

# Price Discrimination along Multiple Dimensions: New Evidence from a Regional Airline\*

Ambarish Chandra<sup>†</sup>  
University of Toronto

January 14, 2019

## Abstract

I examine the case of a firm that practices both second-degree and third-degree price discrimination. I present a model showing conditions under which the premium for higher quality can either rise or fall as the firm implements group pricing. I then use new data from a regional airline to estimate how the two kinds of price discrimination interact, and how each is affected by changes in competition. I establish three key results, all new to the literature. First, in different markets, the two kinds of price discrimination can either offset or reinforce each other. Second, inter-temporal differences in prices are purely driven by price discrimination, rather than by scarcity pricing. Third, competition increases the extent of both kinds of price discrimination.

---

\*I thank Nate Baum-Snow, 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.

<sup>†</sup>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

Price Discrimination is a central and long-studied topic in Microeconomics. With the advent of large new datasets on consumer characteristics, the ability of firms to price discriminate has grown tremendously. As a result, new settings and opportunities have arisen for economists to examine price discrimination. Thus far, the Industrial Organization literature has mostly distinguished between second-degree and third-degree price discrimination strategies. But firms often practice both at the same time, especially as their information on consumer demographics and behavior has improved. Yet, theoretical studies tend to focus on just one of these types, most likely due to the complexity inherent in such models ([Armstrong and Rochet, 1999](#)). Empirical studies also tend to focus on just one type, likely due to data constraints. However, it is unlikely 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 use new data from the airline industry, which is a setting where firms have long used multiple, sophisticated methods to extract consumer surplus using both types of discriminatory pricing. I examine price discounts for advance purchases, as well as price menus for various quality levels, to simultaneously examine both second- and third-degree price discrimination. I then examine how competition affects each of these practices and establish key results that I compare to a model of a firm that practices both kinds of price discrimination. To my knowledge, this is the first paper to directly use booking data directly from an airline to establish how firms price discriminate according to advance purchases. Moreover, this is the first empirical study of how a firm simultaneously practices multiple types of price discrimination in an oligopoly setting.<sup>1</sup>

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

To generate intuition, I first present a simple, stylized model in which I assume that

---

<sup>1</sup>Prior authors have examined multiple kinds of price discrimination in monopoly, for example [Leslie \(2004\)](#). Some other studies have examined settings where firms may implement multiple kinds of discriminatory pricing, but focus on just one kind in their analysis; for example, [Lin and Wang \(2015\)](#). [Borzekowski et al. \(2009\)](#) examine how competition affects the type of price discrimination that firms use.

a firm encounters different types of consumers arriving in each period. It will always be optimal for consumers to buy in the current period rather than wait for future periods, when prices will be higher, but they must self-select into different quality levels. In each period, the firm must choose prices for each quality level. Thus, the change in prices over time—corresponding to advance purchase discounts—represents third degree price discrimination, while the variation in prices for different quality levels at a given point in time represents second-degree price discrimination.<sup>2</sup> I show that the premium for high quality service can either grow or shrink over time, as it depends on the composition of consumers that arrive in each period. For example, if consumers with a high marginal utility of quality are over-represented in later periods then the quality premium will grow over time.

I then turn to the empirical analysis and show that prices for a given flight-day clearly rise over time, for both service quality levels. In other words I show that there is indeed an advance purchase gradient. I then establish three, key, empirical results, all of which are new to the literature. First, I show that in some markets in my sample the two kinds of price discrimination reinforce each other, whereas in others they offset each other, in a manner consistent with the predictions of the model. Specifically, I show that on domestic routes within Canada the premium for high quality service shrinks over time, implying a convergence between low- and high-priced fares as the date of travel approaches. The opposite is true of transborder travel to the U.S. I argue that this is because of different arrival rates of high willingness-to-pay consumers in these markets. It is more complicated to make a last-minute international trip, and therefore later-arriving customers in this market reveal a higher willingness to pay, not just for travel but for quality as well.

Second, I show that the existence of advance purchase gradients is not due to scarcity pricing, but is most likely because the firm implements inter-temporal price discrimination. This is both because the airline that I study is rarely inventory constrained, and because I show that even routes or flight-days that have unusually high demand do not exhibit price gradients that are any steeper than those with normal or low demand. This is an important result because it is established using unambiguous measures of both fares and remaining inventory, as I have these variables directly from the airline, in contrast to previous work which has had to construct or infer measures of remaining inventory.

Third, I establish new results regarding the effect of competition on second-degree and third-degree price discrimination. The results indicate that the firm increases the extent of both kinds of price discrimination when it faces greater competition. Specifically, in markets

---

<sup>2</sup>It is often assumed in the literature that advance purchase discounts are a form of second-degree price discrimination, but I argue that this practice represents group pricing more accurately than menu pricing. I discuss this issue in more detail below.

with a lower Herfindahl Hirschmann Index, or a higher number of rivals, the advance purchase gradient is steeper and the premium for high quality service is greater.

While airlines have long been the setting for studying price discrimination, the available data have limited the extent of our understanding of this phenomenon. Specifically, the most common data source for analyzing prices in this industry has been the DB1B dataset provided by the U.S. Department of Transportation. This dataset has many useful properties, such as its large and representative sample, which has led to its widespread use in studies of price discrimination and price dispersion (Borenstein and Rose, 1994; Gerardi and Shapiro, 2009). However, one characteristic that is unavailable is 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 had to infer price dispersion or discrimination from the empirical distribution of ticket prices, with no knowledge of how much is due to advance purchase discounts versus other factors. It is only recently that some authors have been able to obtain, or construct, information on how fares vary over time for a given flight.<sup>3</sup>

My data source is a private airline based in Toronto, Ontario, which I will refer to as North Air.<sup>4</sup> Though small compared to most well known airlines, North Air is an important player on the set of routes that it operates, which include domestic routes within Canada and transborder routes between Canada and the United States. North Air provided me with a 10% random sample of all its bookings, with information on the exact fares, date of travel, class of service and the number of passengers. Crucially, I also obtained the date of the booking, which allows me to calculate the length of the advance purchase period, and therefore examine price discrimination due to advance purchase.

To my knowledge, this is the first time that research in Economics has obtained booking data directly from an airline, with the resulting ability to examine advance purchase price discrimination. While some recent studies have inferred this information, my data source is superior in two important respects. First, it allows me to accurately calculate the premiums or discounts associated with different advance purchase periods, enabling me to construct fare gradients associated with price discrimination.<sup>5</sup> Second, it allows me to calculate the

---

<sup>3</sup>For example, Alderighi et al. (2015) scrape airfares from the website of Ryanair, Williams (2017) uses data from a search engine and Lazarev (2013) obtains airlines' published fares from a Global Distribution System.

<sup>4</sup>The airline prefers not to be named in print, though its identity should be easy to infer.

<sup>5</sup>Some recent studies have inferred discounts related to advance purchases by scraping the websites of airlines or of online travel agents; for example, Alderighi et al. (2015). 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, say, one day to the next, it is not 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.

remaining inventory of seats available at any point in time, which can potentially be an important constraint for any airline.

Knowing the remaining inventory of seats at any point in time is crucial for accurately understanding price discrimination. 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. Thus, 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.<sup>6</sup> 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.

Thus, the inventory management problem can contribute to price differences over time, that are distinct from price discrimination, and can make the analysis of price differences over time 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.<sup>7</sup> However, in this paper I can address this problem quite simply. A key reason that I approached North Air for its data was because it is well known that this airline is rarely inventory constrained. North Air is often cited by analysts as having one of the lowest load factors in the industry. My own examination of the data shows that its load factors are rarely above 60%, even on its most popular routes. Moreover, my empirical results strongly imply that inventory is not a consideration for North Air. Controlling for the number of remaining seats has no discernible effect on the price discrimination premiums that I estimate, and I show that the price discrimination gradients that I estimate are no different on routes, or route-days, with unusually high or low load-factors.

Therefore, by isolating price discrimination as the cause of dynamic pricing, I show

---

<sup>6</sup>[Sweeting \(2012\)](#) analyzes baseball ticket pricing, where inventory management is the key problem and price discrimination is absent, and shows that prices generally fall over time for a given game. Past research has examined the interaction of the inventory management and price discrimination problems; for example, [Dana and Williams \(2016\)](#) provide a theoretical treatment and [Alderighi et al. \(2015\)](#) conduct an empirical analysis.

<sup>7</sup>[Clark and Vincent \(2012\)](#) use seat maps to construct the remaining inventory, at any point in time, for a given flight, 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.

the presence of an advance purchase gradient even in the absence of scarcity pricing. This is an important finding because a number of papers in the literature have argued 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 Jr (1999), Möller and Watanabe (2010)). My results, however, show that even for a firm which is not inventory constrained, advance purchase discounts are a ubiquitous phenomenon, lending more support to the notion that advance purchase discounts are driven by price discrimination and not solely by optimal inventory management.

## 1.1 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. In this section I will briefly describe the most relevant papers for this study.

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.

Recent empirical studies of second-degree price discrimination include Lin and Wang (2015) who show that the curvature of the hourly price schedule for parking garages decreases with competition. Specifically, short-term parking prices fall by more than long-term prices, making the price schedule more linear, and implying that competition constrains a firm's ability to practice second-degree price discrimination. This result is in contrast to Busse and Rysman (2005), for example, who find that competition increases the curvature of the price schedule in Yellow Pages advertising. Directories that face more competition see greater reductions in the prices of large ads than of small ads. Seim and Viard (2011) also show that, among wireless carriers, quantity discounts are larger in markets that see more entry.

It is worth noting that the empirical literature on second-degree price discrimination has focused mainly on quantity discounts (nonlinear pricing). While it is well known that 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 sell both full-service and self-service gasoline are able to implement discriminatory pricing in a way that cannot be done by single-product stations. As I discuss below, 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 second-degree discrimination according to quality levels, as these do not change in response to demand or competition.

Within the empirical literature on third-degree price discrimination, there are well-known disagreements regarding the effects of competition on price discrimination.<sup>8</sup> [Borenstein and Rose \(1994\)](#) and [Asplund et al. \(2008\)](#) find that firms in more competitive markets offer a wider range of prices, while [Gerardi and Shapiro \(2009\)](#) find the opposite.

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, a number of prior studies argue that similar price dynamics can be achieved solely through the firm's optimal management of its inventory. Specifically, when a firm faces uncertain demand, or when demand is likely to exceed capacity, it can be optimal for prices to rise over time, even in the absence of a discriminatory pricing mechanism. This result is found by [Gale and Holmes \(1993\)](#), [Dana \(1998\)](#), [Dana Jr \(1999\)](#) 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 temporal price discrimination is generally impossible. This apparent contradiction may be due to the fact that baseball games are less likely to be sold out.

Prior research has often assumed that advance purchase discounts are a form of menu pricing, i.e. second-degree price discrimination, since they reason that consumers make a decision of when to purchase tickets and therefore self-select into more or less expensive fares ([Dana, 1998](#)). However, [Chandra and Lederman \(2017\)](#) argue that this assumption should

---

<sup>8</sup>Moreover, a number of studies tend to conflate price *dispersion* and price *discrimination*, such as [Borenstein and Rose \(1994\)](#) and [Gerardi and Shapiro \(2009\)](#), both of which use the observed variation in prices as a measure of price discrimination.

be reevaluated, especially in light of well known equilibrium features of the airline industry. In reality, 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 their travel plans at the last minute never have the option to travel back in time and 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. Those who do so 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. Thus, advance purchase discounts should correctly be understood as a form of third-degree price discrimination.

## 2 Data and Setting

This section describes the data and establishes some simple, but 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. 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. Many prior papers have relied on the U.S. Department of Transportation’s DB1B dataset to calculate measures of price dispersion. However, despite the wealth of information contained in the DB1B, it lacks a key piece of information that would enable researchers to study price discrimination comprehensively—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.<sup>9</sup>

---

<sup>9</sup>See [Gaggero and Piga \(2011\)](#) and [Alderighi et al. \(2015\)](#) for examples of using scraped data on airline prices. While this represents a considerable advance in the use of airline data, the main issue with using

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.<sup>10</sup> 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 carried out by separate teams (Lazarev, 2013). This can make their pricing opaque, and makes it difficult to compare tickets in different buckets. Finally, the airline 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 available to purchase at any time. 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 fee, usually around \$75, 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 prior to the trip.<sup>11</sup> As I will show below, North Air sells very few Refundable tickets, and it makes no difference whether I ignore these, or combine them with Adjustable tickets. I will refer to two main quality levels in my empirical analysis: Low and High fares. 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.

The dataset obtained from North Air contains information on over 900,000 itineraries, during 2008–2014. Each itinerary contains the following information: the date of the booking, the date of travel, the number of passengers on the itinerary, the start and end points of the journey, the connections on the way (if any), the fare group, the currency in which the ticket was purchased, and the fare paid. Crucially, I use the difference between the booking and travel dates to construct the *daysout*, the number of days in advance of the journey that the reservation was made.

---

such data is that it only provides information on the prices offered by firms, not the equilibrium transacted prices. Moreover, the link between the distribution of offered prices and that of transacted prices depends on the frequency and accuracy with which these websites are scraped.

<sup>10</sup>In this regard, the airline is similar to most low-cost carriers. However North Air is not exactly a low cost carrier. It offers certain premium features—to all passengers—such as free snacks in lounges and on board, and leather seats.

<sup>11</sup>There are a few other differences, such as with bag fees.

I restricted the initial 10% random sample in the following ways. First, North Air offers service to and from 19 airports in Canada and the United States, during my sample period. The data contain information on journeys between all of these, as well as some journeys that began or terminated at around 10 other airports. The latter are not standard tickets for travel on North Air, and probably represent situations where travel agents combined tickets for travel on North Air with some other airline. I dropped these itineraries, which accounted for just 122 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.

As the data contain a 10% random sample of all flown itineraries, I can use this information to work out measures of the remaining inventory available to the airline when each seat is sold. This is important in order to distinguish capacity-based and discriminatory reasons for changing fares over time. I use two measures designed to capture the airline’s remaining capacity. The first is the *booking order* which is the order in which each seat is sold, for travel on a given day for a given route. The idea is simply that tickets, for a given travel date, bought earlier are less likely to encounter a capacity constraint on the part of the airline. Although, as I will show below, this measure is correlated with the *daysout* variable, variation in booking order controlling for *daysout* indicates flights that are more or less full, holding constant the days in advance that the reservation was made.

The second measure of the airline’s remaining capacity is created by first constructing North Air’s available seat inventory, for each route and travel day. I do this using data on flights and seats from OAG. I then divide this inventory by 10 (to account for the fact that I only observe 10% of the data) and subtract from it the number of seats that were already purchased on the route-date prior to the given itinerary being purchased.<sup>12</sup>

Table 1 presents summary statistics for the data. Each observation is an itinerary for travel on a certain date for a given route. The average reservation is made 29 days in advance of travel, has 1.2 passengers and costs around \$148 Canadian dollars. 85% of tickets

---

<sup>12</sup>Note that both capacity variables are measured with error, given that they are formed from the 10% random sample rather than from the full distribution.

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
Advance Purchase (Days)	29.2	30.7	0	19	473
Booking Order	11.90	12.1	1	7	86
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

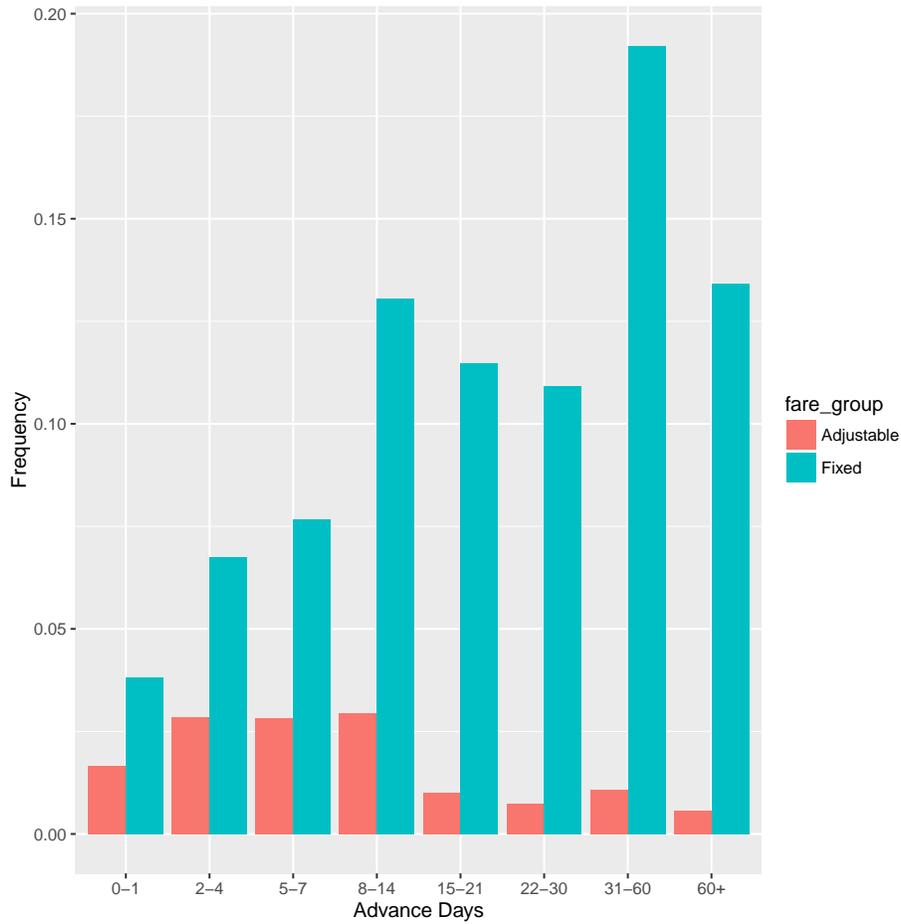
Note: An observation is an itinerary. N=865,492

are in the Fixed fare group, with 13% in the Adjustable group. Over 80% of tickets are part of return journeys, with the rest being one-way trips. The booking order for the average itinerary is 11.9.

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.

In order to calculate measures of competition, as well as accurate measures of remaining inventory, I used a second source of data from OAG, which is a well-known provider of data on the airline industry. For each route that North Air flies, I obtained monthly data on the number of flights and seats offered by the airline, as well as by all the rival carriers that operated in the same route-month. Table 2 presents summary statistics on the competition variables that I use from the OAG. It shows the number of carriers, seats and flights—separately for North Air and all of its rivals—on each route-month on which North Air provided service. The average number of rivals was 1.89, which ranged from 0 (on certain holiday destination routes) to 5 (Toronto-New York). Rivals operated around 415 flights a month on these routes, with about 37,000 available seats on average. North Air itself operated about 150 flights a month with approximately 10,000 seats, showing that the airline operated about the same number of flights as its rivals, but with smaller planes on average. I also calculated the Herfindahl-Hirschmann Index of competition on these routes, based both on seats and flights. This index had a mean of 0.46, and ranged from 0.21 to 1.

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.

Table 2: Summary Statistics: OAG Data

	Mean	SD	Min.	Median	Max.
Rival Carriers	1.89	1.00	0	2	5
Rival Seats	37461	40220	0	19572	157119
Rival Flights	415.0	385.4	0	236	1682
Own Seats	10444	8727	127	7560	38220
Own Flights	147.6	123.3	2	107	546
HHI Seats	0.46	0.16	0.22	0.45	1.00
HHI Flights	0.47	0.17	0.21	0.44	1.00

Note: An observation is a route-month. N=2144.

### 3 Model

In this section I sketch a simple model of a firm that sets price menus over time, to a changing composition of consumers. Average price differences over time will represent the firm's third-degree price discrimination, based on varying willingness-to-pay of consumers for a fixed quality level. Price differences at a given point in time, for different quality levels will represent second-degree price discrimination.

The questions that I hope to understand from the model are the following: First, how does the practice of third-degree discrimination affect a firm's ability to extract profits through second-degree discrimination and vice versa? One possibility is that consumers' tastes and 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, thus allowing the firm to increase its ability to extract profits from such customers. Conversely, it may be the case that it becomes 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.

The second question that the model can shed light on is: How does competition affect the ability to do each of second- and third-degree price discrimination in settings where firms engage in both? Past empirical work has been unable to resolve this question in partial analysis of each kind of discriminatory pricing. For example, [Busse and Rysman \(2005\)](#) and [Lin and Wang \(2015\)](#) come to opposite conclusions when studying the effect of competition on second-degree price discrimination, as do [Borenstein and Rose \(1994\)](#) and [Gerardi and Shapiro \(2009\)](#) when studying third-degree discrimination. Some theoretical models show that the effects of competition can be non-monotonic, e.g. [Dai et al. \(2014\)](#) and [Chandra and Lederman \(2017\)](#). However, a consideration that has not been noted by past authors is that many firms engage in both kinds of discriminatory pricing simultaneously, so that it may be important to account for each kind when examining how competition affects the other. As [Stole \(2007\)](#) notes in his survey, past theoretical studies—especially those of second-degree discrimination—have not provided consistent answers with respect to the effect of competition.<sup>13</sup>

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

---

<sup>13</sup>For examples, see [Stole \(1995\)](#) and [Rochet and Stole \(2002\)](#).

that potential travellers buy immediately once their personal uncertainty regarding travel is resolved. They then sort into different fare classes according to their types to satisfy second degree price discrimination.

Second, I assume that the firm has unconstrained inventory. It incurs some constant marginal cost of providing service. I assume a continuum of consumers in each period. For now, though, I start with the simplest case with discrete types of consumers.

### 3.1 Model with Discrete types

Consider an airline facing three types of consumers in each period  $t$  with marginal willingness-to-pay for quality:  $\theta_1(t)$ ,  $\theta_2(t)$  and  $\theta_3(t)$ . These change every period, but assume that each type buys in the current period as future prices will be higher with certainty. Therefore, this is a series of one-dimensional problems. Assume there are  $n_i$  consumers of each type. Assume that, in each  $t$ ,  $\theta_1 > \theta_2 > \theta_3$ . Consumers get utility  $\theta_i$  from consuming the low quality product and  $\theta_i\phi$  from consuming the high quality product.

The firm only offers two quality levels, Low and High, and must choose associated prices  $p_H$  and  $p_L$ . In order to satisfy the low types' IR constraint, the firm always sets  $p_L = \theta_3$ .<sup>14</sup> The question then is what the premium for high quality service will be in each period.

Either the firm charges a price to attract Type 2's to the high quality level, or it raises the price to only cater to Type 1 consumers.

To do the former, need to satisfy Type 2's IC constraint, which implies

$$p_H = p_L + \theta_2(\phi - 1) = \theta_3 + \theta_2(\phi - 1) \quad (1)$$

To do the latter:

$$p_H = p_L + \theta_1(\phi - 1) = \theta_3 + \theta_1(\phi - 1) \quad (2)$$

Thus the high price premium—expressed in differences—is either  $\theta_2(\phi - 1)$  or  $\theta_1(\phi - 1)$ .

Let  $\Pi_A$  be the profit from charging the higher price and  $\Pi_B$  be the profit from charging the lower price. Then,

$$\Pi_A - \Pi_B = n_2(c_H - c_L) + (\phi - 1)(n_1\theta_1 - n_1\theta_2 - n_2\theta_2) \quad (3)$$

Examining Equation 3, the model provides some simple predictions. First, the firm is more likely to charge the higher price as  $n_1$  and  $\theta_1$  increase. In other words, the price of

---

<sup>14</sup>It is straightforward to assume that  $n_3$  is high enough that the airline does not just write off the most price sensitive consumers.

high quality service increases as the share of the highest types increases, or as their marginal willingness-to-pay for quality increases. Conversely, the firm is more likely to charge the lower price as  $n_2$  and  $\theta_2$  increase.

Second, the effect of  $\phi$  is ambiguous as it depends on the relative numbers of each type and the strength of demand preferences. Higher quality makes both Type 2 and Type 1 willing to pay more, and depending on how many of each type there are, the firm may either try to attract, or write off, Type 2's from the higher quality.

Third, both competition and the arrival rate of various consumer types will affect the evolution of fares. Routes where relatively more high types arrive later will see a greater divergence between the prices for low and high quality levels. By contrast, routes where fewer high types arrive late may see a shrinking of the gap between adjustable and fixed fares over time.

In equilibrium, the airline will sell the high quality seats to high types and low quality to low types. Therefore, empirically we can examine the evolution of fare differences separately for routes where relatively more versus fewer high quality fares are sold over time. In particular, I will examine how the relative premium for high quality service relates to the share of Adjustable and Refundable seats, as the date of travel approaches.

### 3.2 Model with Continuous types

Assume that there are two periods. Period 1 is the advance purchase period where types are drawn on a distribution,  $F_1 [\underline{\theta}, \bar{\theta}]$ . Period 2 is the last-minute purchase period where types are drawn on a distribution,  $F_2 [\underline{\theta}, \bar{\theta}]$ .

I assume that the mean of  $F_2$  is great than that of  $F_1$ , reflecting higher average willingness to pay by consumers who arrive later. The key question will be the relative weight of higher willingness to pay customers in each period. For example, assume that the Period 1 distribution is right skewed, implying more weight on the low types, and in that the Period 2 distribution is left skewed, with more weight on the high types. Then people who buy later are also more likely to have a higher willingness-to-pay for quality, and therefore we should see an increasing quality premium over time. If the opposite, we should see a decreasing quality premium.

In each period, let  $\theta_1$  denote a consumer who is indifferent between the low quality ticket and not purchasing a ticket at all. Then:

$$\theta_1 - p_1(t) = 0 \tag{4}$$

which implies that  $\theta_1 = p_1(t)$ . Let  $\theta_2$  denote a consumer who is indifferent between the low

quality ticket and the high quality ticket. Then:

$$\phi\theta_2 - p_2(t) = \theta_2 - p_1(t) \tag{5}$$

which implies that

$$\theta_2 = \frac{p_2(t) - p_1(t)}{\phi - 1} \tag{6}$$

Therefore the firm chooses prices  $p_1(t)$  and  $p_2(t)$  to maximize:

$$[F(\theta_2) - F(\theta_1)][p_1 - c_1] + [1 - F(\theta_2)][p_2 - c_2] \tag{7}$$

Taking first order conditions with respect to each of  $p_1$  and  $p_2$ , it is straightforward to show that the relative premium for high quality service grows over time if  $f_2$  is more left skewed than  $f_1$ . If the opposite is true then the premium shrinks over time, and prices for low and high service converge for last-minute buyers.

Thus, this simple model illustrates how the relative premium for high-quality service—i.e. the extent of second-degree price discrimination—may grow or shrink over time, depending on the arrival rate of different consumer types. In principle, it may be possible to observe different markets exhibiting each of these phenomena.

## 4 Results

I now present empirical results. There are three main types of results. First, I will establish basic facts about price discrimination in this setting, and show the existence of clear advance purchase gradients, for both quality levels. I will also show how different markets exhibit strikingly different patterns of second-degree price discrimination, but generally similar patterns of third-degree price discrimination. 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. Third, I will examine how competition affects each kind of price discrimination, holding constant the other kind.

### 4.1 Advance Purchase Gradients and the Quality Premium

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: Regression of Fares on Advance Purchase Days: Full Sample

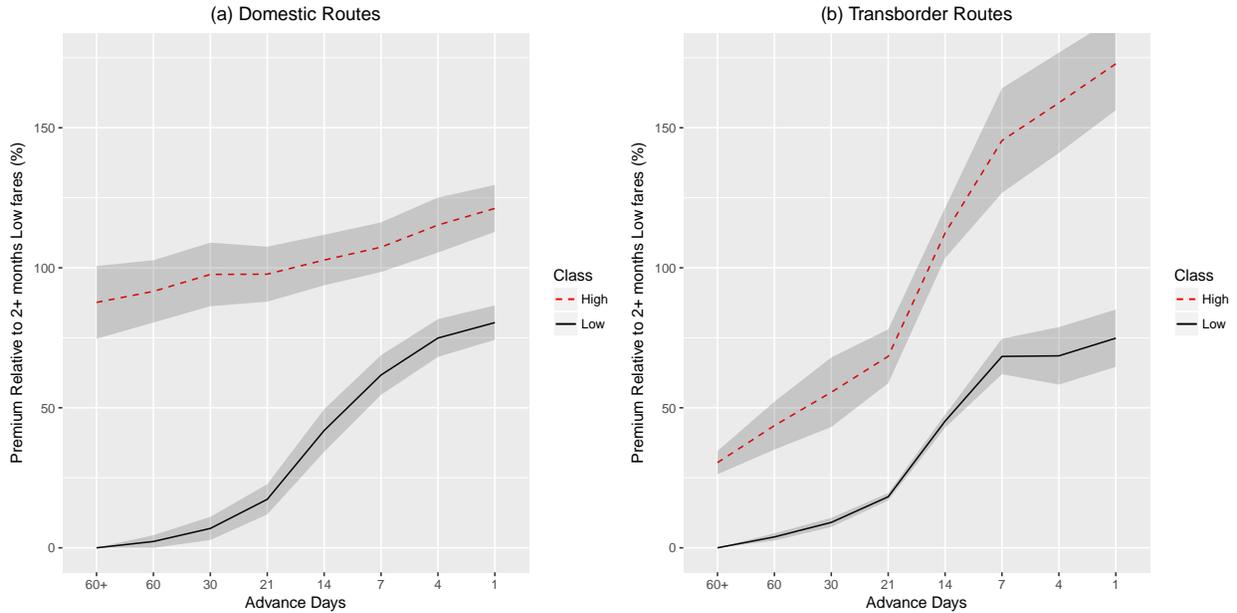
	Low	High-1	High-2
0 to 1 days	0.782*** (0.02)	0.811*** (0.09)	0.868*** (0.10)
2 to 4 days	0.721*** (0.03)	0.772*** (0.09)	0.812*** (0.10)
5 to 7 days	0.605*** (0.03)	0.697*** (0.09)	0.727*** (0.10)
8 to 14 days	0.402*** (0.02)	0.545*** (0.06)	0.570*** (0.07)
15 to 21 days	0.162*** (0.01)	0.308*** (0.02)	0.346*** (0.02)
22 to 30 days	0.066*** (0.01)	0.215*** (0.02)	0.246*** (0.02)
31 to 60 days	0.022*** (0.00)	0.114*** (0.01)	0.129*** (0.01)
Constant	4.345*** (0.04)	4.801*** (0.11)	4.821*** (0.11)
R <sup>2</sup>	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. High-1 refers to using only Adjustable tickets; High-2 refers to using both Adjustable and Refundable 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 Low and High fares. The omitted category is tickets purchased more than 60 days in advance of the flight. High fares include only Adjustable fares in column 2, and include both 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 service levels. Second, including fully Refundable tickets in the High fare classification raises the average fare in any time period, relative to

Figure 2: Advance Purchase Gradients for Domestic and Transborder Routes



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

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 High fares as including both Adjustable and Refundable tickets.

It is perhaps easier to see properties of the advance purchase gradients graphically. In Figure 2 I present these gradients, 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 low and high fares. This is to more easily compare coefficients, and quality premiums, between low and high fares.<sup>15</sup> In each panel, the coefficient for Low 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 2 shows that, as before, both low and high fares rise monotonically as the date of travel approaches. However, the advance purchase gradients are different for low and high 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 Low fares, with tickets purchased within one day of travel being about 110% more

<sup>15</sup>Regression coefficients are presented in Table 10 in the Appendix.

expensive than those bought far in advance.<sup>16</sup> The low 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.<sup>17</sup>

By contrast, the domestic and transborder markets exhibit very different gradients for high fares. High fares for domestic routes exhibit only a slight gradient, with only a 25 percentage point premium for tickets bought just prior to travel versus far in advance. However, high 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 high and low 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.

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  $i$ , consider 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  $t$ ,  $\alpha_i$  is a route-specific premium or discount, and  $\gamma_i$  is the route-specific price elasticity with respect to the remaining days until travel. Thus,  $\gamma$  captures, in a single parameter, the gradient of advance purchase discounts.

I take logs of Equation 8 and estimate elasticities 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 4, I summarize these elasticities using a 60-day window of purchase, but I show in Table 11 in the Appendix that the results are almost identical using a 90-day window.<sup>18</sup> The table presents separate elasticities for low and high fares for some selected routes. For example, the estimated daysout elasticity for

---

<sup>16</sup>This was obtained by calculating  $\exp(0.75) - 1$ .

<sup>17</sup>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.

<sup>18</sup>Tickets purchased within 60 and 90 days of travel account for 88% and 95% of itineraries, respectively.

Table 4: Daysout Elasticities

Route (to and from Toronto)	Low Fares		High Fares	
	$\gamma$	$N$	$\gamma$	$N$
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:  $\gamma$  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.

low fares on the Toronto-Montreal route is -0.25, implying that a 10% decrease in the days until travel is associated with an average price increase of around 2.5% for such fares on this route. 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 4 are mostly very similar to each other, suggesting that the advance purchase gradient for low 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  $\gamma$  is around 0.25, with Boston and Chicago being slight outliers, with estimated elasticities of 0.3 and 0.19, respectively. Overall, though, the average elasticities for domestic and transborder travel are virtually identical, in line with the results in Figure 2.

However, the values in the right column exhibit wide variation, and significant differences between domestic and transborder travel. Specifically, the advance purchase gradient for high fares is considerably flatter for travel between Toronto and the 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 not much inter-temporal price discrimination among buyers of high-quality tickets, suggesting perhaps that there is little room to extract further surplus from such buyers once they self-select into the more expensive, flexible tickets.

On the other hand, the elasticity for high 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 Low Fares. Once again these results are consistent with those shown in Figure 2. This suggests that those travelers on US-Canada routes who self-select into expensive tickets are also subject to extensive third-degree price discrimination by advance purchase. I will examine this suggestive evidence more deeply in the empirical analysis that follows.

## 4.2 Controlling for Remaining Capacity

I now control for remaining capacity at the time of booking each itinerary, using data on flights and seats from the OAG. 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.<sup>19</sup>

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% sample. However, the error is mean zero, and there is no reason to believe it is correlated with pricing.

Table 5 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, with estimated load factors of around 50%. Note that this is well below average load factors for the industry. 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

---

<sup>19</sup>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 5: 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	154.5	139.1	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 6: Regression of Fares on Advance Purchase Days: Capacity Measures

	Low Fares			High Fares		
	(1)	(2)	(3)	(4)	(5)	(6)
$\gamma$ (Gradient)	-0.236*** (0.01)	-0.236*** (0.01)	-0.237*** (0.01)	-0.220*** (0.04)	-0.220*** (0.04)	-0.221*** (0.04)
Remaining Seats (%)		-0.074*** (0.01)			-0.066*** (0.02)	
Load Factor (%)			0.160*** (0.02)			0.159*** (0.03)
Constant	5.193*** (0.03)	5.248*** (0.03)	5.118*** (0.03)	5.804*** (0.07)	5.852*** (0.08)	5.723*** (0.06)
R <sup>2</sup>	0.409	0.409	0.412	0.328	0.329	0.331
R <sup>2</sup> 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. High Fares include only Adjustable tickets in Column 1, and also include Refundable tickets in Column 2.

in Table 6 and are presented separately for low and high 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  $\gamma$ . For both types of fares, the estimated gradient is negative and similar to the average values presented in Table 4. The results on the capacity measures show that average fares do respond to changes in remaining inventory. The coefficients on remaining 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 6 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  $\gamma$  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 intertemporal 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 6. 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 Low 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 low fares, and from 16% to around 33% for high fares.

As a robustness check, in Table 9 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, with the exception that the gradient for the top quartile—*i.e.* for the fullest flights—is slightly lower, implying that the gradient of fares is slightly flatter for such flights.

In a further robustness check, I show in Table 12 in the Appendix, that the Table 6 results are almost exactly replicated when I replace  $\gamma$ , 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  $\gamma$  alone and, once again, the various capacity measures do not affect the estimated gradient using these discrete measures.

Overall, therefore, the results of Table 6 imply that within-route variation in fares can be explained almost entirely by temporal price discrimination rather than by scarcity pricing.<sup>20</sup> 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.

---

<sup>20</sup>This refers to the portion of variation explained by the regressors. Note that the fit measures in Table 6 can obviously be further improved by adding finer indicators for advance purchase days, but that is not the goal of this paper.

### 4.3 Competition and Price Discrimination

I now examine the role of competition in the implementation of both kinds of price discrimination. To do this, I rely on the fact that North Air faces differing levels of competition on the various routes it serves. I use three measures of competition: the number of rivals offering direct service on the route; and the Herfindahl Hirschmann Indices (HHIs) constructed based on each airline's share of seats and flights on the route.

Since I include route fixed effects, identification is obtained by comparing changes in these competition measures within routes. This alleviates the concern that differences in competition across routes are correlated with prices on those routes. However, it leaves two issues: first there may be insufficient variation in competition within a route over time to meaningfully estimate the parameters of interest. And second, within-route changes in competition may still be correlated with unobserved factors affecting prices.

On the first issue, it is quite apparent that there is little variation in the number of direct rivals that North Air faces on a given route. Although some routes have experienced entry and exit during my sample period, for the most part the number of rivals does not change enough to identify the parameters of interest. I expect to rely more on the HHI measures which do, in fact, vary considerably. This is because airlines—both North Air itself and its rivals—exhibit considerable variation in the number and size of flights that they operate on a given route. This is due to factors such as seasonality, consumer demand, constraints on the size of the airline's fleet, and secular trends in airline expansion. North Air itself has steadily grown over time and its market share has correspondingly increased, thereby directly affecting the HHI measures.

On the second issue, while the route fixed effects do account for some of the endogeneity in market structure, they do not account for the possibility of within-route correlation between competition and unobserved demand side-factors. Note, though that the first-order effect of this remaining endogeneity should be on the average *level* of prices, rather than on the *structure* of prices that arises through price discrimination. It is not easily obvious why changes in market structure should be correlated with unobserved factors that affect the ratio of high and low fares (second-degree price discrimination) or the advance purchase discounts offered by the airline (third-degree). I acknowledge, however, that I cannot fully rule out the possibility of such correlations, and therefore the following results should be understood with this caveat in mind.

I first examine the effect of competition on second-degree price discrimination, holding constant the advance purchase gradient. I regress fares on indicators for the advance purchase periods, an indicator for high fares, and an interaction between high fares and various measures of competition. The results are in Table 7. The competition measures have the

Table 7: Effect of Competition on Second-Degree Price Discrimination

	(1)	(2)	(3)
0 to 1 days	0.824*** (0.03)	0.825*** (0.03)	0.826*** (0.03)
2 to 4 days	0.761*** (0.03)	0.762*** (0.03)	0.762*** (0.03)
5 to 7 days	0.655*** (0.03)	0.656*** (0.03)	0.656*** (0.03)
8 to 14 days	0.455*** (0.03)	0.456*** (0.03)	0.455*** (0.03)
15 to 21 days	0.191*** (0.02)	0.191*** (0.02)	0.191*** (0.02)
22 to 30 days	0.089*** (0.01)	0.089*** (0.01)	0.088*** (0.01)
31 to 60 days	0.035*** (0.01)	0.035*** (0.01)	0.034*** (0.01)
High Fares=1	0.742*** (0.06)	0.693*** (0.07)	0.493*** (0.05)
HHI (Seats)	0.427 (0.31)		
High Fares=1 × HHI (Seats)	-0.451** (0.17)		
HHI (Flights)		0.250 (0.36)	
High Fares=1 × HHI (Flights)		-0.342* (0.19)	
Rivals			-0.096*** (0.02)
High Fares=1 × Rivals			0.028* (0.01)
Constant	4.053*** (0.16)	4.143*** (0.17)	4.464*** (0.04)
R <sup>2</sup>	0.568	0.567	0.569
Obs	738504	738504	738504

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

expected signs—positive effects for the HHI measures and a negative effect on the number of rivals—but the HHI measures are imprecisely estimated. This suggests some weak evidence that competition may slightly reduce average fares. As discussed above, however, the endogenous effect of competition will be experienced most directly on the level of fares, suggesting that these coefficients should be interpreted with caution.

The interaction terms in Table 7 suggest a clearer story. These interactions are clearly negative for the HHI measures, and positive for the number of rivals, suggesting in each case that greater competition has the effect of increasing the quality premium that the airline charges for flexible tickets, in other words increasing the extent of third-degree price discrimination.

I now examine the effect of competition on third-degree price discrimination, holding constant the quality premium for high fares. In Table 8 I show the results of regressions of fares on the interaction between the advance purchase elasticity ( $\gamma$ ) and the competition measures, with a separate intercept for High Fares. The results for the number of rivals are inconclusive, but the interactions with the HHI measures suggest that an increase in competition makes the advance purchase gradient steeper, in other words increasing the extent of third-degree price discrimination.

Overall, the results of this section suggest that the airline increases the extent of both second- and third-degree price discrimination when it faces more competition.

## 5 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, the 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 obtain new booking data directly from an airline that, to my knowledge,

Table 8: Effect of Competition on Third-Degree Price Discrimination

	(1)	(2)	(3)
High Fares	0.590*** (0.03)	0.589*** (0.03)	0.590*** (0.03)
$\gamma$	-0.334*** (0.03)	-0.338*** (0.03)	-0.228*** (0.02)
HHI (Seats)	-0.249 (0.44)		
$\gamma \times$ HHI (Seats)	0.218*** (0.07)		
HHI (Flights)		-0.488 (0.48)	
$\gamma \times$ HHI (Flights)		0.235*** (0.08)	
Rivals			-0.068** (0.03)
$\gamma \times$ Rivals			-0.008 (0.01)
Constant	5.309*** (0.21)	5.412*** (0.21)	5.374*** (0.07)
R <sup>2</sup>	0.547	0.547	0.547
Obs	688879	688879	688879

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

are the first of their kind to be used in academic research. I use these to examine both third-degree price discrimination according to advance purchase behaviour, and second-degree price discrimination according to the quality of service that travellers select. I estimate the empirical relationship between the two practices, and then examine how competition affects each. I find that advance purchase gradients clearly exist, for both low and high quality levels, 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 some markets, but a complementarity between them in others. Finally, I find that competition increases the effect of both kinds of discriminatory pricing.

## 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.
- Asplund, M., R. Eriksson, and N. Strand (2008). Price discrimination in oligopoly: Evidence from regional newspapers\*. *The Journal of Industrial Economics* 56(2), 333–346.
- Borenstein, S. and N. L. Rose (1994). Competition and Price Dispersion in the US Airline Industry. *Journal of Political Economy* 102(4), 653–683.
- Borzekowski, R., R. Thomadsen, and C. Taragin (2009). Competition and price discrimination in the market for mailing lists. *QME* 7(2), 147–179.
- Busse, M. and M. Rysman (2005). Competition and price discrimination in yellow pages advertising. *RAND Journal of Economics* 36(2), 3784390.
- Chandra, A. and M. Lederman (2017). Revisiting the relationship between competition and price discrimination. *American Economic Journal: Microeconomics* Forthcoming.
- Clark, R. and N. Vincent (2012). Capacity-contingent pricing and competition in the airline industry. *Journal of Air Transport Management* 24, 7–11.
- Dai, M., Q. Liu, and K. Serfes (2014). Is the effect of competition on price dispersion nonmonotonic? evidence from the us airline industry. *Review of Economics and Statistics* 96(1), 161–170.

- Dana, J. and K. Williams (2016). Oligopoly price discrimination: The role of inventory controls.
- Dana, Jr, J. D. (1998). Advance-purchase discounts and price discrimination in competitive markets. *Journal of Political Economy* 106(2), 395–422.
- Dana Jr, J. D. (1999). Equilibrium price dispersion under demand uncertainty: the roles of costly capacity and market structure. *The RAND Journal of Economics*, 632–660.
- Gaggero, A. A. and C. A. Piga (2011). Airline market power and intertemporal price dispersion. *The Journal of Industrial Economics* 59(4), 552–577.
- Gale, I. L. and T. J. Holmes (1993). Advance-purchase discounts and monopoly allocation of capacity. *The American Economic Review*, 135–146.
- Gerardi, K. S. and A. H. Shapiro (2009). Does competition reduce price dispersion? new evidence from the airline industry. *Journal of Political Economy* 117(1), 1–37.
- 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.
- Seim, K. and V. B. Viard (2011). The effect of market structure on cellular technology adoption and pricing. *American Economic Journal: Microeconomics* 3(2), 221–251.
- Shepard, A. (1991). Price discrimination and retail configuration. *Journal of Political Economy* 99(1), 30–53.

- Stole, L. A. (1995). Nonlinear pricing and oligopoly. *Journal of Economics & Management Strategy* 4(4), 529–562.
- Stole, L. A. (2007). Price discrimination and competition. *Handbook of industrial organization* 3, 2221–2299.
- 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. (2017). Dynamic airline pricing and seat availability. Technical report, mimeo.

## 6 Appendix: Additional Tables

Table 9: Daysout Gradients by Load Factor

	Full	Q1	Q2	Q3	Q4
	(1)	(2)	(3)	(4)	(5)
$\gamma$ (Gradient)	-0.247*** (0.01)	-0.258*** (0.01)	-0.265*** (0.01)	-0.248*** (0.01)	-0.215*** (0.01)
Constant	5.188*** (0.03)	5.128*** (0.03)	5.234*** (0.03)	5.267*** (0.04)	5.276*** (0.05)
R <sup>2</sup>	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 Low Fares. All regressions include route, day-of-week, month and year FEs. Standard errors, clustered by route, in parentheses.

Table 10: Regression of Fares on Advance Purchase Days: Domestic and Transborder Routes

	Domestic	Transborder
High: 0 to 1 days	1.212*** (0.04)	1.728*** (0.08)
High: 2 to 4 days	1.153*** (0.05)	1.589*** (0.09)
High: 5 to 7 days	1.074*** (0.05)	1.454*** (0.10)
High: 8 to 14 days	1.028*** (0.05)	1.124*** (0.05)
High: 15 to 21 days	0.977*** (0.05)	0.684*** (0.05)
High: 22 to 30 days	0.976*** (0.06)	0.555*** (0.06)
High: 31 to 60 days	0.916*** (0.06)	0.436*** (0.04)
High: 60+ days	0.876*** (0.07)	0.304*** (0.02)
Low: 0 to 1 days	0.804*** (0.03)	0.749*** (0.05)
Low: 2 to 4 days	0.749*** (0.03)	0.686*** (0.05)
Low: 5 to 7 days	0.617*** (0.04)	0.683*** (0.03)
Low: 8 to 14 days	0.419*** (0.04)	0.452*** (0.01)
Low: 15 to 21 days	0.173*** (0.03)	0.182*** (0.01)
Low: 22 to 30 days	0.069*** (0.02)	0.091*** (0.01)
Low: 31 to 60 days	0.022* (0.01)	0.039*** (0.01)
Constant	4.246*** (0.06)	4.284*** (0.06)
R <sup>2</sup>	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 11: Daysout Elasticities

Route (to and from Toronto)	Low Fares		High Fares	
	$\gamma$	$N$	$\gamma$	$N$
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:  $\gamma$  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 90 days of travel.

Table 12: Regression of Fares on Advance Purchase Days: Capacity Measures

	Low Fares			High Fares		
	(1)	(2)	(3)	(4)	(5)	(6)
0 to 1 days	0.812*** (0.02)	0.813*** (0.02)	0.816*** (0.02)	0.919*** (0.11)	0.920*** (0.11)	0.924*** (0.11)
2 to 4 days	0.756*** (0.03)	0.757*** (0.03)	0.761*** (0.03)	0.869*** (0.11)	0.871*** (0.11)	0.876*** (0.11)
5 to 7 days	0.639*** (0.03)	0.640*** (0.03)	0.643*** (0.02)	0.778*** (0.11)	0.780*** (0.11)	0.785*** (0.11)
8 to 14 days	0.433*** (0.03)	0.435*** (0.03)	0.438*** (0.02)	0.618*** (0.07)	0.619*** (0.07)	0.623*** (0.07)
15 to 21 days	0.180*** (0.02)	0.181*** (0.02)	0.185*** (0.02)	0.381*** (0.02)	0.382*** (0.02)	0.385*** (0.02)
22 to 30 days	0.080*** (0.01)	0.081*** (0.01)	0.084*** (0.01)	0.278*** (0.02)	0.279*** (0.02)	0.283*** (0.02)
31 to 60 days	0.030*** (0.01)	0.031*** (0.01)	0.033*** (0.01)	0.141*** (0.02)	0.142*** (0.02)	0.143*** (0.02)
Remaining Seats (%)		-0.066*** (0.01)			-0.074*** (0.02)	
Load Factor (%)			0.145*** (0.02)			0.176*** (0.03)
Constant	4.237*** (0.04)	4.285*** (0.05)	4.163*** (0.04)	4.743*** (0.12)	4.795*** (0.12)	4.646*** (0.11)
R <sup>2</sup>	0.434	0.435	0.437	0.348	0.349	0.351
R <sup>2</sup> 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. High Fares include only Adjustable tickets in Column 1, and also include Refundable tickets in Column 2.