Price Discrimination along Multiple Dimensions: Theory and New Evidence*

Ambarish Chandra†
University of Toronto

April 5, 2021

Abstract

This paper examines firms that simultaneously practice quality-based and inter-temporal price discrimination. A model establishes conditions under which these can operate in the same, or opposite, direction, depending on the arrival rate of various consumer types. New data provided by a regional airline reveal that inter-temporal price differences are purely driven by price discrimination, rather than by scarcity pricing. In the domestic market the quality premium increases with inter-temporal price discrimination but the opposite is true in the transborder market. New survey data illustrate how the result is driven by differences in the motivation and characteristics of travelers in each market, which fully reconciles both the predictions of the model and the empirical results.

*I thank Gaurab Aryal, Nate Baum-Snow, Rob Clark, Ricard Gil, Julie Mortimer, Eugene Orlov, Claudio Piga, Steve Puller, Henry Schneider and Alan Sorensen for helpful suggestions. I am especially grateful to Jim Dana for very useful discussions and insights. This research is funded by a SSHRC Insight Grant.

†Department of Management, University of Toronto at Scarborough and Rotman School of Management, 105 St. George Street, Toronto, Ontario, M5S 3E6, Canada. Email: ambarishchandra@gmail.com

1 Introduction

With the advent of large new datasets on consumer characteristics, and the proliferation of consumer tracking technologies, the ability of firms to price discriminate has grown tremendously. Firms, in many industries, now use sophisticated methods to target consumers, often using different kinds of price discrimination strategies simultaneously. It is increasingly common to see firms—in industries as varied as banking, publishing and online retailing—use both menu pricing and group pricing, at the same time, to maximize profits.

Price Discrimination is, of course, a central and long-studied topic in Microeconomics. But although the body of research in this area is vast, most prior work generally models firms as practicing a single kind of price discrimination at a time—typically either second- or third-degree price discrimination strategies. Theoretical studies tend to focus on just one of these types in isolation, most likely due to the complexity inherent in models of multidimensional heterogeneity (Armstrong and Rochet, 1999). Empirical studies also tend to focus on just one type or the other, likely due to data constraints. However, there is no reason to believe that the results from partial analysis of each kind of discriminatory pricing will extend to environments where both kinds are implemented.

In this paper I present a stylized model, augmented with new data, to examine a setting where firms practice two different kinds of price discrimination simultaneously. In the model, a firm encounters different types of consumers arriving in each period who self-select into different quality levels. Therefore, the firm must choose prices for each quality level in multiple periods. The model shows that, while average prices increase over time, the premium for high quality service can either grow or shrink, depending on the composition of consumers that arrive in each period.

I then turn to the airline industry, which is a setting where firms have long employed multiple types of discriminatory pricing to extract surplus, often in finely targeted ways. I obtain data from a regional airline that operates on routes within Canada, and between Canada and the United States. I examine price menus for various quality levels, as well as price discounts for advance purchases, to simultaneously examine both quality-based and inter-temporal price discrimination, and establish key results that confirm the model's predictions. I believe this is the first paper to use booking data directly from an airline to establish how firms price discriminate according to advance purchases. Moreover, this is among the first empirical studies of how a firm simultaneously practices multiple types of price discrimination, especially in a setting where competition is important.[1]

[1]Prior authors have examined multiple kinds of price discrimination in monopoly, for example Leslie (2004). Some studies have examined competitive settings where firms may implement multiple kinds of discriminatory pricing, but focus on just one kind in their analysis (Lin and Wang, 2015), or examine how

Conceptually, the interaction of different kinds of price discrimination can have ambiguous effects on equilibrium prices or firms' profits. For example, if a firm has information about consumer characteristics through which to practice group pricing, this may enhance its ability to implement menu pricing, such as by targeting options more finely. Thus, the ability to practice one kind of discriminatory pricing may make the other kind more lucrative. Conversely, it can be the case that once a firm has optimally implemented group pricing, there will be less surplus available to extract from consumers through menu pricing, and therefore the two kinds of discriminatory pricing may conceivably offset each other.

The empirical analysis first disentangles the role of price discrimination and scarcity pricing in determining the evolution of fares over time. Scarcity pricing, sometimes referred to as inventory management, is the phenomenon by which optimal management of limited inventory can cause prices to vary over time even in the absence of price discrimination; for example, Sweeting (2012) shows that resellers of baseball tickets reduce prices over time, due to declining opportunity costs of holding tickets as the game approaches. By contrast, firms in the airline industry not only need to dispose of a fixed inventory by a certain time, but can also price discriminate, as the enforcement of passenger identification eliminates resale. As a result, studies of the airline industry show that prices tend to rise over time as the date of the flight approaches. However, many such studies focus purely on monopoly routes due to the complexity that would be introduced by studying strategic interaction; for example, Lazarev (2013), Williams (2020), and Aryal et al. (2018).

I show that the temporal variation in fares for a given flight-date is not due to scarcity pricing, but is almost certainly because of price discrimination. Routes or flight-days with unusually high demand do not exhibit steeper price gradients than those with normal or low demand. This result is established using unambiguous measures of both fares and remaining inventory, as these are obtained from the airline, in contrast to previous work which has had to construct or infer measures of remaining inventory. This is an important finding because some prior studies have hypothesized that advance purchase discounts may not necessarily reflect price discrimination, but may instead be the optimal profit-maximizing decision by a firm that faces uncertain demand, or for which demand is likely to exceed capacity (Gale and Holmes (1993), Dana (1998), Dana (1999), Möller and Watanabe (2010)).

I then turn to the most important contribution of this paper: to empirically establish that quality-based and inter-temporal price discrimination work in the same direction in some markets, but in the opposite direction in others. I first depict this using a model where the arrival rate of consumers with high and low willingness-to-pay for quality varies across markets. I then show that there are indeed distinct airline routes in the data where

competition affects the *type* of price discrimination that firms use (Borzekowski et al., 2009).

the premium for higher-quality service either grows or shrinks as the airline practices intertemporal price discrimination. This contribution is entirely new to the literature.

Next, I provide an explanation for the results. I first verify a key prediction of the theoretical model: that differences across markets in the implementation of each kind of price discrimination must be due to differences in the arrival rate of different consumer types in each market. The results show that price discrimination operates very differently on domestic routes within Canada than on transborder routes to and from the United States. Travelers on transborder routes are significantly more likely to buy higher quality tickets closer to the date of travel than those on domestic routes, consistent with the model. These differences generate the observed differences in price discrimination across the two markets.

Finally, I use new survey data to analyze the motivation of travelers to make different kinds of airline travel. The data convincingly establish that leisure travelers on domestic routes are far more likely to stay with friends or family who also live in Canada. By contrast, leisure travelers on transborder routes are more likely to travel for vacation, on trips that are generally longer and involve more expensive accommodation, and often include costly purchases such as cruises, live performances and theme park visits. Consequently, transborder trips require considerably more advance planning, implying that leisure travelers are less likely to make such trips at short notice, as compared to domestic travel. Therefore, travelers who book close to the date of travel on transborder routes are disproportionately likely to be business travelers, who are also more likely to buy higher-quality, flexible tickets. This fact fully reconciles the theoretical predictions and the observed empirical results.

While airlines have long been a setting for studying price discrimination, the available data limit our understanding of this phenomenon. Most commonly used data sources, including the DB1B dataset provided by the U.S. Department of Transportation, do not provide information on the date at which tickets were purchased. Since fares for a given flight tend to rise over time, it is highly likely that price differences over time for a given flight are due to the varying lengths of time in advance that tickets were purchased. Previous researchers have needed to infer price dispersion or discrimination from the empirical distribution of ticket prices, without knowing how much is due to advance purchase discounts versus other factors. Only recently have some authors been able to obtain, or construct, information on how fares vary over time for a given flight.[2]

My data source is a private airline based in Toronto, which I will refer to as North Air.[3] Though small compared to most well known airlines, North Air is an important player on

[2] For example, Alderighi et al. (2015) scrape airfares from the website of Ryanair, Williams (2020) uses data from a search engine and Lazarev (2013) obtains published fares from a Global Distribution System.

[3] The airline prefers not to be named in print, though its identity should be easy to infer.

the set of routes on which it operates. The airline provided me with a 10% random sample of all its bookings, with information on passengers, fares, dates of travel, class of service and, crucially, the date of the reservation. That last variable allows me to examine price discrimination based on advance purchases.[4]

Additionally, the data allow me to calculate the remaining inventory of seats available at any point in time, which can potentially be an important constraint for any airline. This is important because optimal inventory management can create temporal price differences that are distinct from price discrimination, which can make the analysis of dynamic price changing quite challenging due to the complexity of the interaction between price discrimination and inventory management. Moreover, accurately constructing measures of remaining inventory has been another data challenge for past research.[5] I am able to use the schedule of bookings made with the airline to reconstruct the history of purchases made over time for all flight-days and then combine this with information on the airline's capacity.

This paper is organized as follows. Section 2 discusses the terminology used in this paper, especially as it relates to past work, and also discusses the most relevant prior research. Section 3 presents the new airline data source used in this paper. Section 4 presents a stylized model of price discrimination along multiple dimensions. The empirical results are presented in Section 5. Section 6 explains the mechanism behind the main empirical result by incorporating new survey data. Section 7 concludes.

2 Terminology and Literature

2.1 Terminology

Throughout this paper, I refer to two kinds of price discrimination strategies: quality-based and inter-temporal price discrimination. Quality-based price discrimination is the provision of different quality levels at a given point of time, which consumers can select among. This is also referred to as menu pricing, or as indirect or second-degree price discrimination.

[4]Some recent studies have inferred discounts related to advance purchases by scraping the websites of airlines or of online travel agents. A key difference between these studies and mine is that using scraped data provides information on the distribution of *offered* prices, rather than on *transacted* prices. Moreover, when the lowest offered price changes from one period to the next, it is not always clear whether this is due to a smaller advance purchase window or because of a change in inventory due to the earlier fare being sold. My data distinguishes each kind.

[5]Williams (2020) uses seat maps to construct the remaining inventory, at any point in time, for a given flight. Clark and Vincent (2012) do the same, but point out that this metric suffers from measurement error, and that it would be preferable to have daily load factor data directly from the airlines. Puller et al. (2009) obtain booking data, including fares and remaining inventory, from a Computer Reservation System, but this accounts, on average, for only about a third of bookings made on the flights they examine, leaving open the possibility that their constructed measures of load factors are calculated with error.

Inter-temporal price discrimination, by contrast, refers to advance purchase discounts for a given quality level. Prior research has often assumed that advance purchase discounts are also a form of menu pricing, as consumers make the decision of when to purchase tickets and therefore self-select into more or less expensive fares (Dana, 1998). However, as in Chandra and Lederman (2018), I argue that this practice actually has some elements of third-degree price discrimination, also referred to as direct price discrimination or group pricing. This is because potential travelers are never presented with a menu of fares for advance purchase versus last-minute fares. Travelers who consider buying early are never sure what fares will be if they wait, while travelers who only learn of travel needs at the last minute never had the option to purchase early.

Instead, most consumers are well aware that airfares tend to rise over time, and that it is in their interest to purchase as soon as their travel plans are known. Thus, there is rarely any *strategic* reason to defer purchasing. Those who purchase further in advance are likely travelers with fixed dates of travel, such as families or students, while those who learn at the last minute that they must travel are more likely to be business travelers, or those with a higher willingness-to-pay. Airlines infer, from how far in advance travelers wish to purchase tickets, which type they are, and charge them accordingly. Indeed the model in Section 4 explicitly characterizes advance purchase discounts as a form of third-degree price discrimination, as consumers must purchase in the current period or exit the market. It rules out strategic behaviour regarding the timing of purchase, or capacity constraints which could cause the purchase of tickets in one period to affect the shadow cost of future sales.

Nevertheless, as this terminology is controversial and, in any case, does not change the analysis that follows, I will simply refer to this practice as inter-temporal price discrimination for the rest of this paper. The larger point is that airlines practice at least two, very different, forms of price discrimination—one based on quality differences at a given point in time, and the other based on temporal differences for a given quality level.

2.2 Literature

The literature on price discrimination is vast, and both theoretical and empirical papers have focused on the two main classifications of second-degree and third-degree price discrimination, or menu pricing and group pricing.

Few previous empirical papers have studied settings where firms practice both second-degree and third-degree price discrimination. One exception is Leslie (2004), albeit in a monopoly setting. He examines Broadway theatre and distinguishes between menu pricing—implemented by offering consumers a range of quality levels of seats for a given performance—

and group pricing, which is achieved by offering some consumers coupons, and also offering discounts for consumers willing to stand in line on the day of the event. Notably, Broadway theatres do not practice temporal price discrimination. Leslie estimates a discrete choice model of demand, and shows that firm profits are higher under discriminatory pricing than uniform pricing, while consumer surplus is essentially the same.

Aryal et al. (2018) also limit their analysis to monopoly markets. However, in other respects, the focus of their paper is similar to the setting studied here, as they also examine both inter- and intra-temporal price discrimination by airlines. They use a structural model to estimate a multi-dimensional distribution of consumers' preference heterogeneity, exploiting passengers' stated reasons for travel to distinguish between business and leisure travel. They use their results to estimate the extent to which allocative inefficiency can be attributed to demand uncertainty and asymmetric information.

The empirical literature on menu pricing has focused mainly on quantity discounts, *i.e.* nonlinear pricing. While quantity discounts are theoretically the same as offering varying quality levels, there are in fact, few empirical studies of the latter. This may be because quality can change and therefore is hard to measure or compare across different contexts. One noteworthy study of quality differences in second-degree price discrimination is Shepard (1991), who establishes that gasoline stations that offer both full-service and self-service options are able to implement discriminatory pricing in a way that single-product stations cannot. Quality levels are fixed in my setting too, as the airline offers the same amenities for high quality service across all its routes. This makes it feasible to study menu pricing according to quality levels, as these do not change in response to demand or competition.

Within the theoretical literature, a number of studies have focused exclusively on either second-degree or third-degree price discrimination. Few papers examine both together, probably because of the complexity inherent in models of multi-dimensional screening (Armstrong and Rochet, 1999). The results of Rochet and Stole (2002) suggest that competition reduces quality distortions and reduces the dispersion in prices.

While price discrimination is probably a major reason for observing temporal price gradients, some prior studies argue that similar price dynamics can be achieved solely through the optimal inventory management. Specifically, when a firm faces uncertain demand, or when demand is likely to exceed capacity, prices can rise over time even in the absence of price discrimination. This result is found by Gale and Holmes (1993), Dana (1998), Dana (1999) and Möller and Watanabe (2010). Interestingly, though, Sweeting (2012) shows that prices often *decline* over time, for a given game, in Major League Baseball, which is a setting where firms must optimally manage inventory but where price discrimination is generally impossible.

3 Data and Setting

This section describes the data and establishes important empirical facts regarding both advance purchase price discrimination, and the menu of prices offered to consumers.

I obtained data from a small, private airline based in Toronto, which I refer to as North Air. The airline provided me with a 10% random sample of all its bookings, with information on the exact fares, date of travel, class of service. the number of passengers and, crucially, the date of the booking.[6] I then calculate the length of the advance purchase period, which is the number of days between booking travel and the actual day of travel.

This data source represents a significant advance over prior studies that have tried to infer or construct measures of advance purchase discounts in the airline industry, or other settings. The commonly used DB1B dataset of the U.S. Department of Transportation lacks a key piece of information that would enable a comprehensive study of price discrimination—the date of booking for each itinerary. As a result there is no way to tell how fares vary according to advance purchases, which is likely to be the most common and lucrative form of price discrimination by airlines. While other researchers have used creative way to try to obtain advance purchase pricing information, such as by monitoring and 'scraping' airline websites, accurate sources of information on this key element of firm behaviour are lacking.

North Air has a number of features that make it unlike most 'legacy' airlines but that are desirable from the point of view of my research design. First, it offers a single service cabin, with no distinction between leisure and business classes.[7] Thus, in my analysis I can ignore the common type of second-degree price discrimination—between Coach and Business class cabins—practiced by most legacy carriers. Additionally, it does not follow the usual airline practice of offering different fare 'buckets', corresponding to different sets of restrictions, at varying prices. The legacy carriers implement this practice because prices and the quantity of sales are chosen by separate teams (Lazarev, 2013). This can make their pricing opaque, and makes it difficult to compare tickets in different buckets. Finally, North Air does not offer discounts or premiums for roundtrip or multi-city itineraries. A roundtrip itinerary is priced at exactly the sum of each individual journey, which makes it easy to compare the prices of tickets that are one-way or part of longer itineraries.

North Air offers just three *fare groups* on all of its routes, which I will refer to as Fixed, Adjustable and Refundable. These are distinct types of tickets with particular characteristics. On any given flight, all three options are usually available to purchase at any time.[8]

[6] I supervised the randomization myself, to ensure that the sample was truly representative along all of North Air's routes and itineraries.

[7] In this regard, the airline is similar to most low-cost carriers. However North Air is not exactly a low cost carrier. It offers premium features such as free snacks in lounges and on board, and leather seats.

[8] This is unlike other airlines, Southwest Airlines in particular, where the cheapest fares often sell out in

The main difference between the first two kinds of tickets is that Adjustable ticket holders can switch to earlier or later flights on the day of travel, while Fixed ticket holders must pay a significant fee to do so. Refundable ticket holders can not only switch flights on the day of travel, but can also have their entire purchase refunded at any time.[9] As I will show, North Air sells very few Refundable tickets, and it makes no difference whether I ignore these, or combine them with Adjustable tickets. As all customers are shown a menu of these fares whenever they purchase a ticket, this clearly constitutes a form of second-degree price discrimination, which I will refer to as quality-based price discrimination from here on.

The data obtained from North Air has information on over 900,000 itineraries, during 2008–2014, which is a 10% random sample of the actual itineraries reserved with the airline. I restricted the initial sample in a few ways. First, North Air offers service to and from 19 airports in Canada and the United States, during my sample period. I identified 122 itineraries that began or terminated at around 10 other airports. As these are not standard tickets for travel on North Air, and probably represent situations where travel agents combined itineraries with some other airline, I dropped these observations.

Next, I dropped observations that involved complicated multi-city itineraries with three or more separate journeys, which represented about 0.3% of the data. I then dropped itineraries where the fare group was listed as "Other", i.e. not one of the three standard fare groups. These were about 0.15% of the data. Finally, I also dropped observations where the recorded fare for the journey was less than $20—which represented approximately 0.5% of observations— and are likely to have been either employee discount tickets or frequent-flyer reward tickets. I also dropped seven observations where the recorded fare was greater than $1300. The final dataset consists of 865,492 observations, for travel between 19 airports, during the years 2008–2014.[10]

Table 1 presents summary statistics for the data. Each observation is an itinerary for travel on a certain date for a given route. I use the difference between the booking and travel dates to construct *daysout*, the number of days in advance of the journey that the reservation was made. The average reservation is made 29 days in advance of travel, has 1.2 passengers and costs around $148 Canadian dollars.[11] 85% of tickets are in the Fixed fare group, with 13% in the Adjustable group. Over 80% of tickets are part of return journeys,

advance. North Air does occasionally sell out Fixed fares just prior to travel but that is the exception, not the norm. Figure 1 shows that Fixed Fares are sold at all times including the day of travel.

[9]There are a few other differences, such as with bag fees.

[10]Information on the country of residence of travelers on North Air is not directly available, but the vast majority are likely to be Canadians. The data indicate that over 90% of reservations are made in Canadian dollars. While nothing precludes U.S. residents from purchasing in Canadian dollars or vice versa, the currency of purchase is probably a reasonable proxy for the country of residence.

[11]One Canadian Dollar was worth between 0.66 and 1.09 US Dollars during my sample period.

Table 1: Summary Statistics: Airline Data

	Mean	SD	Min.	Median	Max.
One-way Base Fare (CAD/person)	148	99.6	20	123	1296
Passengers	1.21	0.5	1	1	23
DaysOut	29.2	30.7	0	19	473
Fixed Fare	0.85	0.4	0	1	1
Adjustable Fare	0.13	0.3	0	0	1
Refundable Fare	0.01	0.1	0	0	1
Return Journey	0.81	0.4	0	1	1
One-way Journey	0.19	0.4	0	0	1
Multi-City Journey	0.00	0.1	0	0	1
Nonstop	0.86	0.3	0	1	1

Note: An observation is an itinerary. N=865,492

with the rest being one-way trips.

Figure 1 shows the detailed distribution of daysout, separately for the two main fare classes. I divide daysout into bins corresponding to some standard windows of advance purchase periods. Almost 20% of tickets are sold between one and two months in advance of travel, in the Fixed category. Around 5% of tickets are sold either on the day of travel or one day prior. Fixed tickets tend to be purchased further in advance, while Adjustable tickets tend to be purchased closer to the date of travel.[12]

4 Model

In this section I present a simple model of a firm that sets price menus over time, to a changing composition of consumers. Average price differences over time for a given quality level will reflect differences in the willingness-to-pay of consumers who arrive at different times. Price differences for different quality levels, at a given point in time, represent menu pricing, *i.e.* second-degree price discrimination.

The model can help us understand how the practice of inter-temporal discrimination affects a firm's ability to extract profits through quality-based discrimination and vice versa. One possibility is that consumers' preferences are correlated, so that those who have a greater taste for quality, for example, may also have greater willingness-to-pay for a given quality level. This would allow firms to extract greater surplus from consumers, implying that the

[12]It may seem surprising that Adjustable tickets are bought later, on average, as one might intuitively expect less uncertainty regarding the timing of flights when travel occurs soon after booking. In fact, North Air explains that this behaviour is to be expected; business travelers generally book later, but are also more conscious of the value of their time, and therefore tend to reserve Adjustable tickets in order to remain flexible about their journey times even once they commit to their travel dates.

Figure 1: Distribution of DaysOut

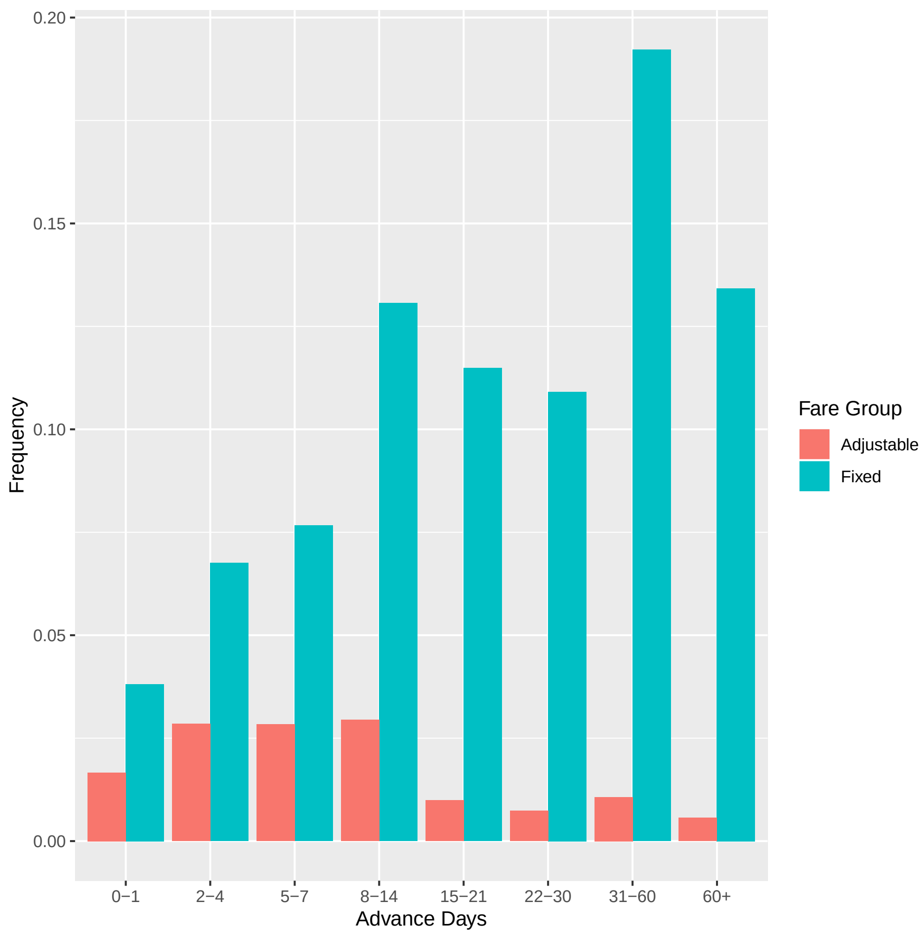

Notes: The figure plots the distribution of advance purchase days, separately for the Fixed and Adjustable Fare classes, across all routes and years in the data.

two kinds of discriminatory pricing may *reinforce* each other. Conversely, it could be that it is harder to extract surplus from, say, group pricing, once the firm has implemented menu pricing, or vice versa. This may be because there may be limits on the total willingness-to-pay of consumers, and so the two types of discriminatory pricing may *offset* each other.

I consider a model of a firm that engages in both kinds of discriminatory pricing. This is related to models of multi-dimensional screening, and it is well known that these models can be difficult to solve. However, I can simplify the model by considering the specific case of the airline industry and relying on two key characteristics of my setting.

First, I assume that the willingness-to-pay of consumers increases, on average, over time. Moreover, this fact is known to consumers themselves, and therefore they correctly understand that equilibrium prices will rise, in expectation, over time as well. As a result, consumers buy immediately once their personal uncertainty regarding travel is resolved. There is no point waiting to buy in a later period, since prices are only expected to increase.

Second, I assume that the firm has unconstrained inventory, though it incurs some constant marginal cost of providing service. Abstracting away from the inventory management problem simplifies the analysis considerably, and allows me to focus on the interaction of the two kinds of price discrimination. Moreover, this assumption also fits well with the empirical results of Section 5, which show that remaining inventory does not affect price discrimination.

Assume that there are two periods. Period 1 is the advance purchase period and Period 2 is the last-minute purchase period. In each period t, consumer types are drawn from a distribution function $F_t(\theta)$, with support $\left[\underline{\theta}(t), \bar{\theta}(t)\right]$. Thus, the distribution functions are allowed to differ across the two periods, reflecting the fact that the composition of traveler types may change over time. In addition, the support of the distribution potentially varies across the two periods, reflecting the possibility that later arriving consumers may have higher average willingness-to-pay. I use $f_1()$ and $f_2()$ to denote the density functions for the distribution functions $F_1()$ and $F_2()$ respectively.

In each period, the firm can offer both low and high quality seats, at a constant marginal cost of c_L and c_H respectively. The firm must choose prices, p_L and p_H, in each period to maximize profits. Given that consumers do not have a strategic reason to delay purchasing, and there are no inventory constraints, the firm's pricing decisions in the two periods are completely independent; in other words we can think of this as a pair of one-period problems.

A consumer of type θ gets a utility of θ from consuming the low-quality ticket, and $\phi\theta$ from consuming the high quality version, where $\phi > 1$. Thus $\phi - 1$ is a measure of the difference in quality between the two goods. A necessary condition for an equilibrium is:

$$c_H - c_L < \phi - 1 \tag{1}$$

In other words, the marginal cost increase of providing high-quality service over low-quality service must be less than the marginal willingness-to-pay for high-quality versus low-quality service.

In each period t, let $\theta_1(t)$ denote a consumer who is indifferent between the low quality ticket and not purchasing a ticket at all. Then the Individual Rationality constraint requires:

$$\theta_1(t) - p_L(t) = 0$$

which implies that $\theta_1(t) = p_L(t)$. Let $\theta_2(t)$ denote a consumer who is indifferent between the low quality ticket and the high quality ticket in each t. Then the Incentive Compatibility constraint requires:

$$\phi\theta_2(t) - p_H(t) = \theta_2(t) - p_L(t)$$

which implies that

$$\theta_2(t) = \frac{p_H(t) - p_L(t)}{\phi - 1}$$

Suppressing time subscripts, we can write the firm's problem in each period as choosing prices to maximize profits:

$$\max_{p_L, p_H} \Big\{ [F(\theta_2) - F(\theta_1)][p_L - c_L] + [1 - F(\theta_2)][p_H - c_H] \Big\}$$

The first order condition with respect to p_L is:

$$[F(\theta_2) - F(\theta_1)] + (p_L - c_L)[\frac{-f(\theta_2)}{\phi - 1} - f(\theta_1)] + (p_H - c_H)\frac{f(\theta_2)}{\phi - 1} = 0$$

which simplifies to:

$$(p_H - c_H)f(\theta_2) + (\phi - 1)[F(\theta_2) - F(\theta_1)] = (p_L - c_L)[f(\theta_2) + f(\theta_1)(\phi - 1)] \tag{2}$$

The first order condition with respect to p_H is:

$$(p_L - c_L)\frac{f(\theta_2)}{\phi - 1} + [1 - F(\theta_2)] = (p_H - c_H)\frac{f(\theta_2)}{\phi - 1}$$

which simplifies to:

$$(p_L - c_L)f(\theta_2) + (\phi - 1)[1 - F(\theta_2)] = (p_H - c_H)f(\theta_2) \tag{3}$$

Combining the two first order conditions, (2) and (3), we get:

$$p_L = c_L + \frac{[1 - F(p_L)]}{f(p_L)} \tag{4}$$

and

$$p_H - p_L = c_H - c_L + \frac{(\phi - 1)[1 - F(\frac{p_H - p_L}{\phi - 1})]}{f(\frac{p_H - p_L}{\phi - 1})} \tag{5}$$

As a benchmark, consider the case where $F(\theta)$ is uniform in each period. Then, equations 4 and 5 imply the following optimal prices:

$$p_L(t) = \frac{\bar{\theta}(t) + \max[\underline{\theta}(t), c_L]}{2} \tag{6}$$

$$p_H(t) = \frac{\phi\bar{\theta}(t) + \max[\phi\underline{\theta}(t), c_H]}{2} \tag{7}$$

The interesting question is whether the premium for high-quality service changes over time. Examining equations 6 and 7 we can see that the change in the premium depends on the evolution of the support of consumer types. For example, if $\bar{\theta}$ is increasing in t but $\underline{\theta}$ is not, then p_H increases faster than p_L and the quality premium grows over time.

By contrast, it is possible for the quality premium to shrink over time. This happens if, for example, if $\underline{\theta}$ is increasing in t but $\bar{\theta}$ is not. In such a case, for parameter values that satisfy $c_L < \underline{\theta} < c_H$, the low-quality price increases over time while the high-quality price does not.

Now consider a more general distribution of consumer types. In particular, the distribution of θ can be different in Periods 1 and 2. Specifically, the mass of the distribution can be weighted more or less heavily towards higher types in the second period, compared to the first. I state the following proposition:

Proposition 1. *If the hazard rate of the distribution of consumer types increases over time then the premium for high quality decreases over time and vice versa.*

Proof. As the hazard rate of any distribution is defined as the ratio of the probability density function to the survival function, $\frac{f(\theta)}{[1-F(\theta)]}$, the proof of the proposition is evident from examining Equation 5, which shows that the relative premium for higher quality is a decreasing function of the hazard rate of the distribution of consumer types. □

Figure 2: Densities, Survival Functions and Hazard rates

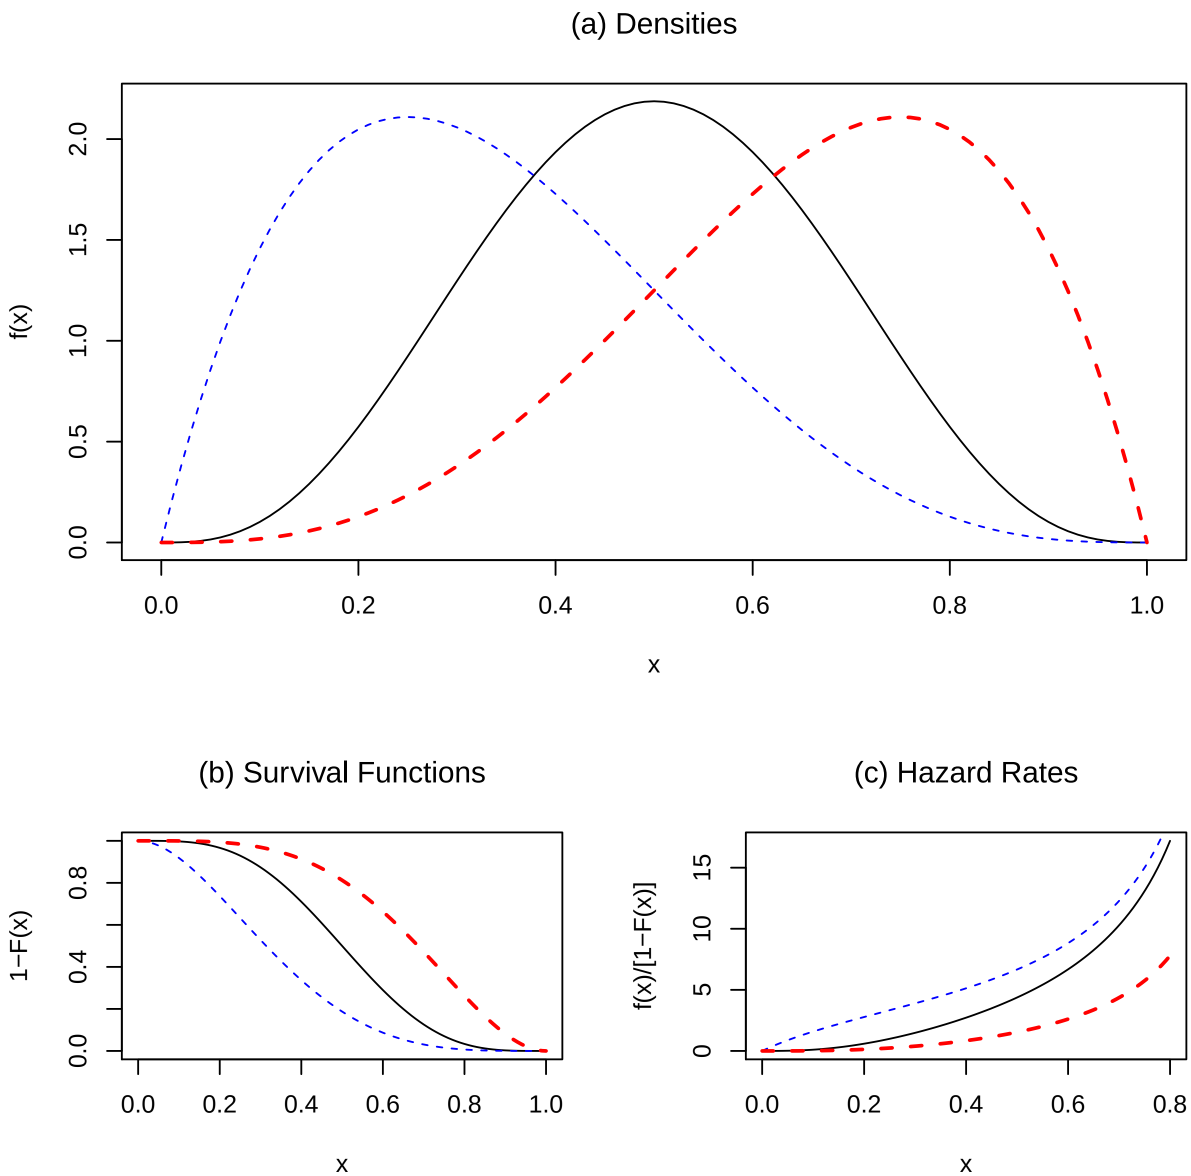

Figure 2 shows the relationship between the hazard rates and densities for three different Beta distributions on the interval $[0, 1]$. Panel (a) shows the probability density functions; the solid black curve depicts a symmetric Beta distribution, while the dashed blue and (thicker) dashed red curves depict right- and left-skewed Beta distributions respectively. A right-skewed (or positive-skewed) Beta distribution, for example, has more weight on lower values. Panel (b) of the figure shows the survival functions for each distribution, which are just the complements of the probability distribution functions, $1 - F(x)$. Finally, Panel (c) shows the hazard functions, which are the ratio of each density function to its survivor function. As can be seen, the right-skewed Beta distribution has a higher hazard rate, for any value along the support, than the symmetric distribution, which in turn has a higher hazard rate than the left-skewed distribution.

Proposition 1, therefore, implies that the evolution of prices over time depends on the arrival rate of consumers in each period and, in particular, whether consumers with a greater taste for quality are over- or under-represented in the last-minute purchase period compared to the advance purchase period. For example, if later arriving consumers are more likely than earlier arrivals to have a greater taste for quality, then we should expect to see the premium for higher quality to grow over time. By contrast, if later arriving consumers are have a relatively lower taste for quality than their counterparts in the first period, the quality premium will shrink over time.

To see this more clearly, and to also understand how equilibrium market shares depend on the arrival rate of consumers, it is helpful to consider numerical simulations, allowing for the willingness-to-pay of consumers to increase over time.

I assume that in period 1, consumer types are drawn on $[0, 1]$, while in period 2, consumer types are drawn on $[1, 2]$. Thus, the support of the Period 2 distribution is just shifted to the right by a constant, reflecting the higher average willingness to pay by consumers who arrive later. Given the cost condition in Equation 1, I assume that $c_H - c_L = \frac{\phi-1}{2}$. Finally, I assume that consumer types in each period follow a Beta distribution, with parameters α and β. I consider three cases:

Case 1: Symmetric Beta distribution in both periods: $\theta \sim \beta\,(2, 2)$. This implies that the density is given by $f(\theta) = 6\theta(1 - \theta)$ and the distribution function is $F(\theta) = 3\theta^2 - 2\theta^3$. Then, equations 4 and 5 imply:

$$8p_L^3 - 9p_L^2 + 1 = 0 \tag{8}$$

which implies that, in period 1, $p_L = 0.42$ and $p_H = 0.68(\phi - 1) + p_L$.

In period 2, the distribution shifts to $[1, 2]$. The analysis is identical to Period 1 if we simply define $\tilde{\theta} = \theta - 1$. Then the period 2 prices are $p_L = 0.42 + 1$ and $p_H = 0.68(\phi - 1) + p_L$. So the premium for high quality, $p_H - p_L$ stays constant over time and equals $0.68(\phi - 1)$ in

Table 2: Numerical Solutions for Period 2 shares and quality premium

Case	θ_1	θ_2	$\frac{1-\theta_2}{1-\theta_1}$	$p_H - p_L$
1	0.42	0.68	0.55	$0.68(\phi - 1)$
2	0.49	0.70	0.58	$0.70(\phi - 1)$
3	0.33	0.64	0.53	$0.64(\phi - 1)$

each period.

Case 2: I continue to assume a symmetric Beta distribution in Period 1, $\theta \sim \beta(2, 2)$, but now a negative-skewed (left-skewed) Beta distribution in Period 2: $\theta \sim \beta(3, 2)$. This implies that the density is given by $f(\theta) = 12\theta^2(1 - \theta)$ and the distribution function is $F(\theta) = 4\theta^3 - 3\theta^4$.

In other words, the distribution shifts to having relatively more high types in period 2. In this case, the premium for high fares increases over time, from $0.68(\phi - 1)$ to $0.70(\phi - 1)$.

Case 3: As before, I assume a symmetric Beta distribution in Period 1, $\theta \sim \beta(2, 2)$, but now a positive-skewed (right-skewed) Beta distribution in Period 2: $\theta \sim \beta(2, 3)$. This implies that the density is given by $f(\theta) = 12\theta(1 - \theta)^2$ and the distribution function is $F(\theta) = 6\theta^2 + 3\theta^4 - 8\theta^3$.

In other words, the distribution shifts to having relatively more low types in period 2. In this case, the premium for high fares decreases over time, from $0.68(\phi - 1)$ to $0.64(\phi - 1)$.

The density functions for each of the three cases, and the corresponding price paths are shown in Figure 3 for a given value of ϕ. The top panel depicts Case 1, where the distribution of consumer types is symmetric in each period, and therefore the premium for the high-quality version is constant over time. The middle panel depicts Case 2, where relatively more high-types arrive in Period 2, which causes the premium for the high-quality version to increase over time. The bottom panel depicts Case 3, where relatively fewer high-types arrive in the second period, causing the high-quality premium to shrink over time.

It is also useful to examine the equilibrium share of consumers who purchase the low- versus high-quality good in each case. Table 2 presents the Period 2 equilibrium price premiums and cutoff values for θ_1 and θ_2. It also shows the share of consumers, in each case, who purchase the high-quality version in the second period, which is calculated as $\frac{1-\theta_2}{1-\theta_1}$.

As the table shows, the second period in Case 1 has the same price premium for high-quality as in Period 1, and 55% of consumers purchase the high-quality version. In Case 2, the premium increases relative to Period 1, due to the relatively larger number of high types arriving in the second period. Despite this, the equilibrium share of consumers who

Figure 3: The arrival rate of consumer types and the evolution of the price premium over time

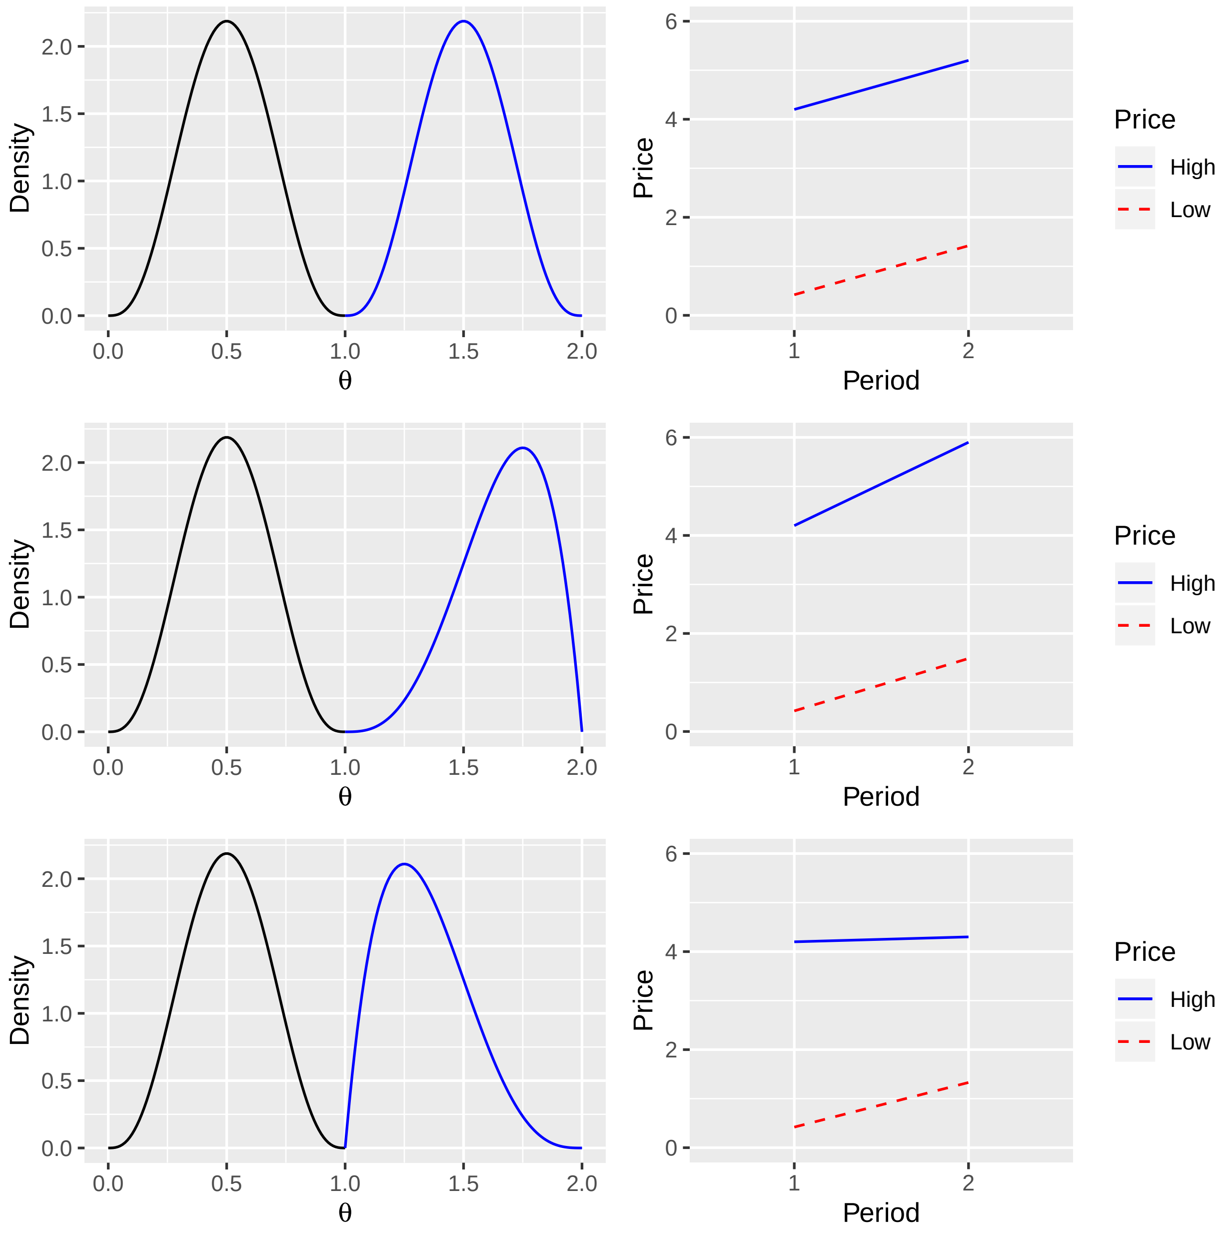

purchase the high-quality version rises, to 58%.[13] Analogously, in Case 3, both the relative premium and the share of consumers purchasing the high-quality version drop, compared to Period 1.

To summarize, this model has implications for the evolution of fares over time, which will depend on the arrival rate of various consumer types. Markets where relatively more high types arrive later will see a greater divergence between the prices for low and high quality levels. By contrast, markets where fewer high types arrive late will see a shrinking of the gap between low and high quality prices over time. In equilibrium, the firm will sell relatively more high-quality seats in the later period in markets where relatively more high-types arrive in that period, and vice versa.

With these predictions from the model, I now turn to empirical results to demonstrate that there exist airline routes where the equilibrium fares and market shares show similar patterns to the cases described above.

5 Results

I now present three main types of results. First, I will show the existence of clear advance purchase gradients, for both of the main quality levels. Second, I will incorporate data on capacity to show that inter-temporal price differences are driven almost entirely by price discrimination, and not by scarcity pricing, thus resolving an open question in the literature. Finally, I will present the main empirical result: that different markets exhibit strikingly different patterns of quality-based and inter-temporal price discrimination.

5.1 Advance Purchase Gradients

In this subsection I will quantify the advance purchase gradient, which is the extent to which fares vary based on how far in advance travelers purchased their tickets. I will also quantify the premium associated with the higher quality level on North Air.

Table 3 presents basic results. I regress the log of the fare, for each itinerary, on a set of indicators for how long in advance the ticket was purchased, separately for the two main fare types. The omitted category in each regression consists of tickets purchased more than 60 days in advance of the flight. Fixed fares are in column 1, Adjustable fares in column 2, and the combination of Refundable and Adjustable fares in column 3. I include year, month, day-of-week and route fixed effects in all specifications, and cluster standard errors by route.

[13]The intuition for this is a simple demand shift, causing both prices and quantity to rise due to the higher willingness-to-pay.

Table 3: Regression of Fares on Advance Purchase Days: Full Sample

	Fixed	Adjustable	Adj+Refundable
0 to 1 days	0.782***	0.811***	0.868***
	(0.02)	(0.09)	(0.10)
2 to 4 days	0.721***	0.772***	0.812***
	(0.03)	(0.09)	(0.10)
5 to 7 days	0.605***	0.697***	0.727***
	(0.03)	(0.09)	(0.10)
8 to 14 days	0.402***	0.545***	0.570***
	(0.02)	(0.06)	(0.07)
15 to 21 days	0.162***	0.308***	0.346***
	(0.01)	(0.02)	(0.02)
22 to 30 days	0.066***	0.215***	0.246***
	(0.01)	(0.02)	(0.02)
31 to 60 days	0.022***	0.114***	0.129***
	(0.00)	(0.01)	(0.01)
Constant	4.345***	4.801***	4.821***
	(0.04)	(0.11)	(0.11)
R^2	0.447	0.379	0.343
Obs	736409	116486	129083

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions include route, day-of-week, month and year FEs. Standard errors, clustered by route, in parentheses. Adj+Refundable refers to using both Adjustable and Refundable tickets.

Two results are apparent from Table 3. First, there are clear advance purchase gradients, for both low and high quality levels. Relative to the omitted category containing itineraries purchased more than 60 days in advance, fares rise monotonically as the date of travel approaches. On average, tickets purchased within one day of travel are about 80% more expensive than tickets of the same quality level purchased more than two months in advance, whether examining low or high service levels. Second, combining Refundable and Adjustable tickets raises the average fare in any time period, relative to only using Adjustable tickets, but does not materially change the estimated advance purchase gradient, as can be seen by comparing coefficients across columns 2 and 3. Therefore, for all the empirical exercises that follow, I define Adjustable fares as including the Refundable tickets as well.[14]

For what follows, it will also be convenient to summarize the advance purchase gradient—captured in Table 3 by the various indicators for each purchase period—in a single coefficient. One simple way to do this is to estimate the price elasticity with respect to the number of

[14]Recall that Refundable Fares are just 1% of total ticket sales, and nest all of the attributes of Adjustable fares in addition to being fully refundable.

days remaining until travel. For any route i, I estimate the following relationship:

$$P_{it} = \alpha_i D_t^{\gamma_i} \tag{9}$$

Here, D_t denotes the days remaining until travel for a fare purchased at time t, α_i is a route-specific premium or discount, and γ_i is the route-specific price elasticity with respect to the remaining days until travel. Thus, γ captures, in a single parameter, the gradient of advance purchase discounts.

I take logs of Equation 9 and estimate gradients for each route, or for groups of routes. Over all of the routes flown by North Air, the average estimated gradient for Fixed fares is -0.258. This implies that, on average, a 10% decrease in the number of days until travel is associated with an average price increase of around 2.5% for such fares, which turns out to be remarkably stable across routes. The average gradient across all routes for Adjustable fares is -0.215 but, as I will show, this masks a high degree of heterogeneity across routes.

5.2 Price Discrimination versus Scarcity Pricing

In this subsection I will disentangle the role of price discrimination and scarcity pricing in producing temporal price variation, by exploiting information about the airline's remaining capacity at the time bookings are made.

Scarcity pricing, sometimes referred to as inventory management, is the phenomenon by which prices may rise or fall over time due solely to a firm's optimal management of its inventory, in a manner that is distinct from price discrimination over time. This is because firms that practice price discrimination according to advance purchases also generally have to contend with the management of a fixed inventory. For example, in theatre management, seats that remain unsold by the time of a given performance represent a lost revenue opportunity, and suggest that prices may have been too high. Conversely, a fully sold-out show eliminates the possibility of making last-minute sales to buyers with a high willingness-to-pay, implying that prices may have been too low.[15] Indeed, it is common to observe temporal price changes for any kind of perishable good, even when price discrimination is absent; for example, discounts for food items just prior to their expiration, or steadily declining prices for baseball tickets as illustrated in Sweeting (2012).

Knowing the remaining inventory of seats at any point in time is therefore crucial for accurately understanding price discrimination. I incorporate data on North Air's capacity at

[15] Past research has examined the interaction of the inventory management and price discrimination problems; for example, Dana and Williams (2020) provide a theoretical treatment and Alderighi et al. (2015) conduct an empirical analysis.

the time of booking each itinerary, using data on flights and seats from the OAG, which is a well-known provider of data on the airline industry. I restrict the sample to include only non-stop itineraries, since it is not straightforward to construct measures of remaining seats for tickets that involve a change of plane. This reduces the sample to 718,691 observations, which is approximately 83% of the number of observations in the full sample. I then construct, for each flight-day, a measure of the total number of available seats. I do this by obtaining OAG data on the number of available seats for each route-month that North Air operates. I then divide this measure by 30 to obtain the approximate daily number of available seats.[16]

Next, I sum up the number of tickets that North Air sold for each flight-day, and multiply these by 10 to account for the fact that I have a 10% sample of tickets. I can then calculate two measures of inventory or capacity. First, I construct the fraction of *remaining seats* for any purchased itinerary, which is the share of the flight's seats that are available for purchase after each ticket is purchased. Second, I construct the *load factor* for each flight-day, which is simply the share of available seats that are in fact occupied. Clearly, both measures are constructed with error, both because I average monthly OAG data to construct daily measures of available seats, and because of the use of a 10% random sample rather than from the full distribution. Importantly, however, the error is mean zero, and there is no reason to believe that it is correlated with the variables of interest, such as fares and the quality choices made by consumers, due to the randomization procedure that was used to generate the sample.

Table 4 presents summary statistics on the number of available, occupied and remaining seats, the total number of passengers on a flight, and the load factor. Note that the fraction of remaining seats varies across purchased tickets on a given flight-day, whereas load factors are constant within a flight-day. The Table shows that the average load factor for North Air flights is 48%, which fits very well with industry observations about North Air. The average itinerary is purchased when about 75% of seats on a flight remain available, and this varies from zero to one across the distribution of purchased tickets.

I now employ the two constructed measures of capacity—the fraction of remaining seats and the load factor—in regressions of fares on advance purchase indicators. The results are in Table 5 and are presented separately for Fixed and Adjustable fares. I first present, in Columns 1 and 4, regressions analogous to those in Table 3, except that I replace the discrete measures of *daysout* with the elasticity γ that was defined in Equation 9. For both types of fares, the estimated gradient is negative. The results on the capacity measures show that average fares do respond to changes in remaining inventory. The coefficients on remaining

[16]There is likely to be day-to-day variation as, for example, the airline may operate fewer flights on weekend days. However, these differences will be captured by the day-of-week fixed effects.

Table 4: Summary Statistics: Capacity Measures

	Mean	SD	Min.	Median	Max.
Total Passengers	297.2	189.2	10	270	1030
Available Seats	624.3	331.6	12	628	1274
Occupied Seats	154.4	139.0	10	110	1030
Remaining Seats (%)	0.75	0.18	0.00	0.78	0.99
Load Factor (%)	0.48	0.19	0.01	0.47	1.00

Note: An observation is an itinerary. N=718691.

Table 5: Regression of Fares on Advance Purchase Days: Capacity Measures

	Fixed Fares			Adjustable Fares		
	(1)	(2)	(3)	(4)	(5)	(6)
γ (Gradient)	-0.236***	-0.236***	-0.237***	-0.220***	-0.220***	-0.221***
	(0.01)	(0.01)	(0.01)	(0.04)	(0.04)	(0.04)
Remaining Seats (%)		-0.074***			-0.066***	
		(0.01)			(0.02)	
Load Factor (%)			0.160***			0.159***
			(0.02)			(0.03)
Constant	5.193***	5.248***	5.118***	5.804***	5.852***	5.723***
	(0.03)	(0.03)	(0.03)	(0.07)	(0.08)	(0.06)
R^2	0.409	0.409	0.412	0.328	0.329	0.331
R^2 excl. γ	0.153	0.153	0.155	0.160	0.161	0.162
Obs	589704	589704	589704	108207	108207	108207

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions include route, day-of-week, month and year FEs. Standard errors, clustered by route, in parentheses.

seats in Columns 2 and 5 are negative, suggesting that fares are lower for itineraries purchased when a greater fraction of seats on the flight-day remain available. Additionally, the results in Columns 3 and 6 imply that fares are higher for flight-days with higher load factors, i.e. for flight-days that have a higher *ex post* likelihood of being sold out.

However, the interesting result of Table 5 is that temporal variation in prices appears to be driven entirely by price discrimination, and not by variation in remaining capacity. Controlling for the fraction of remaining seats, in columns 2 and 5, or for the load factor, in columns 3 and 6, has no discernible effect on the advance purchase elasticity. This can be seen by the small change in the magnitude of γ when going from column 1 to columns 2 and 3, or from column 4 to columns 5 and 6. Thus, these results suggest that, while remaining capacity does affect *average* fares, it has almost no effect in determining the temporal variation in fares of either high or quality tickets. The existence of an almost identical advance purchase gradient, even when controlling for remaining capacity or load factors, suggests that price discrimination, rather than scarcity pricing, is the main driver of inter-temporal price variation.

Additionally, the advance purchase gradient explains a considerable portion of the variation in fares, while capacity measures do not. This can be seen by examining the R^2 measures of fit at the bottom of Table 5. I present the overall R^2, as well as the R^2 obtained from each regression *excluding* the advance purchase gradient, *i.e.* the portion of variation in the dependent variable explained purely by the fixed effects. The fixed-effects alone explain about 15% of the variation in the Fixed fares, as seen in column 1, and this is virtually unchanged when adding the two capacity measures in columns 2 and 3. By contrast, adding the daysout elasticity raises the R^2 from 15% to 41% for Fixed fares, and from 16% to around 33% for Adjustable fares.

As a robustness check, in Table 12 in the Appendix, I divide the sample into quartiles of the load factor, and estimate the daysout elasticity for each quartile. The estimated elasticities are generally very similar to the one estimated for the full sample.

In a further robustness check, I show in Table 15 in the Appendix, that the Table 5 results are almost exactly replicated when I replace γ, with the multiple daysout indicator variables that were originally employed in Table 3. The various indicator variables explain about the same amount of variation as γ alone and, once again, the various capacity measures do not affect the estimated gradient using these discrete measures.

Overall, therefore, the results of Table 5 imply that within-route variation in fares can be explained almost entirely by temporal price discrimination rather than by scarcity pricing.[17] While measures of remaining capacity exhibit the correct sign on their estimated coefficients—implying that fuller flights are likely to be more expensive on average—they explain only a very small portion of the variation in fares, and do not change the estimated gradient with regard to advance purchases.

5.3 Price Discrimination on Domestic and Transborder Routes

In this subsection I will present the main empirical result, namely that domestic and transborder airline routes exhibit strikingly different patterns in the interaction of inter-temporal and quality-based price discrimination.

It is perhaps easier to see properties of the advance purchase gradients graphically. I present these gradients in Figure 4, separately for domestic travel and for transborder travel to the United States. These gradients are estimated, for each group of routes, by a single, pooled regression for Fixed and Adjustable fares. This is to more easily compare coefficients,

[17] This refers to the portion of variation explained by the regressors. Note that the fit measures in Table 5 can obviously be further improved by adding finer indicators for advance purchase days, but that is not the goal of this paper.

Figure 4: Advance Purchase Gradients for Domestic and Transborder Routes

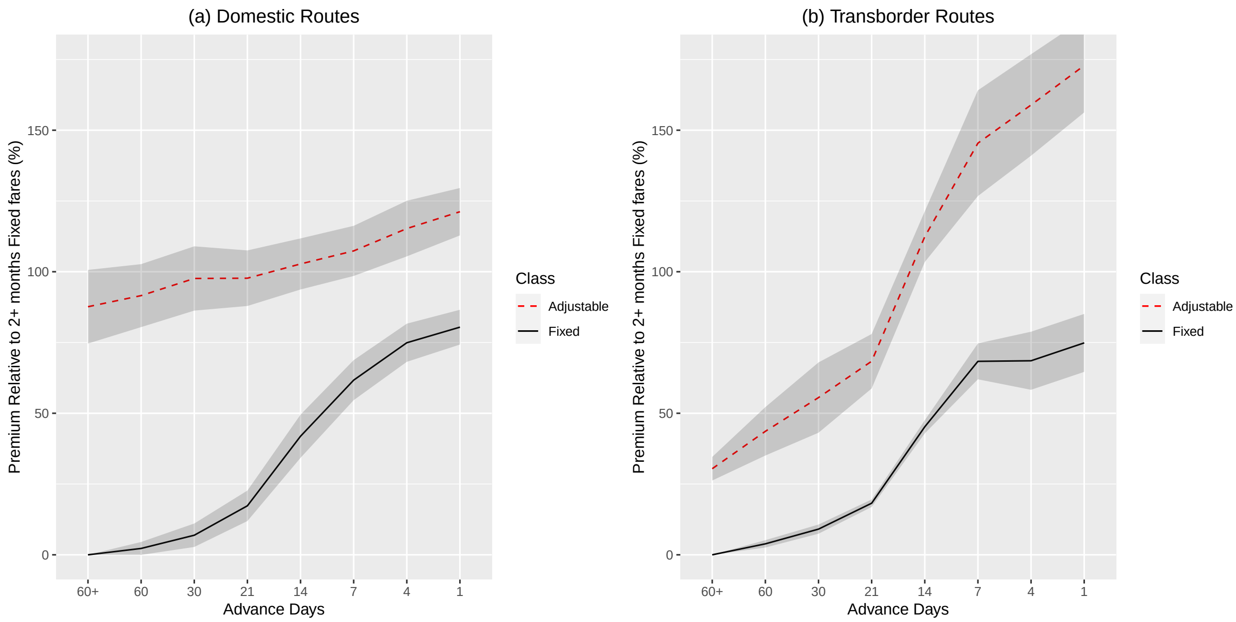

Notes: The figure plots advance purchase gradients, separately for Fixed and Adjustable fares, for domestic and transborder routes. Shaded areas represent 95% confidence intervals, using standard errors clustered by route.

and quality premiums, between these groups.[18] In each panel, the coefficient for Fixed fares purchased more than two months is advanced is normalized to zero, and all other coefficients are expressed as relative premiums to those fares.

Figure 4 shows that, as before, both Fixed and Adjustable fares rise monotonically as the date of travel approaches. However, the advance purchase gradients are different for Fixed and Adjustable fares, and also different based on whether we examine domestic or transborder routes. For both domestic and transborder routes, the advance purchase gradient is fairly steep for Fixed fares, with tickets purchased within one day of travel being about 110% more expensive than those bought far in advance.[19] The Fixed fare gradients in each market follow a relatively similar shape. Even though the coefficients for tickets bought in the last 2–3 days before travel appear to look dissimilar, we cannot reject the hypothesis that they are the same, due to the wide confidence intervals for coefficients on late purchases, especially in the transborder market.[20]

By contrast, domestic and transborder markets exhibit very different gradients for

[18]Regression coefficients are presented in Table 14 in the Appendix.

[19]This was obtained by calculating $\exp(0.75) - 1$.

[20]These confidence intervals widen because relatively few tickets are sold at the last minute for low fares. This is especially true for transborder travel, likely because international travelers generally need to make plans further in advance.

Table 6: Daysout Elasticities

Route (to and from Toronto)	Fixed Fares		Adjustable Fares	
	γ	N	γ	N
Montreal	-0.252	117905	-0.055	20253
Ottawa	-0.272	125364	-0.078	20213
Thunder Bay	-0.207	32357	-0.082	2598
New York	-0.259	80285	-0.250	25326
Boston	-0.307	26336	-0.519	11020
Chicago	-0.189	22506	-0.522	11416
All Domestic	-0.257	389738	-0.065	58098
All Transborder	-0.253	136899	-0.387	49109
All Routes	-0.258	526637	-0.215	107207

Note: γ is the estimated price elasticity with regard to days remaining until travel, for the corresponding route(s) and fare type. Estimating regressions include route, day-of-week, month and year FEs. Sample restricted to purchases made within 60 days of travel.

Adjustable fares. On domestic routes, Adjustable fares exhibit only a slight gradient, with a 25 percentage point premium for tickets bought just prior to travel versus far in advance. However, Adjustable fares on transborder routes exhibit a much steeper gradient, with last-minute purchases costing well over double the price of tickets bought 60 days in advance.

These results imply that, for the domestic market, there is a convergence between Adjustable and Fixed fares over time, whereas these fares diverge—quite sharply—in the transborder market. It is likely that this difference is related to the composition of travelers in each market. Before, turning to an explanation, I first demonstrate that these differences are very robust, and are apparent on a route-by-route basis across the two markets.

I estimate advance purchase gradients for each route, or for groups of routes. In order to perform the estimation, I restrict the sample to tickets purchased within a certain period of travel, to prevent the results being distorted by outlier itineraries that are purchased many months in advance. In Table 6, I summarize these elasticities using a 60-day window of purchase, but I show in Table 13 in the Appendix that the results are almost identical using a 90-day window.[21] The table presents separate elasticities for Fixed and Adjustable fares for some selected routes. I show the estimated elasticities for the top three domestic routes, the top three transborder routes, and then average elasticities for domestic, transborder and all routes.

The values in the left column of Table 6 are mostly very similar to each other, suggesting

[21]Tickets purchased within 60 and 90 days of travel account for 88% and 95% of itineraries, respectively.

that the advance purchase gradient for Fixed fares is the same across the major routes, as it is for the full sample or for subsamples of domestic and transborder flights. The average value of γ is around 0.25, with Boston and Chicago being slight outliers, with estimated elasticities of 0.31 and 0.19, respectively. Overall, though, the average elasticities for domestic and transborder travel are very similar, in line with the results in Figure 4.

However, the values in the right column exhibit wide variation, and significant differences between domestic and transborder travel. Specifically, the advance purchase gradient for Adjustable fares is considerably flatter for travel between Toronto and other major Canadian destinations, as well as for domestic routes in the aggregate, with an elasticity of around -0.06 to -0.08. This implies that, for such routes, there is little inter-temporal price discrimination among buyers of high-quality tickets, suggesting perhaps that it is difficult to extract further surplus from such buyers once they self-select into such tickets.

On the other hand, the elasticity for Adjustable fares on the three busiest transborder routes, and for transborder routes in general, is considerably greater; as high as, or even higher than, the elasticity for Fixed Fares. Once again, these results are consistent with those shown in Figure 4. This suggests that travelers on US–Canada routes who self-select into expensive tickets are also subject to extensive price discrimination by advance purchase. I examine this suggestive evidence more deeply in the empirical analysis that follows.

I now show evidence that the share of high-quality tickets sold is significantly higher closer to the travel date on transborder routes than on domestic routes. Figure 5 presents histograms of the share of tickets sold in each fare group across domestic and transborder routes. Fixed fare tickets comprise the vast majority of sales in general, and are most of the tickets sold in each advance date grouping on domestic routes. However, on transborder routes, Adjustable tickets make up a steadily larger share of tickets sold as the date of travel approaches, and actually outnumber the low-quality tickets in the last week of sales prior to travel. Thus, Figure 5 indicates that transborder routes see relatively more high-quality tickets sold in later periods.

To see this more formally I present, in Table 7, probit regressions of the probability that a ticket is high-quality on measures of advance purchase duration, separately for domestic and transborder routes. These regressions control for a full-set of year, month, day-of-week and route fixed-effects. Comparing columns 1 and 3, we can see that transborder markets are more likely than domestic routes to sell Adjustable tickets in the two weeks prior to travel, but that the reverse is true further in advance. Comparing columns 2 and 4, we see that the likelihood of a ticket being high-quality drops off much faster with the days until travel for transborder routes than for domestic ones. Overall, these results clearly show that high-quality tickets are disproportionately sold in later periods on Transborder routes.

Figure 5: The share of fixed and adjustable tickets on Domestic and Transborder Routes

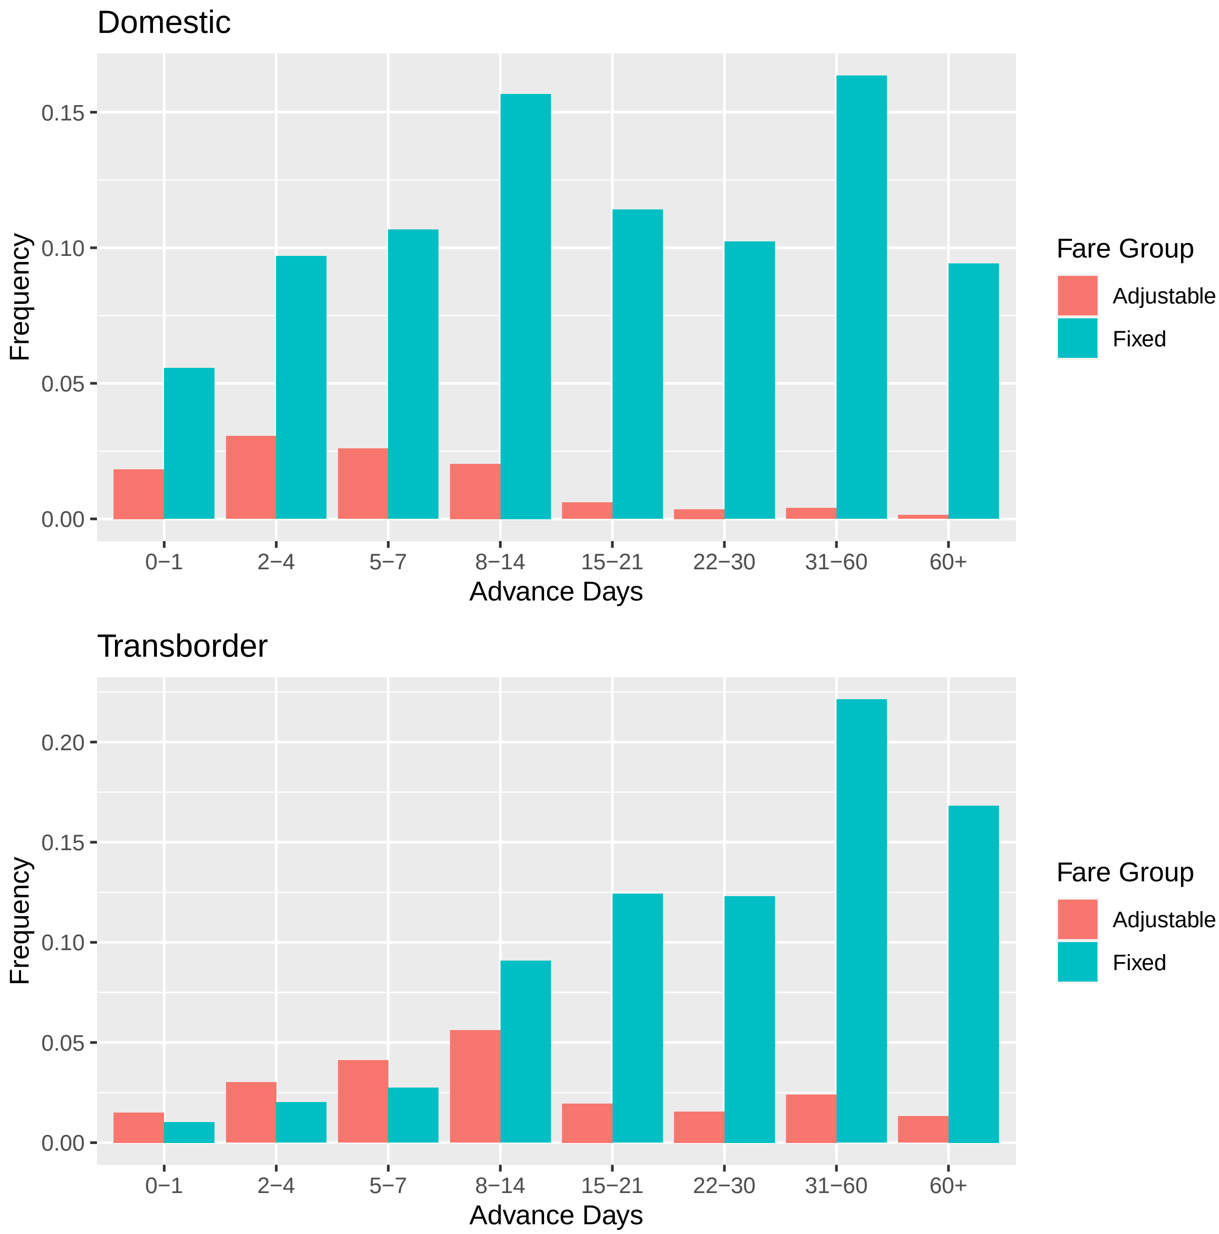

Table 7: Probit Regressions of the Share of Adjustable Fare Tickets sold

	Transborder		Domestic	
	(1)	(2)	(3)	(4)
0 to 1 days	2.327***		1.790***	
	(0.02)		(0.02)	
2 to 4 days	2.257***		1.737***	
	(0.02)		(0.02)	
5 to 7 days	2.176***		1.543***	
	(0.02)		(0.02)	
8 to 14 days	1.425***		1.154***	
	(0.01)		(0.02)	
15 to 21 days	0.433***		0.668***	
	(0.01)		(0.02)	
22 to 30 days	0.286***		0.416***	
	(0.02)		(0.02)	
31 to 60 days	0.182***		0.244***	
	(0.01)		(0.02)	
γ		-0.751***		-0.468***
		(0.00)		(0.00)
Constant	-1.179***	1.827***	-2.384***	-0.240***
	(0.03)	(0.03)	(0.03)	(0.03)
R^2				
Obs	223491	223491	493646	493646

Note: An observation is an itinerary. All regressions include year, month, day-of-week and route fixed-effects.

Table 8: Regression of Fares on Both Kinds of Price Discrimination

	All Routes			Transborder	Domestic
	(1)	(2)	(3)	(4)	(5)
γ (Gradient)	-0.303***		-0.241***	-0.210***	-0.242***
	(0.01)		(0.01)	(0.01)	(0.01)
Adjustable Fares		0.832***	0.574***	1.155***	0.272***
		(0.04)	(0.11)	(0.16)	(0.04)
$\gamma * High$			0.008	-0.190***	0.161***
			(0.05)	(0.05)	(0.02)
Constant	5.461***	4.528***	5.209***	5.135***	5.187***
	(0.04)	(0.03)	(0.03)	(0.08)	(0.02)
R^2	0.424	0.357	0.546	0.614	0.536
R^2 excl. γ	0.087				
Obs	697911	697911	697911	216314	481597

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions include route, day-of-week, month and year FEs. Standard errors, clustered by route, in parentheses.

Taken together, Tables 6 and 7 show results fully consistent with the model presented in Section 4: that it is possible for the low- and high-quality prices to either diverge or converge over time, *i.e.* that temporal price discrimination can either offset or reinforce quality-based price discrimination. Specifically, on domestic routes the premium for high-quality service shrinks as the date of travel approaches, while on transborder routes the opposite is true. Moreover, domestic routes are relatively less likely than transborder routes to sell high-quality tickets close to the date of travel, which is consistent with a relatively larger share of high-type consumers arriving at the last minute on transborder routes.

One final exercise shows how pooling together all markets can lead to incorrect conclusions. In Table 8 I regress fares on both measures of discriminatory pricing together—the gradient of fares measuring temporal price discrimination, and an indicator for Adjustable fare tickets, as well as the interaction of these. The gradient, by itself, explains about 42% of the variation in fares, while the Adjustable fare dummy explains about 35%, with both coefficients having the expected sign. Column 3 suggests that the coefficient on the interaction of these variables is not statistically significant, implying that neither measure of price discrimination is affected by the other. However, these results bely the true relationship; when I break up the sample into transborder and domestic routes, I estimate a negative interaction in the former and a positive interaction in the latter, both of which are highly statistically significant. The results suggest that the negative gradient with respect to advance purchase is particularly pronounced for Adjustable tickets on transborder routes, but mitigated for such tickets on domestic routes, completely in line with Table 6 and Figure 4.

Thus far, the results have shown that domestic and transborder routes are very different in their patterns of converging and diverging premiums for high quality over time, and that these patterns are consistent with relatively more high willingness-to-pay travelers arriving later on transborder routes. This difference in the composition of travelers gives rise to the price discrimination patterns laid out in this subsection.

6 Explanation for the results

I now turn to the mechanism behind the main result, which is the finding that price discrimination appears to work differently in the domestic and transborder markets. The explanation hinges on the composition of travelers in each market. In order to examine this hypothesis, I make use of new survey data on the motivation of travelers to make airline trips. The data reveal that leisure travelers are less likely to make last-minute transborder journeys, because of their underlying motivation for such trips. Transborder trips involve considerably more planning, and potentially more expense, than domestic trips, as I now show.

The data for this exercise are drawn from the National Travel Survey (NTS), operated by Statistics Canada. This is a monthly survey of Canadians, which began in January 2018. I use information from the 2018 and 2019 calendar years. While these years do not overlap with the sample period of North Air's travel data, the general travel habits of Canadians are unlikely to fluctuate greatly from year to year.[22] The NTS is a voluntary survey, although responses are carefully weighted to account for response bias, in order to form a complete picture of travel habits. The survey asks about all travel by members of each surveyed household, whether domestic or international. To my knowledge, this is the first use of this survey in Economics research.

In the NTS data, I restricted the sample to all trips made by Canadian residents for which airline travel was the primary mode of transport. I also restricted the sample to trips that were either entirely within Canada, or international trips that were solely to the United States and back. The resulting sample consists of 18,095 observations which represent around 27 million trips across the two years, once the appropriate survey weights are applied. I use these data to examine the characteristics of domestic and transborder trips.

Table 9 shows the main reason for taking domestic and transborder trips, using data from the NTS.[23] The results show that domestic travel is somewhat more likely (37%) to be

[22] The NTS replaced two earlier surveys, the Travel Survey of Residents of Canada and the International Travel Survey which covered domestic and international travel, respectively. While data from those surveys is available for years corresponding to my sample period, the methodology of each varied considerably, making it impossible to compare domestic and foreign trips for those years. The NTS has the advantage of using exactly the same sampling methodology and survey questions for domestic and international travelers.

[23] There are ten categories in the survey for reporting the motivation for trips, which I have aggregated into these four categories.

Table 9: National Travel Survey: Reasons for Travel

Reason (%)	Domestic	Transborder
Holiday, Recreation, Leisure	16.0	47.2
Visit Friends or Relatives	41.1	20.3
Other Non-Business	6.3	4.5
Business	36.7	28.0
Observations	9,641	8,454
Population-weighted Trips	15.2M	11.5M

Note: Author's calculations from the National Travel Survey for 2018 and 2019.

Table 10: National Travel Survey: Selected Activities During Leisure Travel

Activity during trip (%)	Domestic	Transborder
Spectator at Sporting Event	5.9	11.4
Attend Performance: Play or Concert	6.7	16.5
Visit Theme or Amusement Park	2.6	9.0
Observations	6,436	6,609
Population-weighted Trips	9.5M	8.1M

Note: Author's calculations from the National Travel Survey for 2018 and 2019.

for business than is international travel (28%). Within the categories of non-business travel, however, Canadians are far more likely to take visit friends or family when taking domestic trips (41%) rather than on transborder trips (20%), which is unsurprising since Canadians generally have relatives and friends who also live in Canada rather than abroad.

Another difference between domestic and transborder travel emerges from examining the duration of travel. In the NTS data I restricted the sample to non-business trips with a maximum trip length of one month. Within these, domestic trips lasted for an average of 5.7 days, while transborder trips were, on average, 7.7 days long. Thus, travelers make somewhat shorter trips when flying domestically, and are more likely to stay with friends or family when doing so, as compared to when making transborder trips.

There is also evidence that transborder travelers engage in activities that require more advance planning, or for which ticket prices can rise steeply at the last minute. Table 10 shows selected activities of air travelers who reported making non-business trips. Transborder travelers are significantly more likely to attend sporting events, plays, concerts, or visit theme parks than are domestic travelers.[24] Moreover, while the survey measures activities for all domestic and transborder flyers, it is quite possible that North Air's passengers are even more likely to engage in costly leisure activities on transborder trips than the average flyer.[25] This is yet another reason why transborder leisure trips on North Air are more likely to be planned well in advance.

These differences are not surprising, but also reveal a lot about the nature of planned versus last-minute travel. It is complex, and expensive, for Canadians to make last-minute leisure travel to the U.S. because of the other expenses involved— such as the cost of hotel accommodation, tickets for performances or theme parks, and accommodation on cruises. Therefore, transborder travelers who buy relatively close to the date of travel are disproportionately likely to be business rather than leisure travelers, and these are the ones who are more likely to buy the higher quality tickets.

To be sure, even for domestic travel the share of high-quality tickets rises closer to the travel date. But a relatively higher fraction of later domestic purchases are of lower-quality tickets, likely purchased by leisure travelers, as shown in Figure 5. This makes intuitive sense—a last-minute decision to make a domestic leisure trip is likely to be expensive only due to the cost of the airline ticket, since such travelers will most likely visit friends or family, and their other expenses are also likely to be lower.

Indeed, another piece of evidence supporting this theory comes from examining the accommodation expenses of domestic versus transborder travelers. The NTS asks respondents for the amounts spent on various categories while on domestic and international trips. My calculations show that 23% of domestic trips, but just 14% of transborder trips, reported zero expenses on accommodation. These numbers are sensible given that travelers who stay with friends or family typically do not incur accommodation expenses, and that this is far more likely for domestic travel.

Even restricting the sample to trips with positive accommodation expenses lends further support to this account. Table 11 presents the results from regressing the log of accommodation expenses on a dummy variable for whether the trip was domestic. The regression sample was restricted to the 7,256 domestic or transborder trips by air that were not for business reasons and for which reported spending on accommodation spending was strictly positive (as otherwise the log would be undefined). After applying survey weights, these

[24]The NTS asks about 34 possible activities that travelers may have engaged in on their trip. Many of the activities have similar percentages. Categories that were more common for domestic leisure travelers than transborder ones include visiting friends or family, hiking and shopping.

[25]For example, a common activity of tourists to New York, which is served by North Air, is to watch a Broadway play, which often involves an advance purchase. By contrast, going to the beach is a common holiday activity that does not require advance planning, but North Air does not serve destinations like Florida and California, though these are busy routes for other airlines.

Table 11: Regression of Log(Accommodation Spending)

	(1)	(2)	(3)
Domestic	-0.538***	-0.549***	-0.352***
	(0.04)	(0.04)	(0.04)
Log(Duration)			0.645***
			(0.03)
Constant	6.719***	6.732***	5.614***
	(0.03)	(0.04)	(0.06)
Year FE	No	Yes	Yes
Quarter FE	No	Yes	Yes
R^2	0.074	0.078	0.318
Obs	7256	7256	7256
Population Size	8713813	8713813	8713813

Sample restricted to domestic and transborder non-business trips with positive spending on accommodation.

represent around 8.7 million trips in the overall population.

The results of Table 11 indicate that accommodation expenses are as much as 55% lower for travelers making domestic trips compared to transborder trips. Even when controlling for the fact that transborder trips tend to be longer, domestic trips are still around 35% cheaper. This is at least partially driven by generally lower hotel and resort prices in Canada than in the U.S., but perhaps also because Canadians making transborder leisure trips are probably more likely to 'splurge' by, for example, staying at upscale or luxury resorts or at theme parks. Either way, this provides further evidence that transborder trips are more costly and involve more planning, and are therefore less likely to be undertaken at short notice.[26]

To summarize, this section has established the following: first, Canadians are far more likely to take domestic airline trips in order to visit friends and family than for any other reason, while transborder airline trips are unlikely to involve such personal visits. Second, domestic trips are shorter than transborder trips, on average. Third, domestic trips tend to be less expensive, mainly because domestic travelers tend to stay with friends or family, but also because accommodation in Canada is likely cheaper than in the U.S. Finally, transborder travelers are considerably more likely to engage in activities that require advance planning.

Taken together, this implies that transborder travel involves significantly more organization and expense, and therefore is less likely be undertaken at short notice, compared to domestic travel. This lines up well with the results of Figure 5, showing that Adjustable

[26]Remember that the focus of this analysis is on travel *by air*. Road travel between the U.S. and Canada is common, and often occurs for spur of the minute shopping trips that are usually same-day trips across the border (Chandra et al., 2014).

tickets, which are most likely to be bought by business travelers, are disproportionately more likely to be sold at the last minute on transborder trips rather than on domestic trips.

Overall, the evidence from the NTS explains the mechanism driving the main empirical result of the paper, namely that the arrival rate of consumer types is very different on domestic and transborder routes. Specifically, higher willingness-to-pay consumers are overrepresented among later purchases on transborder routes, compared to domestic routes, as shown in the North Air data. This then results in the patterns of price discrimination that were predicted by the model in Section 4, and empirically documented in Section 5: a convergence of high- and low-quality fares in the periods leading up to travel on domestic routes, but a divergence on transborder routes.

7 Discussion and Conclusion

Price Discrimination is a central and heavily researched topic in Industrial Organization, including both theoretical and empirical lines of inquiry, and focusing on the two major strands of second-degree and third-degree discrimination. For reasons related to complexity and data availability, however, prior authors in the literature have generally focused on just one or the other of these two kinds of practices. This may be adequate for situations where firms exclusively, or even predominantly, practice just one type of discriminatory pricing. However, with the availability of large new data sources on the behaviour, habits and preferences of consumers, it is increasingly likely that firms will practice both simultaneously. It is far from obvious that the results of partial analysis of each type of price discrimination will extend to environments where firms practice both. Past research provides no guide to whether we should expect the two types to be substitutes or complements for each other, or whether they may be completely independent.

In this paper I examine an industry that has long used multiple price discrimination practices—airlines. I develop a model which shows how the arrival rate of different consumer types in different markets can lead to different forms of price discrimination either reinforcing or offsetting each other. I then obtain new booking data directly from an airline that, to my knowledge, are the first of their kind to be used in academic research. I use these to examine both inter-temporal price discrimination according to advance purchase behaviour, and menu-based price discrimination according to the quality of service that travelers select. I find that advance purchase gradients clearly exist, for both low and high quality tickets, and that these are due to price discrimination rather than simply the optimal management of inventory. I show that there appears to be a tradeoff between the two kinds of price discrimination in domestic airline markets markets, but a complementarity between them in transborder markets.

I further show that the market shares of high- and low- quality tickets sold in equilibrium line up with the model's predictions regarding different kinds of markets. Finally, I use new survey data to examine the motivation of airline travellers to make certain kinds of trips. I show that leisure travelers are make more complex, and expensive trips when traveling transborder, compared to domestically, which makes last-minute purchases on these routes unlikely. This fact explains why high-quality tickets are over-represented in later periods on transborder routes, and fully reconciles both the predictions of the model and the observed empirical results.

There are, naturally, some caveats to my results. Mainly, that I use data from a single airline, which is by no means representative of the entire airline industry.[27] Nevertheless, the airline operates on competitive routes and therefore its prices, and the dynamic evolution of its fares, are likely to be similar to those of its rivals in equilibrium.

The main result of this paper—regarding the interaction of two different forms of price discrimination, is novel and has not been established before in other settings. This naturally raises the question of whether the results will extend to other industries that also practice multiple forms of discriminatory pricing. A few examples come to mind: book publishers, pharmaceuticals, and retail banks. In these examples, firms may offer different prices to different groups of consumers, such as due to country-specific pricing, or discounts for seniors or students. At the same time, firms in these industries offer a range of service levels corresponding to different qualities. It would be interesting to study whether the findings of this paper extend to such settings.

References

Alderighi, M., M. Nicolini, and C. A. Piga (2015). Combined effects of capacity and time on fares: insights from the yield management of a low-cost airline. *Review of Economics and Statistics 97*(4), 900–915.

Armstrong, M. and J.-C. Rochet (1999). Multi-dimensional screening:: A user's guide. *European Economic Review 43*(4), 959–979.

Aryal, G., C. Murry, and J. W. Williams (2018). Price discrimination in international airline markets. *Available at SSRN*.

[27]Of course, it is common for research in this area to make similar allowances. For example, Alderighi et al. (2015) also use data from a single airline—Ryanair; Williams (2020) focuses mainly on JetBlue, and that too on monopoly routes. Note that a companion paper, Chandra (2021) examines the effect of competition in this setting.

Borzekowski, R., R. Thomadsen, and C. Taragin (2009). Competition and price discrimination in the market for mailing lists. *QME 7*(2), 147–179.

Chandra, A. (2021). The effects of competition on multiple forms of price discrimination. *Available at SSRN*.

Chandra, A., K. Head, and M. Tappata (2014). The economics of cross-border travel. *Review of Economics and Statistics 96*(4), 648–661.

Chandra, A. and M. Lederman (2018). Revisiting the relationship between competition and price discrimination. *American Economic Journal: Microeconomics 10*(2), 190–224.

Clark, R. and N. Vincent (2012). Capacity-contingent pricing and competition in the airline industry. *Journal of Air Transport Management 24*, 7–11.

Dana, J. and K. Williams (2020). Intertemporal price discrimination in sequential quantity-price games.

Dana, J. D. (1998). Advance-purchase discounts and price discrimination in competitive markets. *Journal of Political Economy 106*(2), 395–422.

Dana, J. D. (1999). Equilibrium price dispersion under demand uncertainty: the roles of costly capacity and market structure. *The RAND Journal of Economics*, 632–660.

Gale, I. L. and T. J. Holmes (1993). Advance-purchase discounts and monopoly allocation of capacity. *The American Economic Review*, 135–146.

Lazarev, J. (2013). The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in us monopoly markets. *New York University*.

Leslie, P. (2004). Price discrimination in broadway theater. *RAND Journal of Economics*, 520–541.

Lin, H. and I. Y. Wang (2015). Competition and price discrimination: evidence from the parking garage industry. *The Journal of Industrial Economics 63*(3), 522–548.

Möller, M. and M. Watanabe (2010). Advance purchase discounts versus clearance sales. *The Economic Journal 120*(547), 1125–1148.

Puller, S. L., A. Sengupta, and S. N. Wiggins (2009). Testing theories of scarcity pricing in the airline industry. Technical report, National Bureau of Economic Research.

Rochet, J.-C. and L. A. Stole (2002). Nonlinear pricing with random participation. *The Review of Economic Studies 69*(1), 277–311.

Shepard, A. (1991). Price discrimination and retail configuration. *Journal of Political Economy 99*(1), 30–53.

Sweeting, A. (2012). Dynamic pricing behavior in perishable goods markets: Evidence from secondary markets for major league baseball tickets. *Journal of Political Economy 120*(6), 1133–1172.

Williams, K. R. (2020). Dynamic airline pricing and seat availability. Technical report, mimeo.

8 Appendix: Additional Tables

Table 12: Daysout Gradients by Load Factor

	Full	Q1	Q2	Q3	Q4
	(1)	(2)	(3)	(4)	(5)
γ (Gradient)	-0.247***	-0.258***	-0.265***	-0.248***	-0.215***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Constant	5.188***	5.128***	5.234***	5.267***	5.276***
	(0.03)	(0.03)	(0.03)	(0.04)	(0.05)
R^2	0.428	0.444	0.460	0.441	0.376
Obs	518530	130008	128254	130201	130067

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Sample restricted to Fixed Fares. All regressions include route, day-of-week, month and year FEs. Standard errors, clustered by route, in parentheses.

Table 13: Daysout Elasticities

Route (to and from Toronto)	Fixed Fares		Adjustable Fares	
	γ	N	γ	N
Montreal	-0.251	122634	-0.054	20359
Ottawa	-0.270	130315	-0.076	20334
Thunder Bay	-0.193	36192	-0.083	2648
New York	-0.235	91906	-0.250	25722
Boston	-0.254	31067	-0.489	11717
Chicago	-0.169	26630	-0.483	12281
All Domestic	-0.246	420246	-0.064	58656
All Transborder	-0.226	158962	-0.379	51093
All Routes	-0.244	579208	-0.217	109749

Note: γ is the estimated price elasticity with regard to days remaining until travel, for the corresponding route(s) and fare type. Estimating regressions include route, day-of-week, month and year FEs. Sample restricted to purchases made within 90 days of travel.

Table 14: Regression of Fares on Advance Purchase Days: Domestic and Transborder Routes

	Domestic	Transborder
	(1)	(2)
Adjustable: 0 to 1 days	1.212***	1.728***
	(0.04)	(0.08)
Adjustable: 2 to 4 days	1.153***	1.589***
	(0.05)	(0.09)
Adjustable: 5 to 7 days	1.074***	1.454***
	(0.05)	(0.10)
Adjustable: 8 to 14 days	1.028***	1.124***
	(0.05)	(0.05)
Adjustable: 15 to 21 days	0.977***	0.684***
	(0.05)	(0.05)
Adjustable: 22 to 30 days	0.976***	0.555***
	(0.06)	(0.06)
Adjustable: 31 to 60 days	0.916***	0.436***
	(0.06)	(0.04)
Adjustable: 60+ days	0.876***	0.304***
	(0.07)	(0.02)
Fixed: 0 to 1 days	0.804***	0.749***
	(0.03)	(0.05)
Fixed: 2 to 4 days	0.749***	0.686***
	(0.03)	(0.05)
Fixed: 5 to 7 days	0.617***	0.683***
	(0.04)	(0.03)
Fixed: 8 to 14 days	0.419***	0.452***
	(0.04)	(0.01)
Fixed: 15 to 21 days	0.173***	0.182***
	(0.03)	(0.01)
Fixed: 22 to 30 days	0.069***	0.091***
	(0.02)	(0.01)
Fixed: 31 to 60 days	0.022*	0.039***
	(0.01)	(0.01)
Constant	4.246***	4.284***
	(0.06)	(0.06)
R^2	0.554	0.636
Obs	508931	229573

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions include route, day-of-week, month and year FEs. Standard errors, clustered by route, in parentheses.

Table 15: Regression of Fares on Advance Purchase Days: Capacity Measures

	Fixed Fares			Adjustable Fares		
	(1)	(2)	(3)	(4)	(5)	(6)
0 to 1 days	0.812***	0.813***	0.816***	0.919***	0.920***	0.924***
	(0.02)	(0.02)	(0.02)	(0.11)	(0.11)	(0.11)
2 to 4 days	0.756***	0.757***	0.761***	0.869***	0.871***	0.876***
	(0.03)	(0.03)	(0.03)	(0.11)	(0.11)	(0.11)
5 to 7 days	0.639***	0.640***	0.643***	0.778***	0.780***	0.785***
	(0.03)	(0.03)	(0.02)	(0.11)	(0.11)	(0.11)
8 to 14 days	0.433***	0.435***	0.438***	0.618***	0.619***	0.623***
	(0.03)	(0.03)	(0.02)	(0.07)	(0.07)	(0.07)
15 to 21 days	0.180***	0.181***	0.185***	0.381***	0.382***	0.385***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
22 to 30 days	0.080***	0.081***	0.084***	0.278***	0.279***	0.283***
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
31 to 60 days	0.030***	0.031***	0.033***	0.141***	0.142***	0.143***
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
Remaining Seats (%)		-0.066***			-0.074***	
		(0.01)			(0.02)	
Load Factor (%)			0.145***			0.176***
			(0.02)			(0.03)
Constant	4.237***	4.285***	4.163***	4.743***	4.795***	4.646***
	(0.04)	(0.05)	(0.04)	(0.12)	(0.12)	(0.11)
R^2	0.434	0.435	0.437	0.348	0.349	0.351
R^2 excl. Daysout	0.153	0.153	0.155	0.160	0.161	0.162
Obs	589704	589704	589704	108207	108207	108207

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions include route, day-of-week, month and year FEs. Standard errors, clustered by route, in parentheses.