Price Discrimination in International Airline Markets.

Gaurab Aryal[y]   Charles Murry[z]   Jonathan W. Williams[x]

November 26, 2018

Abstract

We develop a model of inter-temporal and intra-temporal price discrimination by airlines to study the ability of di erent discriminatory mechanisms to remove sources of ine ciency and the associated distributional implications. To estimate the model's multi-dimensional distribution of preference heterogeneity, we use unique data from international airline markets with ight-level variation in prices across time and cabins, and information on passengers' reason for travel. We nd that current pricing practices

1 Introduction

Discriminatory pricing enables firms with market power to increase their profits, but it can lead to inefficient outcomes and have adverse distributional effects in some markets. This makes the welfare implications of price discrimination theoretically ambiguous, and a subject of substantial interest for empirical studies. The implications are even less clear when demand is revealed gradually and firms can use inter-temporal and intra-temporal price discrimination to screen consumers based on preference for quality and time of purchase.[1] In such settings, different types of allocative inefficiencies can arise due to asymmetric information about preferences and stochastic demand.

In this paper, we use unique data on international air-travel to estimate a model of demand and supply for air travel. We use the estimates to measure allocative inefficiency and identify the portion attributable to each source: asymmetric information and stochastic demand. To achieve this, like Bergemann, Brooks, and Morris (2015), we progressively increase the information the seller has about preferences and measure the increase in welfare under various discriminatory pricing strategies (i.e., second, third and first) and other screening mechanisms (e.g., VCG auction). The inefficiency remaining after removing all forms of asymmetric information is then attributable to stochastic demand. The particular counterfactual strategies that we consider are motivated by recent airline practices (e.g., solicit passengers' reason-to-travel and auctions for upgrades) intended to raise profits by reducing allocative inefficiencies, and our estimates provide insight into the ability of these discriminatory mechanisms to remove each source of inefficiency, as well as any distributional implications across passengers.

At the core of our paper is data from the Department of Commerce's Survey of International Air Travelers (SIAT). The SIAT includes information on passengers traveling on routes that connect U.S. and international markets from more than 70 participating U.S. and international airlines. The data contain detailed information on the purpose of the trip (business or leisure), ticket class (e.g., economy or first class), date of purchase in advance of the flight, and fare paid. Crucially for our purpose, and in contrast to a sample like the Department of Transportation's DB1B survey, the flight-based sampling and richness of passenger information in the SIAT provide a rare opportunity to gain insight into the factors that contribute to the variability in fares and changing composition of passengers (e.g., reason for travel and willingness to pay) as the flight date approaches.

Given the novelty of our data, we first present a rich set of descriptive statistics that reveal

---

[1]There is an extensive theoretical literature studying the implications of dynamic pricing. See, for example, Prescott (1975); Stokey (1979); Dana (1998, 1999); Eden (1990); Courty and Li (2000); Deneckere and Peck (2012); Board and Skrzypacz (2016); Garrett (2016); and Ely, Garrett, and Hinnosaar (2017).

the importance of passenger heterogeneity and stochastic demand in determining pricing dynamics within and across flights. We observe that approximately 15% of passengers travel for business, but this proportion differs substantially across flights, depending on the origin and destination. Compared to leisure passengers, business passengers tend to purchase closer to the departure date, pay substantially more for the same cabin, and are more likely to buy a first-class seat. In markets with a greater proportion of business passengers, fares for both cabins are greater on average in each period, and fares increase substantially as the flight date approaches. Despite the upward pressure from late-arriving business passengers that leads to monotonically-increasing average prices, we observe many flights with non-monotonic price paths.

We then use the model estimates for several counterfactual exercises that quantify the magnitude of inefficiencies due to demand uncertainty and asymmetric information. We begin by considering the first-best allocation when the airline has perfect foresight and observes preferences for all arriving passengers. In this case, the airline allocates seats to passengers with the highest valuations, and the division of welfare then depends upon the prices. Like Bergemann, Brooks, and Morris (2015), this defines an efficient frontier on the set of possible welfare outcomes. We find that current pricing practices, i.e., second-degree discrimination, result in 81% of the first-best welfare. To decompose the source of

(2012); Board and Skrzypacz (2016); Stamatopoulos and Tzamos (2016); and Garrett (2016). There is also an extensive literature in operations research on revenue management. See van Ryzin and Talluri (2005) for a review of this literature. In terms of the empirical application of dynamic pricing for perishable goods, our paper complements the ndings in Sweeting (2010), Williams (2017), Sanders (2017) and Cho et al. (2018). Our analysis is also related to Nair (2007) and Hendel and Nevo (2013) that study inter-temporal price discrimination with storable goods.

include only those flights for which we observe at least 10 nonstop tickets.

Lastly, we define a market as a monopoly market if the carrier operating the flight provides at least 50% of all capacity offered on nonstop flights between the origin and destination. We maintain that such carriers have market power and are likely to price discriminate. To minimize the chance of including non-monopoly markets, we exclude markets that have multiple U.S. carriers. After these sample selection criteria are applied, we are left with 45,473 passenger records across 2,552 flights in 398 markets on 85 unique carriers.

Using this sample, we classify each respondent's reason for travel into one of two categories, business or leisure. Business includes business, conference, and government/military, while leisure includes visiting family, vacation, religious purposes, study/teaching, health, and other. Further, like Borenstein (1989) and others, we make one-way and round-trip fares comparable by dividing round-trip fares by two. In the remainder of this section, we provide descriptive analysis of our sample that motivates the modeling assumptions we make in Section 3.

Table 1: Summary Statistics from SIAT, Ticket Characteristics

| | Proportion | Fare | |
Ticket Class	of Sample	Mean	SD
First	9.00	818.48	839.77
Economy	91.00	425.16	357.98
Advance Purchase			
0-7 Days	9.68	575.79	606.80
8-21 Days	14.56	534.27	536.67
22-35 Days	16.90	471.21	432.95
36-85 Days	21.60	439.99	402.79
85 Days	37.27	408.92	348.60
Travel Purpose			
Leisure	86.55	426.70	371.75
Business	13.45	678.28	699.12

Note: Data from the Survey of International Air Travelers. Sample described in the text.

2.1 Descriptive Statistics

In Table 1 we display some key statistics for relevant ticket characteristics in our sample. From the top panel, in our sample, 91.00% of the passengers purchased economy-class tickets, and their average fare was $425.16. This is in contrast to the 9.00% of the sample who purchased business-class or first-class tickets (henceforth, we use first-class to refer to either

8

In Figure 1(a) we plot the average price for economy fares as a function of the number of days in advance of the ight that a ticket was purchased. Both business and leisure travelers pay more if they buy the ticket closer to the ight date, but the increase is more substantial for the business travelers. The solid line in Figure 1(a) re ects the average price across both reasons for travel. At earlier dates, the total average price is closer to the average price paid by leisure travelers, while it gets closer to the average price paid by the business travelers as the date of the ight nears. In Figure 1(b), we display the proportion of business to leisure travelers across all ights, by the advance purchase categories. In the last week before the ight, the share of passengers traveling for leisure is approximately 60%, while that share is nearly 100% two months earlier. Taken together, Figures 1(a) and 1(b) show that business travelers purchase nearer the ight date, and those markets with a greater proportion of business travelers experience a greater increase in fares as the ight date approaches.

Figure 2: Histogram of Percent of Business Passengers by Market

Note: Histogram of business-travel index (BTI). The business-traveler index is the market-speci c ratio of self-reported business travelers to leisure travelers across the entire sample.

Observing the purpose of travel plays an important role in our empirical analysis, re ecting substantial di erences in the behavior and preferences of business and leisure passengers. This passenger heterogeneity across markets drives variation in pricing, and this covariation permits us to estimate a model with richer consumer heterogeneity than the existing literature like Berry, Carnall, and Spiller (2006); Berry and Jia (2010), and Ciliberto and Williams (2014). Further, a clean taxonomy of passenger types allows a straightforward exploration of the role of asymmetric information in determining ine ciencies and the distribution of surplus that arises from discriminatory pricing of di erent forms.

To further explore the in uence that this source of observable passenger heterogeneity has on fares, we now present some statistics on across-market variation in the dynamics of fares. Speci cally, we rst calculate the proportion of business travelers in each market, i.e., across all ights with the same origin and destination. Like Borenstein (2010), we call this market-speci c ratio the business-traveler index (BTI). In Figure 2, we present the histogram of the BTI across markets in our data. If airlines know of this across-market heterogeneity and use it as a basis to discriminate both intra-temporally (across cabins) and inter-temporally (across time before a ight departs), di erent within- ight temporal patterns in fares should arise for di erent values of the BTI.

Figure 3: Proportion of Business Travelers by Ticket Class

Figure 4: Across-Market Variation in Fares

0	20	40	60	80	100	120

(a) Economy   (b) First-Class

Note:

Figure 5: Flight-Level Dispersion in Fares

in the temporal patterns in fares across ights is attributable to both the across-market heterogeneity in the mix of passengers, and how airlines respond to the gradual realization of stochastic demand for a ight.

Airlines' fares, and their responsiveness to realized demand, depend on the number of unsold seats. In Figure 6(a) we display the joint density of initial capacity of rst and economy class in our sample, adjusting for the fraction of non-stop passengers we observe in the data. The median capacity is 116 economy seats and 15 rst-class seats, and the mode is 138 economy and 16 rst-class seats. The three most common equipment types in our sample are a Boeing 777, 747, and 737 (36% of ights in our sample). The 777 and 747 are wide-body jets. The 777 has a typical seating of around 350 seats (not adjusting for non-stop versus connecting passengers) and the 737 has a typical seating of around 160 total seats. The most common Airbus equipment is the A330, which is about 4% of the ights in our sample. Across all ights, capacity is 88% economy class on average. We merge the SIAT data with the Department of Transportation's T-100 segment data to get a measure of the load factor for our SIAT ights. From the T100, we know the average load factor across a month for a particular route own by a particular type of equipment. In Figure 6(b) we display the density of load factor across ights in our sample. The median load factor is 82%, but there is substantial heterogeneity across ights.

Overall, our descriptive analysis reveals a number of salient features that we capture in our model. We or T3i327passe3(detype)eieasure)-3y
4(thenrices)-5 [(in)3633(the)-405imusing ofpur8(c)28(hure,)-358ilirlg(eness)-405oofpplaF and pfe(ar

Figure 6: Initial Capacity and Load Factor

(a) Density of Initial Capacities

(b) Histogram of Load Factor
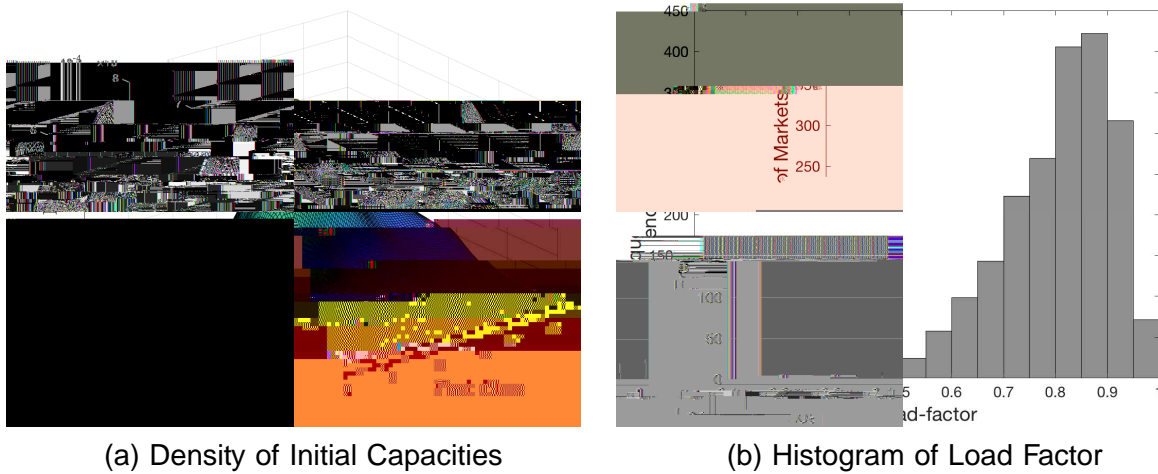
Note: In part (a) this figure presents the Parzen-Rosenblatt Kernel density estimate of the joint-density of initial capacities available for nonstop travel. In part (b) this figure presents the histogram of the passenger load factor across our sample.

Further, we find substantial heterogeneity in the mix of passengers (i.e., business/leisure) across markets, which airlines are aware of and responsive to, creating variation in both the level and temporal patterns of fares across markets. Finally, across flights we observe considerable heterogeneity in whether fares decrease or increase as the flight date approaches. Together, these features motivate our model of the non-stationary and stochastic demand and dynamic pricing by airlines that we present in Section 3, and the flexible estimation approach in Section 4.

3 Model

In this section we present a model of dynamic pricing by a profit-maximizing multi-product monopoly airline that sells a fixed number of economy ($0 \leq K^e < 1$) and first-class ($0 \leq K^f < 1$) seats. We assume passengers with heterogeneous and privately known preferences arrive over time before the date of departure $t (\in \{0, \ldots, T\})$ for a nonstop flight. Every period the airline has to choose the ticket prices and the maximum number of unsold seats to sell at those prices before demand (for that period) is realized.

Our data indicate important sources of heterogeneity in preferences that differ by reason for travel: willingness to pay, valuation of quality, and timing of purchase. Further, variability and non-monotonicity in fares suggest a role for uncertain demand. The demand side of our model seeks to flexibly capture this multi-dimensional heterogeneity and uncertainty

14

that serves as an input into the airline's dynamic-pricing problem. The supply side of our model seeks to capture the inter-temporal and intra-temporal tradeo s faced by an airline in choosing its optimal policy.

3.1 Demand

Let N_t denote the number of individuals thatarrive in period $t \in \{1, \ldots, T\}$

Figure 7: Realization of Demand

Airline Chooses: $(p_t^e; q_t^e); (p_t^f; q_t^f)$

Individuals Arrive : $N_t \sim P(\lambda_t)$

Business: θ_t Leisure: $1 - \theta_t$

Business Arrivals: N_t^b Leisure Arrivals: N_t^l

$(v_i; \epsilon_i) \sim F_v^b \times F_\epsilon; i = 1; \ldots; N_t^b$ $(v_i; \epsilon_i) \sim F_v^l \times F_\epsilon; i = 1; \ldots; N_t^l$

Buy: $q_t^e; q_t^f$; Not Buy: q_t^o

$\bar{q}_t^e \geq q_t^e$
$\bar{q}_t^f \geq q_t^f$

$\bar{q}_t^e < q_t^e$
$\bar{q}_t^f \geq q_t^f$

$\bar{q}_t^e \geq q_t^e$
$\bar{q}_t^f < q_t^f$

Figure 8: Illustration of Random Rationing Rule.

Capacity: $K^f = 1$; $K^e = 2$

$p^f = 2000$; $q^f = 1$

seats in each cabin. One of the defining characteristics of this market is that the airline must commit to policies this period before current and future demand is realized. The airline does not observe a passenger's reason to fly or valuations (ν, ς); however, the airline knows the underlying stochastic process that governs demand, and uses the information to price discriminate, both within period and across periods.[8]

Let c^e and c^f denote the constant cost of servicing a passenger in the respective cabins. These marginal costs, or so-called "peanut costs," capture variable costs like food and beverage service that do not vary with the timing of the purchase but may vary with the different levels of service in the two cabins. Let $\Theta := (\{F_v^b; F_v^l; F_\varsigma; c^f; c^e\}; \{\lambda_t; \tau_t\}_{t=1}^T)$ denote the vector of demand and cost primitives.

The airline maximizes the sum of discounted expected profits by choosing price and seat-release policies for each cabin $\sigma_t = \left(p_t^e; p_t^f; \bar{q}_t^e; \bar{q}_t^f\right)$, in each period $t = 1; \ldots; T$ given ω_t. The optimal policy is a vector $\{\sigma_t : t = 1; \ldots; T\}$ that maximizes expected profit

$$\sum_{t=1}^T E_t \{\pi(\sigma_t; \omega_t; \Theta_t)\};\tag{3}$$

where

because the firm no longer faces any inter-temporal tradeoffs. The dynamic programming characterization of the airline's problem is useful, both for identifying the tradeoffs faced by the airline, and identifying useful sources of variation in our data.[9]

The optimal pricing strategy includes both inter-temporal and intra-temporal price discrimination. First, given the limited capacity, the airline must weigh allocating a seat to a passenger today versus to a passenger tomorrow, who may have higher mean willingness to pay because the fraction of for-business passengers increases as it gets closer to the flight date. This decision is complicated by the fact that both the volume (λ_t) and composition (γ_t) of demand changes as the date of departure nears. Thus, the perishable nature of the good does not necessarily generate declining price paths like Sweeting (2010). Second, simultaneously, every period the airline must allocate passengers across the two cabins by choosing δ_t such that the price and supply restriction-induced selection into cabins is optimal.

To illustrate the problem further, consider the trade-off faced by an airline from increasing the price for economy seats today: (i) decreases the expected number of economy seat purchases but increases the revenue associated with each purchase; (ii) increases the expected number of first-class seat purchases but no change to revenue associated with each purchase; (iii) increases the expected number of economy seats and decreases the expected number of first-class seats available to sell in future periods. Effects (i) and (ii) capture the multi-product tradeoff faced by the firm, while (iii) captures the inter-temporal tradeoff. More generally, differentiating Equation 4 with respect to the two prices gives two first-order conditions that characterize optimal prices given a particular seat-release policy:

$$\begin{pmatrix} E_t(q^e;\, \cdot_t) \\ E_t(q^f;\, \cdot_t) \end{pmatrix} + \begin{pmatrix} \frac{\partial E_t(q^e;\, \cdot_t)}{\partial p^e} & \frac{\partial E_t(q^f;\, \cdot_t)}{\partial} \end{pmatrix}$$

future demand (i.e., variation in volume of passengers and business/leisure mix as flight date nears), and the number of seats remaining in each cabin (i.e. K_t^f and K_t^e). The stochastic nature of demand drives variation in the shadow costs, which can lead to equilibrium price paths that are non-monotonic, and that increase or decrease on average. This flexibility is crucial given the variation observed in our data (see Figure 5).[9]

The airline can use its seat-release policy to dampen both intra-temporal and inter-temporal tradeoffs associated with altering prices. For example, the airline can force everyone to buy economy by not releasing first-class seats in a period, and then appropriately adjust prices to capture rents from consumers.[11] Consider the problem of choosing the number of seats to release at each period $\bar{q}_t := (\bar{q}_t^e; \bar{q}_t^f) \in \Omega_t$. For a choice of $\bar{q}_t$ in period t, let $p_t(\bar{q}_t) := \{ p_t^e(\bar{q}_t); p_t^f(\bar{q}_t) \}$ denote the optimal pricing functions as a function of the number of seats released. Then, the value function can be expressed recursively as

$$
V_t(\omega_t; \theta) = \max_{\bar{q}_t \in \Omega_t} \left\{ \pi_t((p_t(\bar{q}_t); \bar{q}_t); \omega_t; \theta) + \sum_{\omega_{t+1} \in \Omega} V_{t+1}(\omega_{t+1}; \theta) \, Q_t(\omega_{t+1} \mid (p_t(\bar{q}_t); \bar{q}_t); \omega_t; \theta) \right\}.
$$

The profit function is bounded, so this recursive formula is well defined, and under some regularity conditions we can show that it is has a unique optimal policy. We present these regularity conditions and the proof of uniqueness in Appendix A.1.

4 Estimation and Identification

In this section we discuss the parametrization of the model, estimation methodology, and identification strategy. The parametrization of the model is meant to balance the dimensionality of the parameters and the desired richness of the demand structure. The model,

4.1 Model Parametrization and Solution

To retain flexibility in our specification of market heterogeneity while limiting the computational burden of solving the model a large number of times, we parameterize $\Theta = (f F_b; F_l; F_\gamma; c^f; c^e g; f \lambda_t; \pi_t g_{t=1}^T)$ to capture the salient features of our data.

There are two demand primitives, λ_t and π_t, that vary as the flight date approaches, such that a fully non-parametric specification for both would result in a prohibitive $2T$ parameters. Motivated by our data, we choose $\bar{T} = 5$ to capture temporal trends in fares and passenger's reason for travel, where each period is defined as in Table 1. To permit flexibility in the relationship between time before departure and these parameters, we let $\lambda_t := \min \{\alpha(t - 1); 1\}$ and $\pi_t := \kappa + \rho(t - 1)$, where $\alpha; \kappa$, and ρ are scalar constants. This parametrization of the arrival process permits the volume (and λ) and composition (π) of demand to change as the flight date approaches.

There are three distributions $(F_b; F_l; F_\gamma)$ that determine preferences. We assume that for-business and for-leisure valuations are drawn from normal distributions F_b and F_l, respectively, that are left-truncated at zero. Given the disparity in average fares paid by business and leisure passengers, we assume $\mu^b \geq \mu^l$, which we model by letting $\mu^b = \mu^b \mu^l$ with $\mu^b \geq 1$. We assume that the quality premium, γ, equals one plus an Exponential random variable with mean η to ensure that the two cabins are vertically differentiated and that all passengers weakly prefer first class.

Finally, we fix "peanut" costs for first-class and economy, c^f and c^e, respectively, to equal industry estimates of marginal costs for servicing passengers. Specifically, we set $c^f = 40$ and $c^e = 14$ based on information from the International Civil Aviation Organization, Association of Asia Pacific Airlines, and Doganis (2002).[12] Our estimates, and the counterfactual implications of the estimates, are robust to alternative values for these 5.248(one)-247(plus)-23e]TJ -4

Turner, and Williams (2016), we posit that empirical moments can be expressed as a mixture of theoretical moments, with a mixing distribution known up to a nite-dimensional vector of

the S draws of θ taken from $h(\theta_j | \omega_1; \gamma; \sigma)$.

The dimensionality of the integral we approximate through simulation requires a large number of draws. After some experimentation to ensure simulation error is limited for a wide range of parameter values, we let $S = 10{,}000$. Thus, the most straightforward approach to optimization of Equation 8 would require solving the model $S = 10{,}000$ times for each value of $(\gamma; \sigma)$ until a minimum is found. Given the complexity of the model and the dimensionality of the parameter space to search over, such an option is prohibitive. For this reason, we appeal to the importance sampling methodology of Ackerberg (2009).

The integral in Equation 7 can be rewritten as

$$\int_{\underline{\theta}}^{\overline{\theta}} \ell(\omega_1; \theta) \frac{h(\theta_j | \omega_1; \gamma; \sigma)}{g(\theta)} g(\theta) \, d\theta;$$

where $g(\theta)$ is a known well-defined probability density with strictly positive support for $\theta \in [\underline{\theta}; \overline{\theta}]$ and zero elsewhere like ($h(\theta_j | \omega_1; \gamma; \sigma)$)). Recognizing this, one can use importance sampling to approximate this integral with

$$\frac{1}{S} \sum_{j=1}^{S} \ell(\omega_1; \theta_j) \frac{h(\theta_j | \omega_1; \gamma; \sigma)}{g(\theta_j)}$$

where the S draws of θ are taken from $g(\theta)$. Thus, the importance sampling serves to correct the θ_j sampling frequencies so that it is as though the sampling was don32673t9050 h

(ight) re-sampling procedure to account for the dependence structure in our data (Lahiri, 2003).

4.3   Identi cation

In this section, we introduce the moments that we use in (7) to estimate the market hetero-geneity,   h(j! $_1$;   ;   ) and present the identi cation argument that guides our choice. To that end, we present an argument that our moments vary uniquely with each element of
 .

Our identi cation problem is similar to that of Nevo, Turner, and Williams (2016), who study households that optimize their usage of telecommunications services when facing non-linear pricing (i.e., xed fee, allowance, and overage price) and uncertainty about their future usage. This uncertainty introduces a shadow price for current usage that is a function of the overage price and probability of exceeding the usage allowance by the end of a billing cycle. If uncertainty is substantial and varies from month to month, it creates variation in costs useful for identifying a household's preferences.

Similarly for airlines, there is a shadow cost associated with the sale of each seat, which equals the expected revenue from instead selling the seat in a future period. These shadow costs depend on demand and capacity, and can vary substantially across time for a ight due to the stochastic nature of demand. Our model maps these shadow costs to observables like prices, timing of purchase, passenger volumes, and reason for travel. We use this mapping to construct ight-speci c moments for each of these outcomes, which we then pool across ights with similar levels of capacity to construct aggregate moments.[14] This results in a set of empirical moments for each capacity, ($\Uparrow_1$), that we seek to match.

For a given initial capacity ! $_1$ and each period prior to the departure, we use the following moments conditions: (i) the fares for economy and rst-class tickets, for various levels of BTI, which show in Figure 4; and (ii) the distribution of the maximum and minimum di erences in rst-class and economy fares over time, i.e., $\max_{t=1,\ldots,T} (p_t^f - p_t^e)$ and $\min_{t=1,\ldots,T} (p_t^f - p_t^e)$), respectively, which can be derived from Figure 4 by taking the di erence between the two fares; (iii) the proportion of business traveler in each period and the economy/ rst-class fares, as shown in Figure 3; (iv) the joint distribution of ight-BTI and proportion of total arrivals for di erent periods; (v) the quantiles of passenger load factor which is shown in Figure 6(b); (vi) number of tickets, for each class, sold at various levels of BTI, which gives us something similar to Figure 3 with the number of seats on the z-axis; and (vii) overall proportion of business travelers, which is shown in Figure 2.

---

[14]We use a kernel density to apply less weight to ights with less similar capacity when constructing these aggregate moments from the ight-speci c moments.

arrival rate is also the average number of passengers that arrive, the natural way to identify is to use the exact number of seats sold each period. Even though that information is not

Figure 9: Market Heterogeneity: Marginal Distributions of Demand Parameters
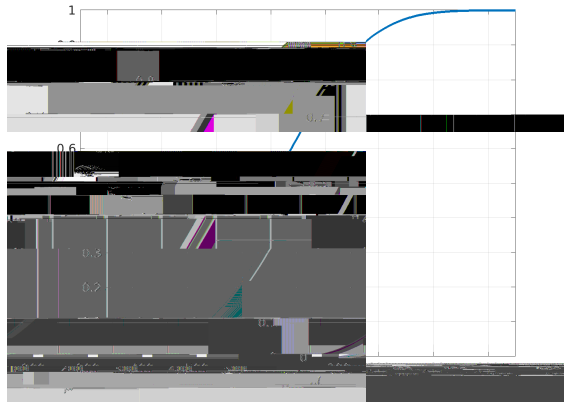
age desirability of the first-class seat is 38% greater than economy, and it varies considerably across markets, which is consistent with the idea that a first-class seat is more desirable in long-haul flights and for business travelers. Finally, in Panel (d) we present the marginal distribution of

Figure 11: Evolution of State for Modal Market and Capacity

Note: The gure displays the contours corresponding to the joint density of unsold seats for every period.

marginal cost of a seat. The total marginal cost of a seat comprises of its \peanut" cost, which is constant, and the opportunity cost that varies over time depending on the evolution of the state{number of unsold economy and rst-class seats. The shadow cost is the right-hand side of Equation 6, the change in expected value for a change in today's price. Or in other words, the shadow cost is the cost of future revenues to the airline of selling an additional seat today. In Figure 11 we present the evolution of the state as is implied by the estimates, and in Figure 12 we present the distributions of the marginal cost for an economy and rst-class seat, for all 5 periods, that are actually realized in equilibrium. This graphically relates the state transitions to the shadow cost of a seat.

In Figure 11 we present the contours corresponding to the joint density of the state. Consider! $_1$, which is the initial capacity for this modal capacity market. So, when we move to the next few periods, we see that the uncertainty increases. But as we get closer to the departure time, the contours move towards the origin, which denotes that with time fewer seats might be left. The contour of the state at the time of departure! ($_{dept}$) denotes the distribution of the state at the time of departure, which is consistent with the load factor observed in our sample (see panel (b) of Figure 6). Thus we can conclude that there is a lot of uncertainty/volatility about the demand.

One of the implications of this demand volatility is the implied volatility in the value of a

6 Quantifying Inefficiency and Welfare

There are two sources of inefficiency in our model, the gradual realization of demand and asymmetric information. The first source leads to inefficient allocations of capacity because the airline chooses policies to allocate capacity without knowledge of future demand. These inter-cabin and inter-temporal inefficiencies represent opportunities for welfare-improving trade. The second source, asymmetric information, has two parts: a passenger's value for a seat and the reason to travel. Airlines' inability to price based on an arrival's reason for travel or even the realizations of valuations can distort the final allocation of seats. In this section, we use our model estimates to simulate counterfactual outcomes to quantify the inefficiencies attributable to each source, and discuss how these inefficiencies manifest as different forms of price dispersion.

6.1 Counterfactual Results

First, consider the first-best allocation when the airline has perfect foresight and observes (v, θ) for all arriving passengers. In that case the first-best allocation is straightforward:

Figure 13: Welfare Triangle

Producer Surplus

A: (1$^{\text{st}}$-best; full extraction)

B: (1$^{\text{st}}$-best; zero price)

Consumer Surplus

that is run every period. Such a division of surplus is denoted by point G in Figure 13.

Figure 14: Discrimination and Price Dispersion for Modal Market and Capacity

Note: For the modal market with modal capacity, this figure displays price dispersion associated with the pricing regimes corresponding to points A, E, C and D in Figure 13. For example, the two lines with circular markers are the 25^{th} and the 75^{th} percentile of the economy fare, in each period, for the pricing regime D. Other pairs of lines with the same markers are defined similarly.

attributable to the stochastic demand. If the airlines had used a VCG mechanism as a way to allocate seats by time-period, their profit would be $73,005, which is almost 85% of the total possible welfare.

Lastly, we calculate the welfare when airlines can third-degree price discriminate based on reason for travel, D in Figure 13. In this scenario, profits increase by more than 15%, and business traveler welfare goes down substantially, from 7,081 to 4,206 (41%), compared to the baseline case. As expected, leisure traveler welfare increases, but only by 1.2%. Airlines' profits increase from $58,996 to $67,573, suggesting that with better ability to price discriminate, total welfare increases but at a cost of lower consumer welfare. In other words, the airlines can capture the informational rents from business travelers.

7   Conclusion

We develop a model of intra-temporal and inter-temporal price discrimination by airlines that sell a fixed number of seats of different quality to heterogenous consumers arriving before a

Table 3: Price Discrimination Counterfactuals

	1st Best B	2nd Degree C	3d Degree D	4t Degree F	VCG G
Producer Surplus	0	58,996	67,573	0	73,005
	(0.00)	(199.59)	(210.77)	(0.00)	(222.87)
Consumer Surplus	92,074	16,215	13,792	86,229	13,224
	(281.02)	(99.26)	(101.81)	(256.41)	(89.49)
{Business	19,372	7,081	4,206	16,893	4,754
	(87.12)	(34.90)	(27.31)	(42.57)	(37.87)
{Leisure	72,702	9,134	9,586	69,336	8,470
	(267.18)	(92.93)	(98.08)	(252.86)	(81.09)
Total Welfare	92,074	75,211	81,365	86,229	86,229
	(281.02)	(222.9)	(234.07)	(256.41)	(240.16)

Note: In this table we present measures of welfare for six di erent outcomes, corresponding to points A-G in Figure 13. These calculations are performed for all market types receiving positive weight for the modal initial capacity. Bootstrapped standard errors are in the parenthesis.

References

Ackerberg, Daniel. 2009. \A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation." Quantitative Marketing and Economics 7 (4):343{376.

Armstrong, Mark. 2006. \Recent Developments in the Economics of Price Discrimination." In Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, edited by Richard Blundell, Whitney K. Newey, and Torsten Persson. 97{141.

Barnhart, Cynthia, Peter Belobaba, and Amedeo R. Odoni. 2003. \Application of Operations Research in the Air Transport Industry."

Chandra, Ambarish and Mara Lederman. 2018. \Revisiting the Relationship between Competition and Price Discrimination." American Economic Journal: Microeconomics 10 (2):190{224.

Cho, Sungjin, Gong Lee, John Rust, and Mengkai Yu. 2018. \Optimal Dynamic Hotel Pricing." Mimeo .

Ciliberto, Federico and Jonathan W. Williams. 2014. \Does Multimarket Contact Facilitate Tacit Collusion? Inference on Conduct Parameters in the Airline Industry."Rand Journal of Economics 45 (4):764{791.

Courty, Pascal and Hao Li. 2000. \Sequential Screening."Review of Economic Studies 67 (4):697{717.

Eden, B. 1990. "Marginal Cost Pricing when Spot Markets Are Complete." Journal of Political Economy 98 (6):1293–1306.

Einav, Liran, Amy Finkelstein, and Mark R. Cullen. 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." Quarterly Journal of Economics 125 (3):877–921.

Einav, Liran, Amy Finkelstein, and Paul Schrimpf. 2010. "Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market." Econometrica 78 (3):1031–1092.

Ely, Jeffrey C, Daniel F. Garrett, and Toomas Hinnosaar. 2017. "Overbooking." Journal of European Economic Association 15 (6):1258–1301.

Fox, Jeremy, K. Kim, and Chenyu Yang. 2016. "A Simple Nonparametric Approach to Estimating the Distribution of Random Coefficients in Structural Models." Journal of Econometrics 195 (2):236–254.

Gale, David and Hukukane Nikaido. 1965. "The Jacobian Matrix and Global Univalence of Mappings." Mathematische Annalen 159 (2):81–93.

Gale, I. and T. Holmes. 1993. "Advance-purchase Discounts and Monopoly Allocation of Capacity." American Economic Review 83 (1):135–146.

Garrett, Daniel F. 2016. "Intertemporal Price Discrimination: Dynamic Arrivals and Changing values." American Economic Review 106 (11):3275–3299.

Gerardi, Kristopher S. and Adam Hale Shapiro. 2009. "Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry." Journal of Political Economy 117 (1):1–37.

Goettler, Ronald L. and Karen Clay. 2011. "Tariff Choice with Consumer Learning and Switching Costs." Journal of Marketing Research 48 (4):633–652.

Hartman, Wes. 2006. "Intertemporal Effects of Consumption and Their IMplications for Demand Elasticity Estimates." Quantitative Marketing and Economics 4 (4):325–349.

Hendel, Igal and Aviv Nevo. 2013. "Intertemporal Price Discrimination in Storable Goods Markets." American Economic Review 103 (7):2722–2751.

Hirshleifer, J. 1971. "The Private and Social Value of Information and the Reward to Incentive Activity." American Economic Review 61 (4):561–574.

|||. 1980. \Privacy: Its Origin, Function, and Future." The Journal of Legal Studies 9 (4):649{664.

Johnson, Oliver. 2007. \Log-concavity and the Maximum Entropy Property of the Poisson Distribution." Stochastic Processes and their Applications 17 (6):791{802.

Klemperer, Paul D. and Margaret A. Meyer. 1989. \Supply Function Equilibria in Oligopoly under Uncertainty." Econometrica 57 (6):1243{1277.

Lahiri, Soumendra N. 2003. Resampling Methods for Dependent Data, Springer Series in Statistics. Springer.

Lazarev, John. 2013. \The Welfare E ects of Intertemporal Price Discrimination: An Empirical Analysis of Airline Pricing in U.S. Monopoly Markets." Working Paper .

Leslie, Phillip. 2004. \Price Discrimination in Broadway Theater." Rand Journal of Economics 35 (3):520{541.

McAfee, R. P. and V. Te. Velde. 2006. \Dynamic Pricing in the Airline Industry." In Handbook on Economics and Information Systems, Handbooks in Information Systems, vol. 1, edited by T. J. Hendershott. Elsevier, 527{570.

McCartney, Scott. 2016. \Going Once, Going Twice: Airlines Auction Seat Upgrades." The Wall Street Journal .

McManus, Brian. 2007. \Nonlinear Pricing in an Oligopoly Market: The Case of Specialty Co ee." Rand Journal of Economics 38 (2):512{532.

Murphy, R. S. 1996. \Property Rights in Personal Information: An Economic Defense of Privacy." Georgetown Law Journal 84:2381{2417.

Mussa, M. and S. Rosen. 1978. \Monopoly and Product Quality." Journal of Economic Theory 18 (2):301{317.

Nair, Harikesh. 2007. \Intertemporal Price Discrimination with Forward-Looking Consumers: Application to the US Market for Console Video-Games." Quantitative Marketing and Economics 5 (3):239{292.

Nesterak, Max. 2017. \How Delta Masters the Game of Overbooking Flights." PBS News Hour .

Nevo, Aviv, J. Turner, and Jonathan W. Williams. 2016. \Usage-Based Pricing and Demand for Residential Broadband." Econometrica 84 (2):411{443.

Odlyzko, A. 2003. \Privacy, Economics, and Price Discrimination on the Internet." In *ICEC2003: Fifth International Conference on Electronic Commerce*, 355{366.

Posner, Richard A. 1981. \The Economics of Privacy." *American Economic Review* 71 (2):405{409.

Prescott, Edward. 1975. \E ciency of the Natural Rate." *Journal of Political Economy* 83 (6):1229{1236.

Puller, Steven L., Anirban Sengupta, and Steven N. Wiggins. 2012. \Does Scarcity Drive Intra-route Price Dispersion in Airlines?" NBER Working Paper 15555.

Roberts, James W. and Andrew Sweeting. 2012. \The Potential Entry Defense in Airline Mergers." Working Paper .

|||. 2016. \Bailouts and the Preservation of Competition: The Case of the Federal Timber Contract Payment Modi cation Act." *American Economic Journal: Microeconomics* 8 (3):257{288.

Sanders, Robert E. 2017. \Reducing Retailer Food Waste through Revenue Management." Working Paper .

Snider, Connan and Jonathan W. Williams. 2015. \Barriers to Entry in the Airline Industry: A Multidimensional Regression-Discontinuity Analysis of AIR-21." *Review of Economics and Statistics* 97 (5):1002{1022.

Spence, Michael. 1977. \Nonlinear Pricing and Welfare." *Journal of Public Economics* 8 (1):1{18.

Stamatopoulos, Ioannis and Christos Tzamos. 2016. \Optimal Seat Pricing with Multiple Classes and Why Economy Seats Feel So Small." *Mimeo* .

Stigler, George J. 1980. \An Introduction to Privacy in Economics and Politics." *The Journal of Legal Studies* 9 (4):623{644.

Stiglitz, Joseph E. 1977. \Monopoly, Non-linear Pricing, and Imperfect Information: The Insurance Market." *Review of Economic Studies* 44 (3):407{430.

Stokey, Nancy. 1979. \Intertemporal Price Discrimination." *Quarterly Journal of Economics* 93 (3):355{371.

Sweeting, Andrew. 2010. \Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets." *Journal of Political Economy* 120 (6):1133{1172.

Taylor, C. 2004. \Consumer Privacy and the Market for Customer Information." *Rand Journal of Economics* 35 (4):631{650.

Tully, Shawn. 2015. \Southwest Bets Big on Business Travelers." *Fortune* .

van Ryzin, Garrett and Kalyan Talluri. 2005. \An Introduction to Revenue Management." In *Emerging Theory, Methods, and Applications*, chap. 6. INFORMS, hap6.927S94

Appendix

A.1   Uniqueness of the Optimal Policy

In this section we show that the optimal policy is unique under some regularity conditions. These regularity conditions are widely used in the literature and ensure that demand is decreasing in its own and cross price, and that the demand for each seat class is concave. We begin by presenting these conditions below, but for notational ease suppress the time index.

Assumption 1.   1. (Downward Demand   ): $\frac{\partial \mathbb{E} q^e(p^e; p^f)}{\partial p}$ 0, and $\frac{\partial \mathbb{E} q^f(p^e; p^f)}{}$

Step 1

Here, $T = 1$ and for notational ease suppress the time index. The airline solves:

$$V(\cdot) = \max_{p^e;p^f} \left(\sum_{k=e;f} (p_t^k - c^k) \int q^k(p^e;p^f)g^k(q^k(p^e;p^f);p^e;p^f)dq^k \right)$$

$$= \max_{p^e;p^f} \sum_{k=e;f} (p_t^k - c^k)Eq^k(p^e;p^f) \tag{A.1}$$

Then the equilibrium prices $(p^e;p^f)$ solve the following system of equations:

$$\left[\begin{array}{c} Eq^e(p^e;p^f) + (p^e - c^e)\frac{@Eq^e(p^e;p^f)}{@p^e} + (p^f - c^f)\frac{@Eq^f(p^e;p^f)}{@p^e} \\ Eq^f(p^e;p^f) + (p^f - c^f)\frac{@Eq^f(p^e;p^f)}{@p^f} + (p^e - c^e)\frac{@Eq^e(p^e;p^f)}{@p^f} \end{array}\right] = \left[\begin{array}{c} 0 \\ 0 \end{array}\right] : \tag{A.2}$$

The above system has a unique solution $(p^e;p^f)$ if the negative of the Jacobian corresponding to the above system is a P-matrix (Gale and Nikaido, 1965). In other words, all principal minors of the Jacobian matrix are non-positive, which follows from Assumption 1.
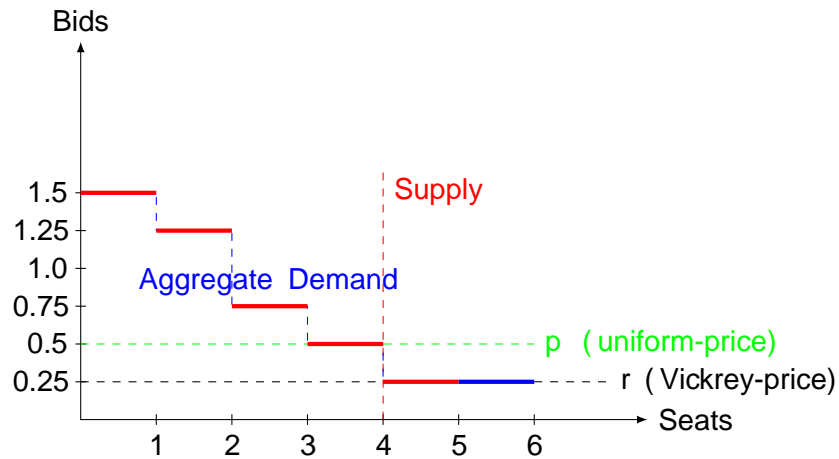
Step 2

Suppose we have a unique solution when $T = t$ and all nite pair $f K^e; K^f g$. Now we want to show that the solution is still unique if we have one additional period, i.e. $T = t + 1$. Consider the value function

$$V(\pi^t) := \max_{f \pi_t g_{t=0}^t} E_0 \left[\sum_{t=0}^t \sum_{k=e;f} (p_t^k - c^k)q_t^k(\pi_t) \middle| K_0^f; K_0^e \right] \tag{A.3}$$

where $\pi^t := (\pi_1;\ldots;\pi_t)$ is the unique optimal policy. Now, suppose we have $t + 1$ periods to consider. So the maximization problem faced by the airline becomes

$$\max_{f \pi_t g_{t=0}^{t+1}} E_0 \left[\sum_{t=0}^{t+1} \sum_{k=e;f} (p_t^k - c^k)q_t^k(\pi_t) \middle| K_0^f; K_0^e \right]$$

$$= \max_{f \pi_t g_{t=0}^t} E_0 \left[\sum_{k=e;f} (p_t^k - c^k)q_t^k(\pi_t) \middle| K_0^f; K_0^e \right]$$

$$+ \max_{\pi_{t+1}} \sum_{!_{t+1}} \pi_{t+1} E_{t+1} \left[\sum_{k=e;f} (p_t^k - c^k)q_t^k(\pi_t) !_{t+1} \right] Pr(!_{t+1} j \pi^t)$$

47

Figure 15: Multi-Unit Auction

Note. 4 seats are auctioned among 6 bidders, arranged in a descending order of their bids (thick red line). All bids are aggregated to get the aggregate demand (the thick-red and dotted blue line). The supply is ﬁxed at 4 seats. Market-clearing price (or the uniform price) is $p = 0.5$. Bidders 5 and 6 submit the same bid (thick red and blue part of the aggregate demand).

$n < d$ is zero. Therefore, we see that the transition probability is a Poisson distribution with parameter $\lambda_t(1 - \bar{F}_t(p))$.

□

A.2 Vickrey-Groves-Clarke Auction

In this section, we explain the steps we take to implement VCG. We can explain the basic idea behind VCG by considering an example of where we auction four seats of the same class to potential passengers who want to buy at most one seat; see Figure 15. Each passenger submits a bid (the price she is willing to pay for a ticket) and these bids are ranked and aggregated to generate an aggregate demand function. The price at which the demand equals four (or the supply) is called the cut-oﬀ price (p). The airline then allocates the four seats to those bidders who bid at least p. Under the uniform price auction these four passengers pay the same (uniform) price p, while under discriminatory price auction they pay their own bid. Under VCG, however, each bidder pays the opportunity cost of each seat, which is the highest rejected bid r.

Now, we explain the auction procedure in detail. Let $K^f + K^e = \{1, \ldots, K\}$ be the total number of seats available for sale in a given period. For notational simplicity, we suppress the time index. Both the airline and the passengers consider the two classes as weak substitutes, but within each classes each seat is a perfect substitute. Moreover, a passenger only wants

49

to buy one ticket. Let $S_i \in \{0, 1\} \times \{0, 1\}$ denote an allocation to passenger $i \in N$. For example, $S_i = (1; 0)$ denotes passenger i gets only a first-class ticket, similarly $S_i = (1; 1)$ denotes passenger gets a ticket for both classes. Since each passenger wants at most one ticket, if at all, the passenger i's value as a vector is

$$v_i := (v_i(0; 0); v_i(0; 1); v_i(1; 0); v_i(1; 1)) = (0; v; v; v): \quad \text{(A.4)}$$

Thus it is never efficient to give one passenger more than one ticket, if at all. To reflect this, we restrict the allocation rule to be $S_i \in \{(0; 0); (1; 0); (0; 1)\}$.

An allocation is then is an ordered collection $(S_1; \ldots; S_N)$ of seats among N passengers, such that $\bigcup_i S_i = K$, i.e., each seat is allocated to at least one passenger, and no seat is allocated to more than one passenger. Let V denote the space of all value vector $v = (v_1; \ldots; v_N)$ of passengers, and let $: V \to K$ denote an allocation

$$(v) = \langle S_1(v); \ldots; S_N(v) \rangle;$$

which is the "externality" passenger i imposes on everybody else by being present in the auction. Fix the value of everybody else other than i at some value v_{-i}. Under VCG passenger i with value v_i if he reports his value is z_i is $v_i(\pi_i(z_i; v_{-i}))$

Step 4. From the solution vector $\{a_i((0;0); a_i(1;0); a_i(0;1)\}_{i=1}^N$ determine the total number K^f of rst-class seats sold, and K^e economy-class seats sold at cost $C = K^f c^f + K^e c^e$.

Step 7. Then the net e ciency loss due to dynamic demand is $W(v)$ C utility from the data .

Step 8. The pro t for the airline is E (auction, given seats) := $\sum_{i=1}^{N_t} P_i$ C.