Price Discrimination along Multiple Dimensions: Theory and New Evidence*

Ambarish Chandra†
University of Toronto

May 18, 2023

Abstract

This paper examines firms that simultaneously practice quality-based and intertemporal price discrimination. A model illustrates how the quality premium charged by a price discriminating firm can increase, decrease or remain unchanged as time elapses, depending on the distribution of consumer types across time. Data from a regional airline reveal that, in some markets, the quality premium increases as time to departure falls but the opposite is true in others. New survey data illustrate how the result is driven by differences in the motivation and characteristics of travelers in each market, which fully reconciles the predictions of the model and the empirical results.

Conflict of Interest: The author declares no conflicts or competing interests.

*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

In most industries, consumers vary in their willingness-to-pay for goods in both observed and unobserved ways. Firms commonly exploit this heterogeneity using sophisticated screening methods to practice price discrimination. But consider what happens if a firm has multiple instruments to screen consumers. Does one instrument dominate the others, which is then used exclusively? Or might it be profitable to use multiple screens simultaneously, such as by offering different quality levels while also charging a premium to certain groups?

Price discrimination is, of course, a central topic in both Economics and Marketing. 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. 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. 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.

This paper examines a setting where consumers' willingness-to-pay varies along two dimensions—quality and time—which permits firms to practice two distinct forms of price discrimination. A stylized model depicts firms that implement both kinds simultaneously. Depending on the distribution of consumer types, relative price differences within a menu can either increase or decrease over time. Specifically, as prices rise due to intertemporal price discrimination, some markets can see the premium for higher quality grow, while others will see it shrink. The paper then empirically establishes that there are indeed distinct markets where each phenomenon is observed, and these differences are consistent with the type and behavior of consumers.

It is well established that many individual industries price discriminate in multiple ways. One example is the market for books, which exhibits quality discrimination based on paperback and hardback versions, as well as price discrimination across countries for identical versions of the same book.[1] The market for prescription drugs is characterized by both cross-country price differences as well as differences between generics and branded drugs within the same country.[2] As a final example, the market for new cars price discriminates based on observable characteristics such as race and gender as well as country pricing.[3] It also exhibits quality discrimination, based on different trim levels for the same model.

[1] Clerides (2002) documents the former and Cabolis et al. (2007) and Boivin et al. (2012) show the latter.

[2] Schut and Van Bergeijk (1986) and Frank and Salkever (1997), are some examples of these phenomena.

[3] Verboven (1996) documents cross-country price discrimination, and Ayres and Siegelman (1995) and Chandra et al. (2017) establish discrimination based on race, gender and age.

In none of these industries, or any others, has prior research examined whether there is a relationship between multiple forms of price discrimination. It would be interesting to know, for example, if countries with higher average prices for cars, books and drugs also exhibit a greater quality premium for such goods. Knowing this would help explain whether it is profitable for firms to employ multiple consumer screens simultaneously, or whether a single price discrimination instrument can suffice.

My analysis models a firm which 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 key intuition of the model is that if later-arriving consumers are also more willing to pay for quality on average, then the quality premium will increase in conjunction with intertemporal price discrimination. Moreover, in such a situation the equilibrium market share of the high-quality version will increase over time. However, the opposite will be true, for both quantities and prices, if later-arriving consumers have a lower taste for quality.

I then obtain booking data from a regional airline based in Canada that practices both intertemporal and quality-based price discrimination, and which operates on both domestic routes as well as transborder routes to the United States. The results show that the airline price discriminates intertemporally in a remarkably similar manner across all routes in the data. However, as prices rise due to intertemporal price discrimination, the premium for higher-quality service increases on some routes but decreases on others, a result that is entirely consistent with the model.

Next, I 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 distribution of 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 provide an explanation for the results. I obtain new survey data on the characteristics of travelers and their motivation to make different kinds of airline travel. The data convincingly establish that leisure travelers on domestic routes are far more likely to visit friends or family, and incur lower costs of travel. 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.

Few previous empirical papers have studied settings where firms practice multiple forms of price discrimination. One example is Leslie (2004), who 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. However, Leslie does not examine whether each kind of discriminatory pricing is affected by the other.[4]

Aryal et al. (2022) also examine inter- and intra-temporal price discrimination by airlines. They 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. Aryal et al's main focus is on estimating welfare effects from various kinds of price discrimination, while the emphasis of this paper is on relative price differences within and across menus offered on different airline routes.

The empirical literature on menu pricing has focused mainly on quantity discounts. While quantity discounts are theoretically the same as offering varying quality levels, there are 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 price discrimination based on quality differences is by 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, as these amenities do not change in response to demand or competition.

Though 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 optimal inventory management. Specifically, when a firm faces uncertain demand, or when

[4] Some studies have examined 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 competition affects the *type* of price discrimination that firms use (Borzekowski et al., 2009).

demand is likely to exceed capacity, prices can rise over time even in the absence of price discrimination. This question is investigated more thoroughly in a companion paper, Chandra (2023), which finds that the inter-temporal variation in fares appears to be almost entirely explained by price discrimination rather than by the limited inventory of seats.

The airline industry is a setting where firms have long employed multiple types of discriminatory pricing to extract surplus. However, the available data limit our understanding of the relationship between multiple types of price discrimination. 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 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.[5]

My data source is a private airline based in Toronto, which I will refer to as North Air.[6] Though small compared to most well known airlines, North Air is an important player on the set of routes on which it operates, which are mainly to large cities in the Eastern half of Canada and the United States. 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.[7]

This paper is organized as follows. Section 2 presents a stylized model of price discrimination along multiple dimensions. Section 3 presents the new airline data source used in this paper. The empirical results are presented in Section 4. Section 5 explains the mechanism behind the main empirical result by incorporating new survey data. Section 6 concludes.

[5]For example, Chandra and Lederman (2018) examine price dispersion in airlines based solely on the observed distribution of fares. Among papers that calculate the temporal variation in fares, Alderighi et al. (2015) scrape airfares from the website of Ryanair, Williams (2022) uses data from a search engine and Lazarev (2013) obtains published fares from a Global Distribution System.

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

[7]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.

2 Model

In this section I present a model of a firm that sets price menus over time, to a changing composition of consumers. The model will make predictions about relative prices and quantities in a situation where a firm can screen consumers on the basis of both the time of booking, which is an observable consumer characteristic, and willingness-to-pay for quality, which is an unobserved characteristic.

It is well known models of multi-dimensional screening can be difficult to solve (Armstrong and Rochet, 1999). 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 for a given flight increases, on average, over time. Moreover, this fact is known to consumers themselves, and therefore they correctly anticipate 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 two kinds of price discrimination. Moreover, this assumption also fits well with the empirical results of Section 4, which show that remaining inventory does not affect price discrimination, at least for the airline I study.

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 tickets, 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 versions. I assume that $c_H - c_L < \phi - 1$. In other words, the marginal cost increase of providing high-quality over low-quality must be less than the marginal willingness-to-pay for high-quality versus low-quality.

In each period t, let $\theta_1(t)$ denote a consumer who is indifferent between the low quality seat and not purchasing 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 the argument, 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:

$$(p_H - c_H)f(\theta_2) + (\phi - 1)[F(\theta_2) - F(\theta_1)] = (p_L - c_L)[f(\theta_2) + f(\theta_1)(\phi - 1)] \qquad (1)$$

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:

$$(p_L - c_L)f(\theta_2) + (\phi - 1)[1 - F(\theta_2)] = (p_H - c_H)f(\theta_2) \qquad (2)$$

Combining the two first order conditions, (1) and (2), we get:

$$p_L = c_L + \frac{[1 - F(p_L)]}{f(p_L)} \tag{3}$$

and

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

As a benchmark, consider the case where $F(\theta)$ is uniform in each period. Then, equations 3 and 4 imply the following optimal prices:

$$p_L(t) = \frac{\bar{\theta}(t) + \max[\underline{\theta}(t), c_L]}{2} \tag{5}$$

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

The interesting question is whether the premium for high-quality service changes over time. Examining equations 5 and 6 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. Specifically, I allow the distribution of θ to be different in Periods 1 and 2, so that 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 4, which shows that the relative premium for higher quality is a decreasing function of the hazard rate of the distribution of consumer types. □

Figure 1 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

Figure 1: Densities, Survival Functions and Hazard rates

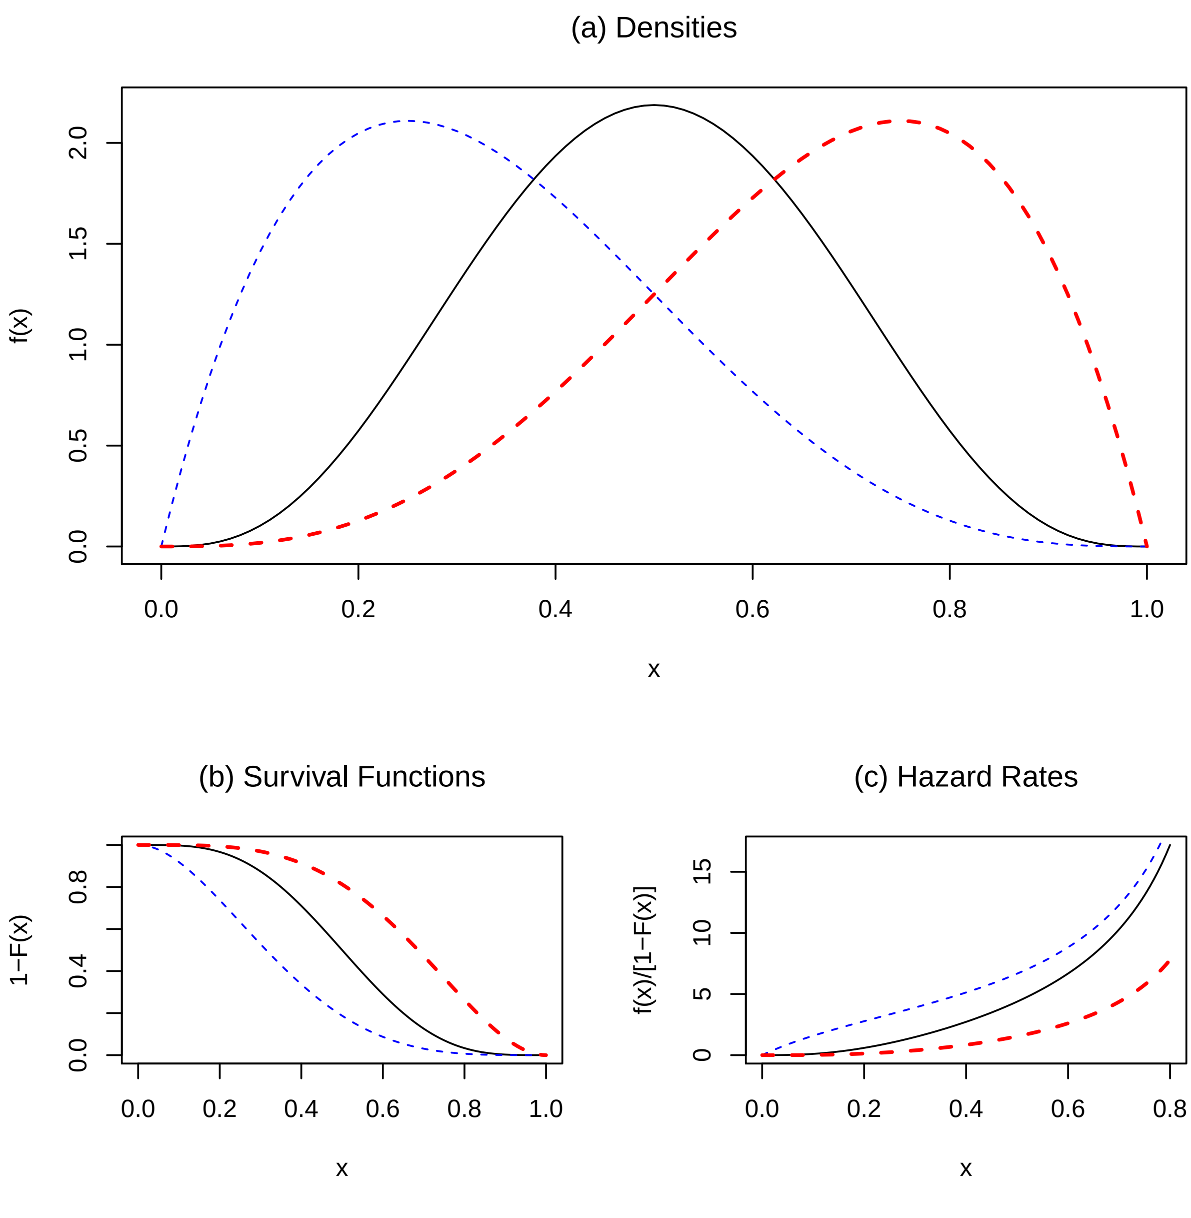

and (thicker) dashed red curves depict right- and left-skewed Beta distributions respectively. A right-skewed (or positive-skewed) Beta distribution, for example, has more weight on lower values. Panel (b) of the figure shows the survival functions for each distribution, which are just the complements 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 distribution of consumer types 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 will see the premium for higher quality grow over time. By contrast, if later arriving consumers 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 distribution of consumer types, 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. 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 3 and 4 imply:

$$8p_L^3 - 9p_L^2 + 1 = 0 \tag{7}$$

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 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

Table 1: 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	$0.64(\phi-1)$

$F(\theta) = 4\theta^3 - 3\theta^4$.

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 2 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 version in each case. Table 1 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 more consumers with a higher willingness-to-pay arriving in the second period. Despite this, the equilibrium share of consumers who purchase the high-quality version rises, to 58%.[8] Analogously, in Case 3, both the relative premium and the share of consumers purchasing the high-quality version drop, compared to Period 1.

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

Figure 2: The distribution of consumer types and the evolution of the price premium over time

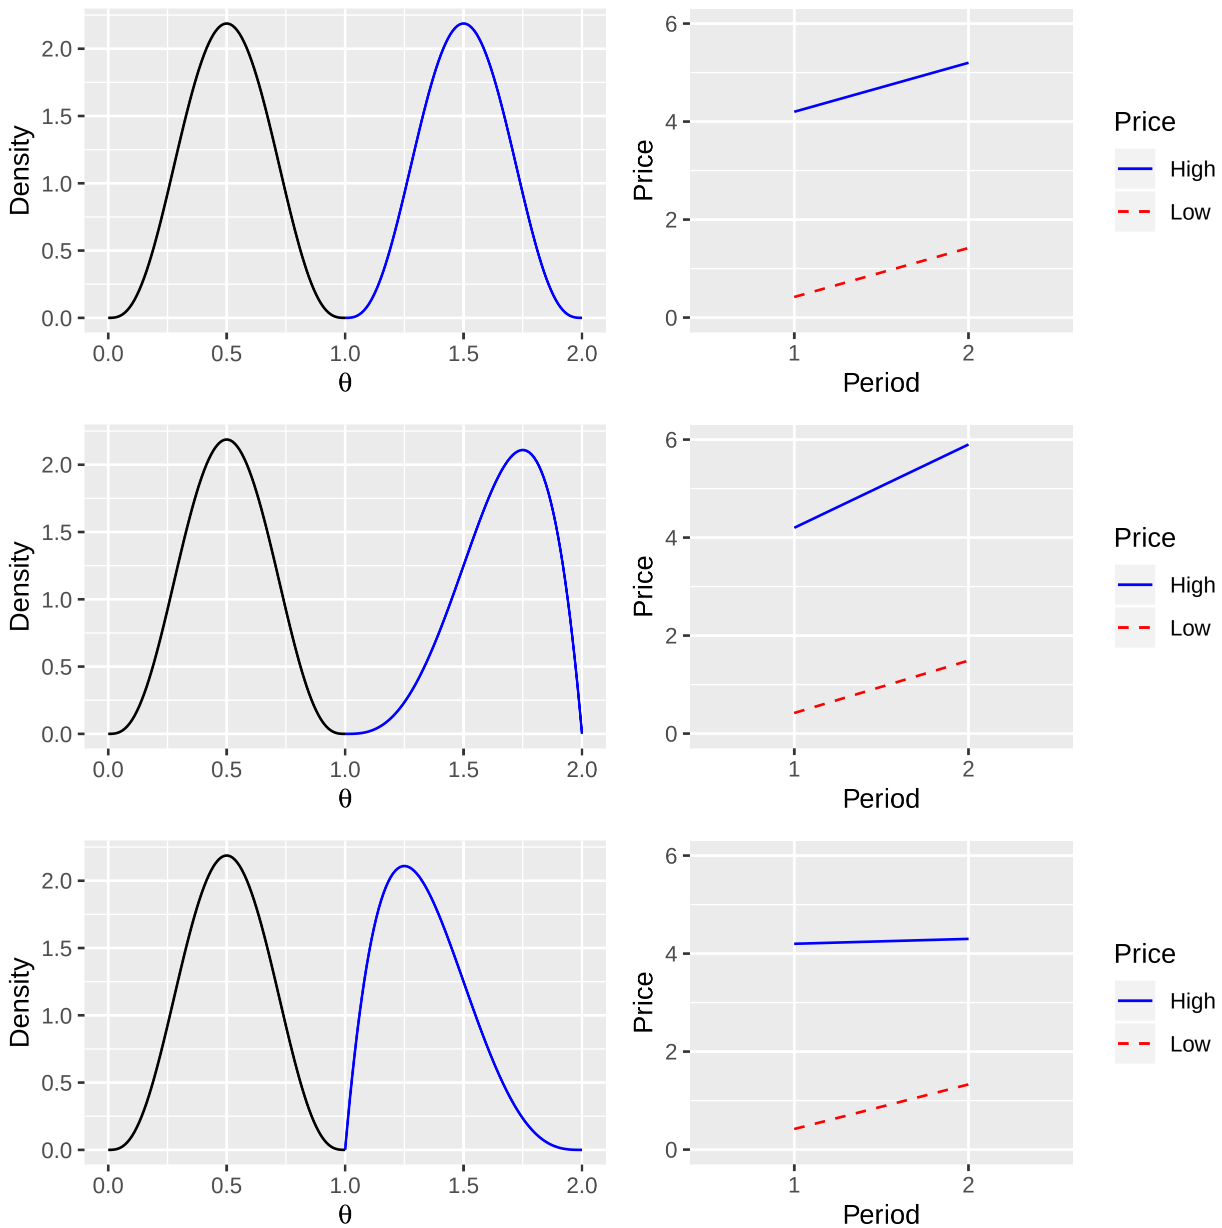

To summarize, this model has implications for the evolution of fares over time, which will depend on the distribution of various consumer types. Markets where relatively higher consumer types arrive later will see a greater divergence between the prices for low and high quality levels. By contrast, markets where relatively lower 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.

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 a 10% random sample of all its bookings, with information on fares, class of service. the number of passengers and the dates of both reservation and travel. [9] This allows me to calculate the length of the advance purchase period.

This data source represents a significant advance over prior studies that infer or construct measures of advance purchase discounts. The commonly used DB1B dataset of the U.S. Department of Transportation lacks a key piece of information on 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 a 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.

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, by dropping observations that involved complicated multi-city itineraries with three or more separate journeys. Second, by dropping itineraries where the fare group was listed as "Other", i.e. not one of the three standard fare groups. Finally, I dropped observations where the recorded fare for the journey was less than $20 or greater than $1300. The final dataset consists of 865,492 observations, for travel

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

Table 2: 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

between 19 airports, during the years 2008–2014.

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.[10] There are clear quality differences between these ticket types. Adjustable ticket-holders can switch to earlier or later flights on the day of travel, unlike Fixed ticket holders who must pay a fee to do so. Adjustable tickets also have other perks such as free seat selection and free checked bags. Refundable tickets have all the characteristics of Adjustable ones, in addition to being fully refundable. 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.

Table 2 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.

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 flies a single type of plane, and offers a single service cabin, with no distinction between leisure and business classes.[12]

[10] This is unlike other airlines, Southwest Airlines in particular, where the cheapest fares often sell out in advance. North Air does occasionally sell out Fixed fares just prior to travel but that is the exception, not the norm. Figure 3 shows that Fixed Fares are sold at all times including the day of travel.

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

[12] 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.

Figure 3: Distribution of DaysOut

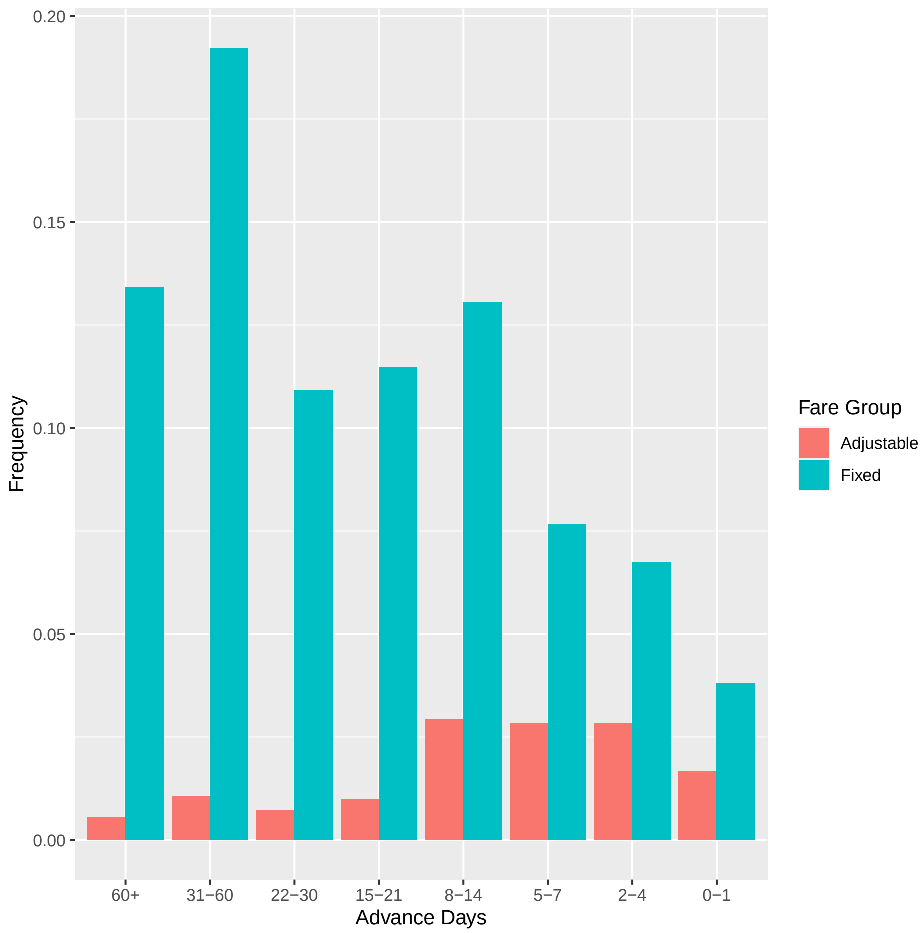

Thus, in my analysis I can ignore the common type of second-degree price discrimination—between Coach and Business class cabins—practiced by most legacy carriers. Additionally, it does not follow the usual airline practice of offering different fare 'buckets', corresponding to different sets of restrictions, at varying prices. The legacy carriers implement this practice because prices and the quantity of sales are chosen by separate teams (Hortaçsu et al., 2021). 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. This also means that the airline does not price discriminate along other dimensions that are commonly used by legacy carriers, such as length of stay, originating airport or so-called 'hidden-city' ticketing.

Figure 3 shows the detailed distribution of daysout, separately for the two main fare groups. 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.[13]

4 Results

This Section will first show the existence of clear advance purchase gradients, for both quality levels. It will then present the main empirical result: that different markets exhibit strikingly different patterns of quality-based and intertemporal price discrimination.

4.1 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. 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.

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 quality tickets. 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]

It is easier to see the estimated gradient visually. Figure 4 depicts advance purchase

[13]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.

[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.

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.

gradients for Fixed and Adjustable fares, for different groups of routes. These gradients are estimated, for each group of routes, by a single, pooled regression for all fares.[15] In each panel, the coefficient for Fixed fares purchased more than two months in advance is normalized to zero, and all other coefficients are expressed relative to those fares.

The top panel of Figure 4 plots the coefficients for the full sample of routes. The figure shows that fares rise monotonically over time, and prices for the two fare classes appear to rise in tandem suggesting that the relative premium for higher quality does not change over time.

However, this masks underlying heterogeneity. Panels (b) and (c) of Figure 4 plot coefficients separately for domestic travel and for transborder travel to the United States. These figures tell a different story from the one shown in Panel (a). As before, both Fixed and Adjustable fares rise monotonically as the date of travel approaches. Moreover, for both domestic and transborder routes, the Fixed fare gradients follow a relatively similar shape,

[15] Regression coefficients are presented in Table 9 in the Appendix.

Figure 4: Advance Purchase Gradients

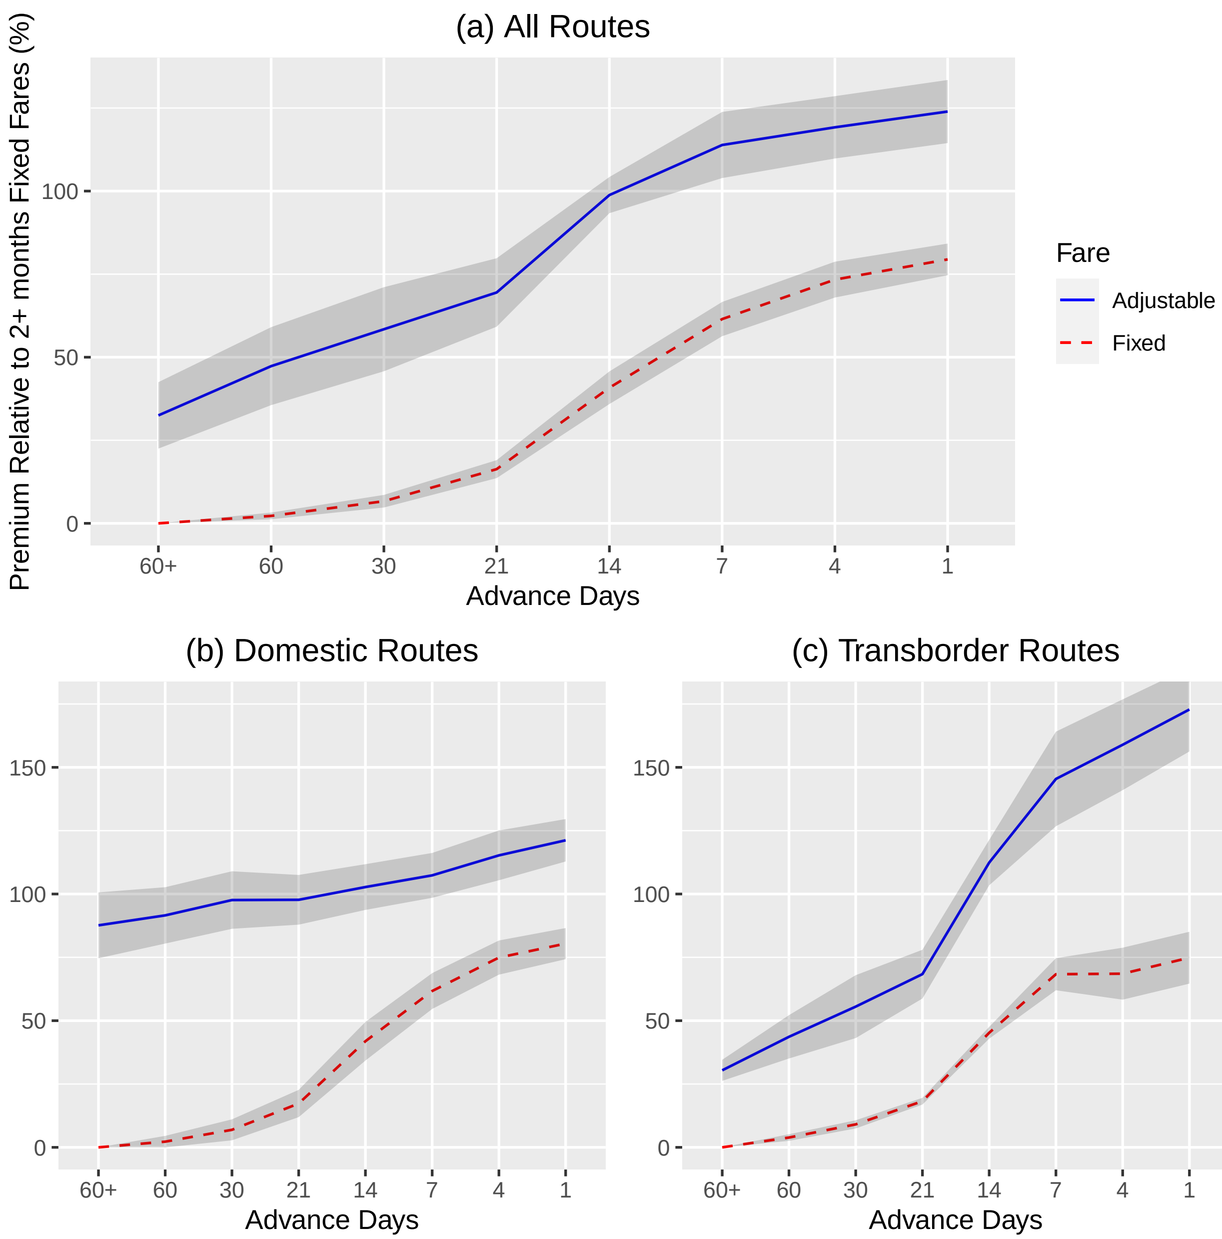

Notes: Shaded areas represent 95% confidence intervals, using standard errors clustered by route.

with tickets purchased within one day of travel being about 110% more expensive than those bought far in advance.[16] Even though the coefficients for tickets bought in the last 2–3 days before travel appear to look dissimilar across panels (b) and (c), 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.[17]

By contrast, the gradients for Adjustable fares differ markedly across the two bottom panels. On domestic routes, the high quality premium appears to shrink as the travel date approaches, whereas it appears to grow on transborder routes. As a result, Fixed and Adjustable fares converge in the former, but diverge in the latter. I now turn to a deeper explanation of this finding, which constitutes the main empirical contribution of the paper.

4.2 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 relationships between intertemporal and quality-based price discrimination.

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

$$P_{it} = \alpha_i D_t^{\gamma_i} \tag{8}$$

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 8 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.

In Table 4, I regress fares on both measures of discriminatory pricing together—the

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

[17] These confidence intervals widen because, on transborder routes, relatively few Fixed tickets are sold in the last few days, likely because international travelers generally need to make plans further in advance.

Table 4: Regression of Fares on Both Kinds of Price Discrimination

	All Routes			Transborder	Domestic
	(1)	(2)	(3)	(4)	(5)
Ln(Days)	-0.303***		-0.241***	-0.210***	-0.242***
	(0.01)		(0.01)	(0.01)	(0.01)
Adj. Fares		0.832***	0.574***	1.155***	0.272***
		(0.04)	(0.11)	(0.16)	(0.04)
Ln(Days)*Adj.			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
Base R^2	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.

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 based on the R^2 in column 1, while the Adjustable fare dummy on its own 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 decompose 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 Figure 4.

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. 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 5, I summarize these elasticities using a 60-day window of purchase, but I show in Table 10 in the Appendix that the results are almost identical using

Table 5: Daysout Elasticities

Route (to and from Toronto)	Fixed Fares		Adjustable Fares	
	$\hat{\gamma}$	N	$\hat{\gamma}$	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: $\hat{\gamma}$ 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.

a 90-day window.[18] 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 5 are mostly very similar to each other, suggesting 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 $\hat{\gamma}$ 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. The 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 intertemporal 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. By contrast, the elasticity for Adjustable fares on the three busiest transborder routes, and for transborder routes in general,

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

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

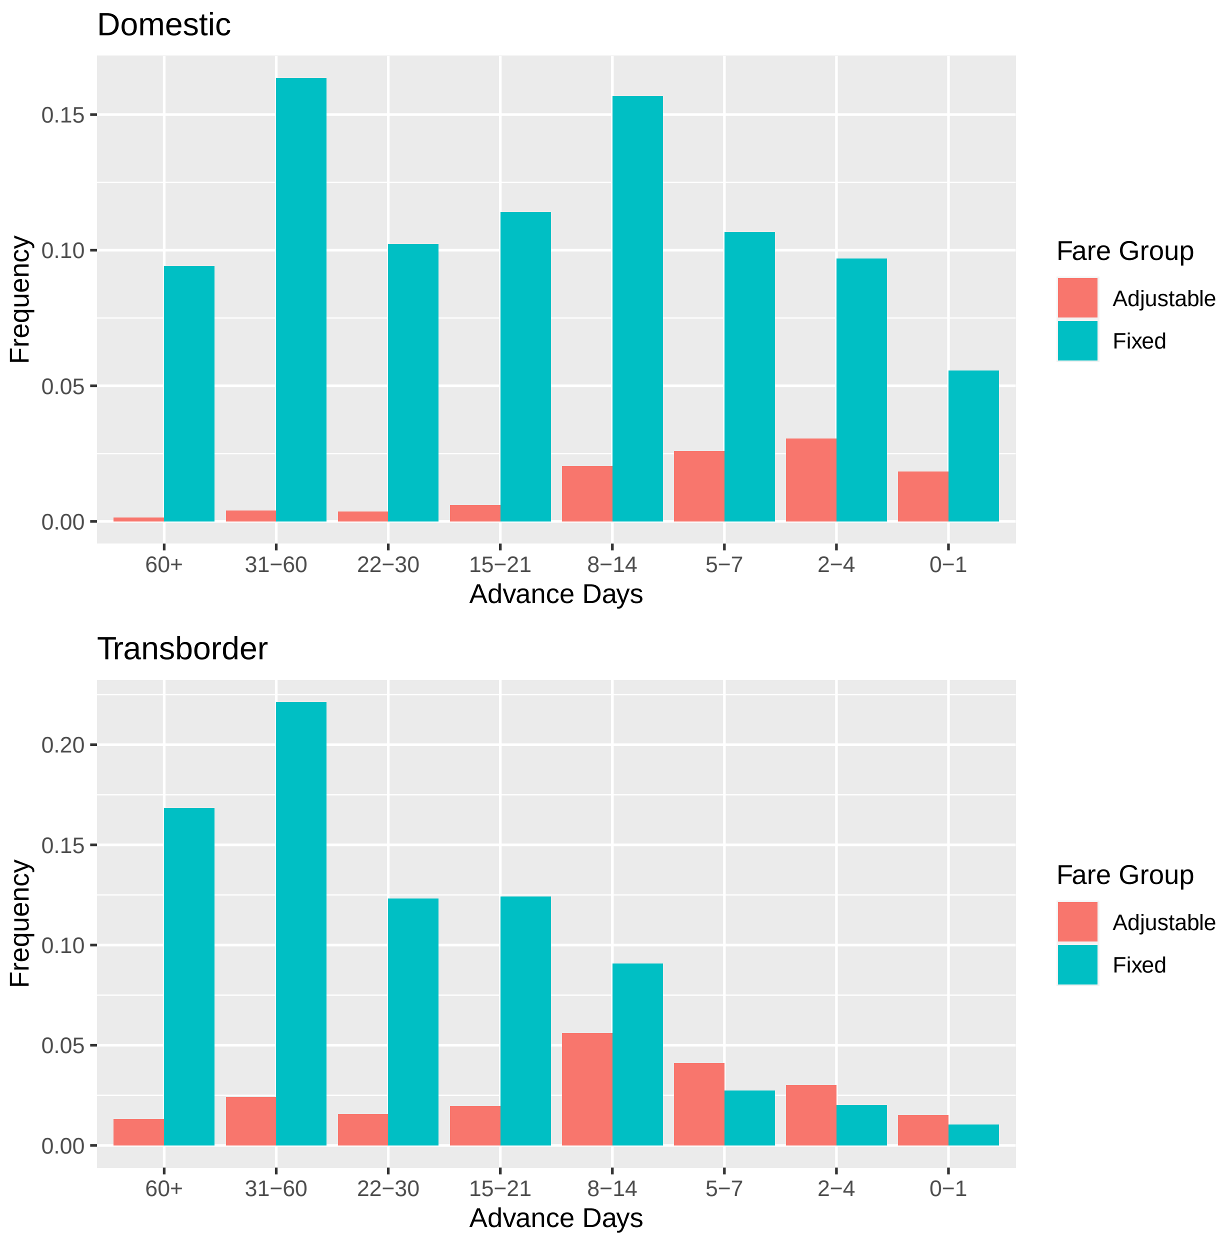

is considerably greater; as high as, or even higher than, the elasticity for Fixed Fares. Once again, these results are consistent with those shown in Figure 4 and suggest that travelers on US–Canada routes who self-select into expensive tickets are also subject to extensive price discrimination by time of booking.

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.

This result is shown more formally in Appendix Table 11, which presents probit regressions of the probability that a ticket is high-quality on measures of advance purchase duration, separately for domestic and transborder routes. It shows 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. Moreover, 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, Table 5, Figure 5 and Appendix Table 11 show results fully consistent with the model presented in Section 2: that it is possible for the low- and high-quality prices to either diverge or converge over time, *i.e.* that intertemporal and quality-based price discrimination can operate in either the same or opposite direction. 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.

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.

5 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. One hypothesis that may explain this pattern is that the composition of travelers differs in the two markets. 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

Table 6: 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 2018 and 2019 National Travel Survey.

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.[19] 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 6 shows the main reason for taking domestic and transborder trips, using data from the NTS.[20] The results show that domestic travel is somewhat more likely (37%) to be for business than is international travel (28%). Within the categories of non-business travel,

[19]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.

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

Table 7: 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.

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 7 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.[21] 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.[22] 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

[21] 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.

[22] 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 8: 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 non-business trips with positive spending on accommodation.

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 8 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 represent around 8.7 million trips in the overall population.

The results of Table 8 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.[23]

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 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 2, and empirically documented in Section 4: a convergence of high- and low-quality fares in the periods leading up to travel on domestic routes, but a divergence on transborder routes.

[23] Remember that this analysis is of 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).

5.1 Other explanations

While the explanation presented above appears to represent the most likely mechanism for the results, I briefly consider other hypotheses. One potential criticism of my analysis is that domestic travel within Canada may represent a very different kind of market than transborder travel to the U.S., both in terms of industry structure and the composition of travelers.

I first examine whether market structure plays a role. The Herfindahl-Hirschmann Index is, on average, higher on transborder routes, though this is largely due to the extremely competitive Toronto–New York route. However, controlling for the number of rivals that North Air faces on a route, or for the HHI (as measured by either seats or flights) does not change the coefficients of interest on either the domestic or transborder markets. This can be seen in Appendix Table 12, the results of which are very similar to those in Columns 5 and 6 of Table 4.[24] Of course, the results of Appendix Table 12 do not have a causal interpretation, as market structure is endogenously determined, but this is nevertheless evidence against the notion that the main results are driven by differences in competition across routes.

I then examine whether travelers are different on domestic versus transborder routes. Of course, this is partly the explanation, as I argue that Canadians travelling to the U.S. behave differently from those travelling domestically. But could the result be driven by a greater proportion of Americans on transborder routes, who perhaps behave differently from Canadians traveling domestically? The answer is no, for two reasons. First, even on transborder routes, 84% of itineraries are purchased in Canadian dollars, suggesting that the vast majority of these flyers are, in fact, Canadians.[25]

Second, I can restrict attention to roundtrip itineraries, and compare itineraries that originate in the U.S. with those that originate in Canada. For example, a flyer who travels from Toronto to New York and back is likely to be a resident of Canada, while one who flies from New York to Toronto and back is likely to be a U.S. resident. It turns out that roundtrips that originate in one country are no different from those originating in the other. This is evident from Table 13 in the Appendix, which shows that the two kinds of price discrimination, as well as their interaction, are very similar for roundtrip transborder itineraries that originate or terminate in Toronto, and also similar to the results of the full sample of transborder trips, as shown in Column 4 of Table 4.

[24]While the number of rivals on a route tends not vary much over time, the HHI measures vary considerably, even within a route, over time. This is because of airline-specific responses to demand and seasonality, airline specific fleet and capacity constraints, and North Air's own growth over time.

[25]I do not have data on the country of residence of travelers, but the currency of purchase should be a very good proxy for this as the airline's website defaults to the local currency of the booking computer.

6 Discussion and Conclusion

Price Discrimination is a central and heavily researched topic, including both theoretical and empirical lines of inquiry, and focusing on both self-selection by consumers as well as group-based pricing. For reasons related to complexity and data availability, however, the prior literature has 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.

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 operating in either the same or opposite direction. 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 intertemporal price discrimination according to how far in advance purchases are made, 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. I show that, in domestic airline markets, the change in the quality premium operates in the same direction as inter-temporal price discrimination, but in transborder markets the opposite is true.

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 make more complex and expensive trips when traveling transborder, compared to domestically, which makes last-minute purchases on these routes less likely. 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.

One caveat to these results is that I use data from a single airline, which is by no means representative of the entire airline industry.[26] 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.

[26] 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 (2022) focuses mainly on JetBlue.

The main result of this paper—regarding the relationship between 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; for example, book publishing, 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 while also offering 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 (2022). Price discrimination in international airline markets. *Available at SSRN*.

Ayres, I. and P. Siegelman (1995). Race and gender discrimination in bargaining for a new car. *The American Economic Review*, 304–321.

Boivin, J., R. Clark, and N. Vincent (2012). Virtual borders. *Journal of International Economics 86*(2), 327–335.

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

Cabolis, C., S. Clerides, I. Ioannou, and D. Senft (2007). A textbook example of international price discrimination. *Economics Letters 95*(1), 91–95.

Chandra, A. (2023). Why Do Airline Prices Rise over Time? Working Paper, University of Toronto.

Chandra, A., S. Gulati, and J. M. Sallee (2017). Who loses when prices are negotiated? an analysis of the new car market. *The Journal of Industrial Economics 65*(2), 235–274.

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.

Clerides, S. K. (2002). Book value: intertemporal pricing and quality discrimination in the us market for books. *International Journal of Industrial Organization 20*(10), 1385–1408.

Frank, R. G. and D. S. Salkever (1997). Generic entry and the pricing of pharmaceuticals. *Journal of Economics & Management Strategy 6*(1), 75–90.

Hortaçsu, A., O. R. Natan, H. Parsley, T. Schwieg, and K. R. Williams (2021, November). Organizational structure and pricing: Evidence from a large u.s. airline. Working Paper 29508, National Bureau of Economic Research.

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.

Schut, F. T. and P. A. Van Bergeijk (1986). International price discrimination: the pharmaceutical industry. *World Development 14*(9), 1141–1150.

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

Verboven, F. (1996). International price discrimination in the european car market. *The RAND Journal of Economics*, 240–268.

Williams, K. R. (2022). The welfare effects of dynamic pricing: Evidence from airline markets. *Econometrica 90*(2), 831–858.

7 Appendix: Additional Tables

Table 9: 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 10: 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 11: 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)	
Ln (Days)		-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)
Obs	223491	223491	493646	493646

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

Table 12: Regression of Fares on Both Kinds of Price Discrimination: Market Structure measures

	Transborder			Domestic		
	(1)	(2)	(3)	(4)	(5)	(6)
Ln(Days)	-0.213***	-0.210***	-0.210***	-0.242***	-0.242***	-0.241***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Adj. Fares	1.125***	1.141***	1.148***	0.272***	0.272***	0.273***
	(0.17)	(0.17)	(0.17)	(0.04)	(0.04)	(0.04)
Ln(Days)*Adj.	-0.171**	-0.180**	-0.184***	0.161***	0.161***	0.161***
	(0.06)	(0.06)	(0.06)	(0.02)	(0.02)	(0.02)
Rival Carriers	-0.159***			-0.032***		
	(0.04)			(0.01)		
HHI (Seats)		1.473**			-0.451	
		(0.52)			(0.29)	
HHI (Flights)			1.027**			-0.598
			(0.38)			(0.40)
Constant	5.611***	4.623***	4.809***	5.244***	5.420***	5.471***
	(0.15)	(0.15)	(0.09)	(0.03)	(0.16)	(0.20)
R^2	0.623	0.618	0.617	0.536	0.537	0.537
Obs	216314	216314	216314	481597	481597	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.

Table 13: Regression of Roundtrip Transborder Itineraries: To and From Toronto

	From Toronto	To Toronto
	(1)	(2)
Ln(Days)	-0.222***	-0.227***
	(0.00)	(0.00)
Adj. Fares	1.153***	1.124***
	(0.01)	(0.01)
Ln(Days)*Adj.	-0.205***	-0.186***
	(0.00)	(0.00)
Constant	5.056***	5.309***
	(0.01)	(0.01)
R^2	0.618	0.605
Obs	91655	91735

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