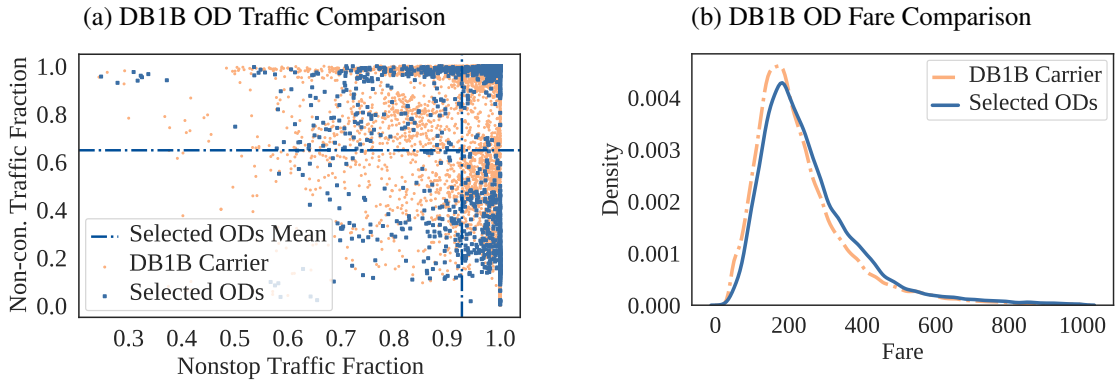
Figure 11: Route Selection Using Bureau of Transportation Statistics Data



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.

significant variation across markets, with the average being close to 50%.

Figure 12: 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.

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 markets.

In Figure 12-(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 12-(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.

C.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 13 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.

Table 5: Estimation Routes 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 13: 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.

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. 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 (2021).

- 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:

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

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

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

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

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

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

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

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

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

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

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

E Additional Counterfactuals

EMSR-b is a heuristic and is itself biased (Wollmer, 1992) because it does not consider substitute products. To account for substitutes, we also consider counterfactuals where prices are determined by solving a dynamic pricing problem. We follow the dynamic pricing (DP) problem in Williams (2021), where a firm selects a price for each flight from a discrete set of prices that maximizes its current and expected future profits. We assume that the firm solves

$$V_t(C_t, p_t) = \max_{p \in P_t} (R_t^e(C_t, p_t) + EV_{t+1}(C_{t+1}, p_{t+1} | C_t, p_t)),$$

where C_t is the vector of remaining capacity for each flight offered in that time period, p_t is the vector of prices the firm selects, and $R_t^e(C_t, p_t)$ is the firm's expected flow revenue. These value functions are specific to a route and departure date.

We consider two versions of the DP. We first simulate pricing for each flight indepen-

dently, assuming other flights will be priced at the lowest priced fare. This is analogous to how we proceed with EMSR-b. We then consider a multi-product DP and limit ourselves to $|J| = 2$ due to the dimensionality of the more complicated environments. Our DP results are thus based on a selected set of routes (and departure dates). We use the coordinated fare menus derived under Model B and Model E as inputs. These fares may not be optimal in the multi-product setting.

Table 6: Counterfactual Estimates under an Alternative Pricing System

	Counterfactual	CS_L	CS_B	Q	Rev	W
1)	Pricing heuristic EMSR	100.0	100.0	100.0	100.0	100.0
2)	Single- J Dynamic Pricing	95.0	99.0	98.5	100.6	99.6
3)	Multi- J Dynamic Pricing	93.4	97.2	101.5	100.9	98.8

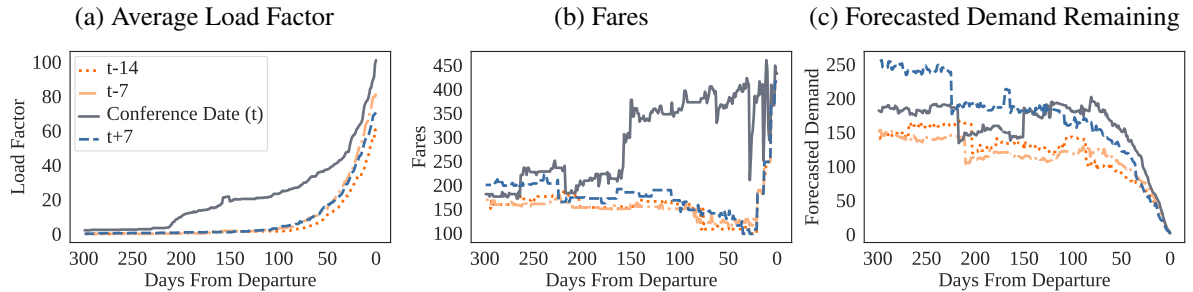
Note: In counterfactual (1) prices are set using EMSR-b with Model E and the coordinated price menus. Counterfactual (2), endogenously sets prices for each flight independently using the DP. Finally, counterfactual (3) we jointly set prices of all products in the same market using the DP.

In Table 6, we compare Row 4 of Table 9 to models of dynamic pricing that also use these inputs. We report two rows after our EMSR-b results corresponding to the situation where the firm optimizes flight prices individually and one in which the firm prices flights jointly. Outcomes are normalized to EMSR-b. We estimate marginally lower consumer surplus and slightly higher revenues under dynamic pricing. These results are not due to the discrete nature of prices—implementing a continuous-price version for single flight markets yields quantitatively similar results. Overall, we find that accounting for cross-price elasticities results in a marginally higher revenues and lower consumer surplus compared to EMSR-b. These effects are relatively small compared to correcting forecasting bias and coordinating fares to the correct demand curves using the pricing heuristic.

One concern with Model E is that it assumes the firm knows preferences and arrivals rates in advance. The forecasting data allow us to explore learning about demand. We find that with the current system, reactions to “surprises” occur too little and too late. In particular, demand forecasts respond to demand surprises with delay, leading to missed opportunities both for the flight in question, but also for future flights which are mistakenly

thought to be over-(or under-) demanded.

Figure 14: Reacting to Surprises: A Conference Example



Note: All plots contain data series for flights departing the date of a large conference (which moves both location and dates each year) and the corresponding flights in the surrounding weeks of the conference. (a) Shows the average load factor across flights (b) Contains the average lowest available fares across flights (c) Has the average total expected seats to sell at a given point in time unconstrained by the remaining capacity.

To demonstrate how the firm updates its beliefs about demand, in Figure 14 we show average load factors, fares, and forecasted demand remaining for a particular route-departure date. This departure date is special because it involves a conference which alternates both date and location each year. In addition to flights on the conference date, we include information for flights on this route one week before and after the conference date for comparison. As shown in panel (a), as soon as the location and date of the conference is announced, around 200 days from departure, there is a sudden jump in load factor. The firm’s revenue management software responds with delay (over a month) to the sudden jump in bookings. Prices eventually increase dramatically as seen in panel (b). Panel (c) shows that the forecasting algorithm, having observed the conference shock, then inflates the forecast of remaining demand for the following week—to higher levels than the conference date. That is, the algorithm incorrectly believes the next week will now also involve a conference. However, in panel (a) we see that the flights a week later contain no surprises—bookings follow a similar pattern as other dates. Consequently, fares are too high for the non-conference flights and too low for a conference flights.

Instead of directly incorporating how the RM department updates its forecasts toward the departure date, we consider an alternative, simpler model where the firm has “persistently average beliefs,” or where the firm prices according to $\text{avg}(\lambda)$ and $\text{avg}(\beta)$. The firm

knows γ and α in this counterfactual. This behavioral model need not provide a good approximation to the actual learning environment but it allows us to investigate if our findings are driven by the firm knowing all preferences in advance. We find that even with average beliefs, results are similar to scenario (4) in Table 3, with a 15% increase in revenues. Essentially, knowing the average change in willingness to pay over time is very important, much more so than variation in preferences and arrival rates across departure dates. Correcting beliefs about the average shape and evolution of demand—when paired with coordinated fare options—is the key driver of the welfare effects in our counterfactuals.