# Essays in Industrial Organization

A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Kevin R Williams

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
Doctor of Philosophy

Thomas J Holmes & Amil Petrin, Advisors

June, 2014

# Acknowledgements

I am deeply indebted to my advisors Thomas J Holmes and Amil Petrin for their constant support and encouragement over the years. Tom and Amil are excellent mentors and they have greatly shaped how I approach research.

I am further indebted to my committee members, Joel Waldfogel and Kyoo-il Kim, for their valuable guidance and insights. I thank Joel for his great enthusiasm in my work and Kyoo-il for helping me develop the tools necessary to answer empirical questions.

This research has benefitted from conversations with Mark Bergen, Konstantin Golyaev, Ellen McGrattan, Thomas Quan, Filippo Rebessi, Imke Reimers, Ethan Singer, Maria Ana Vitorino, Chunying Xie, Motohiro Yogo, Thomas Youle, and the many participants at conferences and seminars I have presented at while at the University of Minnesota. I give special thanks to Brian Adams, who coauthored the second chapter of this dissertation.

I acknowledge support of the Minnesota Supercomputing Institute for providing computational resources. I thank the Department of Economics staff, especially Caty Bach, Kara Kersteter, and Wendy Williamson. Finally, I gratefully acknowledge the financial support of Morton and Artice Silverman and Richard and Ellen Sandor.

# Dedication

To my family and friends. To AJ for his patience and constant support.

# Abstract

This dissertation is comprised of two essays, both of which study how particular market environments affect firms' abilities to price discriminate.

In the first chapter, I analyze the pricing decisions of airlines in monopoly markets. Airline markets are noted for having several key features: (1) airlines have limited capacity and limited time to sell, (2) airlines face uncertainty in the popularity of any given flight, and (3) consumers who purchase tickets close to the departure date are less price sensitive than those who buy well in advance. These forces influence the pricing decision – what I call dynamic adjustment to stochastic demand (1 and 2) and intertemporal price discrimination (3). While the previous literature has emphasized each force in isolation, in this chapter, I estimate a model of dynamic airline pricing taking both into account. I use an original data set of daily fares and seat availabilities at the flight level. With model estimates, I disentangle key interactions between the arrival pattern of consumer types and remaining capacity under stochastic demand. I find dynamic adjustment to stochastic demand is particularly important as a means to secure seats for high-valuing consumers who arrive close to the departure date. It leads to substantial revenue gains compared to pricing policies which depend on date of purchase but not remaining capacity. In aggregate consumers benefit, despite facing higher fares on average, as a result of more efficient capacity allocation. Finally, I show an empirical procedure abstracting from stochastic demand will systematically understate the price insensitivity of consumers who search for tickets close to the date of travel.

In the second chapter, Brian Adams and I develop an empirical analysis of zone pricing under competition. While monopolists can only increase profits by adopting more granular pricing policies, this is not necessarily the case in markets with competition. Using an original data set for the retail drywall industry, we estimate a structural model of supply and demand. We use the model estimates to calculate equilibrium under alternative pricing policies to quantify the welfare implications of zone pricing. We find consumer surplus decreases substantially but firm profits appear to increase with finer pricing. As firms have not adopted these policies, they must face some additional costs. We call these costs "spatial menu costs," and our analysis finds them to be substantial: at least 22.1% of estimated profits, or 2.2% of revenues. Finally, we show that competitive interaction plays an important role in recovering menu costs – failing to account for competitive effects leads to an overestimate of profit gains and implied menu costs by 32.9%.

# Contents

# List of Tables

# List of Figures

# Chapter 1

# Dynamic Airline Pricing and Seat Availability

## 1.1 Introduction

Airlines tend to charge high prices to passengers who search for tickets close to the date of travel. The conventional view is that these are business travelers, and airlines capture their high willingness to pay through intertemporal price discrimination. Airlines also adjust prices on a day-to-day basis as capacity is limited and the future demand for any given flight is uncertain. While fares generally increase as the departure date approaches, prices can actually fall from one day to the next, after a sequence of low demand realizations.

This paper examines pricing in the airline industry taking into account both forces – intertemporal price discrimination and dynamic adjustment to stochastic demand. I use a new flight-level data set to estimate a structural model of dynamic airline pricing where firms face a stochastic arrival of consumers. The mix of consumer types – business and leisure travelers – changes over time, and in the estimated model, late-arriving consumers are significantly more price inelastic than consumers who arrive early on. With model estimates, I simulate the revenue losses associated with a pricing system which allows for intertemporal price discrimination, but not dynamic adjustment. I find these losses to be substantial, suggesting that the addition of dynamic adjustment creates an important complementarity in the pricing channels. My results provide credence as to why airlines have pioneered such complex pricing systems: having prices respond to stochastic demand allows firms to first secure seats for the high-valuing consumers who arrive close to the departure date, and then charge these consumers very high prices.

Existing research demonstrates the importance of intertemporal price discrimination in the airline industry. The view is that business consumers learn of last-minute meetings and are willing to pay a premium in order to reserve a seat, while leisure consumers are more price sensitive and book tickets well in advance. Consistent with the idea of market segmentation, Puller, Sengupta, and Wiggins (2012) find that ticket characteristics, such as advance purchase discount (APD) requirements, explain much of the variation in fares over time. Recently, Lazarev (2013) estimated a model of intertemporal price discrimination and he found a substantial role for this force.

The literature also shows that dynamic adjustment plays an important role in airline pricing. Escobari (2012) and Alderighi, Nicolini, and Piga (2012) find evidence that airlines face stochastic demand and prices respond to remaining capacity. In particular, Escobari (2012) estimates the pricing functions of airlines. He notes that fares decline in the absence of sales while having reduced capacity in any given period results in increased fares.[1] These results support the theoretical

[1]Puller, Sengupta, and Wiggins (2012) find limited support that seat scarcity explains the

predictions of Gallego and Van Ryzin (1994) and a large branch of research in operations management that have studied optimal pricing under uncertain demand, limited capacity, and limited time to sell. This work has been used to inform airline revenue management systems.[2] Such systems allow airlines to respond to stochastic demand by increasing fares when a sellout is likely and fall otherwise, as to not leave as many seats unfilled.

While previous work emphasizes the importance of intertemporal price discrimination and stochastic demand pricing separately, this paper examines both forces together and highlights how they interact. I establish two key points about the interaction. First, intertemporal price discrimination and dynamic adjustment to stochastic demand are *complements* in the airline industry. This follows because inelastic consumers tend to arrive last. In order to be in a position to price discriminate and set high prices to these late-arriving consumers, the firm will want to allow fares to adjust to realizations in demand. Second, in order to estimate how demand elasticity changes over time, which is needed to calculate the welfare effects of intertemporal price discrimination, it is necessary to take stochastic demand into account. The reason is that by ignoring stochastic demand, the opportunity cost of selling a seat is the same regardless of the date of purchase. But with stochastic demand, the opportunity cost changes over time. This matters because inferences regarding demand elasticity come from the firm's first-order condition in choice of price, which relates prices and marginal costs to demand elasticities. If marginal costs are incorrect, then the estimated change in

variation of fares for major US routes. However, Alderighi, Nicolini, and Piga (2012) find evidence that both characteristics and capacity matter. They find the role of capacity is pronounced in less competitive markets.

[2]There is a large literature on stochastic demand pricing (revenue management, yield management, or dynamic pricing). An overview of the dynamic pricing literature can be found in Elmaghraby and Keskinocak (2003) and Talluri and Van Ryzin (2005). Seat inventory control has also been studied; see Dana (1999).

demand elasticity is also incorrect.

In order to investigate dynamic airline pricing, a detailed data set of ticket purchases is required. However, the standard airline data sets used in economic studies (Goolsbee and Syverson (2008); Gerardi and Shapiro (2009); Berry and Jia (2010) for example) are either at the monthly or quarterly level. Recently, papers have been using new data to get to the flight level. McAfee and Te Velde (2006) and Lazarev (2013) create data sets containing high frequency fares. Other papers have obtained high frequency data on prices and a measure of quantities. Puller, Sengupta, and Wiggins (2012) use a unique transaction data set from a single computer reservation system. Escobari (2012) and Clark and Vincent (2012) collect fare and flight availability data, where the available number of seats is derived from publicly available seat maps. I create a similar data set with a key improvement. I use a new data source that allows me to see the same flight availability data that travel agents see. Specifically, the seat maps I collect allow me to distinguish between blocked and occupied seats. Without accounting for blocked seats, I find seat maps overstate load factors (seats occupied / capacity) by 10%. In addition, I provide evidence that seat maps are a useful proxy of bookings.

The sample contains the time path of fares and seat availabilities for over 1,300 flights in US monopoly markets. The structure of the data allows me to capture over one hundred departures of a single flight number, where each flight is tracked for sixty days. Descriptive analysis of the data reveals a strong role of remaining capacity in explaining the variation of daily fares. By investigating the pricing decisions of routes with two flights a day, I find that if one flight option is 30% more full, the flight is roughly 35% more expensive compared to the other option. Also, consistent with the ideas of stochastic demand pricing, 35% of fares change daily and 10% of the itineraries in the sample result in a fire-sale.

While the empirical evidence above is informative, it cannot be used to disentangle the interactions between intertemporal price discrimination and stochastic demand pricing. I proceed by estimating a structural model. The model includes three key ingredients: (i) firms have finite capacity and finite time to sell, (ii) firms face a stochastic arrival of consumers, and (iii) the mix of consumers, corresponding to business and leisure travelers, is allowed to change over time. The firm solves a stochastic dynamic programming problem. For the demand system, I assume a stochastic process brings new consumers to the market. The consumers that arrive know when they want to travel and solve a static problem, choosing to either buy a ticket on an available flight or exit the market permanently. The demand model differs from earlier theoretical work, including Gale and Holmes (1993), where consumers do not know if they wish to fly and waiting provides more information. In my model the only reason to wait is to bet on price, and since prices tend to increase, I show only a small transaction cost is needed to persuade consumers to decide whether to travel in the current period. In addition, I provide empirical evidence that suggests this is a reasonable assumption.[3]

The key identification challenge of the model is to separately identify the parameters associated with the arrival process from the parameters governing the change in elasticity over time. My identification strategy relies on accounting for the firm's pricing choice. By solving the firm's dynamic problem, I back out the firm's beliefs on current and future demand. Variation in sales conditional on seats and time relating inform the arrival process, and the firm's optimality conditions relate the pricing decision to the demand elasticity each period.

Using the model estimates, I compare the allocation of scarce capacity across

[3] Li, Granados, and Netessine (2013) studies dynamic consumer behavior in airline markets. Depending on the specification, they find between 5% and 20% of consumers are dynamic.

time under dynamic pricing with several counterfactual pricing systems. I first shut down the use of dynamic pricing so that the monopolist can only charge a uniform price. I find that uniform pricing results in a significant reallocation of capacity across consumers and time, but the gains in consumer surplus are mitigated as a result of inefficient capacity allocation. Further, uniform pricing results is a significant decline (6.6%) in revenues, more than offsetting the increase in consumer welfare (1.4%). As airlines operate under razor thin margins, the decrease in revenues would likely result in market exit in the long run.[4] Using dynamic and uniform pricing as benchmarks I then allow the firm to use dynamic pricing, but restrict the frequency of price updates. I find that even minor restrictions on the frequency of price adjustments results in significant revenue reductions.

I then single out the use of intertemporal price discrimination alone by considering a pricing system which depends on date of purchase, but not remaining capacity. By comparing uniform pricing to this intermediate case, and this intermediate case with dynamic pricing, I quantify the relative importance of intertemporal price discrimination and adjusting to stochastic demand. I find that roughly half the revenue gains of using dynamic pricing over uniform pricing comes from the intertemporal price discrimination channel, with the remaining half coming from dynamic adjustment. Dynamic pricing substantially increases revenues (3.5%) over the use of intertemporal price discrimination alone as firms are able to allocate more seats to late-arriving business consumers, who are then charged high prices. Additionally, I find that overall, both business and leisure welfare is higher under dynamic pricing compared to intertemporal price discrimination

[4] According to an IATA industry report (Profitability and the Air Transport Profitability and the Air Transport Value Chain), the average fare paid per passenger in 2012 was \$181.91, with an average cost per passenger of \$225.70. After accounting for auxiliary and cargo revenue, they estimate the net profit per passenger to be \$2.56.

alone. Although business consumers are charged higher fares under dynamic pricing, they also benefit from having more seats available. Leisure consumers benefit from lower fares as dynamic adjustment reduces the firm's incentive of holding back capacity in early periods.

Intertemporal price discrimination and dynamic adjustment to stochastic demand are complements in the airline industry because high-valuing consumers arrive late. To highlight this complementarity, I perform two analyses. First, I reverse the arrival process of consumers so that the high-valuing consumers arrive first. In this environment, I find that intertemporal price discrimination accounts for a much larger percentage (25% more) of the value of dynamic pricing. This follows because there is no need for the firm to save seats until close to the departure date.[5] Second, I hold the mix of business and leisure consumers constant over time. This analysis reveals the revenue gains associated with dynamic adjustment are half the gains attained under the estimated arrival process where high-valuing consumers show up late. This result is consistent with the theory of Gallego and Van Ryzin (1994), and emphasizes that stochastic demand pricing is especially valuable in the airline industry because of the particular pattern of consumer arrival.

Finally, I show how estimation approaches that do not take into account stochastic demand will systematically understate the degree to which demand becomes more inelastic as the departure date approaches.

[5]This analysis assumes consumers do not wait to purchase. Stokey (1979) shows an environment in which a monopolist of durable goods that commits to pricing would not use intertemporal price discrimination as consumers with high valuations would strategically wait to purchase. Conlisk, Gerstner, and Sobel (1984) and Board (2008) consider durable goods models with time dependent demand.

## Related Literature

This paper adds to the growing empirical work on intertemporal price discrimination and dynamic adjustment to stochastic demand. Intertemporal price discrimination can be found in many markets, including video games, Broadway theater, and concerts (Nair (2007), Leslie (2004), and Courty and Pagliero (2012), respectively).[6] A closely related paper to mine is Lazarev (2013), who estimates the welfare effects of intertemporal price discrimination in US monopoly airline markets. His model includes dynamic consumers, but abstracts away from aggregate demand uncertainty. There is a large literature in economics and operations research on stochastic demand pricing.[7] Talluri and Van Ryzin (2004) and Vulcano, van Ryzin, and Chaar (2010) provide insights into the estimation of (discretized) continuous time demand models with myopic consumers. Like Vulcano, van Ryzin, and Chaar (2010), I estimate a discrete choice model with Poisson arrival; however, I also allow for two consumer types. I use information on the pricing decision of the firm to aid in the identification of the parameters. Importantly, McAfee and Te Velde (2006) note that stochastic demand models – models which do not incorporate changes in willingness to pay over time – do not match the positive trend in airfares as the departure date approaches. By investigating both forces simultaneously, my model is able to capture both the positive trend in fares as well as the day-to-day variation in fares, including price declines. To the best of

[6] Lambrecht, Seim, Vilcassim, Cheema, Chen, Crawford, Hosanagar, Iyengar, Koenigsberg, Lee, et al. (2012) provide an overview of empirical work on price discrimination. Interestingly, Jones (2012) notes that some theaters are now using the same pricing techniques of airlines.

[7] An overview of the stochastic demand pricing (or also called: dynamic pricing or revenue management depending on the context) literature can be found in Elmaghraby and Keskinocak (2003) and Talluri and Van Ryzin (2005). Sweeting (2012) analyzes ticket resale markets. Pashigian and Bowen (1991) and Soysal and Krishnamurthi (2012) study clearance sales and seasonal goods, respectively. Zhao and Zheng (2000) and Su (2007) discuss extensions to dynamic pricing models, including consumer dynamics.

my knowledge, this is the first paper to quantify the complementarities between intertemporal price discrimination and dynamic adjustment to stochastic demand through a structural model.

The rest of the paper proceeds as follows. Section 1.2 describes the data collected for this study. Section 1.3 presents the model. Section 1.4 discusses the econometric specification and identification of the model parameters. Section 1.5 presents the results of estimation. Section 1.6 presents the counterfactuals. The conclusion follows.

## 1.2 Data

I create an original data set of high frequency airfares and seat availabilities with data collected from two popular online travel services. The first web service used is a travel metasearch engine. I use the web service to obtain daily fares at the itinerary level.[8] I obtain all one way and round trip itinerary fares where the length of stay is less than eight days. The fares recorded correspond to the cheapest ticket available for purchase. The second web service returns flight availabilities by allowing users to query real-time seat maps as well as look up detailed fare information. I compare the time series of seat maps to derive seat availabilities and thus, recover bookings across time. The data set contains fare and flight availability data for ten markets collected over a six month period in 2012. In total, the sample contains 1,328 flight departures and more than 80,000 one-way fare/seat map observations.

In the following subsections, I highlight features of the data. I first discuss route selection. I then confront the issue that seat maps may not accurately

[8] I define an itinerary to be a routing, airline, flight number(s), and departure date(s) combination.

reflect true flight loads. I perform two analyses that suggest the measurement error in seat maps is likely to be small. I then provide summary statistics for the sample. The last subsection documents preliminary evidence in the data. I document that: (i) prices fluctuate across time, but systematic fare increases are common, (ii) remaining capacity is important in explaining the variation in fares, (iii) there is no evidence that consumers anticipate systematic fare hikes.

### 1.2.1 Route Selection

I select markets to study using historical DB1B tables. These publicly available tables contain a 10% of domestic ticket purchases and are at the quarterly level. I define a market in the DB1B as an origin, destination, quarter. I single out markets where:

(i) there is only one carrier operating;[9]

(ii) there is no nearby alternative airport;

(iii) at least 95% of flight traffic is not connecting to other cities;

(iv) total quarterly traffic is greater than 3,000 passengers;

(v) total quarterly traffic is less than 30,000 passengers;

(vi) there is high nonstop traffic.

One important issue with using seat maps is figuring out which itinerary and hence which fare, to attribute to each seat map change. Since airlines offer extensive networks, the disappearance of a single seat could be associated with one of several thousand possible itineraries. This is an important consideration since

[9] 70.12% of routes in the US are monopoly.

the pricing of feeder routes tends to be different than main routes. Criteria (iii) addresses this by selecting routes where most traffic is not connecting to other cities. Criteria (iv) corresponds to routes with less than 75% load factor (seats occupied / capacity) of a daily 50-seat aircraft. This criteria removes routes with irregular service. Criteria (v) removes most routes with greater than three flights a day. I implement this criteria to keep the data collection process manageable. I then look for routes with (vi) high nonstop traffic. This criteria is important for establishing the relevant outside option in the demand model. In the data, (vi) is negatively correlated with distance ($\rho = -.5$). Cities with very high nonstop traffic percentages tend to be short flights. Given such short distances, many consumers may choose to instead drive. At the same time, markets with large distances typically have lower nonstop traffic percentages, meaning more consumers purchase tickets such as one stop.

I select five city pairs, or ten directional routings (markets), given the selection criteria above. All directional routings either originate or end in Boston, MA. The other cities are: Portland, OR, San Diego, CA, Austin, TX, Kansas City, MO, and Jacksonville, FL. The selected markets have close to 100% direct traffic, meaning very few passengers connect to other cities. The percent is nonstop traffic ranges between 40% and 60%. Three of the five city pairs are operated by JetBlue. The other airlines in the sample are Delta Air Lines and Alaska Airlines.[10] The selected markets are all low frequency as at most two flights are operated in either direction daily. Most days see a single flight frequency with double daily service on select routes during peak weeks of the summer or on weekends.[11]

[10]At the time of data collection, flights between Kansas City and Boston were operated by regional carriers on behalf of Delta Air Lines. Since Delta Air Lines determines the fares for this market, I collectively call these regional carriers Delta.

[11]This is not exceptional. The average number of frequencies across routes in the US is 1.95 flights per day. Over 60% of US routes see a single fight a day.

Two other features of the data are worth being noted. First, both JetBlue and Alaska price itineraries at the segment level. Consumers wishing to purchase round-trip tickets on these carriers in fact purchase two one-way tickets. As a consequence, round-trip fares in these markets are exactly equal to the sum of the corresponding one-way fares. Since fares must be attributed to each seat map change, this feature of the data makes it easier to justify the fare involved. Second, JetBlue does not oversell flights, while most other air carriers do.[12] Since most of the markets I study are operated by JetBlue, in the model, I assume firms do not oversell.

### 1.2.2 Inference and Accuracy of Seat Maps

A seat map is a graphical representation of occupied and unoccupied seats for a given flight at a select point in time before the departure date. Many airlines that have assigned seating present seat maps to consumers during the booking process. When a consumer books a ticket and selects a seat, the seat map changes – an unoccupied seat becomes occupied. The next consumer wishing to purchase a ticket on the flight is offered an updated seat map and has a choice amongst the remaining unoccupied seats. By *differencing* seat maps across time – in this case daily – inferences can be made about bookings.

Figure 1.1 presents a sample seat map. The seat map indicates occupied seats in solid blue. The unshaded blocks correspond to unoccupied seats. Seats with a “P” are available seats, but classified as premium. These seats are toward the front of the aircraft or seats located in exit rows. Finally, the seat map indicates seats currently blocked by the airline with “X”s. Seats that are blocked are usually

[12] On the legal section of the JetBlue website, under Passenger Service Plan: “JetBlue does not overbook flights. However some situations, such as flight cancellations and reaccommodation, might create a similar situation.”

**Figure 1.1:** An example seat map. The white blocks are unoccupied seats, the blue blocks are occupied, the blocks with the X's are blocked seats, and the blocks with a "P" are premium, unoccupied seats.

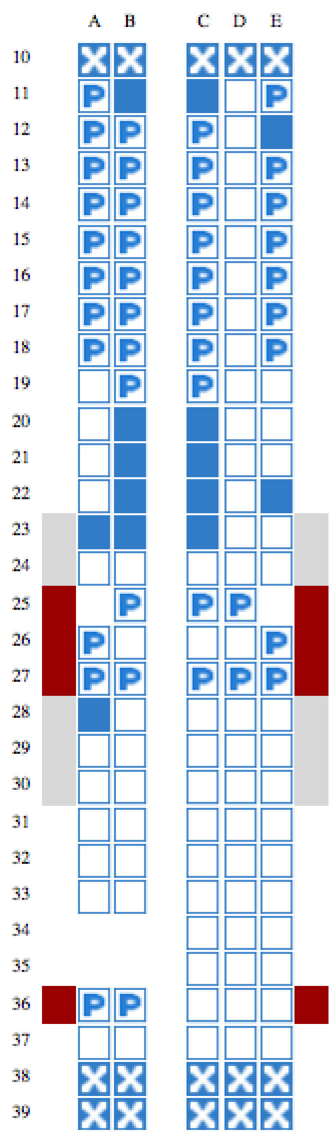

not disclosed on airline websites; however, I am able to capture this data through the web service used. Seats may be blocked due to crew rest, weight and balance, because a seat is broken, or because the airline reserves handicap accessible seats until the day of departure. In addition, seat blocking may be used to encourage consumers to purchase tickets or upgrade as they give the impression the cabin is closer to capacity. However, the data suggests that airlines predominantly block seats in exit rows and at the front and/or back of the cabin until closer to the departure date or when bookings demand additional seats. The decision is dynamic as over 70% of the flights in the sample experience changes in the number of blocked seats.[13] For every seat map collected, I aggregate the number of occupied, unoccupied, and blocked seats. I compare the aggregate counts across days to determine net bookings by day before departure.

Unfortunately, seat maps may not be accurate representations of true flight loads. This is especially problematic if consumers do not select seats at the time of booking. This measurement error would systematically understate sales early on, but then overstate last minute sales when consumers without existing seat assignments are assigned seats. From a modeling perspective, this measurement error would lead to an overstatement of the arrival of business consumers. Ideally, the severity of measurement error of my data can be assessed by matching changes in seat maps with bookings, however this is impossible with the publicly available data. In order to gauge the magnitude of the measurement error in using seat maps, I perform two different data analyses, which are only briefly discussed here,

[13] I do not model the decision to block/unblock seats; however, I do take this information into account when determining bookings. Knowing which seats are blocked is important because it allows me to distinguish between consumers canceling or purchasing tickets and airlines adjusting the supply of seats. For example, if it were not possible to distinguish between blocked and occupied seats, if an airline unblocks six seats, I would erroneously conclude six passengers canceled tickets.

but are detailed in the Appendix.

First, I match monthly enplanements using my seat maps with actual monthly enplanements reported in the T100 Segment tables. By comparing these aggregate measures, I find my seat maps understate true enplanements by 0.81% of load factor at the monthly level. To investigate the measurement error at the observation, or flight-day level, I create a separate data set by collecting information from an airline that provides seat maps as well as reported flight loads. With this data, I find seat maps understate reported load factor by 2.3%, with a range of 0% to 4% by day before departure. These two analyses suggest the measurement error associated with the seat maps in the sample is likely to be small. I proceed by using the capacity of the plane minus the number of occupied and blocked seats as the number of seats available in the proceeding analysis.[14]

### 1.2.3 Summary Statistics

Summary statistics for the data sample appear in Table 1.1. The average oneway ticket in my sample is \$282 whereas the average roundtrip fare is \$528. The discrepancy in one-way and round-trip fares can be attributed to the flights operated by Delta, since Delta does not price at the segment level but both Alaska and JetBlue do.

Reported load factor is the number of occupied seats divided by capacity the day flights leave, and is reported between 0 and 1. In my sample, the average load factor is 85%, ranging from 77% to 89% by market. The booking rate corresponds to the mean difference in occupied seats across consecutive days. I find the average

[14]By treating blocked seats as occupied, I find the monthly load factors for my routes exceed 95%, which is inconsistent with the reported carrier statistics found in the T100. Treating blocked seats as unoccupied results in a 0.81% difference in load factor with reported carrier data.

**Table 1.1:** Summary statistics for the sample.

| Variable | Mean | Std. Dev. | 5th pctile | 95th pctile |
|---|---|---|---|---|
| Oneway fare ($) | 282.23 | 118.03 | 129.80 | 498.80 |
| Roundtrip fare ($) | 528.09 | 202.30 | 279.60 | 917.60 |
| Load factor | 0.85 | 0.08 | 0.69 | 0.98 |
| Daily booking rate | 0.78 | 1.61 | 0 | 4 |
| Daily fare change ($) | 5.77 | 52.79 | -60.00 | 87.00 |
| Daily fare change rate | 0.35 | 0.10 | 0.17 | 0.52 |
| Unique fares (per itin.) | 12.76 | 3.50 | 7 | 19 |

$N_{\text{oneway}} = 80,550$

booking rate to be 0.78 seats per day, per flight. At the 5th percentile, zero seats per flight are booked a day, and at the 95th percentile, four seats per flight are booked a day. Airline markets are associated with *low* daily demand as 56.8% of the seat maps in the sample do not change across consecutive days. The fare change rate is an indicator variable equal to one if the itinerary fare changes across consecutive days. I find the daily rate of fare changes to be 35%, so that the itineraries in my sample typically change price 21 times in 60 days.

Due to institutional details concerning airline pricing practices, only a discrete number of fares are seen in the data. The last row indicates the number of unique fares per itinerary. On average, each itinerary reaches 12.7 unique fares, and given that the average itinerary sees 21 fare changes on average, this implies fares fluctuate up and down usually several times within 60 days.

## 1.3 Preliminary Evidence from the Data

With a description of the main features of the data complete, I now move into documenting preliminary evidence from the data. This analysis provides additional details concerning the data, but it is also meant to motivate features of the model. First, I document the pricing patterns in the data. This descriptive analysis shows airlines commonly implement systematic fare hikes, but fares frequently change daily. Second, I show that remaining capacity is important in explaining the variation in observed fares. Finally, I investigate if there is evidence that consumers anticipate systematic fare hikes.

### Pricing Patterns

Figure 1.2 shows the frequency and magnitude of fare changes across time. The left panel indicates the fraction of itineraries that experience fare hikes versus fare discounts by day before departure, and the right panel indicates the magnitude of these fare changes (i.e. a plot of first differences, conditional on the direction of the fare change). For example, in the left plot, 40 days prior to departure, roughly 20% of fares increase and 20% of fares decrease. The remaining 60% of fares are held constant. Moving to the right panel, the magnitude of fare hikes and declines 40 days out is roughly \$50.

The left panel confirms fares change throughout time, including fare declines. The fraction of fares that decline over time is roughly U-shaped, increasing through roughly three weeks prior to departure, peaking at 20%, and then declining to roughly 10% the day before departure. Note that well before the departure date, the number of fare hikes and declines is roughly split even. The fraction of itineraries that experience fare hikes increases over time. There are four noticeable

**Figure 1.2:** Fraction and magnitude of fare changes by day before departure

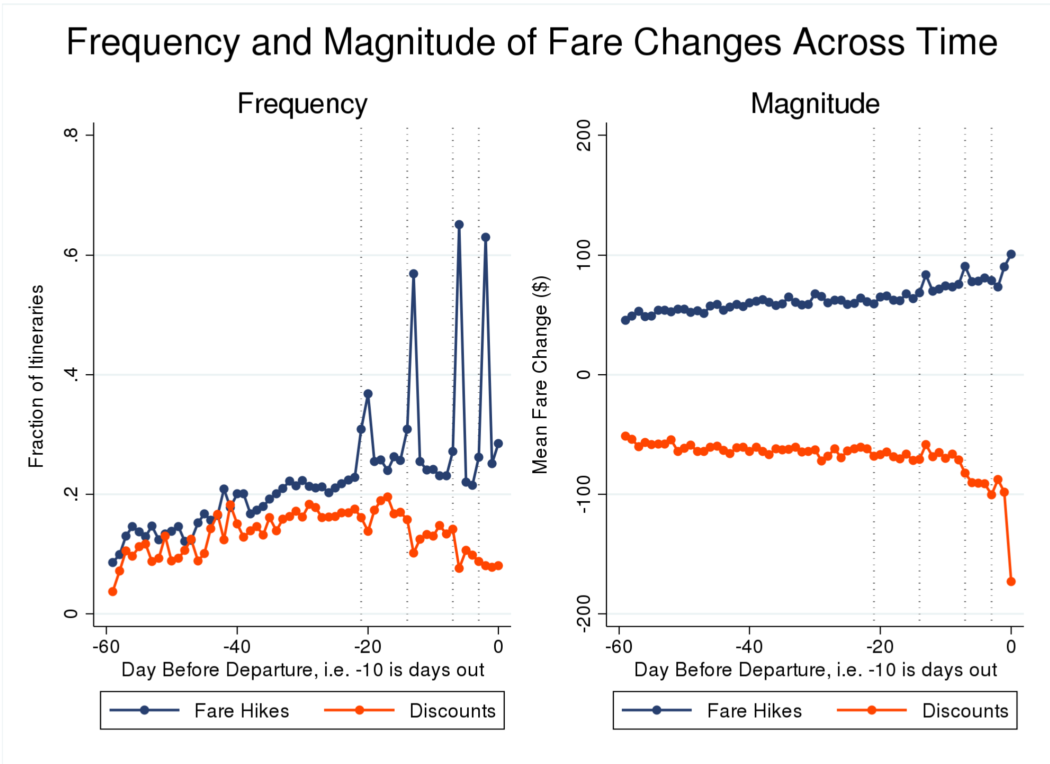


Notes: (left) Fraction of fares that increase/decrease across consecutive days. (right) Magnitude of fare changes across time.

jumps in the line indicating fare hikes. These jumps correspond to crossing 3, 7, 14 and 21 days prior to departure, or when advance purchase discounts placed on many tickets expire.[15] The use of advance purchase discounts (APDs) is consistent with the story of intertemporal price discrimination, since fares increase unconditional on remaining capacity. Surprisingly though, the use of advance purchase discounts is not universal. Only 35% of itineraries experience fare hikes at 21 days and less than 60% increase at 14 days. Just under 70% of itineraries see an increase in fare when crossing the 7-day APD requirement.

[15]Advance purchase discounts are sometimes placed at 4, 10, and 30 days prior to departure, but this is not the case for the data I collect.

The right panel shows the magnitude of fare changes towards the departure date, conditional on the direction of the change (increase/decrease). There are two findings worth mentioning. First is that the magnitude of fare hikes and declines are similar – at around \$50 but increasing in time. Second, the magnitude of fare increases when crossing APD days is similar to the magnitude of fare increases on other days.

Figure 1.2 shows just how dynamic airline pricing is. Fares are constantly increasing and decreasing. Consistent with the theory of stochastic demand pricing, fire-sales do occur – roughly 10% of the itineraries in these sample decrease over \$175 within 24 hours of the departure date. On the other hand, fares systematically increase at routine intervals, which is consistent with airlines segmenting the market between business and leisure customers.

Moving to statistics in levels, Figure 1.5 plots the mean fare and mean load factor (seats occupied/capacity) by day before departure. The plot confirms the overall trend in prices is positive, with fares increasing from roughly \$225 to over \$375 in sixty days. The noticeable jumps in the fare time series occur when crossing the APD fences noted in Figure 1.2. At sixty days prior to departure, roughly 45% of seats are already occupied; consequently, I observe about half the bookings on any given flight. The booking curve for flights in the sample is smooth across time, leveling off roughly three days prior to departure at 85%, which closely matches monthly enplanement totals found in the publicly available T100 Segment data. The fact that fares tend to double but consumers still purchase tickets is suggestive evidence that consumers of different types purchase tickets towards the date of travel.

**The Role of Remaining Capacity**

A key source of the variation in fares can be attributed to the scarcity of seats. I provide descriptive evidence of this two ways. First, I compare fares and load factor when two flights are offered a day. I calculate the difference in fare ($\Delta^{\text{fare}}$) and difference in load factor ($\Delta^{\text{LF}}$) across the two flight options by day before departure. When calculating the difference, I assign the first flight of the pair to be the flight with the greater number of seats occupied. This implies $\Delta^{\text{LF}} > 0$. By comparing fares for the two flights by day before departure, I control for systematic fare changes associated with intertemporal price discrimination.

**Figure 1.3:** The role of capacity in explaining fare variation

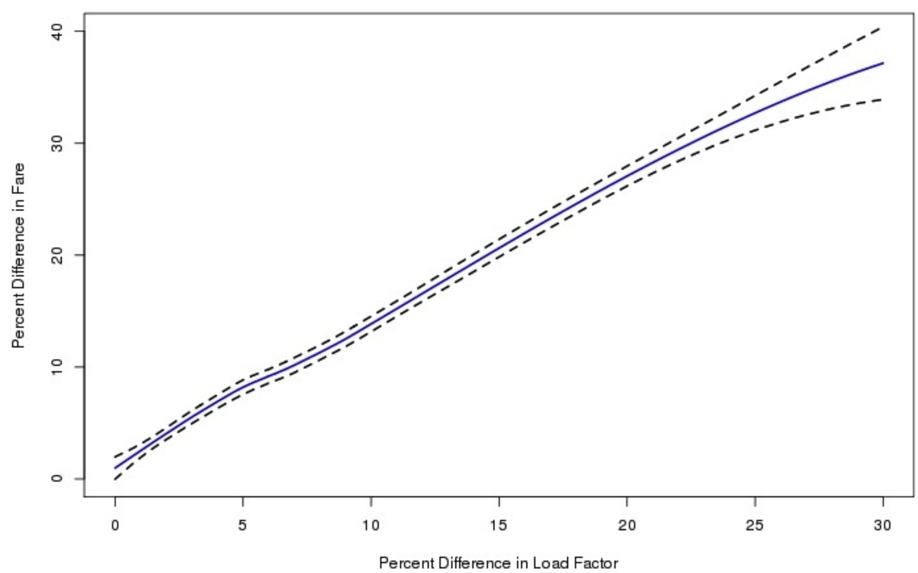


Nonparametric regression of difference in fare by difference in load factor from comparing fares and load factors for one-way itineraries where two flights are available. $n = 20,062$.

Figure 1.3 shows nonparametric fitted values as well as the 95% confidence interval of these calculations in percentage terms. The plot suggests that when the two options have the same number of seats occupied, the average difference in fare is close to zero. If one option is 20% more full, the flight which is more full is also 25%, or roughly $60, more expensive. The line remains upward sloping throughout observed differences in load factor, where at the extreme, a flight that is 30% more full is also 35% more expensive.[16]

I perform a similar analysis using the entire sample by calculating the mean difference in fare and load factor at the flight number, day before departure level. Figure 1.4 shows nonparametric fitted values of this procedure. The line is again upward sloping across differences in load factor, where flights with lower load factor compared to the average are also less expensive. Likewise, flights that that have a higher occupancy compared to the average are also more expensive.

**Figure 1.4:** Lowess of mean differences in fare and load factor.

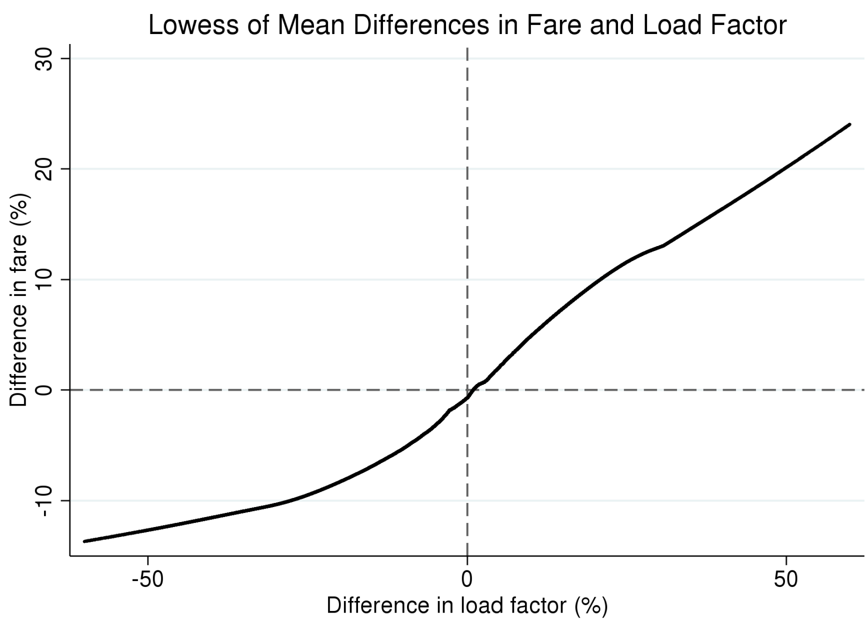


[16]Applying similar methodology to round trip itineraries with two options – a single outbound flight and the choice of two return flights, or two outbound flights and a single return flight – yields similar results.

**Figure 1.5:** Mean fare and load factor by day before departure.

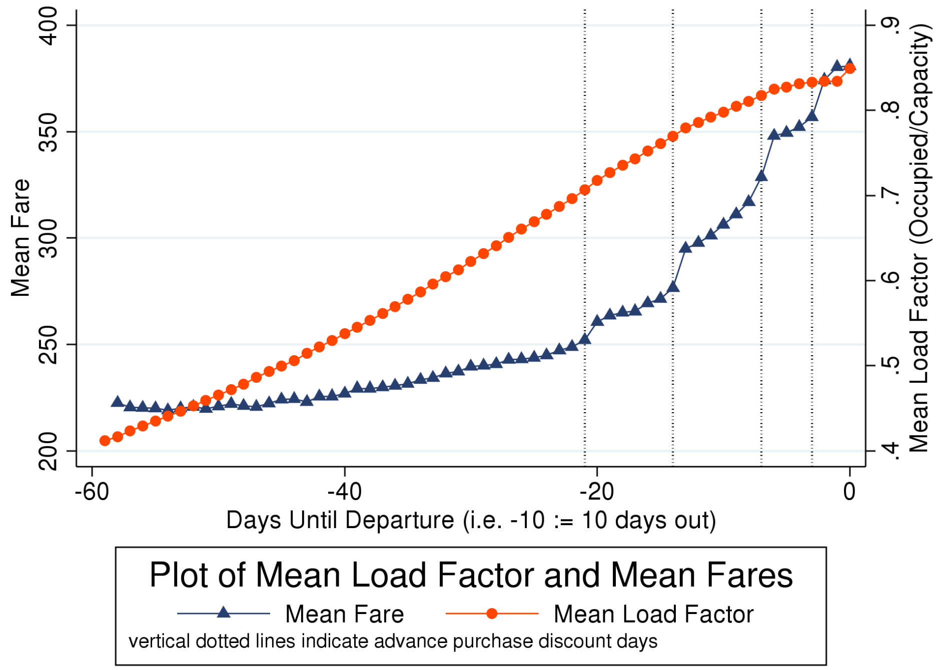


### Do consumers dynamically substitute across booking days?

The booking curve of flights plotted in Figure 1.5 is smooth, even though fares tend to increase by nearly \$50 when crossing the advance purchase discount thresholds. This result is surprising since bunching in sales should be seen before the discounts expire if consumers anticipate systematic fare hikes. This is not the case as the only noticeable jump in load factor appears right before flights leave. I test for discontinuities in the booking curve using regressions of the following form:

$$\text{LF} = \overrightarrow{\text{APD}} + m(t) + u + \varepsilon,$$

where $\overrightarrow{\text{APD}}$ are dummy variables corresponding to the day before advance purchase discounts expire, $m(t)$ is a flexible function in time, and $u$ are other fixed effects. Regression results appear in Table 1.2. Across all specifications, I find that none of the advance purchase discount dummies are significant; moreover, the 21 and 3 day advance purchase discount dummies are negative, which is inconsistent with bunching.

The fact that the there is no evidence of bunching suggests that consumers either do not anticipate the fences or are possibly restricted in some other way from being able to purchase before the advance purchase discounts expire. Alternatively, it could be the case that consumers substitute to a different departure date, however this does not explain why the booking rate is similar after the discounts expire. I use this feature of the data to motivate a demand system where consumers do not dynamically substitute across booking days. Further, I show after estimating the model that since prices tend to increase across time, this provides little incentive for consumers to wait to purchase tickets.

## 1.4 A Model of Dynamic Airline Pricing

In this section, I write down a structural model of dynamic airline pricing where firms face a stochastic arrival of consumers, and the mix of consumer types – corresponding to leisure and business consumers – is allowed to change over time. To make the analysis tractable, I incorporate the following simplifications in the model. First, the model studies the pricing decisions of airlines operating in monopoly markets. In this way, the paper can focus on intertemporal price discrimination and stochastic demand, and avoid the complexities of modeling oligopolistic competition. As noted earlier, a large fraction of airline routes are monopoly. Second, I assume that when consumers first learn about their interest in travel, all travel plan uncertainty is immediately resolved. Consumers pay a

**Table 1.2:** Dynamic Substitution Regressions

| | Model Specification | | |
|---|---|---|---|
| | (1) | (2) | (3) |
| $\text{APD}_3$ | $-0.170$<br>(0.109) | $-0.171$<br>(0.109) | $-0.163$<br>(0.108) |
| $\text{APD}_7$ | 0.032<br>(0.077) | 0.032<br>(0.077) | 0.031<br>(0.077) |
| $\text{APD}_{14}$ | 0.033<br>(0.041) | 0.033<br>(0.041) | 0.032<br>(0.041) |
| $\text{APD}_{21}$ | $-0.069$<br>(0.062) | $-0.069$<br>(0.062) | $-0.071$<br>(0.061) |
| $m(t)$ | spline | spline | spline |
| Flight FEs | | X | X |
| d.o.w. flight FEs | | | X |
| d.o.w. purchase FEs | | | X |
| obs | $80,550$ | $80,550$ | $80,550$ |
| $R^2$ | 0.605 | 0.724 | 0.739 |

Route clustered standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

fixed cost to come back and check on fares. Since fares tend to increase over time, the combined effect of these assumptions reduce the consumer problem to a static choice of either purchasing a ticket the day the consumer's travel plans are realized, or not buying at all. Third, while in the actual airline business two consumers buying on the same day might face different fares (e.g. from different fare categories, or buying at different times of the day, or from different web sites), in the model a single fare is offered to consumers each day. Fourth, I model consumers as purchasing one-way tickets. A consumer interested in a round-trip can be thought of as two consumers interested in one-way tickets. As noted earlier, in the collected data, round-trip fares are very close to the corresponding one-way fares. Fifth, I assume firms utilize a finite set of fares. Firms take the set as given, and choose a single fare to offer for each flight daily. This assumption captures the fact that only a discrete number of fares is seen in the data. Next, I assume firms do not oversell flights since most of the routes studied are operated by an airline that does not oversell flights. If demand exceeds remaining capacity, tickets are rationed. Finally, since most markets see one flight departure a day, I write down the down with a single flight option. The model does extend to multiple flight frequencies.

### 1.4.1 Consumer Problem

A market is defined as an origin, destination, search date, departure date. At time $t$, $\tilde{M}_t \in \mathbb{N}$ consumers arrive interested in traveling between the two cities.[17] For each of these newly-arrived consumers, all uncertainty about travel preferences is resolved at this point. This assumption differs from Lazarev (2013) and earlier theoretical work, including Gale and Holmes (1993), where consumer uncertainty exists and this uncertainty can be resolved by delaying purchase until closer to the departure date. In my model, when the date $t$ consumers arrive, they choose to either purchase a ticket on an available flight, exit the market, or pay a cost $\overline{\phi}$ to search again the following day. Throughout the rest of this section and when estimating the model, I assume $\overline{\phi}$ is sufficiently large so that waiting is never optimal. In Section 1.6.3, I calculate a bound on $\overline{\phi}$ for this assumption to hold, and show that it is relatively small – since fares tend to increase, there is little

[17] More broadly, since $\tilde{M}_t$ is market specific, $\tilde{M}_{t,d}$ consumers look to travel on $d, \tilde{M}_{t,d'}$ consumers look to travel on date $d'$, etc.

incentive to wait.

Preferences of consumers follow the earlier two-type consumer approach to study airline markets (Berry, Carnall, and Spiller (1996) and Berry and Jia (2010)). Consumer $i$ is either a business traveler or a leisure traveler. With probability $\gamma_t$, consumer $i$ is a business traveler. Consumer $i$ receives indirect utility from product characteristics $X_t \in \mathbb{R}^K$ and price $p_t$. Let $\beta_i, \alpha_i$ denote the taste parameters over $X_t$ and $p_t$, respectively. Let 0 denote the outside option. Each consumer $i$ receives idiosyncratic preference shocks for choosing to travel ($\varepsilon_{i1t}$). Let $\varepsilon_{i0t}$ be the taste shock for the outside option. Following the discrete choice literature, consumer $i$ chooses to fly iff

$$U_{i1t}(X, p, \beta, \alpha, \varepsilon) \geq U_{i0t}(X, p, \beta, \alpha, \varepsilon).$$

I assume utility is linear in product characteristics and of the form

$$U_{i1t} = X_t\beta_i - \alpha_i p_t + \epsilon_{i1t}.$$

Let $\boldsymbol{\varepsilon}_{it} = (\varepsilon_{i0t}, \varepsilon_{i1t})$ be the idiosyncratic preference shocks for products in the choice set for consumer $i$. Define $y_t = \left(\alpha_i, \beta_i, \boldsymbol{\varepsilon}_{it}\right)_{i \in 1,..,\tilde{M}_t}$ to be the vector of preferences for the consumers that enter the market. The demand for flight $j$ at $t$ is a mapping given the fare offered ($p_t$) and consumer preferences ($y_t$), defined as

$$Q_t(p, y) := \sum_{i=0}^{\tilde{M}_t} 1\left[U_{i1t} \geq U_{i0t}\right] \longmapsto \{0, ..., \tilde{M}_t\},$$

where $1(\cdot)$ denotes the indicator function.[18]

[18] Here, and for remainder of the paper, I suppress the dependence of $Q$ on $X$ to emphasize the role of price. $X$ may contain a flight characteristics such as whether a particular flight is a

Let $s_t \in \mathbb{N}$ denote the remaining capacity for the flight at time $t$. Demand is integer-valued; however, it may be the case that more consumers want to travel than there are seats remaining, i.e. $Q_t(p, y) > s_t$. Since firms are not allowed to oversell, in these instances, I assume remaining capacity is rationed by random selection. If demand exceeds capacity for the flight, consumers that wanted to travel are randomly shuffled. The first $s_t$ are selected and the rest receive their outside option.

Recall that a market is departure date-specific. Although the model assumes consumers arrive and purchase a single one-way ticket, the model does allow for round-trip ticket purchases in the following way. At time $t$, two consumers arrive. One consumer is interested in leaving on date $d$, and another consumer is interested in returning on date $d'$. The consumers receive idiosyncratic preference shocks for each of the available flights, and choose which tickets to purchase. Since the round-trip fares in the sample are very close to the sum of the corresponding one-ways, there is little measurement error in this approach.

### 1.4.2 Monopoly Pricing Problem

A monopolist sells tickets for a single flight over a finite horizon. Period 0 corresponds to the first period of sale, and $t = T$ is the time at which the flights depart. The flight has an initial capacity constraint of $s_0$ seats, which is exogenous to the model. I assume the cost of operation is sunk, so the only cost facing the firm is the opportunity cost of selling seats. I assume the marginal cost per passenger is zero, which is reasonable as almost all flight costs are not influenced by the number of seats occupied.

morning or afternoon flight. In addition, $X$ may contain an indicator for a particular departure date which would allow the firm to use peak-load pricing.

The firm maximizes expected discounted revenues. The firm knows on average how many business and leisure travelers will search for tickets over time, but is unsure exactly how many consumers will arrive in any given period. Since fares are posted before consumers arrive, the firm forms expectation over present and future revenues. The decision rule depends on the number of seats remaining, number of periods left to sell, and an unobserved error term that the firm sees, but the econometrician does not. This term, which is price specific, reflects the fact that in data, there are flights in which the firm charges different prices given the same number of seats remaining and time left to sell.[19] Let $\omega$ denote this term, which is separable from demand.

Since excess demand is rationed, by charging prices $p_t$ and receiving consumers $y_t$, the firm can sell at most $\min(Q_t(p,y), s_t)$ seats. This implies the law of motion for remaining capacity can be written as

$$s_{t+1} = s_t - \min\Big(Q_t(p,y), s_t\Big).$$

Define the incremental revenue for the firm to be

$$R_t(p,y,s) = \min\Big(Q_t(p,y), s_t\Big) \cdot p_t.$$

I assume the firm is restricted to choosing prices $p_t \in \mathcal{P}$. If the flight is has no seats remaining, the price is infinite.

To write the firm's problem as a dynamic program, define the value function, $V_t(s,\omega)$ to be the discounted expected revenue left to go with remaining capacities

[19]The decision may also depend on other observed (at least to the firm) states $Z_t$. Note that $X \subset Z$ since all product characteristics that enter the consumer problem affect purchasing behavior, and consequently, affect expected revenues. Like $X$, I suppress the notation of $Z_t$ to highlight the importance of remaining capacity $s_t$ in the firm problem.

$s_t$ and $t$ periods to sell. The restrictions on capacity form two boundary conditions on the value function. The first is that with zero seats remaining, the firm cannot capture additional revenue, which is $V_t(0,\omega) = 0$. Second, unsold seats the day the flights leave must be scrapped with zero value implying $V_T(s,\omega) = 0$.

The firm's problem can be written recursively as

$$V_t(s,\omega) = \max_{p_t \in \mathcal{P}} \mathbb{E}\Big[R_t(p,y,s) + \omega_p + \rho V_{t+1}(s',\omega' \,|\, y,s,\omega)\Big]$$

$$\text{s.t.} \quad \begin{cases} s_{t+1} = s_t - \min\Big(Q_t(p,y), s_t\Big) \\ V_T(s,\omega) = 0 \\ V_t(0,\omega) = 0 \\ s_0 \text{ given,} \end{cases}$$

where $\rho \in [0,1]$ is the discount factor (which I set equal to 1).

The value function of the firm illustrates the important interactions between intertemporal price discrimination and dynamic adjustment to stochastic demand. If business consumers are less price sensitive and the proportion of business consumers increases as the departure date approaches, the firm can extract more revenue by increasing fares over time. However, since the arrival of consumers is uncertain, it is possible that a flight may sell out early. This creates an incentive for the firm to save seats until close to the departure date. Looking at the value function, if capacity becomes scarce early on, the firm can increase fares to reduce current period expected sales and revenues, but increase the probability that seats will remain the following day. For example, if the firm sets $p_t = \infty$, then current period expected revenues would be zero, but the probability that $s_t = s_{t+1}$ would be one. Hence, the firm would enter the subsequent period under $V_{t+1}(s,\omega)$ which may be the optimal pricing strategy if the firm expects high-valuing consumers to arrive closer to the departure date. Alternatively, it may be the case that the firm receives a sequence of low demand realizations. In order to not leave as many seats unfilled, the firm may opt to lower prices.

Recall I assume separability in demand and the unobserved error term. The firm is uncertain about demand and forms expectation over $y_t$. Expected sales for the flight in any given period is

$$Q_t^e(s,p) = \int_{y_t} \min\Big(Q_t(p,y), s_t\Big) dF(y_t).$$

By charging a price $p_t$, the firm has a probability distribution over remaining capacity tomorrow – which is integer valued. This means the dynamic program of the firm can be written as

$$\begin{aligned} V_t(s,\omega) &= \max_{p_t \in \mathcal{P}} \Big[p_t Q_t^e(s,p) + \omega_p + \rho \mathbb{E}[V_{t+1}(s'\omega'|s,\omega)]\Big] \\ &= \max_{p_t \in \mathcal{P}} \left[p_t Q_t^e(s,p) + \omega_p + \rho \int_{\omega'} \sum_{j=0}^{s_t} \Pr(s_{t+1} = j \mid p_t, s_t) V_{t+1}(j,\omega') dF\omega'\right]. \end{aligned}$$

such that $s_{t+1} = s_t - \min\Big(Q_t(p,y), s_t\Big)$ and the two boundary conditions hold.

## 1.5 Econometric Specification & Estimation

I first parameterize the demand model and derive analytic expressions for purchase probabilities. The firm's pricing decision can be written as a dynamic discrete choice model. In the last section, I discuss the estimation approach and identification strategy.

### 1.5.1 Consumer Demand

First I derive the purchase probabilities for the consumer demand system. Recall that the preferences of consumers that arrive to the market are

$$y_t = \left(\alpha_i, \beta_i, \boldsymbol{\varepsilon}_i\right)_{i \in 1,..,\tilde{M}_t}.$$

The number of consumers, as well as the relative proportion of each type, that will arrive is not observed by the firm before pricing (or by the econometrician). In order to proceed, I integrate over the distribution of $y_t$.[20]

Assumption 1: Consumer idiosyncratic preferences are distributed Type-1 Extreme Value (T1EV).

I assume the outside option yields a normalized utility $u_{i0t} = \varepsilon_{i0t}$. The distributional assumption on the idiosyncratic preferences leads to the frequently used conditional logit demand system. Since there is only a single product in the choice set,

$$\varsigma_{1t}^i := \Pr(i \text{ wants to purchase ticket } |\text{type} = i) = \frac{1}{1 + \exp(-X_t\beta_i + \alpha_i p_t)}.$$

The discrete choice literature typically does not model capacity constraints. Given my environment, it is possible that consumers wish to purchase a product, but are unable to due to a violation of the capacity constraint. As noted in the previous section, if demand exceeds available capacity, demand is rationed, and consumers who are not selected for travel receive their outside option. This is why the purchase probabilities state "$i$ wants to purchase ticket" instead of

[20] I proceed this way because individual decisions are not observed. Using the EM-algorithm would be another way to address this issue. See: Talluri and Van Ryzin (2004) and Vulcano, van Ryzin, and Chaar (2010).

"$i$ purchases ticket". Let $\varsigma_{0t}^{i}$ denote the share of the outside option.

In practice, appropriate data to enter $X$ would be variables such as if a flight is operated on a holiday or weekend. Preferences over these characteristics could also vary across consumer types and/or routes. While the model allows for these covariates, I assume consumers only care about price, which is not unreasonable as the routes studied are monopoly with a single flight daily.[21]

Assumption 2: There are two consumer types. Let $B$ denote the business type and $L$ to denote the leisure type. The probability of a consumer being type-$B$ is $\gamma_t$. Then $\gamma_t s_t^B$ defines the probability that a consumer is of the business type and wants to purchase a ticket. Consider the market share of the outside good at time $t$. Conditional on $k$ consumers arriving, the probability of zero consumers wishing to travel is

$$\Pr\Big(Q_t = 0 \,|\, k \in \mathbb{N}\Big) = \sum_{i=0}^{k} \binom{k}{i} \Big((1-\gamma_t)\varsigma_{0t}^{L}\Big)^{i} \Big(\gamma_t \varsigma_{0t}^{B}\Big)^{k-i}.$$

For example, let $k = 2$. Then conditional on two consumers arriving, both consumers want to purchase the outside option. Either both consumers are leisure travelers, both consumers are business travelers, or one of each type arrive. The first two situations correspond to $i = 2$ and $i = 0$, respectively since $((1-\gamma)\varsigma_{0t}^{L})^2$ is the probability of two leisure consumers arriving and wanting to choose the outside option, and $(\gamma\varsigma_{0t}^{B})^2$ is the probability two business consumers want to choose the outside option. Lastly, it could be the case that one business and one leisure consumer arrives. There are two possibilities, the first consumer is the business consumer, or vice versa. Hence, $2[(1-\gamma)\varsigma_{0t}^{L}][\gamma\varsigma_{0t}^{B}]$ enters the probability.

[21]When I solve the firm problem, adding additional characteristics makes the problem very computationally demanding. I do allow all parameters to be route-specific.

Next, integrating over the arrival process of consumers yields

$$\Pr\Big(Q_t = 0\Big) = \sum_{k=0}^{\infty} \Pr(Q_t = 0, \tilde{M}_t = k) = \sum_{k=0}^{\infty} \Pr(Q_t = 0 | \tilde{M}_t = k) \Pr(\tilde{M}_t = k).$$

Assumption 3: Consumers arrive according to a Poisson process.[22]

Let $\tilde{M}_t \sim \text{Poisson}(\mu_t)$ denote the Poisson process. The probability mass function for the Poisson distribution is $(\mu_t)^k e^{-\mu_t}/k!, k \in \mathbb{N}$. With this assumption, $\Pr\Big(Q_t = 0\Big)$ has an analytic form and can be written as

$$\Pr\Big(Q_t = 0\Big) = \sum_{k=0}^{\infty} \frac{\mu_t^k e^{-\mu_t}}{k!} \sum_{i=0}^{k} \binom{k}{i} \Big((1-\gamma_t)\varsigma_{0t}^L\Big)^i \Big(\gamma_t \varsigma_{0t}^B\Big)^{k-i}.$$

Implicitly, this depends on both price and capacity remaining. With zero seats remaining, price is infinite so $\varsigma_{0t}^L = \varsigma_{0t}^B = 1$, which implies $\Pr\Big(Q_t = 0\Big) = 1$, as expected. Next, consider the probability of selling a positive number of seats, but not selling out, conditional on $k$. This probability can be written as

$$\begin{aligned} \Pr\Big(Q_t = q \,|\, k \geq q, q < s\Big) = & \binom{k}{q} \sum_{\ell=0}^{q} \binom{q}{\ell} \Big(\gamma_t \varsigma_{1t}^B\Big)^{\ell} \Big((1-\gamma_t)\varsigma_{1t}^L\Big)^{q-\ell} \times \\ & \left[ \sum_{i=0}^{k-q} \binom{k-q}{i} \Big((1-\gamma_t)\varsigma_{0t}^L\Big)^i \Big(\gamma_t \varsigma_{0t}^B\Big)^{k-q-i} \right]. \end{aligned}$$

In the formula above, the terms following the first sum correspond to all the combinations of having $q$ consumers wanting to purchase a ticket. The terms following the second summation (second line) correspond to all the combinations of the remaining $k - q$ consumers wanting to purchase the outside option. Finally, $k\text{C}q$ at the beginning of the equation sorts all the possible combinations of (want

[22]Given this assumption, the demand model closely follows Talluri and Van Ryzin (2004) and Vulcano, van Ryzin, and Chaar (2010), except this model has two consumer types.

to buy, do not want to buy) amongst the $q$ of $k$ consumers that wish to purchase tickets. Of course, selling $q$ seats requires at least $k = q$ consumers to enter so $\Pr(\text{sell } q \mid k < q, q < C, p) = 0$. Integrating over the Poisson arrival process results in an analytic expression for $\Pr(\text{sell } q \mid q < C, p)$, which is also a Poisson-Binomial mixture.

Demand is latent in situations where flights are observed to sell out since it is possible some consumers are forced to the outside option. These probabilities can be constructed based off the fact that at least $s_t$ seats are demanded. For example, if $s_t = 2$ and the flight is observed to sell out, that implies at least 2 seats were demanded, which also has an analytic expression of a Poisson-Binomial mixture. Since capacity is assumed to be monotonically decreasing, all purchase probabilities have been defined, which appear in Equation 1.1 - Equation 1.3 below. Collectively call these probabilities $f(s'|s,p,t)$.[23]

[23] To be clear, $s'$ denotes the number of seats sold in the current period given remaining capacity $s$.

$$\Pr\Big(Q_t \geq s \,|\, s, p\Big) = \sum_{q=s}^{\infty} \sum_{k=q}^{\infty} \frac{\mu_t^k e^{-\mu_t}}{k!} \binom{k}{q} \sum_{\ell=0}^{q} \binom{q}{\ell} \Big(\gamma_t \varsigma_{1t}^B\Big)^{\ell} \Big((1-\gamma_t)\varsigma_{1t}^L\Big)^{q-\ell} \times \tag{1.1}$$
$$\left[\sum_{i=0}^{k-q} \binom{k-q}{i} \Big((1-\gamma_t)\varsigma_{0t}^L\Big)^{i} \Big(\gamma_t \varsigma_{0t}^B\Big)^{k-q-i}\right]$$

$$\Pr\Big(Q_t = q \,|\, 0 < q < s, p\Big) = \sum_{k=q}^{\infty} \frac{\mu_t^k e^{-\mu_t}}{k!} \binom{k}{q} \sum_{\ell=0}^{q} \binom{q}{\ell} \Big(\gamma_t \varsigma_{1t}^B\Big)^{\ell} \Big((1-\gamma_t)\varsigma_{1t}^L\Big)^{q-\ell} \times \tag{1.2}$$
$$\left[\sum_{i=0}^{k-q} \binom{k-q}{i} \Big((1-\gamma_t)\varsigma_{t0}^L\Big)^{i} \Big(\gamma_t \varsigma_{t0}^B\Big)^{k-q-i}\right]$$

$$\Pr\Big(Q_t = 0 \,|\, s > 0, p\Big) = \sum_{k=0}^{\infty} \frac{\mu_t^k e^{-\mu_t}}{k!} \sum_{i=0}^{k} \binom{k}{i} \Big((1-\gamma_t)\varsigma_{t0}^L\Big)^{i} \Big(\gamma_t \varsigma_{t0}^B\Big)^{k-i}. \tag{1.3}$$

### 1.5.2 Solving the Firm's Problem

The firm's pricing decision depends on remaining capacity and time to sell. I assume fares are chosen from a discrete set, which is exogenous to the model. I define this set to be the set of observed fares in the data. This assumption accurately reflects that fares for any given flight tend to fluctuate between relatively few distinct prices. This assumption allows me to write down the firm's problem a dynamic choice choice model.

Following Rust (1987), I make the following assumptions:

Firm Assumptions:

(i) The choice shocks are distributed Type-1 Extreme Value (T1EV);

(ii) The per-period payoff function and choice shocks are separable;

(iii) Conditional independence is satisfied.

These assumptions collectively lead to a dynamic discrete choice model, or more specifically, a dynamic logit model. Due to conditional independence, the transition probabilities correspond to the probabilities derived in the last section.[24] In addition, these functions also can be used to define the expected per period revenues, since $\sum_{s'} f(s'|s,p,t) \cdot s' = Q^e(p,s)$.

Let $V_t(s,\omega)$ be the discounted expected revenue at state $(s_t, \omega_t)$. $V_t(s,\omega)$ solves

$$V_t(s,\omega) = \max \left\{ \begin{array}{c} v_t(p^1, s) + \omega(p_t^1) \\ v_t(p^2, s) + \omega(p_t^2) \\ \vdots \\ v_t(p^{|\mathcal{P}|}, s) + \omega(p_t^{|\mathcal{P}|}) \end{array} \right\},$$

where $\{v_t(p^k, s)\}_{k \in 1..|\mathcal{P}(s_t)|}$ are the choice specific value functions and can be written as

$$\begin{aligned} v_t(p^k, s) &= Q_t^e(s, p^k) p_t^k + \rho \mathbb{E}_{s',\omega'|s,p^k} V(s', \omega') \\ &= ER_t(p^k, s) + EV_t(p^k, s). \end{aligned}$$

The T1EV assumption along with conditional independence imply the expected value functions have a closed form and can be computed as

$$EV_t(p,s) = \int_{s'} \underbrace{\left[ \sigma \ln \left( \sum_{p' \in \mathcal{P}(s')} \exp \left( \frac{ER_{t'}(p', s') + EV_{t'}(p', s')}{\sigma} \right) \right) \right]}_{:= V_{t+1}(s')} f_t(s'|s,p) ds',$$

[24] That is, $f(s', \omega'|s, \omega, p, t) = f(\omega'|s', t) f(s'|s, p, t)$

where $\sigma$ is the scale parameter the choice shocks ($\omega$).[25] Lastly, the assumptions on the firm's problem imply the conditional choice probabilities also have a closed form and can be computed as

$$CP_t(p,s) = \frac{\exp\left(ER(p,s) + EV(p,s)/\sigma\right)}{\sum_{p' \in \mathcal{P}(s_t)} \left(\exp(ER(p',s) + EV(p',s))/\sigma\right)}.$$

### 1.5.3 Estimation Approach

In this section, I discuss estimation of the structural parameters. There are two possible approaches. The first strategy is to estimate the demand parameters without accounting for the optimal choice in price. This strategy would rely solely on the purchase probabilities derived in Section 1.5.1. However, as pointed out in Talluri and Van Ryzin (2004), for a specific price, the quantity of tickets sold in a given period can be higher either because of a higher overall arrival rate, or because consumers become more inelastic for a given arrival rate. The second approach is to estimate the model by accounting for the firm's choice in price. I take this approach and argue this information provides additional identification power of the structural parameters.

To account for the firm's pricing decision, I incorporate the dynamic discrete choice model of the firm (Section 1.5.2). With the assumptions in Section 1.5.2, given a set of flights ($F$) each tracked for ($T$) periods, the likelihood for the data

[25] Additionally,

$$\begin{aligned} v_t(p,s) &= ER_t(p,s) + \beta EV_t(p,s). \\ &= \sum_{s'=0}^{s} f(s'|s,p,t) p_t s' + \sum_{q=0}^{s} f(s'|s,p,t) V_{t-1}(s-s') \\ &= \sum_{s'=0}^{s} \Big( f(s'|s,p,t) \cdot \Big[ p_t s' + \beta V_{t-1}(s-s') \Big] \Big) \end{aligned}$$

is

$$\mathcal{L}(\text{data}|\theta) = \max_{\theta} \prod_F \prod_T CP_t(p, s) f(s'|s, p, t).$$

Note that under the first strategy, the likelihood for the data is $\prod_F \prod_T f(s'|s, p, t)$. By accounting for the firm's pricing decision, the likelihood is disciplined through the conditional choice probabilities, and that Bellman's equation is satisfied. One thing to note is that the parameters in the transition probabilities and the parameters in the per-period payoff function, which enter $CP$, are the same – they are both functions of $\theta$. For this reason, the transition probabilities cannot be estimated in a first stage (which would be just estimating the consumer demand model), as is typically done using this framework, and in alternative methods of estimating dynamic models, including Hotz and Miller (1993) and Bajari, Benkard, and Levin (2007).

To estimate the structural parameters of the model, I solve the dynamic program of the firm using mathematical programming with equilibrium constraints (MPEC) (see Su and Judd (2012)). The estimation procedure equates to solving the constrained maximization problem

$$\max_{\theta, EV} \sum_F \sum_T \ln(CP_t(p, s)) + \ln(f(s'|s, p, t))$$

$$\text{such that} \quad EV = \mathcal{T}(EV, \theta).$$

Recall that the arrival rate ($\boldsymbol{\mu}$) and mixture of consumer types ($\boldsymbol{\gamma}$) is allowed to change over time. I specify both sets of parameters to be continuous functions in time. I assign the Poisson arrival rates ($\boldsymbol{\mu}$) to be a 4th-degree polynomial series.

For the probability on types, I specify $\boldsymbol{\gamma}$ as a logistic function,

$$\gamma_t = \frac{\exp(\gamma^1 + \gamma^2 t)}{1 + \exp(\gamma^1 + \gamma^2 t)}, \gamma^2 \geq 0$$

This functional form assumption implies $\gamma_t \in (0,1), \forall t$ and that the probability of being a business consumer is increasing over time. However, this specification does not require the proportion of business consumers to strictly increase over time as $\gamma^2 = 0$ is allowed.

Since the firm knows the demand elasticities of consumers as well as the arrival process, studying the firm's incentives to set prices given remaining capacity and time to sell relays important information regarding the structural parameters. By solving the firm's dynamic program, I recover the shadow price of capacity across time and the pricing policy functions of the firm, which through the markup rule, informs the demand elasticity across time. Random demand results in variation in sales for a given number of seats remaining and time left to sell, which informs the arrival process parameters.

## 1.6 Results

In this section, I discuss the parameter estimates and the fit of the model. I then discuss how the estimated pricing policies relate to theory on dynamic pricing. Finally, I return to the assumption that the cost for consumers to search again the next period is sufficiently high that waiting is never optimal. I show only a small transactions cost is needed to make this assumption valid. In the next section, I conduct the counterfactual exercises.

### 1.6.1 Model Estimates and Fit

I estimate the model by city pair. I utilize observations for the last 45 days prior to departure, as average prices are relatively constant and bookings are low between 45 and 60 days prior to departure, and it greatly eases the computational burden of solving the dynamic program. Here I discuss the results for one city pair, which I use to conduct the counterfactual exercises in the next section.

**Table 1.3:** Parameter Estimates

| Parameter | | Point Estimate | SE |
|---|---|---|---|
| Price Sens. | $\alpha_L$ | $-0.0128^{***}$ | 0.000141 |
| | $\alpha_B$ | $-0.0048^{***}$ | 0.000165 |
| Pr(business) | $\gamma_1$ | $8.352^{***}$ | 0.344 |
| | $\gamma_2$ | $-0.214^{***}$ | 0.0124 |
| Poisson Process | $\mu_1$ | $8.870^{***}$ | 0.0501 |
| | $\mu_2$ | $0.102^{***}$ | 0.00833 |
| | $\mu_3$ | $0.0037^{***}$ | 0.000404 |
| | $\mu_4$ | $-0.00015^{***}$ | 5.750$e$-05 |
| Pseudo LogLike | | -44398 | |
| $e_L^D$ | | $-4.254$ | |
| $e_B^D$ | | $-1.397$ | |

[a] Standard errors calculated using block bootstrap, $n = 10,440$.

Parameter estimates appear in Table 1.3. All parameters are significant at the 1% level. The parameter estimates imply business consumers are over three times less price sensitive than leisure consumers, and are willing to pay up to 75% more in order to secure a seat. Figure 1.6 plots fitted values of the Poisson arrival rate and probability on types across time. Estimates of the arrival rate start around eight persons per flight 45 days prior to departure. The process peaks roughly three weeks prior to departure with a rate just under 12 persons per flight. From there, the potential market decreases in time with the lowest arrival of consumers appearing the day before departure. Here the arrival rate is just under 8 persons per flight.

**Figure 1.6:** Fitted values of the arrival process: probability on consumer types and Poisson rates.

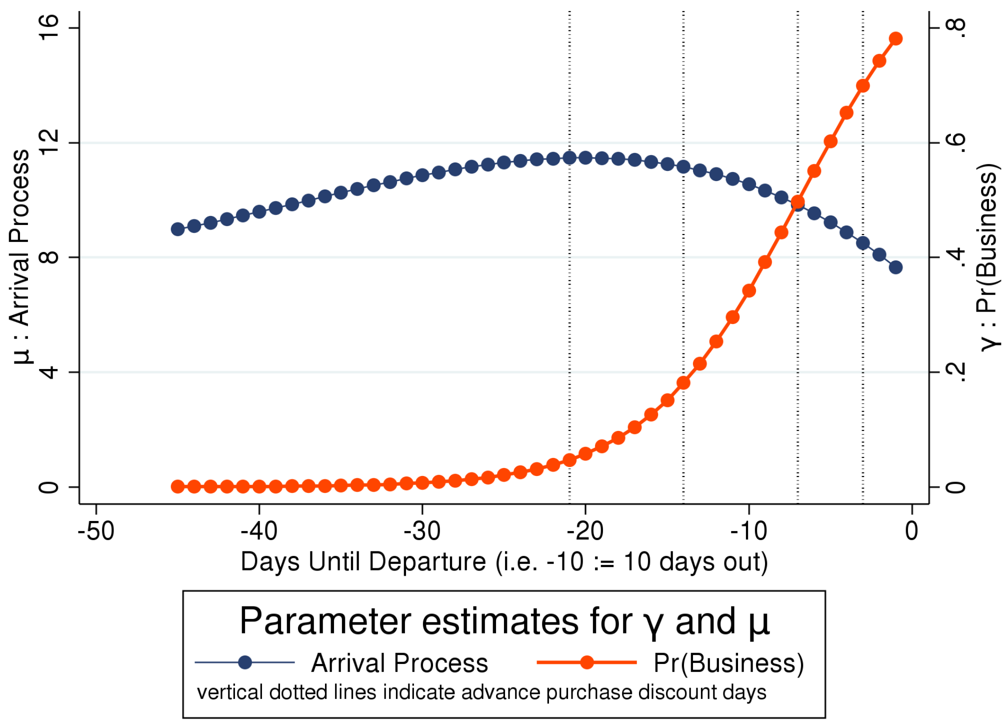


The model estimates suggest a large shift in the makeup of consumers across time. More than a month prior to departure, the share of business travelers is close to zero. Starting at approximately two weeks prior to departure, corresponding to the 14-day advance purchase discount, the share of business consumers in the market increases dramatically from 20% fourteen days to nearly 80% the day before departure.

**Figure 1.7:** Comparison of model prices with observed prices.

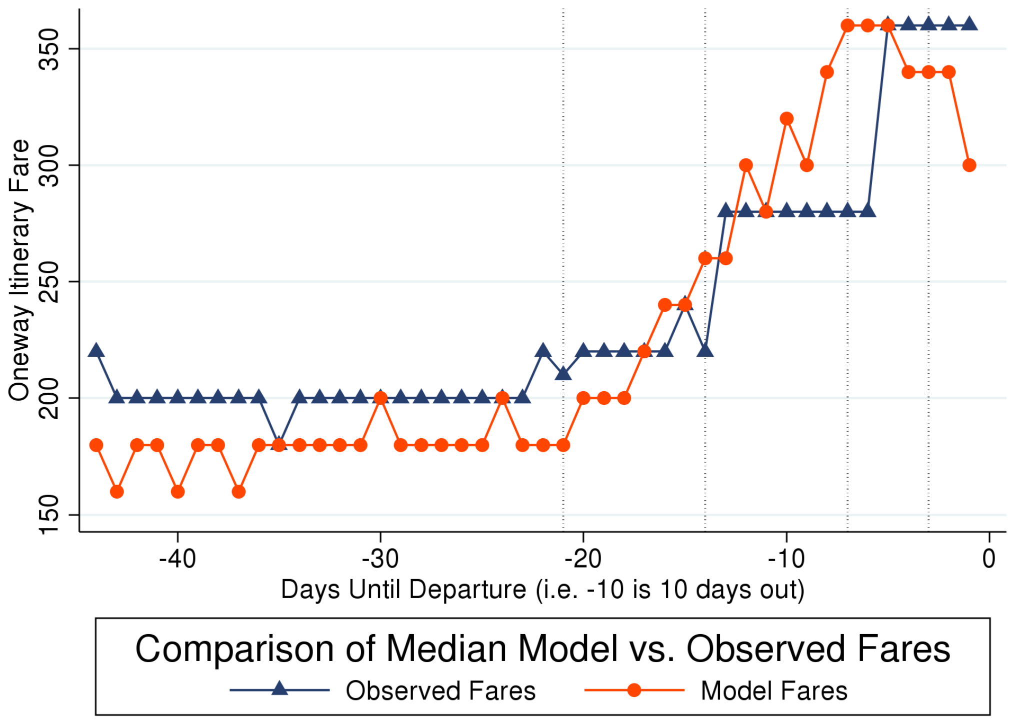


Figure 1.7 plots the median model fares and observed fares by day before departure.[26] The plot shows that the model fares are quite similar to observed fares, with differences of usually less than \$25 between 45 and 10 days prior to departure. The model accurately picks up the increasing pattern of fares within three weeks of the departure date. The model predicts fares to increase to their highest levels between seven to three days prior to departure, whereas in the data, the highest fares occur over the last five days prior to departure. With model prices, flights with excess capacity result in fire-sales, with the 50% percentile decreasing nearly \$75. The reason for this last-minute decline in prices is that

[26] Mean fares appear in the counterfactuals.

there is no value of holding capacity in the last period. Recall in the data fire-sales do occur – roughly 10% of the time, but presumably, fire-sales do not occur more often because otherwise business consumers would learn that delaying purchase results in lower fares.

### 1.6.2 Optimal Pricing

Figure 1.8 plots the optimal pricing policies and expected revenues for the firm as a function of remaining capacity. The left panel plots the optimal price given remaining capacity for three selected periods, corresponding to 15, 30, and 45 days prior to departure. The right plot indicates the expected revenues associated with remaining capacities for these three periods.

**Figure 1.8:** Estimated policy functions and expected revenues

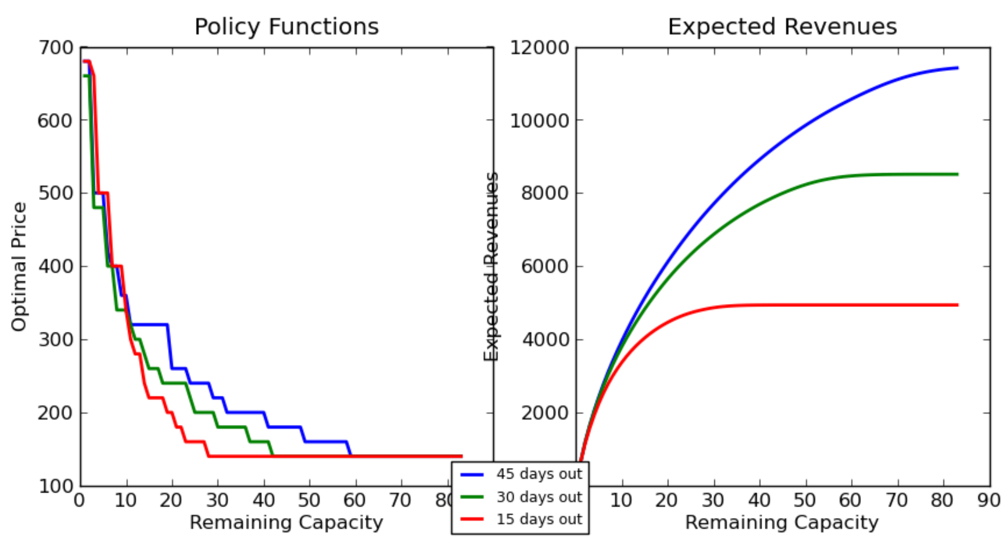


The right panel shows that expected revenues are increasing in capacity for a given period. Second, the panel shows that expected revenues are increasing in time to sell for a given capacity. These results are consistent with the theory on dynamic pricing found in Gallego and Van Ryzin (1994). Expected revenues flatten out for a given period because the firm cannot capture additional revenue when there is sufficiently high capacity remaining. The plot shows that the probability of a sell out is close to zero if the firm has at least 30 seats remaining with 15 days left to sell. With 30 days remaining, there is excess capacity when at least 55 seats are remaining.

The left panel plots the policy function of the firm, $p(s,t)$ for a given unobservable shock $\omega$. As already mentioned, with 30 seats remaining and 15 days left to sell, the probability of a sell out is close to zero. In return, the firm charges a low price. However, the price charged by the firm with 30 seats remaining is increasing in time to sell; that is, $p(30,15) < p(30,30) < p(30,45)$. This result is driven by the complementarity between intertemporal price discrimination and dynamic adjustment to stochastic demand when business consumers arrive close to the departure date. With additional time to sell for a given capacity, the probability of a sell out increases and in particular, there is an increased probability that a sell out would occur under a low price before business consumers become active. By charging a higher price early on when capacity is expected to be scarce, the firm can save seats for the high-valuing consumers who arrive late.

### 1.6.3 Allowing Consumers to Wait

The demand model assumes that waiting is never optimal. This assumption was motivated by the fact that there are no discontinuities in bookings immediately before advance purchase discounts expire. In this section, I study the incentives of consumers to wait in purchasing tickets.

I change the model in the following way: after consumers arrive, each consumer has the option to either buy a ticket, choose not to travel, or wait one additional day to decide. By choosing to wait, each consumer retains her private valuations (the $\varepsilon$'s) for traveling but may be offered a new price tomorrow. Consumers do not have perfect foresight, so they forecast both fares and remaining capacity for the next period. Additionally, each consumer has to pay a transactions cost $\phi_i$ to wait. This cost reflects the disutility consumers incur when needing to return to the market the next period. The goal of this section is to derive a waiting cost $\overline{\phi}$ such that if all consumers have a waiting cost at least as high as $\overline{\phi}$, then no one will wait. I then calculate the waiting cost in the data.[27]

Dropping the $i$ subscript, the choice set of a consumer arriving at time $t$ in a model of waiting is

$$\max\Big\{\varepsilon_0, \beta - \alpha p_t + \varepsilon_1, EU^{\text{wait}}(p,s) - \phi\Big\},$$

where $EU^{\text{wait}}$ is the expected value of waiting and can be written as

$$EU^{\text{wait}}(p,s) = \mathbb{E}_{p'|p,s}\Big[\max\{\varepsilon_0, \beta - \alpha p_{t-1} + \varepsilon_1\}\Big],$$

To derive $\overline{\phi}$, I first investigate the decision to wait for the marginal consumer, or a consumer such that $\varepsilon_0 = \beta - \alpha p_t + \varepsilon_1$. This consumer has no incentive to wait if the price tomorrow is at least as high as today. If price drops, the gain from

[27] For the proceeding analysis, I assume capacity is infinite. This means $\overline{\phi}$ is not the lower bound on waiting costs because the probability of not getting a seat creates an additional incentive to not wait, i.e. there is a positive probability of being offered an infinite price. By ignoring capacity, the consumer does not forecast the arrival process, or make decisions due to possible rationing.

waiting is

$$u_{t-1} - u_t = (\beta - \alpha p_{t-1} + \varepsilon_1) - (\beta - \alpha p_t + \varepsilon_1)$$
$$= \alpha(p_t - p_{t-1}),$$

which implies the expected gains from waiting are $\Pr(p_{t-1} < p_t)\, \mathbb{E}[\alpha(p_t - p_{t-1}) \,|\, p_{t-1} < p_t]$. Hence, an indifferent consumer will not wait if

$$\phi_i > \overline{\phi} := \Pr(p_{t-1} < p_t)\, \mathbb{E}[\alpha(p_t - p_{t-1}) \,|\, p_{t-1} < p_t].$$

**Proposition 1:** With $\overline{\phi} := \Pr(p_{t-1} < p_t)\, \mathbb{E}[\alpha(p_t - p_{t-1})$, then all consumers will choose not to wait.

**Proof:**

Take a consumer such that $\varepsilon_0 < \beta - \alpha p_t + \varepsilon_1$. Then there exists a $\overline{p} > p$ such that $\varepsilon_0 = \beta - \alpha\overline{p} + \varepsilon_1$. The expected gain for this consumer waiting is

$$\Bigg(\underbrace{\Pr(p_{t-1} < p_t)\, \mathbb{E}[\alpha(p_t - p_{t-1}) \,|\, p_{t-1} < p_t]}_{=\overline{\phi}} + \underbrace{\Pr(p_t < p_{t-1} \le \overline{p})\, \mathbb{E}[\alpha(p_t - p_{t-1}) \,|\, p_t < p_{t-1} \le \overline{p}]}_{\le 0} - \overline{\phi}\Bigg) \le 0,$$

which shows that a consumer that prefers to purchase today would not wait under the defined $\overline{\phi}$. Next, consider a consumer such that $\varepsilon_0 > \beta - \alpha p_t + \varepsilon_1$. Then there exists a $\underline{p} < p_t$ such that $\varepsilon_0 = \beta - \alpha\underline{p} + \varepsilon_1$. The expected gain from this consumer waiting is

$$
\begin{aligned}
\mathbb{E}[\text{gain}] = \Big(&\Pr(p_{t-1} < \underline{p})\mathbb{E}[\beta - \alpha p_{t-1} + \varepsilon_1 - \varepsilon_0 \,|\, p_{t-1} < \underline{p}] \\
&\qquad - \underbrace{\Pr\left(p_{t-1} < p_t\right)\mathbb{E}[\alpha(p_t - p_{t-1}) \,|\, p_{t-1} < p_t]}_{=\overline{\phi}}\Bigg) \leq 0, \\
&\leq \Pr(p_{t-1} < p_t)\Big[\mathbb{E}[\beta - \alpha p_{t-1} + \varepsilon_1 - \varepsilon_0 \,|\, p_{t-1} < \underline{p}] - \mathbb{E}[\alpha(p_t - p_{t-1}) \,|\, p_{t-1} < p_t]\Big] \\
&\leq \Pr(p_{t-1} < p_t)\Big[\mathbb{E}[\beta - \alpha p_{t-1} + \varepsilon_1 - \varepsilon_0 \,|\, p_{t-1} < \underline{p}] - \mathbb{E}[\alpha(p_t - p_{t-1}) \,|\, p_{t-1} < \underline{p}]\Big] \\
&= \Pr(p_{t-1} < p_t)\Big[\mathbb{E}[\beta - \alpha p_{t-1} + \varepsilon_1 - \varepsilon_0 - \alpha(p_t - p_{t-1}) \,|\, p_{t-1} < \underline{p}] \\
&= \Pr(p_{t-1} < p_t)\Big[\mathbb{E}[\beta + \varepsilon_1 - \varepsilon_0 - \alpha p_t \,|\, p_{t-1} < \underline{p}] \\
&= \Pr(p_{t-1} < p_t)\Pr(p_{t-1} < \underline{p})\underbrace{[\beta - \alpha p_t + \varepsilon_1 - \varepsilon_0]}_{<0 \text{ by assumption}} \leq 0.
\end{aligned}
$$

Hence, a consumer that prefers the outside option today would also not wait. ■

In monetary terms, $\tilde{\phi} = \overline{\phi}/\alpha = \Pr\left(p_{t-1} < p_t\right)\mathbb{E}[(p_t - p_{t-1}) \,|\, p_{t-1} < p_t]$ defines the the transaction cost a consumer would have to face in order to never wait. These statistics can be calculated in the data. Figure 1.9 plots $\tilde{\phi}$ across time using the model estimates. I find that the average transactions cost required to make consumers not wait to be \$5.49. For many days, the transactions required to make consumers not wait is close to zero. Notably, $\tilde{\phi}$ is higher close to the departure date as there are greater incentives to wait under model prices. In particular, business consumers who arrive the day before departure must incur a high waiting cost to persuade them not to wait. By waiting, they can capture fire-sale prices.

The calculated waiting costs suggest that the assumption that consumers do

**Figure 1.9:** Plot of mean transactions costs that would induce consumers not to wait

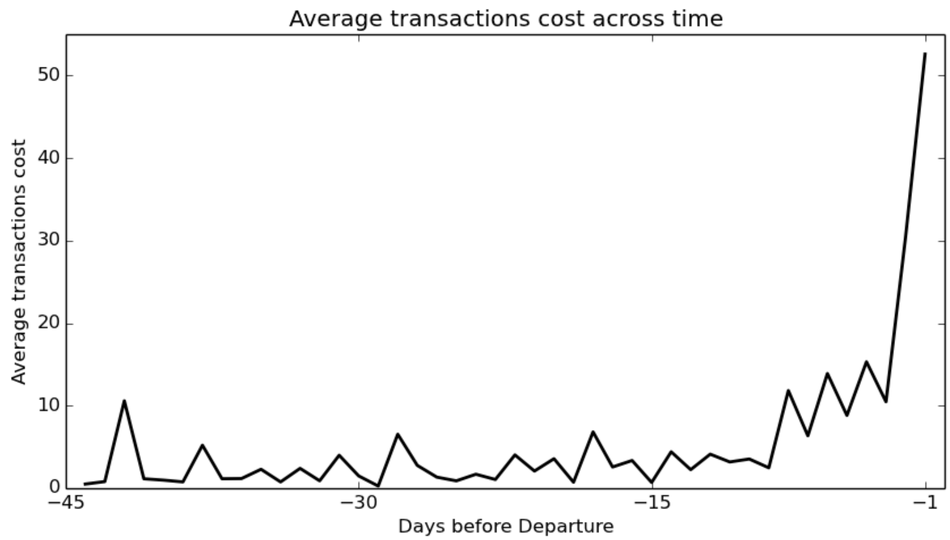


not dynamically substitute is very reasonable for almost all days prior to departure. The result is driven by the upward price trend. Since prices tend to increase, there is little incentive to wait a day to purchase once consumers first learn about their interest in travel. Further, while the required transaction cost is high close to the departure date, this analysis analyzed the indifferent consumer. Consumers who prefer to either purchase or not purchase would require a lower cost to not wait.

## 1.7 Analysis of the Estimated Model

In this section, I use the estimated model to examine three issues. In Section 6.1, I first perform three exercises that reduce the firm's ability to price discriminate. I investigate uniform pricing, dynamic pricing but restrict the number of price changes, and a pricing system which depends on date of purchase, but not remaining capacity. Section 6.2 uses the estimated model to highlight the complementarity between intertemporal price discrimination and stochastic demand pricing. Finally, in Section 6.3, I show how estimation approaches that do not take into account stochastic demand will systematically produce biased estimates of the degree to which demand becomes more inelastic as the departure date approaches.

For each exercise, I use the empirical distribution of remaining capacity 45 days prior to departure as the initial capacity condition.[28] I maintain the assumption that prices are chosen from the price set observed in the data. Relaxing this assumption so that prices are continuous yields qualitatively similar findings. For each counterfactual, I calculate the following benchmarks:

- Fare : overall mean fare for flights that have not sold out;
- LF : mean load factor the day flights leave;
- Sell outs : percent of flights that sell out after the last pricing period;
- Rev : mean revenue across flights;
- $\text{CS}_L$ : mean leisure consumer surplus;
- $\text{CS}_B$ : mean business consumer surplus;
- SW : mean daily welfare across flights (less sunk costs).

[28] I simulate 100,000 flights using this distribution. This provides an initial capacity condition between 6 and 83 seats. This controls for the fact that I model pricing for the last 45 days prior to the departure date. With just six seats remaining, the probability of selling out exceeds 85% under dynamic pricing versus less than 5% for a a flight with 83 seats.

### 1.7.1 The Welfare Effects of Flexible Pricing

In the model, the firm can set prices flexibly over time, to respond to changes in the consumer composition, and in response to random realizations of demand. At the opposite extreme is a pricing system that sets a uniform price over the entire time period. This subsection compares these extreme cases: dynamic pricing to uniform pricing. I also examine an intermediate case, where prices depend upon time to departure, but not on random fluctuations in demand. By comparing uniform to the intermediate case, and the intermediate to the full, the following analysis separates out the gains from intertemporal price discrimination from the gains to adjusting to stochastic demand.

#### Uniform Pricing

I start by removing the firm's ability to adjust prices as the departure date approaches. The firm maximizes expected revenues subject to the constraint that it must charge a uniform price across time. The price is solely dependent on the initial capacity condition.

Under uniform pricing, a high fare ensures the firm can successfully secure seats for business consumers, but doing so prices leisure consumers out of the market. At the same time, going after only the business market results in unused capacity which can be filled by lowering the fare. The optimal pricing strategy is one that balances saving seats for the high-valuing customers who arrive close to the departure date along and filling seats that would otherwise be scrapped.

Results for the counterfactual appear in Table 1.4. The left panel plots mean fare, weighted by initial capacity, across time for flights that have not yet sold out. The plot shows that the uniform price is higher relative to dynamic prices well in advance of the departure date. An optimal uniform fare of nearly \$200

**Table 1.4:** Dynamic vs. Uniform Pricing

| Policy / Mean | Fare | $LF^f$ | Sell outs | $Rev^f$ | $CS_L^i$ | $CS_B^i$ | $SW^t$ |
|---|---|---|---|---|---|---|---|
| Dynamic | \$208.70 | 95.69% | 28.65% | \$9,956.51 | \$52.59 | \$171.97 | \$955.85 |
| Uniform | \$197.62 | 91.87% | 40.46% | \$9,299.37 | \$51.46 | \$184.50 | \$951.35 |
| Difference (%) | -5.30% | -3.82% | 11.81% | -6.60% | -2.15% | 7.29% | – |

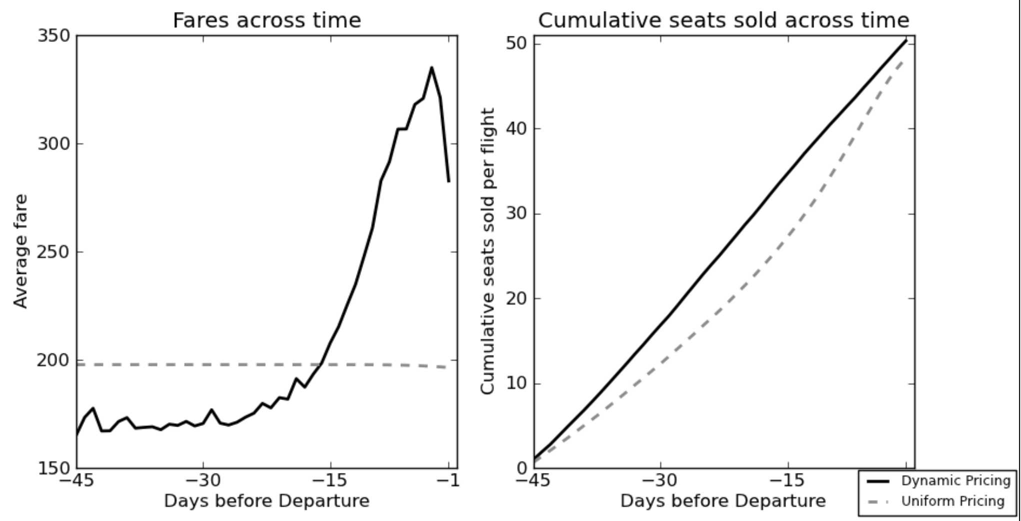


means only leisure consumers with high private valuations actually purchase. The fare is sufficiently high that the firm can save (some) seats for the high-valuing consumers who arrive close to the departure date.

The fact that fares are relatively higher under uniform pricing early on, but relatively lower closer to the departure date results in a significant reallocation of capacity across time. This is shown in the right panel, which plots the booking curve, or mean cumulative seats sold, towards the departure date. The uniform pricing booking curve is bowed out as fewer consumers purchase under the relatively higher fare early on. Relatively lower fares close to the departure date results in a higher booking rate compared to dynamic pricing through the day of departure. However, even with this increase in the booking rate, the overall load factor for flights under uniform pricing is lower than dynamic pricing. I find that 17% fewer seats are purchased by leisure consumers under uniform pricing. At the same time, nearly 17% more seats are booked by business consumers under uniform pricing. Since more leisure consumers purchase tickets than business consumers, the net change in bookings under uniform pricing is negative.

There is also a large reallocation within consumer type. Figure 1.10 plots the fraction of consumer types that purchase conditional on entering the market (i.e. number of consumers that buy / number of consumers that enter). The left panel shows that under dynamic pricing, the purchase rate amongst leisure consumers is close to 11% well before the departure date. As prices increase under dynamic pricing, very few leisure consumers purchase. On the other hand, under uniform pricing, fares are high early on resulting in a 33% decline in purchases by leisure consumers. The take up rate is constant, but also declines towards the departure date. This is because under uniform pricing, many flights sell out in advance. For business consumers, the lower fares offered by uniform pricing results in higher purchasing rates across all periods. Again, the uniform pricing rate declines towards the departure date since many flights sell out in advance.

The change in allocation under uniform pricing is mitigated because of two forces. First, as discussed above, uniform pricing implies leisure consumers with high private valuations purchase throughout time, whereas under dynamic pricing essentially no leisure consumers purchase when fares go over $250. While leisure consumers are made worse off due to high prices early on, it also means some successfully purchase closer to the departure date and overall, leisure consumer surplus decreases over 2% under uniform pricing. Second, while business

**Figure 1.10:** Purchasing rates across consumer types and time

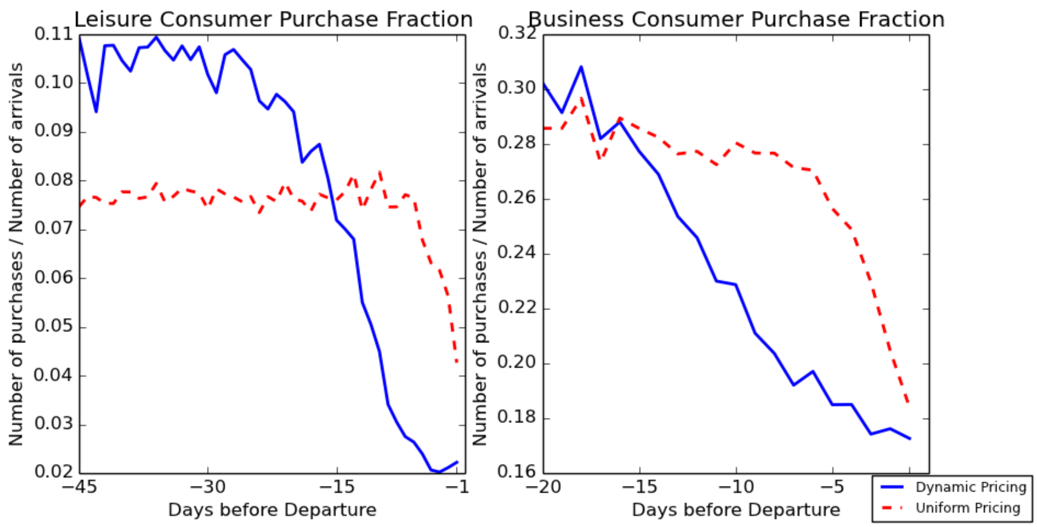


consumers benefit from lower prices, the increase in business consumer surplus is mitigated because of the increased number of early sell outs. An important consequence of uniform pricing is that the firm cannot control the booking rate of flights that turn out to be popular. Uniform pricing increases sell outs by 12%, which consequently forces more late-arriving business consumers to the outside option. Hence, business consumer surplus increases 7% under uniform pricing.

On average, leisure consumer welfare declines and business consumer welfare increases under uniform pricing. Overall, I find consumer welfare is 1.36% higher under uniform pricing compared to dynamic pricing. The consumer welfare gains in the absence of price discrimination are relatively modest as a result of inefficient capacity allocation. Further, compared to dynamic pricing, revenues fall 6.6% under uniform pricing. As airlines operate razor thin margins, the decline in revenues is significant, and suggests the firm would probability choose to exit the market in the long run. Moreover, total welfare is lower under uniform pricing compared to dynamic pricing.

### The Role of Frequent Price Adjustments

The previous exercise compared the extremes in pricing capabilities of the firm – either the firm maintains a single price across time, or the firm can update prices daily. Now I allow the firm to use dynamic pricing, with the restriction that prices must be maintained for $k$ days. I conduct four counterfactuals, corresponding to $k = 3, 5, 9, 15$. The idea here is that dynamic pricing is clearly valuable to the firm, but it is not necessarily true that daily price adjustments are needed to obtain the revenues observed under (daily) dynamic pricing.

**Figure 1.11:** Comparing different degrees of dynamic pricing

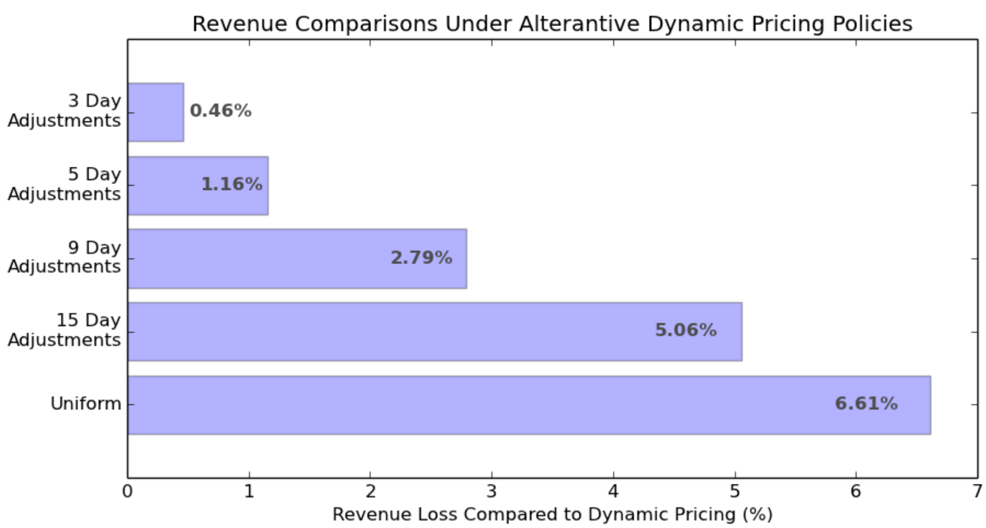


Figure 1.11 plots the revenue loss compared to the baseline case of daily price adjustments, for the four counterfactuals. For example, while uniform pricing reduces revenue by 6.6% compared to daily dynamic pricing, the ability to update prices every 15 days increases revenue by over 1%. The difference in revenues

between 15 day adjustments and 9 day adjustments is quite large. I find that revenues are 1.1% lower under 5-day adjustments compared to daily dynamic pricing. Under 3-day adjustments, revenues are 0.5% less than under daily adjustments. There are two ways to interpret these results. The first is that while 9 and 15 day adjustments result in significant revenue declines, adjustments made at 3 and 5 day intervals result in revenues similar to daily adjustments. At the same time, even under 3-day adjustments, revenues are nearly half a percent lower compared to daily adjustments. As reported by the IATA (2013), the margins for airlines are very small (around 1%), suggesting that even the losses associated with 3 and 5 day adjustments would lead to a significant decline in profits in percentage terms.

### The Use of Intertemporal Price Discrimination Alone

I next single out the intertemporal price discrimination force by investigating pricing policies which depend on day until departure, but not on the scarcity of seats. Viewed from the dynamic pricing perspective, this counterfactual investigates pricing with the dynamic adjustment force shut down and thus, quantifies the complementarity between the two pricing forces. Specifically, this counterfactual quantifies the additional revenue gains possible by dynamically adjusting fares.

Under intertemporal price discrimination alone the optimal sequence of fares, $\mathbf{p}^*$, solves

$$\max_{\mathbf{p} \in \times_{t=0}^{T} \mathcal{P}} \mathbb{E}_{\mathbf{y}} \left[ \sum_t \min \left\{ Q_t(p, y) \, , \, s_{t+1} - \min \Big( Q_{t+1}(p, y), s_{t+1} \Big) \right\} p_t \, \Big| \, s_T \right],$$

where $s_{t+1} - \min \Big( Q_{t+1}(p, y), s_{t+1} \Big) \equiv s_t$ is the capacity remaining after the previous period's demand is realized and $s_T$ is the initial capacity condition. The decision space of the firm's problem has cardinality $|\mathcal{P}|^T$. I simplify the problem by adding the restriction that the firm only adjusts fares on the usual advance purchase discount days – 3, 7, 14, and 21 days prior to departure.[29]

**Table 1.5:** Dynamic vs. Intertemporal Price Discrimination

| Policy / Mean | Fare | LF | Sell outs | Rev | $CS_L$ | $CS_B$ | SW |
|---|---|---|---|---|---|---|---|
| Dynamic | \$208.70 | 95.69% | 28.65% | \$9,956.51 | \$52.59 | \$171.97 | \$955.85 |
| Discrim. Only | \$205.94 | 90.01% | 31.43% | \$9,608.13 | \$52.01 | \$171.54 | \$942.29 |
| Difference (%) | -1.32% | -5.68% | 2.78% | -3.49% | -1.01% | -.25% | – |

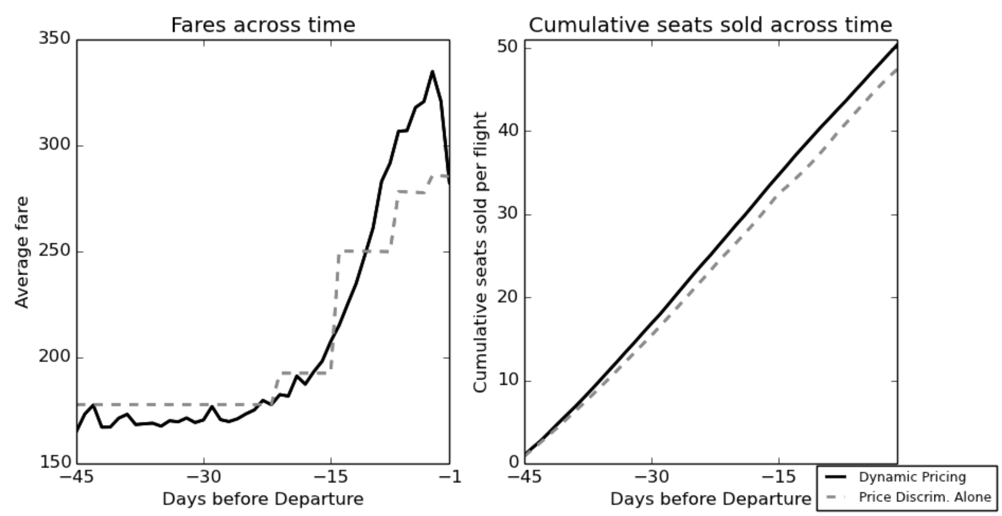


Results for this counterfactual appear in Table 1.5. The time path of prices under intertemporal price discrimination alone are monotonically increasing, with a substantial increase in fares when crossing the 14-day advance purchase discount. This corresponds to the increase in proportion of business consumers given

[29] I have also investigated the pricing decision over $|\mathcal{P}|^T$ using simulated annealing. I obtain similar findings compared to this setup.

the model estimates. The fares offered to leisure consumers well before the departure date are slightly higher under intertemporal price discrimination alone. The reason for this is because without dynamic adjustment, there is an additional incentive for the firm to reserve capacity for business consumers. In order to do this, the firm charges higher prices early on, which decreases the number of seats purchased by leisure consumers compared to dynamic pricing.

Importantly, as the pricing becomes more flexible, the fares offered early on are lower. Fares under intertemporal price discrimination alone are substantially lower than uniform pricing well in advance of the date of travel. Under dynamic pricing, fares are even lower. The point where leisure consumers stop purchasing is similar under both intertemporal price discrimination and dynamic pricing, and as a result, overall leisure consumer surplus has the following order: $CS_L^{\text{dynamic}} > CS_L^{\text{discrim.}} > CS_L^{\text{uniform}}$. That is overall, but within consumer type there is also a reallocation, and the ordering is reversed for the leisure consumers that arrive closer to the departure date. Business consumer surplus is essentially unchanged as prices are either higher or lower under dynamic pricing depending on when they arrive. Moreover, without dynamic adjustment, firms have a reduced ability to save seats for late-arriving business consumers, which results in a 2.8% increase in sell outs.

Not only is consumer surplus (marginally) higher under dynamic pricing, revenues are substantially higher. Revenues under intertemporal price discrimination alone are 3.5% lower compared to dynamic pricing.[30] These results demonstrate there is a significant complementarity between the pricing channels. Allowing firms to adjust prices dynamically results in additional revenues from early-arriving leisure consumers. At the same time, dynamic adjust allows airlines to

[30]Revenue Management Overview states revenue management systems have increased airline revenues by 3-9%.

secure seats on flights that are realized to have high demand. When the business consumers arrive, they receive higher fares compared to pricing under intertemporal price discrimination alone. In terms of value, these results suggest 50% of the revenue gains associated with dynamic pricing over uniform pricing comes from dynamic adjustment.

### 1.7.2 The Complementarity of Intertemporal Price Discrimination and Stochastic Demand Pricing

This subsection shows how intertemporal price discrimination and stochastic demand pricing are complements in the airline industry. It arises because of the particular pattern of consumer arrival. The consumers who arrive last are the ones with the highest willingness to pay, which creates an incentive for firms to save seats until close to the departure date. If seats become scarce early on, the firm's optimal pricing strategy is to sharply increase fares, which reduces the rate at which seats are sold. With capacity reserved, the firm then price discriminates toward the late-arriving consumers by charging them high fares.

The point about this complementarity can be made with a simple example. Consider a reversal of the arrival process, where high-valuing consumers arrive first. In particular, suppose there are three periods and a flight has two seats. In each period, a consumer arrives with a 50/50 probability. In the first period ($t = 1$), if the consumer arrives, the reservation price is \$1000. For the remaining two periods ($t = 2, 3$) the reservation price is \$200. In this environment, the profit maximizing policy is to set $p_1 = 1000$ and $p_2 = 200$, and $p_3 = 200$. There is no need for price to respond to demand since the high-valuing consumer is guaranteed a seat. In this case, there are no gains to having a pricing system that reacts to demand realizations.

Now reverse the order of arrival so that in the first two periods, the reservation price is \$200, and in the last period, the reservation price is \$1,000. Under dynamic pricing, the firm will charge a price of \$200 in the first period. If a seat sells, the firm will set a price $p_2 > 200$, and finally, will set a price of \$1,000 in the last period. If the seat does not sell, the firm will charge $p_2 = 200$ and $p_3 = 1000$, which yields an expected profit of \$650.[31]. However, if prices are not allowed to respond to demand, the firm will set up a fare schedule that reserves the seat for the potential business consumer. By charging $(p_1, p_2, p_3) = (\infty, 200, 1000)$, the firm has expected revenues of \$600. Hence, the airline would be willing to pay up to \$50 to utilize a pricing system that responds to demand realizations.

This simple example demonstrates the importance of dynamic adjustment when high-valuing consumers arrive last. In particular, it is important that airlines price to keep seats available until close to the departure date. To get a sense of the magnitude of this force for the airline industry, I perform the same exercise above using the estimated model. I reverse the arrival process of consumers and compare revenues under dynamic pricing with a pricing system that only depends on the date of purchase. Note that if the arrival process was constant, and the mix of consumer types did not change over time, reversing the order would have no effect on revenues. Thus the magnitude of the difference will depend upon how stochastic demand is, and the extent to which elasticity varies. Both of these are pinned down in the estimation exercise.

[31]If the seat sells, the expected revenues for the remaining period are \$500. If the seat does not sell, the expected revenues are $1/2 \cdot 200 + 1/2 \cdot 1000$. Hence, $ER = 1/2(200) + 1/2(500) + 1/2(1/2 \cdot 200 + 1/2 \cdot 1000) = 650$

**Reversal of the Arrival Process**

Table 1.6 presents the counterfactual results of dynamic pricing, intertemporal price discrimination alone, and uniform pricing when the arrival process is reversed. Compared to when business consumers arrive late (observed arrival process), a reversal of the arrival process brings: lower fares, higher load factors, more sell outs, and increased revenues (see Table 1.7 for comparisons). Overall consumer surpluses are also higher under a reversal of the arrival process, with the exception of leisure consumer surplus under uniform pricing, which is lower, but very close to the observed arrival process.

**Table 1.6:** Reversal of the arrival process

| Policy / Mean | Fare | LF | Sell outs | Rev | $CS_L$ | $CS_B$ | SW |
|---|---|---|---|---|---|---|---|
| Dynamic | \$205.61 | 96.69% | 49.67% | \$10,086.01 | \$52.86 | \$172.60 | \$959.78 |
| Discrim. Only | \$204.89 | 93.72% | 48.79% | \$9,851.89 | \$52.55 | \$172.46 | \$951.69 |
| Uniform | \$196.53 | 91.96% | 40.83% | \$9,304.30 | \$51.34 | \$190.20 | \$957.15 |

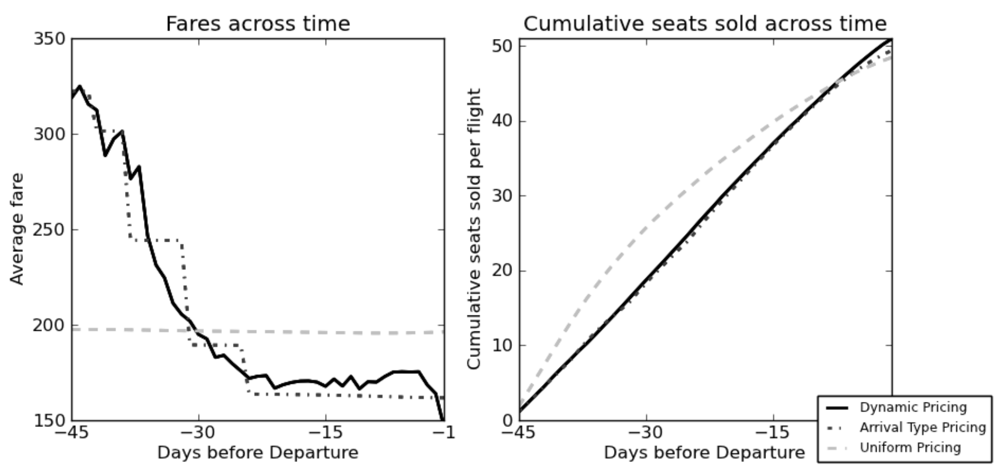


With a reversal of the arrival process, firms have no incentive to hold remaining capacity. This is particularly noticeable under intertemporal price discrimination alone. Under the reversed arrival process, leisure consumers who arrive within 21 days of the departure date receive fares that are $10-$20 lower on average. Fares increase slightly under dynamic pricing closer to the departure date for flights with scarce capacity, driving up mean fares. Flights that have not already sold out result in fare-sales the day before departure as the firm tries to fill any remaining open seats. The lower prices offered with a reversed arrival process result in substantially higher sell outs. As shown in Table 1.7, the ordering of percent of sell outs is reversed under the two arrival processes. The percent of sell outs under dynamic pricing nearly doubles. Under intertemporal price discrimination alone, sell outs are 17% higher.

**Table 1.7:** Comparing Arrival Processes

| Policy | Arrival / Mean | Fare | LF | Sell outs | Rev |
|---|---|---|---|---|---|
| Dynamic | Reversed | \$205.61 | 96.69% | 49.67% | \$10,086.01 |
| | Observed | \$208.70 | 95.69% | 28.65% | \$9,956.51 |
| Discrim. Only | Reversed | \$204.89 | 93.72% | 48.79% | \$9,851.89 |
| | Observed | \$205.94 | 90.01% | 31.43% | \$9,608.13 |
| Uniform | Reversed | \$196.53 | 91.96% | 40.83% | \$9,304.30 |
| | Observed | \$197.62 | 91.87% | 40.46% | \$9,299.37 |

The most important feature of the reversed arrival process is the role of intertemporal price discrimination. While revenues are $100 greater per flight with dynamic pricing across arrival processes and only $5 per flight more under uniform pricing, the difference in revenues under intertemporal price discrimination alone is $250 across arrival processes. The other relevant measure is the revenue across policies within arrival process. As previously discussed, 50% of the revenue gains of using dynamic pricing over uniform pricing can be attributed to intertemporal price discrimination. The remaining can be attributed to dynamic adjustment. However, when the arrival process is reversed, 75% of the revenue gains can be attributed to intertemporal price discrimination. The reason for this is because the role of dynamic adjustment is unique to the arrival process. When business consumers arrive late, dynamic adjustment is used to reserve capacity across time. With business consumers arriving early, firms can capture their willingness to pay solely through the intertemporal price discrimination channel. The role of dynamic adjustment is simply to fill remaining seats after the business market is served.

### Constant Arrival Process

The last counterfactual demonstrated the importance of dynamic adjustment in *complementing* intertemporal price discrimination when business consumers arrive late. The counterfactual also established that the role of dynamic adjustment is unique to the arrival process. In this counterfactual, I investigate pricing under a constant arrival process. Demand is still stochastic, but the mix of consumers that arrive over time is held constant. This shuts down the use of the intertemporal price discrimination channel. In this environment, prices adjust across time solely due to the scarcity of seats under stochastic demand.

Figure 1.12 plots the mean price under dynamic pricing and uniform pricing under constant arrival.[32] Notably, instead of prices increasing over time, prices simply fluctuate around the levels offered under uniform pricing. Fares increase

[32] Under constant arrival, I take the probability of a consumer being of the business type to be the mean of $\hat{\boldsymbol{\gamma}}$ for all periods

**Figure 1.12:** Pricing under a constant arrival process

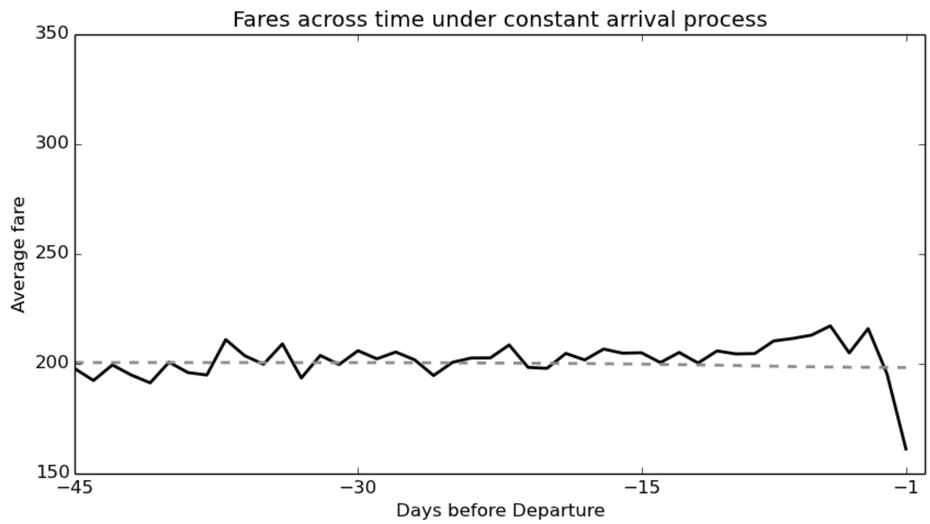


for flights that are realized to be scarce close to the departure date in a similar fashion to the dynamic pricing policies under the reversed arrival process. Flights are then offered under fire-sale prices as to not leave as many seats unfilled. One of the predictions found in Gallego and Van Ryzin (1994) is that the value of dynamic pricing is lower under constant arrival. My empirical findings support this theory. While revenues increase 6.6% by using dynamic pricing over uniform pricing under observed arrival, the revenue gains under constant arrival are half of that.

### 1.7.3 Consequences of Abstracting from Stochastic Demand

In this section, I illustrate that in order to conduct welfare analysis in markets where both the intertemporal price discrimination and dynamic adjustment forces operate, it is important to take into account the uncertainty about demand. In particular, by ignoring stochastic demand, an empirical analysis will fail to take into account the opportunity cost of holding back capacity tends to fall over time, which leads to a systematic bias in estimating demand elasticities.

Consider the pricing of a single flight. Suppose demand is stochastic but the empirical approach ignores the uncertainty about demand. In this case, the firm faces a static problem which can be written as

$$\max_{\mathbf{p}} \sum_t Q_t(p_t)p_t \quad \text{s.t.} \quad \sum_t Q_t(p_t) \leq s.$$

Letting $c(s)$ be the shadow price of capacity, the firm problem can be written as the following unconstrained problem:[33]

$$\begin{aligned} &\max_{\mathbf{p}} \sum_t Q_t(p_t)p_t - c(s)\left(\sum_t Q_t(p_t) - s\right) \\ \Leftrightarrow &\max_{\mathbf{p}} \sum_t Q_t(p_t)\Big(p_t - c(s)\Big) + c(s)s. \end{aligned}$$

Letting $c$ be the opportunity cost at the optimum, Lerner's index reveals

$$\frac{p_t - c}{p_t} = \frac{1}{e_t^D(p_t)},$$

where $e_t^D(p_t) \in \mathbb{R}_+$. Rearranging terms to solve for price yields the markup rule, $p_t = \frac{e_t^D(p_t)}{e_t^D(p_t)-1}c$. Then, taking the ratio of prices over time, in this case with the

[33] Here I assume price is not restricted to the discrete set of fares, otherwise the firm problem is an integer programming problem.

first period, reveals

$$\frac{p_t}{p_T} = \frac{\frac{e_t^D(p_t)}{e_t^D(p_t)-1}c}{\frac{e_T^D(p_T)}{e_T^D(p_T)-1}c} = \frac{\frac{e_t^D(p_t)}{e_t^D(p_t)-1}}{\frac{e_T^D(p_T)}{e_T^D(p_T)-1}}. \qquad \text{(Ignoring stochastic demand)}$$

In the equation above, the opportunity costs cancel which reveals that information on prices directly relate to elasticity ratios. On the other hand, when accounting for stochastic demand, the opportunity cost changes over time and the terms do not cancel. Instead,

$$\frac{p_t}{p_T} = \frac{\frac{\tilde{e}_t^D(p_t)}{\tilde{e}_t^D(p_t)-1}c_t}{\frac{\tilde{e}_T^D(p_T)}{\tilde{e}_T^D(p_T)-1}c_T}. \qquad \text{(Accounting for stochastic demand)}$$

In particular, the opportunity cost of selling a given seat tends to be lower closer to the departure date ($c_t \leq c_T$), because if a seat is not sold in the current period there is less of a chance it will be sold in the future. As a result, an empirical analysis ignoring stochastic demand will systematically understate the degree to which the late arriving consumers are less price sensitive.

Figure 1.13 demonstrates this result, by plotting the relative elasticity ratios when accounting for stochastic demand and when ignoring stochastic demand using the model estimates. Notably, as the opportunity cost falls closer to the departure date, the bias becomes significant.

## 1.8 Concluding Remarks

There are two broad rationales for product prices that change over time: segmentation of consumers who differ in their willingness to pay, and changes in scarcity – or shadow costs – arising from stochastic demand. In this paper, I study the

**Figure 1.13:** Comparing elasticity ratios across time

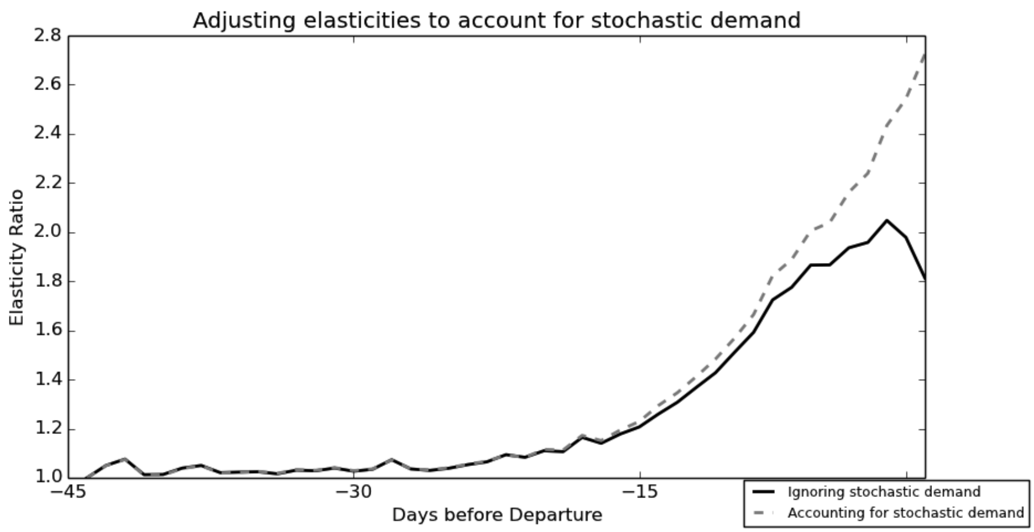


interactions of these forces by investigating the pricing decisions of airlines in US monopoly markets. I create a novel data set of high frequency fares and seat availabilities to estimate a structural model of dynamic airline pricing. In the model, firms face a stochastic arrival of consumers. The mix of consumer types, corresponding to leisure and business travelers, changes over time.

I show dynamic adjustment to stochastic demand *complements* intertemporal price discrimination in the airline industry. The complementarity arises because price inelastic consumers tend to arrive close to the date of travel. I find there are significant revenue losses associated with a pricing system that depends on the date of purchase, but not on the random realizations of demand. There are two reasons for this. First, dynamic adjustment allows firms to secure seats for business consumers who arrive close to the departure date. These consumers are then charged high prices. Second, dynamic adjustment reduces the incentive to hold back capacity well before the date of travel. Firms offer lower fares early on which results in additional revenue from leisure consumers who would otherwise not purchase.

Compared to pricing policies that depend only on the date of purchase, I find dynamic adjustment increases consumer surplus. Leisure consumers benefit from the lower fares offered early on. While business consumers face higher prices under dynamic adjustment, the ability of firms to save seats reduces the probability of an early sell out. This allows some business consumers to receive seats on flights that would have been sold out otherwise.

One of the limitations of this study is that I only examine the pricing decisions in monopoly markets. Recent research has suggested the ability of airlines to respond to scarcity may be reduced in more competitive markets. An open question would be to examine the complementarities in the pricing channels in oligopoly markets. In addition to presence of multiple equilibria, this research would have to address what information airlines know about their competitors. Airlines surely keep track of competitor fares, but it is not clear that airlines keep track of their competitors' seat availabilities.

# Chapter 2

# Zone Pricing and Spatial Menu Costs: Evidence from Drywall

## 2.1 Introduction

In the empirical industrial organization literature, the standard approach models prices as being set at the *market level*. Competing firms in a particular market each set a specific price for that market, taking as given the prices set by competing firms. In practice, retail firms commonly set prices at the *zone level*. These pricing zones, although usually geographically contiguous, often combine distinct markets that may be hundreds of miles apart and that differ in significant ways. A firm's pricing zone might include urban and rural markets, markets with different degrees of competition, and markets where input costs vary substantially. With these differences across markets, we might expect the firm to set individual prices in each market rather than a common price throughout the zone.

In this article, we develop an empirical analysis of zone pricing under competition. While monopolists can only increase profits by adopting more granular pricing policies, price discrimination theory has shown this is not necessarily the case in markets with competition. When competitors are present, a commitment to not use more granular pricing may allow firms to obtain higher profits. We explore this ambiguity by examining the zone pricing practices of the major home improvement retailers. We explain a number of features of drywall retailing that make it an ideal industry for such a study, and further we are able to construct a unique, rich data set for this industry. We estimate a structural model of supply and demand, which we use to estimate equilibria under alternative pricing regimes. We find profits increase if firms adopt more granular pricing. Since the major retailers have chosen not to adopt these policies, they must face some additional costs. We call these costs "spatial menu costs", and our analysis finds them to be substantial.

The spatial menu costs considered here relate in spirit to the concept of menu costs prominent in the macroeconomics literature. The macro literature documents how prices change infrequently over time, and this inter-temporal price rigidity has potentially significant implications for the macroeconomy. There are also studies that examine the menu costs associated with price changes at the micro level, such as for products within stores (see Levy, Bergen, Dutta, and Venable (1997) for example). Here, the price rigidity is across space – a pricing zone. Many of the key issues from the macro literature apply in our context.

Retail drywall markets have several features that aid analysis of zone pricing. In the mainstream retail sector for drywall, competition is between a few large chains, all of which practice zone pricing. Drywall, also known as wallboard or gypsum board, is costly to transport, so consumers buy from local stores and retailer costs vary geographically. The distribution centers for each major retail chain have known locations, making costs estimable. Some observed pricing zones span multiple, diverse markets, making advertising an unlikely reason for setting such large pricing zones. Some pricing zones also include monopoly markets and markets with multiple stores from each chain. The variation of costs, competition, and demand variables within a pricing region allow us to uncover the role of spatial menu costs in forcing pricing zones.

Our work contributes to the literature on price discrimination, as spatial menu costs are an interesting impediment to price discrimination. Firms have chosen not to set completely uniform prices and so engage in limited price discrimination. We show that without spatial menu costs, drywall retailers would set a discriminatory price in each market. In the standard monopoly setting, limiting the firm to a zone system can only lower profits. However, Holmes (1989) and Corts (1998) show that in environments where firms compete, the effect of being able to price discriminate has an ambiguous effect on profit; impediments to price discrimination limit what a particular firm can do, but it also affects what its rivals are doing. In more recent work, Aguirre, Cowan, and Vickers (2010) and Cowan (2012) explore sufficient conditions on demand for a third-degree price discriminating monopolist to either increase or decrease social welfare and consumer surplus, respectively. Chen and Schwartz (2013) examine the welfare implications of differential pricing in a monopoly setting where there are differences in marginal costs of serving consumer groups. Our empirical model allows for both demand and cost-based rationales for differential pricing, and we quantify the impact of zone pricing to both consumers and firms.

In order to examine the welfare implications of zone pricing in retail drywall, we estimate a structural model of demand and supply using a new data set for the industry. The two largest retail chains report prices and up-to-date inventory levels; we difference reported inventory levels to derive daily sales. The sales data allow estimation of a discrete-choice demand model. Consumers select between all available drywall products at each nearby store from either chain. We find drywall to be a highly substitutable product, but overall industry elasticity is very low. We estimate the marginal cost of each product in each store by accounting for transportation costs from the warehouse. With these demand and costs estimates, we can estimate profits under the current pricing zones and compare them to pre-menu cost profits in counterfactual equilibria where one or both of the firms instead use market level pricing. Multiple equilibria exist for a given zone configuration. We compare the current regime with a small adjustment in zones to yield unique counterfactual equilibria. Aggregating across markets, we find the spatial menu costs to be 2.2% of current revenues. For the 128 stores in the sample, this equates to roughly $4.6 million in additional profits for retail drywall annually. Applying a selection mechanism across multiple equilibria on the meta game of zone pricing yields a menu cost of nearly $2.3 million. The spatial menu costs that would rationalize the chains' decision not to separate stores into their own pricing districts are substantial, though small enough that managerial effort costs (as in Zbaracki, Ritson, Levy, Dutta, and Bergen (2004)) are a likely explanation.

Previous work on zone pricing is relatively sparse in both the economics and marketing literature. Montgomery (1997), Chintagunta, Dubé, and Singh (2003), and Khan and Jain (2005) are important exceptions. They all examine a supermarket chain that practices zone pricing and ask how profits and consumer surplus would change if the chain switched from zone pricing to store-by-store pricing. Marginal costs were assumed to be the same at all stores. Although this is an acceptable abstraction for supermarkets within a city, such an assumption cannot be maintained with drywall. An important difference between our work and the previous literature is that instead of analyzing what is happening to one firm in isolation, holding fixed the environment of the firm as it changes from zone to store-by-store, we take into account how switching regimes can affect the entire competitive interaction. We show the profit gains to using market level pricing are greatly overstated when abstracting from effects of competition, causing the menu costs to be overstated by as much as 32% or over $2.1 million annually for the stores examined.

This article is organized as follows: In Section 2, we document zone pricing for the home improvement retail industry and describe the data used for this study. In Section 3, we introduce the supply and demand system. In Section 4, we present estimation results and in Section 5, we conduct analysis on alternative pricing regimes.

## 2.2 Data and Descriptive Evidence

We create two new data sets for products sold at the largest home improvement warehouse retailers in the United States. In Section 2.2.1, we describe the detailed price and quantity data we use to estimate the empirical model in Section 2.4. We construct a panel of daily prices and sales quantities for all drywall products offered at 128 Home Depot and Lowe's stores in the Intermountain West. To obtain sales quantities, we monitor daily inventory levels reported on the retailers' websites; by differencing daily inventory levels, we obtain sales for each product-store pair.

In Section 2.2.2, we document the use of zone pricing throughout the retail home improvement industry. We obtain a cross section of prices at all Home Depot, Menards, and Lowe's locations in the United States for several product categories. While zone pricing is used in most of the product categories examined, the size and character of the zones varies for different products. We describe what features of the drywall data generalize to other product categories.

### 2.2.1 Drywall Data

Drywall prices posted on company websites reveal the use of zone pricing by both Home Depot and Lowe's. For both chains, prices on a given product vary considerably at a national level, but are exactly the same price at all of the chain's stores within contiguous areas. These areas are largely the same for different drywall products. We define a pricing zone as an area in which stores have exactly the same price on all products offered. We show that although most pricing zones are small, some contain a large number of stores and span a diverse set of markets.

Price levels for regular 5/8" x 4' x 8' gypsum board are mapped at every Lowe's and at every Home Depot in Figure 2.1. Each dot on the map represents a store, and its color represents its price intensity. Nationally this product has considerable price variation, at between \$5.98 a sheet to \$20.35 a sheet for Home Depot. In large geographical areas in the United States, such as every store located in Idaho, there is no price variation. A similar pattern can be seen for Lowe's stores. For this particular sheet of drywall, Lowe's charges 126 distinct prices across its 1,714 stores in the United States. Surprisingly in some areas, such as the Carolina Piedmont, a sharp boundary separates two zones with very different prices. Elsewhere, such as the upper Midwest, prices are similar over a wide area.

Figure 2.2 plots unique prices for a regular 5/8" x 4' x 8' gypsum board product available at Home Depot stores in the Western United States. Each dot again represents a store; stores with exactly the same price are the same color. The figure shows geographically contiguous pricing zones. For example, there is a unique price for this drywall product for all Home Depot stores in Oregon. All locations in Utah and Southern Idaho have the same price, although stores in eastern and western Washington have different prices. The unique pricing regions for different products within a chain are often, but not always, the same. The prices in Figure 2.2 are only for one product, and there is some variation in the pricing patterns across products. For example, prices for 1/2" x 4' x 12' drywall board exhibit three prices in the state of Washington instead of two, but the two pricing zones in western Washington combine to correspond exactly into the pricing region for 5/8" x 4' x 8' drywall. For comparison, Figure 2.3 in the Appendix plots unique prices for a regular 5/8" x 4' x 8' gypsum board at all national Lowe's locations.

**Figure 2.1:** Map of US Lowe's and Home Depot stores. Each point indicates a store location, the color of the point corresponds to a pricing intensity for 4' x 8' x 5./8" drywall.

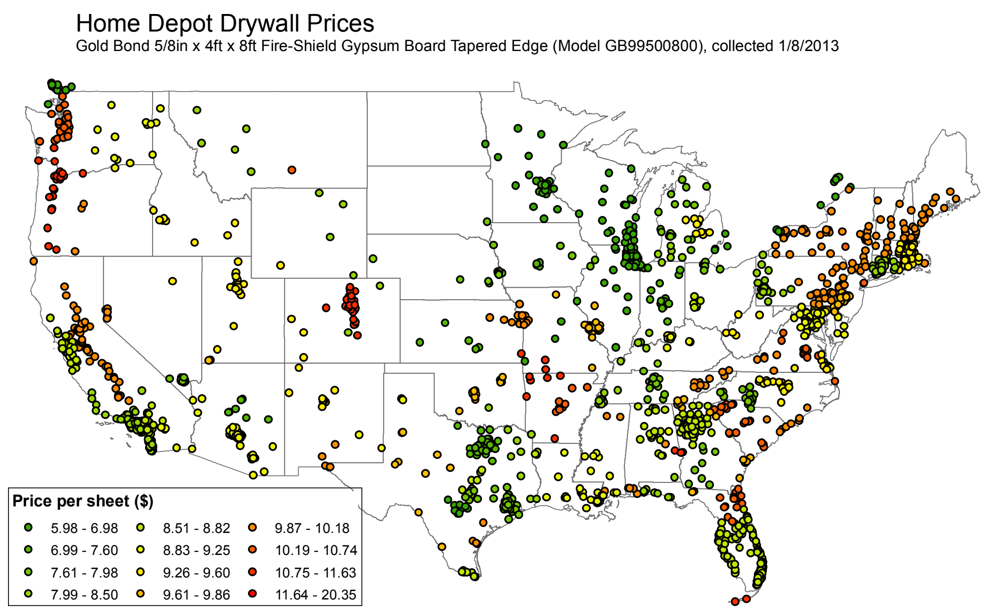


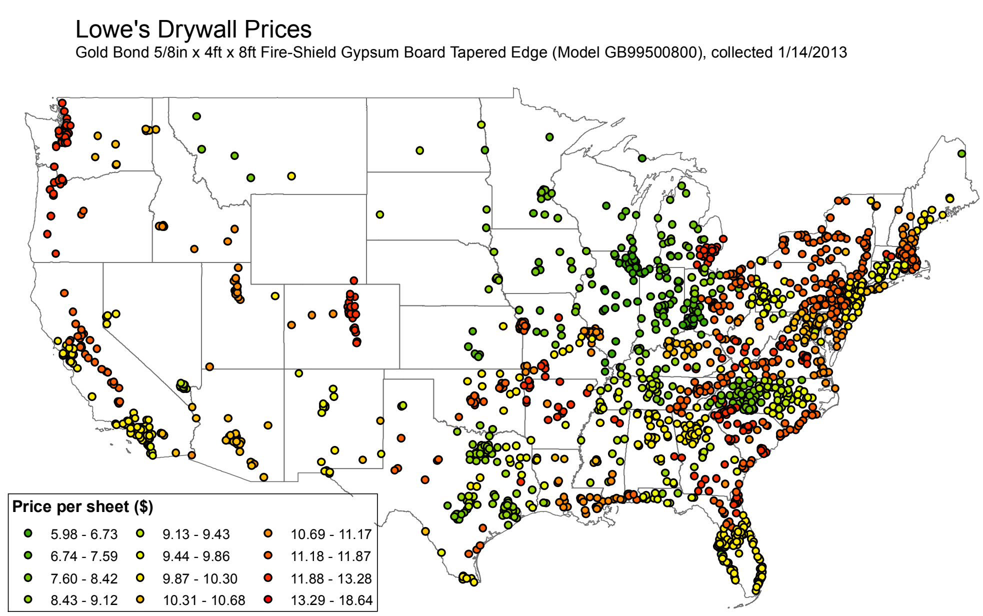

We take into account all products within a category when assigning zones. By definition, these pricing zones are no larger than the uniform price region for any product. Using our definition of a zone, we determine Home Depot has 165 drywall pricing zones for its 1,979 stores and Lowe's has 129 drywall pricing zones for its 1,714 stores in the US.[1]

Many pricing zones contain only one metropolitan area. Such zone might be justified for marketing reasons or because costs and competition are similar within a city. This would lead to profit-maximizing prices being the same across stores. The drywall pricing zones with the most stores for both Lowe's and Home Depot are in Southern California, both extending from Los Angeles to San Luis Obispo. Lowe's has 32 pricing zones that span more than 200 miles; Home Depot has 16 such zones. Pricing zones of such size represent multiple consumer markets. Drywall is bulky, making it unlikely that consumers would substitute to stores a great distance away.

A notable feature about the decision to use zone pricing in retail drywall is that costs and market structure vary considerably within a zone. Table 2.1 presents an example from a large drywall pricing zone based around Salt Lake City, Utah. The Home Depot stores in Logan, Utah; Rock Springs, Wyoming; and Elko, Nevada are all in this pricing zone, and hence, the prices for drywall within these stores are the same – the 5/8" x 4' x 8' drywall board is $10.98. The Home Depot in Logan faces competition from Lowe's, located a mile away. The nearest Lowe's to the stores in Rock Springs and Elko are 107 and 168 miles away, respectively. Further, at around 50 pounds per sheet of drywall, distance should play an important role in costs. The distance to the nearest distribution center and the distance to the nearest factory both vary by hundreds of miles. Profit maximizing prices for each of these stores should differ substantially, yet Home Depot places all three stores

[1] Three pricing zones had the prices for regular 5/8" x 4' x 8' sheets that were seen in other zones, hence the 126 distinct prices for that product.

in the same zone and assigns identical prices.

**Figure 2.2:** Map of unique prices for Home Depot 4’ x 8’ x 5/8” drywall.

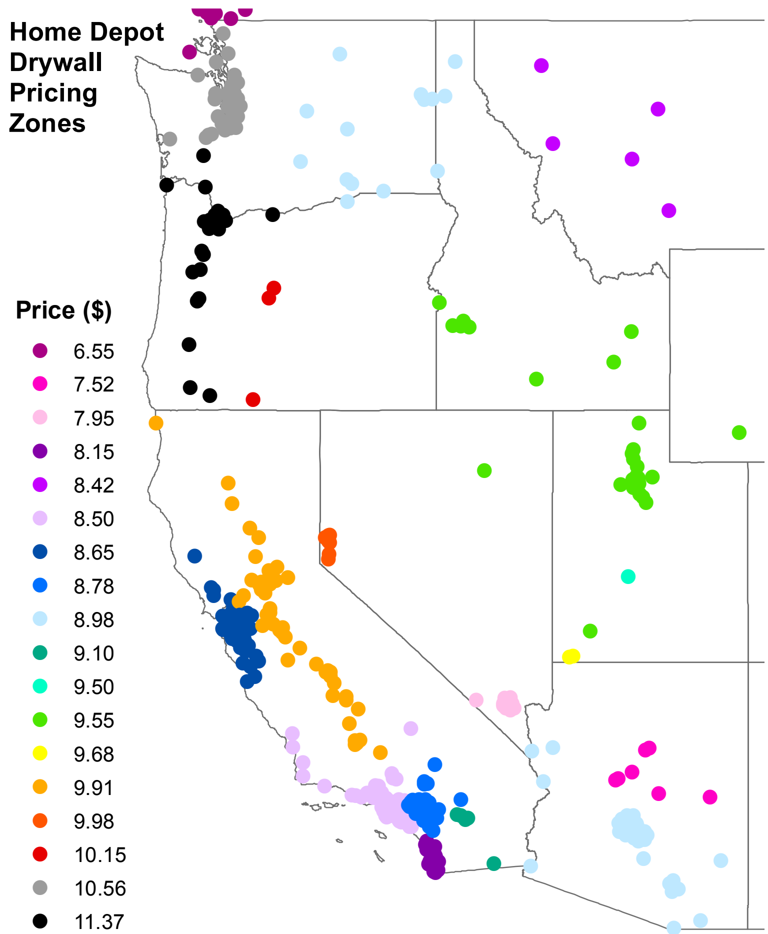

**Figure 2.3:** Zone pricing map for Lowe's 4' x 8' x 5./8" drywall in the Western United States. Each color denotes a unique price.

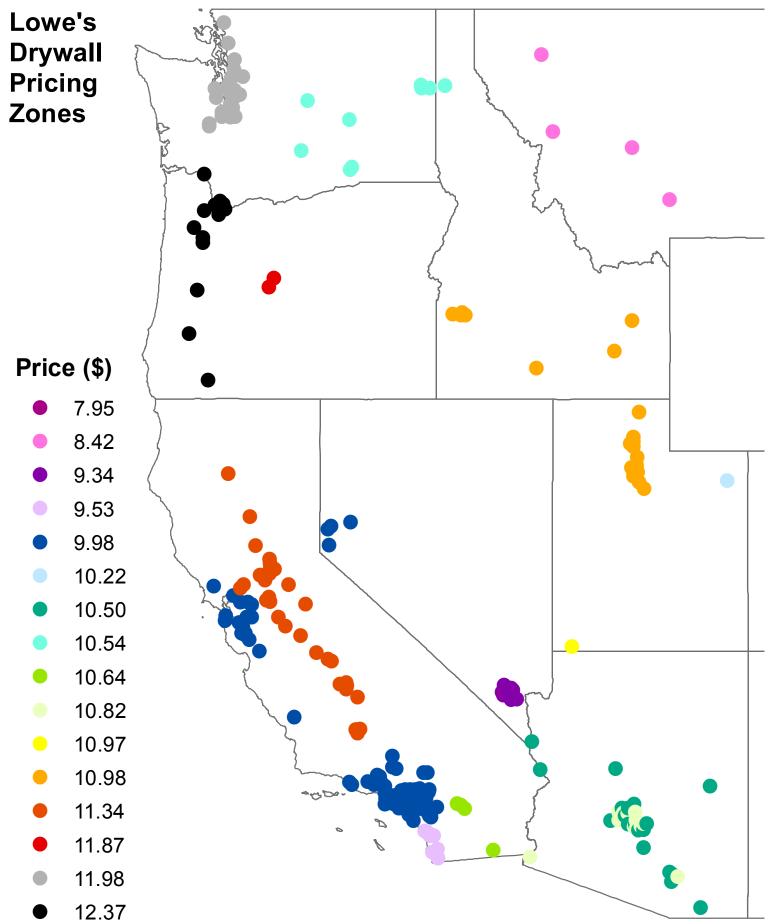


The drywall pricing zones for Lowe's and Home Depot in the United States often have the same boundaries. In 84.9% of ZIP codes with both Lowe's and Home Depot stores, prices for this product match exactly. The success of the law of one price is consistent with theories ranging from Bertrand competition to full collusion, but it is not consistent with Bertrand competition if products are strongly differentiated and costs differ between competing stores. However, drywall is not a highly differentiated product – and in estimating the model, we find that drywall products are highly substitutable within markets.

**Table 2.1:** Example documenting differences in costs and competition within a zone.

| | Home Depot Stores | | |
|---|---|---|---|
| | Logan, UT | Rock Springs, WY | Elko, NV |
| **Drywall Prices** | | | |
| regular 8’x4’x5/8” | \$10.98 | \$10.98 | \$10.98 |
| mold resistant 8’x4’x1/2” | \$11.47 | \$11.47 | \$11.47 |
| **Distances - nearest, miles** | | | |
| Lowe’s | 1 | 107 | 168 |
| American Gypsum factory | 743 | 821 | 491 |
| HD Distribution Center | 58 | 177 | 251 |

To estimate an empirical model of zone pricing with competition, we create an original data set of prices, sales quantities, and product characteristics for all drywall products available at 75 Home Depot stores and 53 Lowe’s stores in the Intermountain West. We select this region of the United States as it allows us to capture considerable variation in competition and costs, while keeping the data collection manageable. While pricing strategies do vary across product categories, drywall is the focus of this study for several reasons. Consumer markets are small and relatively well-defined, because buyers are unlikely to transport something as bulky and fragile as drywall far. Because of transportation difficulties, costs will vary predictably across stores. Drywall is rarely used in price promotions or as a loss leader, so category profit maximization is reasonable. Finally, drywall pricing zones are large enough to be economically interesting, but small enough that dozens of zones can be studied with a limited number of stores. We comment on the differences between drywall and other product categories in the next section.

Figure 2.4 maps the stores for which we have obtained quantity data. Our data set includes all stores in Idaho, Montana, New Mexico, Utah, western Colorado, eastern Washington, and stores in adjacent states needed to complete pricing zones. This region includes locations where only Home Depot operates (for example, Elko, Nevada) and locations where only Lowe's operates (for example, Vernal, Utah). Menards, the third largest home improvement warehouse, operates no stores in this region and so is omitted from the analysis.

**Figure 2.4:** Home Depot and Lowe's locations where detailed pricing and sales information was obtained

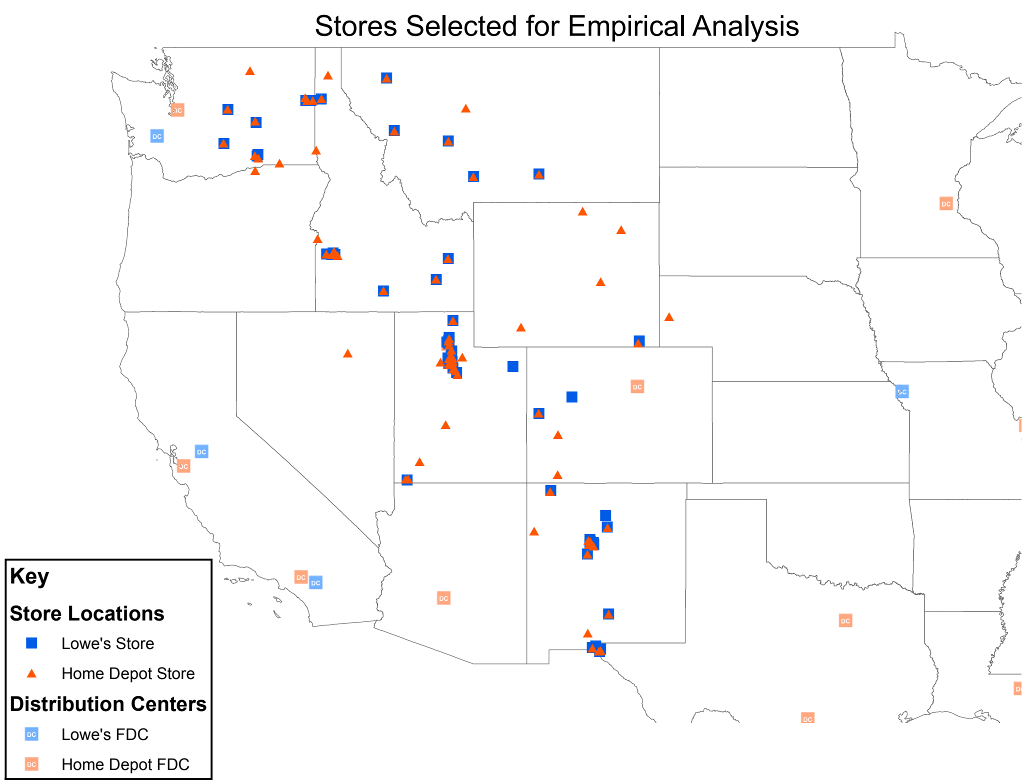


Our sample includes 14 complete Home Depot pricing zones and 11 complete Lowe's pricing zones. The area considered contains several single store pricing zones as well as one of the largest pricing zones in the nation. The zone boundaries largely match between chains. As in the national sample, the chains often charge the same price in locations where they compete.

We download prices and inventory levels for individual store stock keeping units (skus) and match these to products. For several products, Lowe's lists several brands on their website as different products, but those skus have identical prices and inventory levels that (with a one day lag) coincide perfectly. We eliminate these duplicates. Using manufacturer model numbers, we match products offered by both chains. In all, we identify 31 distinct products. We record the thickness, width, length, mold resistance, and moisture resistance for each product. We do not use brand identifiers because our site visits found brands frequently mislabeled at both chains.

Net changes in daily inventory levels for each product at each store are used to calculate sales quantities. Decreases in inventory levels give sales quantities. Inventory level increases of more than 20 sheets are classified as deliveries. When deliveries occur, we take the net change in inventory for the day as the volume of the shipment, meaning we assume no sales take place on delivery days. This systematically under-reports sales, but deliveries occur only every 16 days on average.[2] Smaller net increases in inventory levels are counted as returns, or negative sales.

Table 2.2 provides summary statistics for the data sample, which was collected between 12 February 2013 and 29 July 2013. The sample includes $N = 155,184$ observations. On average, daily inventory decreases by 6.2 and 7.0 sheets per product-store for Lowe's and Home Depot, respectively. Because the sales volume is low, we aggregate to the fortnight level in model estimation. The observed sales quantity would represent a small fraction of drywall used in new construction. We interpret Lowe's and Home Depot to be supplying the market for smaller consumer projects, such as wall repair or room remodeling. Construction contractors and

[2] We provide evidence in the estimation section that suggests our results are not sensitive to the possible measurement error in sales.

their supply networks are participating in a separate market. Consumers in our model will always have an outside option which will include buying from contractor suppliers.

**Table 2.2:** Summary statistics for the sample

| **Means** | Lowe's | Home Depot |
|---|---|---|
| Sales (per product, per day) | 6.19<br>(24.67) | 6.97<br>(19.77) |
| Delivery size (per product) | 200.90<br>(294.76) | 194.11<br>(216.05) |
| # Products (per store) | 8.60<br>(1.13) | 8.97<br>(1.15) |
| Revenue (per store, per day) | $531.87<br>(381.42) | $651.15<br>(305.15) |
| Price (per product) | $10.90<br>(3.04) | $12.04<br>(3.78) |
| Observations | 58,100 | 85,787 |

On delivery days, typically around a hundred sheets are delivered per product-store. The two chains have similar drywall product selection, offering around eight products per store. The price of drywall within a market ranges from just over $5 to over $20 per sheet, depending on the dimensions and features. The total drywall sales revenue for the 128 stores we study sums to $25.4 million per year.

Both Home Depot and Lowe's operate flatbed distribution centers. These distribution centers provide store locations with lumber and board products.[3]

[3]A Lowe's public document states: "FLATBED DISTRIBUTION CENTERS (FDC) - The purpose is to serve Lowe's stores with lumber, plywood, boards, and other building materials that can be forklift loaded onto flatbed trailers."

**Figure 2.5:** Histogram of distance for each store-product to the closest firm distribution center.

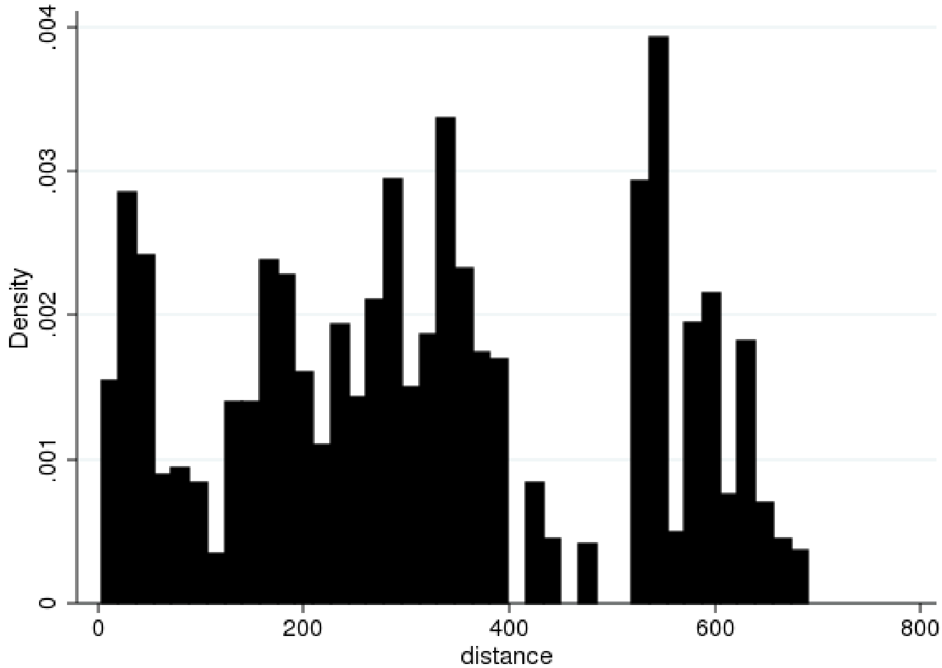


Using distribution center locations, we calculate the closest distribution center to each store, which we utilize in cost estimation. Figure 2.5 displays a histogram of the distances calculated. We find the average distance to stores from distribution centers is 318 miles, with a standard deviation of 190. At the extremes, the closest store to a distribution center is three miles, whereas the greatest distance is just over 690 miles. Distribution centers are usually near large markets and often Lowe's and Home Depot distribution centers are near each other. In our sample region, one big difference is Home Depot's placement of a distribution center in northern Utah, whereas Lowe's nearest distribution center is in southern Nevada. Since a sheet of drywall exceeds fifty pounds, we expect distance to be an important driver of costs. Labor costs may also be important. To measure them, we use ZIP code level wages for home improvement retailers from the Quarterly

Census of Employment and Wages.

### 2.2.2 Pricing Patterns in Retail Home Improvement

Home improvement warehouses chains use zone pricing in many, but not all of their product categories. The pricing zones for different products vary considerably in their size and coherence. Drywall, like other heavy products, has many small pricing zones. Other product categories have large regional pricing zones.

To explore the heterogeneity in zones across products, we collect national cross sections of prices in nine product categories: Drywall, Hardwood Plywood, Roof Underlayment, Stone Pavers, Window Film, Insulation Panels, Mosiac Glass Tiles, Phillips Screwdrivers, and LED Light Bulbs. Some of these categories are heavy and costly to transport (such as Stone Pavers and Drywall), while others are light or have a high value/weight ratio (such as LED Light Bulbs or Mosiac Glass Tiles). The products also vary in whether demand is seasonal or steady, whether demand is regional or national, and whether they are transported on flatbed trucks or in containers.

We collect prices for these categories for all store locations at Home Depot, Lowe's, and Menards, as each site offers store specific product selection and price quotes. In total, over 600,000 prices are recorded. Prices for most categories were collected in March 2014. The drywall and plywood prices are exceptions and are from April 2013. In our spot checks, the listed price on the website exactly matched the posted price inside the store.

Table 2.3 contains descriptive statistics for each category and chain. The table shows there is considerable variation in the number of products offered across stores and in the number of prices per product nationally. For example, Home Depot offers 45 unique products in the plywood category, and each store carries roughly 24 plywood products. On average, each product has 5.44 different prices nationally. The distribution is skewed; the median number of prices per product is one, suggesting that plywood at Home Depot is largely uniformly priced. These facts do not hold for plywood sold at Lowe's, where the median number of prices if 21. On the other hand, mosaic glass tile is a category where most products have a single, uniform price at all the stores within any of the three chains. House brand paint is an example of a product category that has no variation in price nationally for both Lowe's and Home Depot.

We define a pricing zone for a product category as before: a set of stores with exactly the same price on all products offered in the category (see Figure 2.6 for a clear example of zone pricing for a specific product offered at Menards). The pricing areas for individual products within a category usually align (such as drywall). There are occasional exceptions where the pricing areas for products within a category do not align (Window films at Lowe's is the most notable exception). This would occur, for example, if a product has two prices along a North-South line and for a second product to have two prices along an East-West line.

Since products within a category may have pricing areas that do not align perfectly, we employ an algorithm to classify the zone structures across chains and categories. Some ambiguity is created by differences in product availability. Consider, for example, three stores with common prices for most of their products. If one product is offered at a high price in the first store, offered at a low price in the second store, and not carried at all by the third store, then the third store could be in the pricing zone of either the first or second store. Our algorithm to classify zones starts with the westernmost store in a chain, which is declared to be in the first zone. If prices at the second westernmost store match for all the products at the first store, then this store is also in the first zone. If not, it is designated as belonging to a new pricing zone. This is repeated store by store from east to west. Other ways of sorting the stores yield slightly different zones and zone counts. We find pricing zones for the product categories at all three

**Table 2.3:** Data collected for descriptive evidence

| Category | Statistic | HD | Lowe's | Menards |
|---|---|---|---|---|
| Plywood | | | | |
| | $N$ | 47,462 | 77,417 | 5,559 |
| | Unique Prods. | 45 | 148 | 22 |
| | Prods. per store | 23.98 | 45.20 | 19.57 |
| | Prices per prod. | 5.44 / 1 | 23.65 / 21 | 5.19 / 4 |
| Roof Underlayment | | | | |
| | $N$ | 11,120 | 10,510 | 2,418 |
| | Unique Prods. | 28 | 25 | 11 |
| | Prods. per store | 5.65 | 8.34 | 8.45 |
| | Prices per prod. | 33.82 / 32 | 52.14 / 38 | 10.15 / 3 |
| Stone Pavers | | | | |
| | $N$ | 31,132 | 28,334 | 6,011 |
| | Unique Prods. | 130 | 303 | 37 |
| | Prods. per store | 15.84 | 22.49 | 21.02 |
| | Prices per prod. | 8.65 / 2 | 12.43 / 2 | 6.00 / 2 |
| Window Film | | | | |
| | $N$ | 19,680 | 40,320 | 4,839 |
| | Unique Prods. | 10 | 42 | 19 |
| | Prods. per store | 10.00 | 32.00 | 16.92 |
| | Prices per prod. | 2.1 / 1.5 | 4.53 / 2 | 3.42 / 3 |
| Insulation Panels | | | | |
| | $N$ | 29,868 | 14,242 | 5,093 |
| | Unique Prods. | 64 | 68 | 31 |
| | Prods. per store | 15.17 | 11.30 | 17.81 |
| | Prices per prod. | 14.48 / 7 | 22.50 / 22 | 16.53 / 19 |
| Mosaic Glass Tile | | | | |
| | $N$ | 30,710 | 40,272 | 12,716 |
| | Unique Prods. | 53 | 95 | 124 |
| | Prods. per store | 15.59 | 31.96 | 44.46 |
| | Prices per prod. | 2.08 / 1 | 1.66 / 1 | 2.07 / 2 |
| Phillips Screwdrivers | | | | |
| | $N$ | 29,613 | 40,232 | 30,939 |
| | Unique Prods. | 18 | 37 | 113 |
| | Prods. per store | 15.03 | 32 | 108.18 |
| | Prices per prod. | 4.93 / 5 | 1.31 / 1 | 1.01 / 1 |
| LED Light Bulbs | | | | |
| | $N$ | 47,280 | 40,320 | 26,976 |
| | Unique Prods. | 40 | 35 | 135 |
| | Prods. per store | 24.00 | 32.00 | 94.32 |
| | Prices per prod. | 6.23 / 4 | 3.02 / 3 | 1.62 / 2 |

Notes: $N$ denotes the number of product-store observations obtained. Unique Prods. is the number of unique products across all stores using both internal SKU and brand ID numbers. Prods. per store is the average store assortment size. Prices per prod. reports the mean and median number of prices across products, frequency weighted by the number of stores carrying each product.

**Figure 2.6:** Example of zone pricing at Menards.

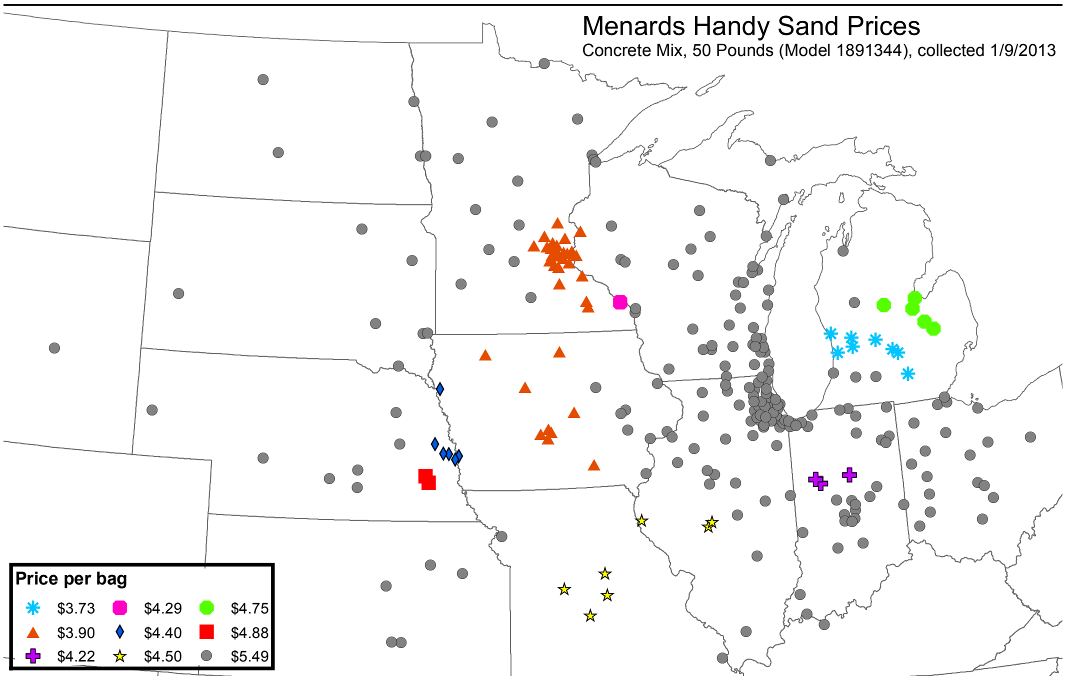


chains.

The pricing zones are not the same for different product categories. Figure 2.7 maps the pricing areas for a particular LED light bulb and sheet of drywall at Home Depot in the western United States. While drywall (right) has well-defined pricing regions that span hundreds of miles, the prices for the LED light bulb are more concentrated around a single price, with lower prices available in Southern California and select stores in Washington and Oregon. This variation in pricing strategy across categories is reflected in the difference in zone size and in the number of zones for categories.

Table 2.4 reports the number of zones classified using the algorithm described above. There is considerable variation in the number of zones across categories as well as the number of zones across firms within a category. Drywall, plywood, and roofing underfilm have the most pricing zones. On the other hand, there are

**Figure 2.7:** Comparison in zone structure for drywall and LED light bulbs.

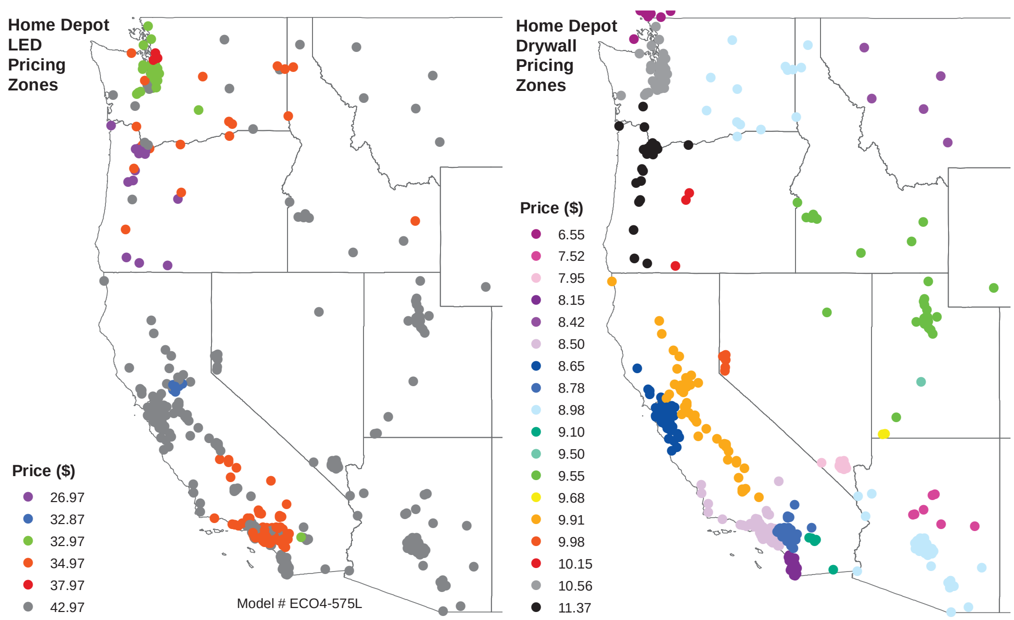


only a few national mosaic tile zones at all three retailers. While both Lowe's and Home Depot use zone pricing for power sanders, Menards has adopted a uniform pricing policy. Products with low weight-to-value ratios tend to have finer pricing zones, reflecting the importance of transportation costs.

For drywall, the pricing zones for Home Depot and Lowe's are quite similar in size and in their boundaries. This is not generally the case. In insulation and windows films, for example, Lowe's has much finer pricing zones.

Table 2.5 shows the number of stores in the largest price zone. Some product categories have expansive zones with a substantial fraction of a chain's stores sharing a common price on all products. In heavy, bulky products, shipping costs vary too much over such broad areas for the same set of prices to be maintained nationally.

The smaller zones are not simply subsets of the zones for other products. Zone boundaries for different product categories intersect, with a zone in one category

**Table 2.4:** Number of pricing zones, by product category

| | Number of zones | | |
|---|---|---|---|
| Category | Home Depot | Lowe's | Menards |
| Drywall | 165 | 129 | 50 |
| Hardwood Plywood | 69 | 247 | 47 |
| Store-brand interior paint | 14 | 1 | 1 |
| Insulation | 9 | 170 | 91 |
| Mosaic | | 9 | 4 |
| Power Sanders | 24 | 22 | 1 |
| Roofing | 51 | 213 | 75 |
| Window Films | 9 | 108 | 28 |

**Table 2.5:** Number of stores in largest pricing zone, by product category

| | Number of stores | | |
|---|---|---|---|
| Category | Home Depot | Lowe's | Menards |
| Drywall | 80 | 94 | |
| Hardwood Plywood | 232 | 43 | |
| Store-brand interior paint | 743 | 1714 | |
| Insulation | 1759 | 170 | 21 |
| Window Films | 805 | 616 | 38 |
| Roofing | 498 | 32 | 24 |
| Power Sanders | 569 | 548 | |

overlapping multiple zones in another category and vis versa. Figure 2.7 compares pricing regions for a single pair of products in the LED light bulbs (on the left) and drywall (duplicating Figure 2.2 on the right) categories. The drywall product has different prices in the San Francisco Bay Area, California's Central Valley, and Utah's Wasatch Front, but all three regions pay the same price for the light bulb. All of Western Oregon is part of one drywall pricing zone, but three different prices of light bulb are found in the area. Pricing zones for whole pricing categories exhibit the same complexity as the individual products of Figure 2.7.

## 2.3 Model

In this section we introduce the structural model of supply under zone pricing and demand. Competing firms operate stores in multiple markets, and each store sells multiple products. Consumers choose between all the products at stores in their market according to a standard discrete choice setup. Firms set prices in two-stages. First, they partition their stores into pricing zones. Next, they simultaneously chose the price levels of each zone to maximize profits subject to the zones they have chosen. To start, we introduce some notation and our definition of a pricing zone. Next, we describe the consumer's problem. Finally, we detail the pricing game the firms play and show how spatial menu costs determine the selection of pricing zones.

### 2.3.1 Products, Stores, Markets, and Zones

Each firm operates a networks of stores. Stores, indexed by $s$, each have a location $\ell$. Firms may operate more than one store at $\ell$. Let $S_\ell^f$ be the set of stores operated by firm $f$ at location $\ell$, and let $S_\ell := \bigcup_f S_\ell^f$ be the set of all stores in location $\ell$.

Each firm partitions its stores into pricing zones. Let these partitions be denoted by $Z^f$. An element in the partition is comprised of all the stores in the same pricing zone. Conceptually, if firm $f$ (superscript suppressed) operates four stores across two zones, with a single store in the first zone,

$$\begin{aligned} Z &:= \{z_1, z_2\} \\ &:= \Big\{ \underbrace{\{s_1\}}_{z_1}, \underbrace{\{s_2, s_3, s_4\}}_{z_2} \Big\}. \end{aligned}$$

If $s_1$ is in location 1, $s_2$ is in location 2, and $s_3$ and $s_4$ are both in location 3.

The same partition could also be described by

$$Z = \Big\{ \underbrace{\{S_1\}}_{z_1}, \underbrace{\{S_2, S_3\}}_{z_2} \Big\}.$$

Zone pricing implies that for every product $j$ in every period $t$

$$p_{jst} = p_{js't}, \, \forall s, s' \in z$$

With the example above, the firm uses store level pricing for $z_1$. However, for all $j \in \bigcap_{i \in 2,3,4} J_{s_i}$, the price is constant for $j$ across $z_2$. Because the product set is allowed to be different across stores, the definition of a zone implies that if a product is offered at at least two stores within a zone, the price is the same across the stores.

Compared to the observed zone structure, alternative pricing regimes are associated with a spatial menu cost $\mu^{Z^f;Z^{-f}}$ that the firm must pay. These menu costs encompass all costs related to changing the zone structure – including reevaluating profits.

In total, there are $J$ differentiated drywall products, where each store offers a subset of these products each period. Let $J_{s,t}$ be the set of products offered at store $s$ in period $t$. The product set may change over time due to inventory or the discontinuation or introduction of a product. The discontinuation, introduction, and overall selection of products is not modeled; however, the product set is mostly constant within a zone so there is little evidence of strategic product placement. Given products, product characteristics, and prices, consumers at each location solve for demands.

A market is a location in time period, $(\ell, t)$. It has a market size of $M_{\ell,t}$.

### 2.3.2 Demand

Consumers solve a nested-logit discrete choice utility maximization problem. Consider a consumer living at location $\ell$. The choice set facing this consumer is the set of products sold by all stores at $\ell$, that is $\bigcup_{s \in S_\ell} J_{s,t}$ or an outside option. The decision to not purchase a good yields a normalized utility, $\mathcal{U}_{i0t} = \epsilon_{i0t}$. By purchasing product $j$, consumer $i$ receives indirect utility

$$\mathcal{U}_{ijt} = \mathbf{x}_j \boldsymbol{\beta} - \alpha p_{jt} + \xi_{jt} + \zeta_{igt}(\sigma) + (1-\sigma)\epsilon_{ijt}, \tag{2.1}$$

where $\beta$ measures preferences over a vector of product characteristics $\mathbf{x}_j$, $p_{jt}$ is price, $\alpha$ is the marginal utility to income, and $\xi_{jt}$ is unobserved (to the econometrician) product quality. The composite taste shock, $\zeta_{ig}(\sigma) + (1-\sigma)\epsilon_{ij}$, follows a Type-1 Extreme Value distribution among group $g$ – the nesting variable. The outside good is in its own nest. Note when $\sigma = 0$, the composite error term simplifies to just $\epsilon_{ij}$, which yields the standard logit demand system. As $\sigma \to 1$, products within nests are increasingly close substitutes, and in the limit, when $\sigma = 1$, there is no substitution outside of the nest. Each period consumers purchase the good that maximizes their individual indirect utility $U_{ijt}$ or select the outside good if $\mathcal{U}_{i0t} > \mathcal{U}_{ijt}$ for all $j \in \bigcup_{s \in S_\ell} J_{s,t}$.

As shown in Berry (1994), given the logit structure of demand, the log difference in market share of good $j$ compared with outside good $j = 0$ equals

$$\ln(\varsigma_{jt}) - \ln(\varsigma_{0t}) = \mathbf{x}_j \boldsymbol{\beta} - \alpha p_{jt} + \sigma \ln(\varsigma_{jt}/g) + \xi_{jt}.$$

Here, $\varsigma_j/g$ is the market share of product $j$ within group $g$.[4] The demand parameters to be estimated are $\boldsymbol{\theta}^D = (\boldsymbol{\beta}, \alpha, \sigma)$. We address the endogeneity of prices in Section 2.4.1.

[4] We use $\varsigma_j$ to denote the purchase probability (market share) instead of the typically seen $s_j$ because we use $s$ to denote a store.

### 2.3.3 Supply

Prices are set in two stages. First, firms simultaneously partition their stores into zones, selecting $Z^f$ and paying spatial menu cost $\mu_{Z^f}$. After zone partitions are publicly known, firms select zone prices for each product to maximize total firm profits. The second stage price decisions differ from standard multiproduct Bertrand competition only in that firms are constrained for a given product to set identical prices within a zone.

The first stage is a simultaneous move game in which every possible zone partition is a possible action. Payoffs in the first stage depend on the spatial menu costs for the zone partition selected and the pre-menu cost profits that emerge from Bertrand competition in the second stage.

Let $c_{js}$ be the constant marginal cost associated with offering $j$ at store $s$. Given a zone structure chosen in the first stage, the profits accrued to a firm for selling product $j$ in period $t$ are

$$\pi_j^f := \sum_{z \in Z} \sum_{s \in z} (p_{jz} - c_{js}) q_{js}, \tag{2.2}$$

where $q_{js} := M_\ell \varsigma_{js}$ and $M_\ell$ is the market size corresponding to the location of store $s$. Implicitly, only zones and stores that offer $j$ are included in the sum, and $\varsigma_j := \varsigma_j(\mathbf{X}, \mathbf{p}, \boldsymbol{\xi}; Z, \boldsymbol{\theta}^D)$. Lastly, we assume there are no further fixed costs associated with offering products.

Firms maximize total profits. Total profits are the summation of profits over the products offered by the firm minus the spatial menu costs. Once zones are set, the second stage profit maximization problem involves selecting a price for every zone-product:

$$\max_{\mathbf{p}^f} \sum_{z \in Z^f} \sum_{s \in z} \sum_{j \in J_s} (p_{jz} - c_{js})\, q_{js} - \mu(Z^f; Z^{-f}). \tag{2.3}$$

Each market share $\varsigma_{jst}$ is a function of the prices in market $(\ell, t)$, including all competitor prices.

The model is quite general and encompasses several cases which are commonly seen in retailing. In the special case of uniform pricing, there is only one zone for the firm and it contains the entire network. Since, $Z \equiv z := \{S\}$, the profit maximization takes simpler form. The first sum in Equation 2.3 disappears entirely so

$$\max_{\mathbf{p}^f} \sum_{s \in S} \sum_{j \in J_s} \Big(p_{jZ} - c_{js}\Big) q_{js} - \mu^{\text{uniform}}.$$

Another possibility is that firms operate store-level pricing, or market-level pricing Finally, in the second stage, firms take price zones as given and set prices to maximize profits.

An equilibrium for the game depends on zones and prices chosen amongst the players. Formally, an equilibrium is a set of pricing zones $Z^*$, prices $\mathbf{p}^* \in \mathbb{R}_+^{|Z^*|}$, and market shares $\boldsymbol{\varsigma}^* \in \mathbb{R}^{|J_S|}$ such that

1. Given pricing zones $(Z^*)$ and competitor prices $(\mathbf{p}^{*-f})$, $\mathbf{p}^{*f}$ solves Equation 2.3

2. Given competitor pricing zones $Z^{*-f}$, $Z^{*f}$ is chosen such that

$$\pi^{*f}\Big(Z^{*f}; Z^{*-f}\Big) \geq \pi^{*f}\Big(Z'^{f}; Z^{*-f}\Big) \ \forall \ Z'^{f}$$

3. Given prices $\mathbf{p}^*$, $\boldsymbol{\varsigma}^*$ follows from consumers solving Equation 2.1

We assume firms play a game of perfect information in pure strategies.

Unfortunately, the equilibria defined for a given zone structure $Z$ are not in general unique. We found dozens of distinct equilibria for each system of zone partitions we examined. Caplin and Nalebuff (1991) proves uniqueness for competition within multinomial logit demand systems for single product firms, but its result does not generalize to the multiproduct firms we see in our data. In particular, there are equilibria in which firms assign high prices to some products to take advantage of the tail consumers with particularly high $\epsilon_{ij}$ draws while shepherding the rest of the consumers into moderately priced products. Different equilibria assign the role of the moderately priced mainstream alternative to different products and vary as to which store engage in this version of price discrimination.

In order to evaluate the menu costs associated with alternative pricing policies we must ensure that differences we find are due to the policies themselves, and not due to switching between vastly different equilibria. We therefore investigate small deviations from the observed equilibrium, allowing firms to switch to market level pricing one product at a time. By only allowing for adjustments in prices of a specific good in a single market, the result of Caplin and Nalebuff (1991) does guarantee unique equilibria. Hence, in our definition of a Bertrand-Nash Equilibrium, we condition on the zone structure instead of having firms choose zones. We also explore a move to alternative pricing regimes for all products simultaneously. We discuss a metagame of zone choice and pricing, and calculate menu costs after implementing a selection mechanism on equilibrium. All experiments are reported in Section 5.

## 2.4 Estimation and Results

We proceed by estimating the demand parameters which enter shares $\varsigma$ as well as marginal costs. We run our estimates in two stages. First, we estimate the demand parameters of the nested logit model. Given estimates of the demand parameters, we solve for marginal costs assuming firms are competing in a Bertrand game of zone pricing. As zone pricing is a consequence of solving a constrained optimization problem, we cannot invert the first-order conditions to back out marginal costs. Instead, we parametrize the cost function, and simultaneously recover marginal costs and cost parameters using mathematical programming with equilibrium constraints (MPEC). With our estimates, we calculate observed profits of the current zone structure. Given the presence of multiple equilibria, we use observed prices to calculate current zone profits.[5]

### 2.4.1 Demand

We invert market shares, as shown in Berry (1994), to obtain $\ln(\varsigma_j) - \ln(\varsigma_0) = \delta_j$, where $\delta_j$ is the mean utility from purchasing product $j$. To account for the endogeneity of unobserved product quality being correlated with price, we pursue both instrumental variable and fixed effects approaches. We separate unobserved product quality as

$$\xi_{jt} := \xi_j + \xi_t + \Delta\xi_{jt},$$

where we assume $\Delta\xi_{jt}$ is uncorrelated with price and observed product characteristics. For product characteristics, we use the dimensions of the gypsum board (length, width, height), and whether the drywall is mold resistant and/or moisture resistant. We estimate $\xi_t$ as a market-time fixed effect. We compare this method with an instrument variables approach, where we instrument price by using a Hausman instrument – average prices for a given product in other markets where the product is offered.

We use a product-store hierarchy for nests within markets. The top level of nests denotes the various product types available in the market $(\ell, t)$. We define product types as the grouping of product dimensions. The second nest comprises the various stores in the market that sell that particular product type. The interpretation of $\sigma$ in this nesting structure is the degree of substitutability

[5] By solving for equilibrium zone prices given observed prices as starting values, we obtain equilibria prices quite close to observed prices. The geometric fit between the two is 93%. This is discussed further in Section 5.

of a product type across the various stores at location $\ell$ at time $t$. We specify the nesting structure this way because of conversations we have had with home builders, who say that by far, the most important characteristic of drywall is the size, particularly the thickness. Specifically, almost all walls use 1/2-inch boards, but 5/8-inch is necessary on fire walls, such as the walls separating the interior from a garage. If size is the most important characteristic, we would expect that consumers substitute to other stores in the market for a particular drywall type instead of substituting to a different size sheet at the same store. Hence, our nests are at the product type level instead of the store level. We expect to (and do) estimate $\sigma$ close to one, which suggests that a particular product type is highly substitutable across stores within a market.

Identification for parameters in this stage results from the observed purchases of consumers given their choice set, following the standard revealed preference identification for discrete choice demand systems. In every market, all products and their associated prices and characteristics are known. Product indicators and characteristics are constant across all markets. The other products offered vary by market and relative prices vary by zone. The response of sales quantities to these different relative prices and product offerings identifies the price coefficient $\alpha$. Preferences for observed and unobserved product characteristics are revealed through market shares.

We estimate several demand specifications. First, we set $\sigma = 0$ in the nested logit model so that the nests do not matter. This results in the classic logit demand system. We estimate this specification assuming prices are exogenous (using Ordinary Least Squares) and then use a fixed-effects approach to address the endogeneity of prices. We then estimate $\sigma$ along with the demand parameters, again assuming prices are exogenous, and then accounting for endogeneity. Given the nesting structure, all observed product characteristics within a nest are identical, except price which may vary across the stores in the market. Hence, for the nested logit model using instrumental variables, we use store fixed effects to address the endogeneity problem on the group shares. For the nested logit model using fixed effects, the endogeneity problem on group shares is already addressed by having the fixed effects be product-store ("$j$") specific.

We aggregate daily sales so that $t$ denotes a two week period because many drywall products have few daily sales. Observations with zero sales must be dropped, because $\ln(\varsigma_{jt})$ would be undefined[6] At the biweekly level, 5.70% of products exhibit zero sales. Aggregating data across time does not introduce as much measurement error as it might in other applications, because product characteristics are all time-invariant and prices rarely change. Over 127 days of data collection, only 9.4% of product-store combinations exhibit price changes. Of the products to see price adjustments, 88% (77 products) of them experience a single change and 12% (11 products) see two price adjustments.

In order to complete the demand system, we need to specify market size. We define a market to be a Core Based Statistical Area (CBSA) over a two week period. For stores not located within CBSAs, we set the market to be the county in which the store resides. With this interpretation of markets, each location $\ell$ typically has several stores from both Home Depot and Lowe's. Further, given the structure of pricing regions by both firms, regions overlap into several markets. We take market size to be proportional to the 2010 CSBA population.[7]

The results of the demand estimation appear in Table 2.6. Across all specifications, all coefficients have expected signs and are significant across specifications. We estimate that consumers prefer larger drywall sheets and mold resistant panels. The unreported coefficients on drywall thickness are reasonable and show that industry standard 1/2-inch panels are much more desirable than all other

[6] See Gandhi, Lu, and Shi (2013) for estimating discrete choice demand systems with products that exhibit zero sales.

[7] For observations not within CBSAs, we take the population to be proportional to the 2010 Census county population.

**Table 2.6:** Demand estimation results

| | (1) Logit | (2) Nested logit | (3) IV nested logit | (4) FE nested logit |
|---|---|---|---|---|
| Price | -0.594*** | -0.650*** | -0.831*** | -0.359*** |
| | (0.0172) | (0.0108) | (0.0129) | (0.0141) |
| Area | 0.130*** | 0.158*** | 0.194*** | 0.107*** |
| | (0.00404) | (0.00269) | (0.00308) | (0.00639) |
| Mold resist. | 1.167*** | 1.385*** | 1.838*** | 1.555*** |
| | (0.0667) | (0.0441) | (0.0457) | (0.260) |
| Chain | 0.938*** | 0.295*** | 1.719*** | 0.292* |
| | (0.0334) | (0.0231) | (0.176) | (0.116) |
| $\sigma$ | | 0.918*** | 0.601*** | 0.831*** |
| | | (0.00821) | (0.0119) | (0.00887) |
| Thickness FE | Yes | Yes | Yes | Yes |
| Store FE | No | No | Yes | Yes |
| Product FE | No | No | No | Yes |
| elasmean | -6.620 | -55.99 | -17.58 | -15.79 |
| elasmin | -12.92 | -136.7 | -35.87 | -36.57 |
| elasmax | -2.665 | -2.894 | -2.869 | -1.609 |

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

thicknesses. We estimate that consumers are price sensitive, with the marginal utility of income $-0.375$ in the fixed effects, nested-logit model (Model 4). Our estimates on price sensitivity do not become more negative when accounting for the potential endogeneity between prices and unobserved drywall quality, which is not what we would expect as high priced items are typically assumed to be positively correlated with unobserved quality. This is true for both the fixed effects and instrumental variables approach. Our interpretation on the price coefficient across specifications is that given any two drywall products with identical (observed) characteristics but different prices, consumers would gravitate towards the cheaper good as the other (unobserved) characteristics are not worth the additional expense. Hence, in our setting, the correlation between unobserved quality and price is negative leading to price sensitivities closer to zero after accounting for endogeneity. We estimate the mean own-price elasticities to be -16 and -18 for the fixed effects and instrumental variables nested logit models, respectively. These values are large in magnitude but not unreasonable given the high substitutability of drywall products, especially within nests. For the final specification (4), we obtain industry elasticities of $-0.03$ to $-0.04$ depending on the market. Finally, we estimate the coefficient reflecting substitutability within nests to be high, at 0.830 in the last specification. This suggests that consumers would rather drive to another store within the market to buy a particular drywall panel than substitute to a different size.[8]

In the following analysis, we use Model (4) – the nested logit model with fixed effects – as our model of consumer demand. Our results are not sensitive to this choice as the nested logit model with instrumental variables yields quantitatively similar answers. Aggregating the data to just the week level also yields similar results.

### 2.4.2 Recovering Marginal Costs

Marginal costs can typically be recovered using the demand estimates and the first order conditions of each profit maximizing firm. In the single good, market level pricing case, Lerner's index is inversely proportional to the own-price elasticity. With observed prices and an estimated demand elasticity, the marginal cost is

[8] Our model of demand assumes it is costless for consumers to travel to stores within a market.

identified. A multi-product analog, as seen in Nevo (2001) and Petrin (2002), can be used when firms set prices for all products at the market level. Since firms here set a uniform zone price, the first order condition based on Equation 2.4 differs from the condition on which the standard approach is based. Each price $p_{jz}$ is obtained from solving

$$\frac{\partial \pi}{\partial p_{jz}} = \sum_{s \in z} \sum_{i \in J_s} (p_{iz} - c_{is}) M_\ell \frac{\partial \varsigma_{is}}{\partial p_{jz}} + \sum_{s \in z} M_\ell \varsigma_{js} = 0. \tag{2.4}$$

The first-order condition for each price contains marginal costs for all stores within its pricing zone. The supply system yields $|J \times Z|$ first-order conditions of the form in Equation 2.4. However, there are $|J \times S|$ marginal costs to identify and $|S| > |Z|$. As there are more marginal costs than first-order conditions, no set of first-order conditions can be directly solved to recover marginal costs.

To make progress in recovering marginal costs, we first parametrize costs as

$$c_{jst} = a_j + \kappa d_s + \nu_{jzt},$$

where $a_j$ is a fixed effect for product $j$, and $d_s$ is the distance from store $s$ to the closest flatbed distribution center. The cost shock $\nu_{jzt}$ is an unobserved (to the econometrician) and enters at the product, zone, time level. Let $\theta^S = (\mathbf{a}_j, \kappa)$. Because the cost shock is at the product level, we cannot manipulate Equation 2.4 to recover costs directly; however, given an objective function on $\nu$, we can simultaneously recover marginal costs and the parameters governing costs. Instead of using a nested fixed-point approach, we proceed with using mathematical programming with equilibrium constraints (MPEC) as seen in Su and Judd (2012). Forming moment conditions on $\nu$ directly, along with the optimality conditions from the firms' zone pricing problems, completes the mathematical program.

Firms want high prices at stores with high costs and at stores with monopoly power. An unconstrained optimum price for a competitive, low-cost store would be lower. When a firm has both types of stores in the same price zone, the optimal price balances these considerations using Equation 2.4. Only the zone price, market power (through estimated price elasticities), market size, and a few cost variables are observed, but the zone price reveals information on the marginal costs for its stores. Identification effectively comes from how weighted averages of store cost variables are correlated with the zone price. For example, if zones full of competitive stores far removed from a distribution center have a high zone price, then distribution center distance is an important driver of costs.

Our simplification in making the unobserved error term $\nu$ be product, zone, time specific instead of product, store, time specific results in the dimension of the error term being equal to the number of equilibrium conditions. Without this assumption on the cost shock, we would have an unidentified system. Although restrictive, we do account for transportation costs and wages at the store level.

The objective comes from moment conditions on $\boldsymbol{\nu}$. Let $\mathbf{W} := [\mathbf{a}_j, \mathbf{d}]$ be the matrix of covariates on costs. The method of moments estimator is derived from $\mathbb{E}[\mathbf{W}'\boldsymbol{\nu}] = \mathbf{0}$, leading to the sample analogue

$$g_j(\mathbf{W}, \boldsymbol{\theta}^S) = \frac{1}{N} \sum_{i=1}^{N} w_{ji} \nu_i = 0.$$

Letting $\text{FOC}(\boldsymbol{\theta}^S, \boldsymbol{\nu})$ denote the set of equilibrium conditions characterized by the first-order conditions of the firms' problems, the MPEC program is to solve

$$\begin{aligned} &\min_{\boldsymbol{\theta}^S, \boldsymbol{\nu}} g(\mathbf{W}, \boldsymbol{\theta}^S)' g(\mathbf{W}, \boldsymbol{\theta}^S) \\ &\text{s.t. } \text{FOC}(\boldsymbol{\theta}^S, \boldsymbol{\nu}) = \mathbf{0}. \end{aligned}$$

In estimating costs, we obtain a negative, but insignificant coefficient on wage.

We drop wage from the model and proceed with estimating product fixed effects and the coefficient on distance. The remaining coefficients are very similar to the model with wages included. Cost estimations are reported in Table 2.7. We estimate the parameter on distance (per mile) to be \$0.00052. On average, transportation from the distribution center contributes \$0.20 to the cost of a drywall sheet. We find transportation costs for different stores range from \$0.002 to \$0.46. The coefficient on distance is lower than other estimates of transportation costs in similar settings. Miller and Osborune (2011) estimate a transportation cost \$0.30/ton mile for Portland cement. The equivalent cost for a fifty pound sheet of drywall would be \$0.0075/mile, which is over eleven times larger than we find. If, however, other products shipped to stores on flatbed trucks need frequent deliveries, perhaps much of the drywall inventory is shipped on trucks with spare capacity. Indeed, we find that the deliveries for drywall are less than the full capacity of a flatbed trailer.

### 2.4.3 Observed Pricing Regime

With observed prices and the marginal costs we estimate, we calculate the sales weighted average markup on a sheet of drywall to be \$1.11. With an average price of a sheet of drywall at \$10.22, we find the margin on drywall to be around 11.0%. We estimate profits on drywall for the stores of interest to be about \$29 million annually. Table 2.8 details equilibrium zone pricing profits by chain and competition type. Only 3 of the 53 Lowe's stores in our region are in markets where Home Depot is absent. Home Depot on the other hand has 20 of their 75 stores in markets without competition from Lowe's. These twenty account for 38% of Home Depot's revenue. Interestingly, 15 of the 20 Home Depot monopoly stores are in pricing zones with stores that do face competition from Lowe's. A higher price that would extract the most profit in the monopoly markets must

**Table 2.7:** Cost Estimates

| | Point Estimates | Std Error |
|---|---|---|
| dist ($\kappa$) | 0.00052 | (0.00016)$^{***}$ |
| | *Product Fixed Effects* | |
| $a_1$ : 8.384 | $a_{11}$ : 11.631 | $a_{21}$ : 12.609 |
| $a_2$ : 7.561 | $a_{12}$ : 12.154 | $a_{22}$ : 11.323 |
| $a_3$ : 2.088 | $a_{13}$ : 10.670 | $a_{23}$ : 14.052 |
| $a_4$ : 7.137 | $a_{14}$ : 10.900 | $a_{24}$ : 18.613 |
| $a_5$ : 7.947 | $a_{15}$ : 12.395 | $a_{25}$ : 13.860 |
| $a_6$ : 7.948 | $a_{16}$ : 12.063 | $a_{26}$ : 9.760 |
| $a_7$ : 9.307 | $a_{17}$ : 8.921 | $a_{27}$ : 9.989 |
| $a_8$ : 14.592 | $a_{18}$ : 11.042 | $a_{28}$ : 15.298 |
| $a_9$ : 7.137 | $a_{19}$ : 11.180 | $a_{29}$ : 10.236 |
| $a_{10}$ : 12.990 | $a_{20}$ : 13.269 | $a_{30}$ : 14.363 |
| | | $a_{31}$ : 14.873 |

Zone clustered standard errors. All FEs significant at 1%
$^{*}$ $p < 0.05$, $^{**}$ $p < 0.01$, $^{***}$ $p < 0.001$

be balanced by a lower price needed to maintain market share in competitive markets. Because several of the monopoly store are in large, lower cost zones, average prices for monopoly markets are slightly below average prices overall. As a consequence, the current zone structure greatly limits Home Depot's effective market power. Indeed, we estimate Home Depot only obtains 16.48% of its profits from monopoly stores

Figure 2.8 plots a histogram of observed profits, by chain, aggregated by store. The histograms show there is considerable variation in profits across stores. In this region, Home Depot distribution centers are closer to more stores, so our estimates generally find Home Depot stores to have lower costs and higher profits than Lowe's stores. We estimate a majority of the Lowe's stores have less than

\$1 million in annual profits for drywall, with the maximum profits being nearly \$2 million annually. On the other hand, Home Depot operates a few stores that exceed \$2 million in annual profits. Both chains operate stores with nearly zero profits from drywall sales.

## 2.5 Alternative Pricing Regimes

Here we calculate optimal pricing policies by changing the zone structure of firms. Given the presence of multiple equilibria, we investigate moving one product from its current zone pricing scheme to market level pricing. In this exercise we hold all other product prices constant at their observed levels. This guarantees a unique equilibrium for the exercise. We first consider holding competitor prices fixed. In a second exercise, we allow both chains to switch to market level pricing for a single product. Previous work on zone pricing has not taken into account this competitor response, and we highlight that this leads to an overstatement in profit gains by 33% in retail drywall. We discuss how profit changes in either experiment relate to various interpretations of spatial menu costs.

**Table 2.8:** Current Profits by Chain and Market Type

| | Lowe's | | Home Depot | |
|---|---|---|---|---|
| | % of $\pi$ | Annual $\pi$ | % of $\pi$ | Annual $\pi$ |
| Monopoly | 3.28% | \$273,462 | 16.48% | \$3,356,029 |
| Duopoly | 96.72% | \$8,057,492 | 83.52% | \$17,007,531 |
| Annual $\pi$ | | \$8,330,954 | | \$20,363,560 |

Duopoly means there is a competitor store in the market (CSBA). Profits are annualized.

**Figure 2.8:** Histogram of observed profits, by chain, and aggregated to the store level.

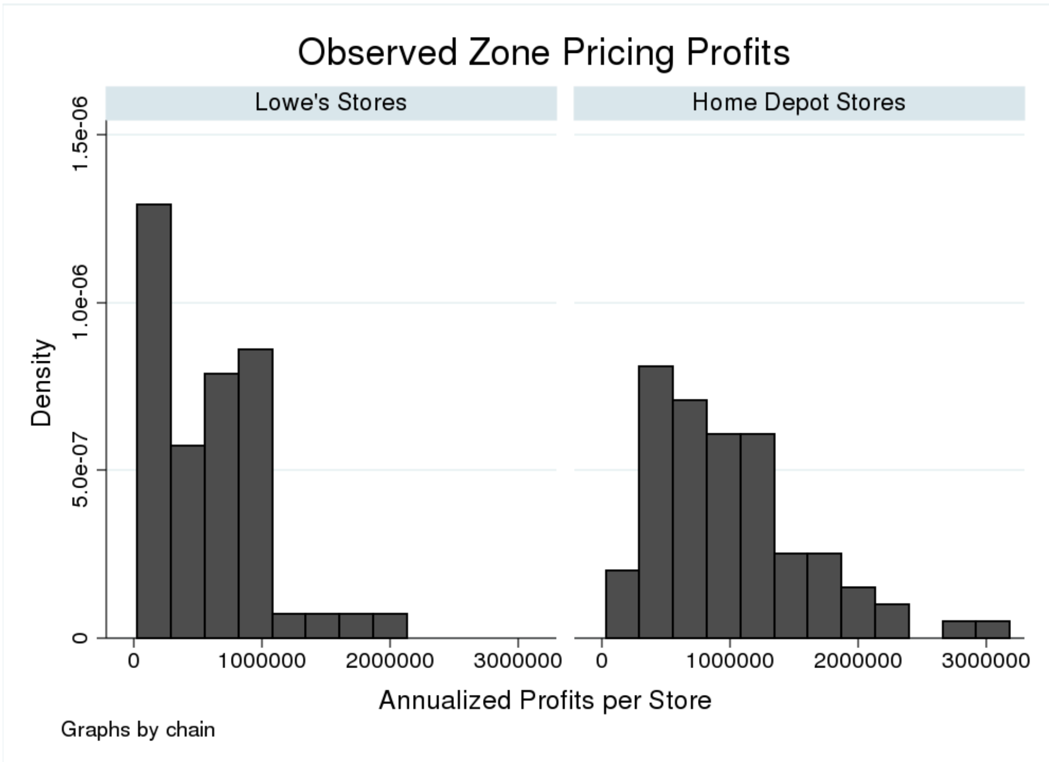


### 2.5.1 Unilateral Single Product Deviations

We first let one chain switch to market level pricing for one product while holding fixed prices on all of its other products and prices on all of its competitor's products. Market by market we calculate new prices that maximize total firm profits. For this exercise, profit gains must be nonnegative as the zone price is in the choice set and all other prices are unchanged. In some sense, the profit gains here represent a lower bound on menu costs. If post-menu cost profits were higher with market level pricing given competitor prices, then the observed zone pricing would not be a best response to those competitor prices and hence not an equilibrium strategy. This also parallels the approach of the prior literature, which estimate profit gains from store level pricing implicitly keeping competitor prices fixed by not modeling competition.

A zone price must balance the profit gains available by raising the price at stores facing high residual demand or high marginal cost with the losses that would follow from overcharging at stores in the zone where marginal costs are lowest or residual demand is most elastic. The optimal market-specific prices are spread both above and below the zone price. In our sample, when one firm can adjust one product-market price, we find prices in 62% of product-store observations increase and decrease for 38% of product-store observations. Together the sales weighted average price increases by $0.12. As expected, prices in monopoly markets increase. Figure 2.9 provides a density plot of price deviations by market competition for Home Depot stores. As expected, the average price in monopoly markets increases $1.63. For some products in monopoly markets, prices increase by over $3 or by nearly 25%. In contested duopoly markets, the average price decreases by $0.08. While the average deviation in duopoly markets is negative, Figure 2.9 shows that for some products Home Depot increases prices in duopoly markets by over $2. Some of these price increases occur because Lowe's usually offers the same product but does not stock it in that particular market, giving the Home Depot store some market power for that product.

For monopoly stores in their own pricing zone, the zone pricing problem and this exercise coincide. However, since costs are not perfectly measured, observed zone prices are not necessarily equal to optimal zone prices. Hence, when solving for optimal prices for products in monopoly markets and own zones, we find small price adjustments to solve the first-order conditions on the firm's problem. Therefore, we report nonzero profit gains for these products. These errors are lumped together with the menu cost. An alternative exercise would begin with a simulated equilibrium instead of observed prices. There are multiple pricing equilibria for the current zone structure, and while some are very close to observed

**Figure 2.9:** Density of Home Depot price changes when a chain adopts market level pricing holding competitor prices fixed. The plot shows that in monopoly markets, prices only increase whereas in duopoly markets, prices both increase and decrease.

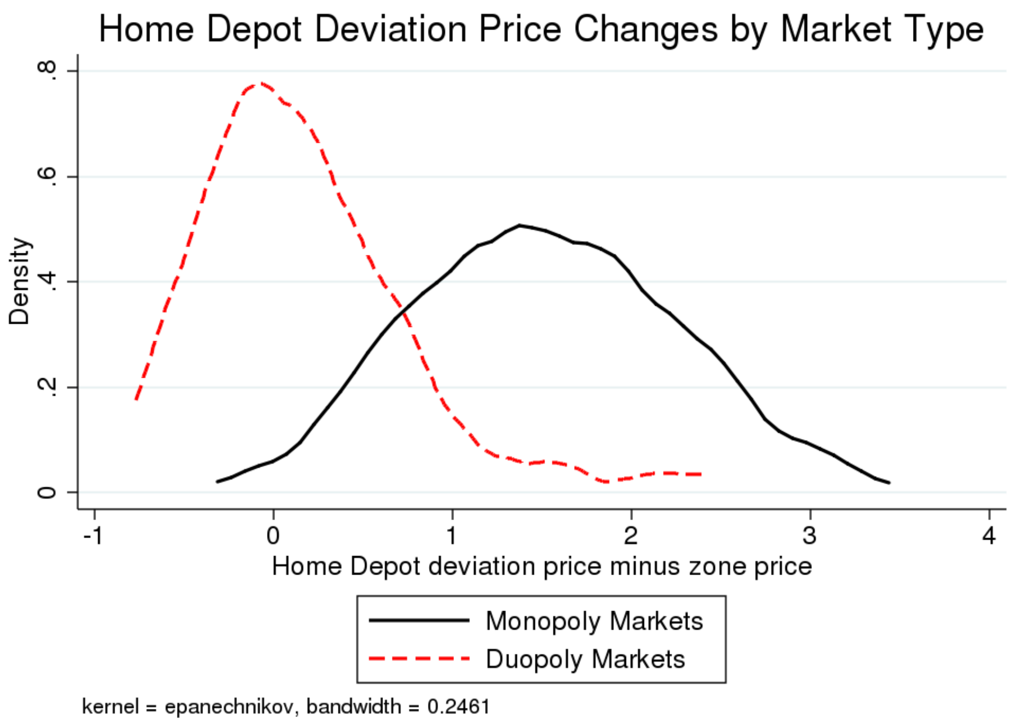


prices and quantities, we would need to dictate an equilibrium selection rule. Here we begin with the equilibrium that has been observed and proceed with small changes that would force the selection of a new equilibrium.

When summing the profit gains from all the single product deviations, we find market level pricing to give profit gains of around 3.3% of revenue, or 33% of current profits. This equates to nearly $6.7 million in additional industry profits annually for the 128 stores of study. We calculate the average annual gain to be $10,430 per store for Home Depot and $11,617 per store for Lowe's.

Our findings are consistent with the previous empirical studies of zone pricing and with elementary price discrimination theory. Without competitive effects, offering separate prices in each market will increase profits, while lowering prices for some consumers and raising them for others.

### 2.5.2 Single Product Market Equilibria

We next allow competitors to respond with market level pricing of their own, again in one product. For that product, we find an equilibria in which both chains set a market-specific price that are best responses to the others price. The equilibrium prices we find have the same general features as in the first exercise. In duopoly markets prices decrease \$0.09 from the observed zone price. The price increase in monopoly markets are the same, since there the equilibrium price and the unilateral deviation price are identical. In 55% of product-store observations prices increase.

Prices that were lowered in the unilateral deviation are decreased further in equilibrium in 74.4% of observations in duopoly markets. In the Salt Lake City market, for example, 4'x8'x5/8" regular drywall is sold for \$10.98 when firms use zone prices. Both Lowe's and Home Depot have monopoly stores and distant, higher cost stores in the same price zone, and their zone prices reflects this. If all other prices were fixed, Home Depot would set a market level price of \$10.70. Lowe's would charge \$10.74. In equilibrium, Home Depot has to compete not against a fixed \$10.98 Lowe's prices, but against a more competitive Lowe's price. Likewise, Lowe's must respond to a lower Home Depot price. Therefore, Home Depot's equilibrium price is reduced to \$10.63.

In a few markets, one firm's zone price falls below their optimal price, while the other firm's is above their optimal price. This can come from cost differences stemming from the dissimilar placement of distribution centers or from one firm's zone containing more monopoly stores that push up the optimal zone price. In these markets, equilibrium considerations moderate the price movements seen in the unilateral deviation experiment. In some high cost markets, a firm in the unilateral deviation experiment will double its price, forfeiting the mainstream market to capture consumers with strong product-specific preferences. If the rival chain raises its price, its equilibrium price will be lower than the unilateral deviation, as returning to something closer to normal competition becomes more appealing. Under certain conditions, to ignore competition would cause overstatement of the effect of moving to market level pricing, but in our region understatement was more common.

Because competitive effects are absent in monopoly markets, the effects that competition magnifies are mostly price decreases in duopoly markets. Because industry demand is so inelastic, these price decreases reduce profits. In total, equilibrium profit gains are 32.9% lower than in the move to market level pricing holding competitor prices fixed. The sum over all products of moving from a zone pricing equilibrium to a market pricing equilibrium are 2.2% of current revenue. The profit gains are almost evenly split between Home Depot and Lowe's in total; per store Lowe's gains more: $8,879 to Home Depot's $6,665.

In the price discrimination theory literature, the profit effects of increased ability to price discriminate depend on the industry and cross-price elasticities of demand. Drywall industry demand is nearly inelastic and many products are close substitutes, bringing it close to the conditions in Holmes (1989) where discrimination yields lower industry profits than uniform pricing. Yet, in our counterfactual experiment, industry profits increase. Some of this is driven by the presence of monopoly markets, but even in duopoly markets profits increase on net. Furthermore, in every zone industry profits increase under market level pricing, although Lowe's profits decrease in a few zones.

If firms can credibly commit to their pricing zones, the (pre-menu cost) profit gains calculated here are lower bounds on the relevant menu costs. If the menu costs for market level pricing are more than a chain's profit gains, it would prefer to stay in the zone pricing equilibrium rather than invest in market level pricing. The magnitude of bounds we calculate here could be consistent with a managerial source to the spatial menu costs. Retail segments with more elastic industry demand or with less substitubility between competitors should have the higher profit gains. If those profit gain are enough to overcome spatial menu costs, retailers will use market level pricing or smaller pricing zones.

### 2.5.3 Metagame Analysis

#### Pricing and Profits

In this section we explore the metagame in which firms choose to adopt zone pricing or market level pricing. There are millions of zone combinations so we only explore the option of selecting the current zone structure for each firm, or a move to market level pricing for the entire network. We calculate the Bertrand-Nash Equilibrium (BNE) for four pricing regimes: Lowe's and Home Depot keep their current zone structure, Lowe's moves to market level pricing and Home Depot keeps its current zone regime, Home Depot moves to market level pricing and Lowe's keeps its current zone structure, and finally, both firms move to market level pricing. Prices adjust for all products in all periods.

As previously noted, there are multiple equilibria for each of these pricing regimes. We utilize a selection mechanism on the number of products with low sales. Due to the logit error term, high prices yield marginal sales, but the firm may choose to set very high prices so that consumers substitute to other products with better margins. Of the equilibria found, we select the equilibrium for each scenario that has the lowest number of products priced sufficiently high as to yield marginal sales. For example, in solving for equilibria based on the current zone structure, there are equilibria in which the price of a product is such that sales are close to $10^{-6}$ sheets per week. We sum up the number of product-store combinations in which this occurs and select the equilibrium with the lowest number. We gauge the performance of this selection mechanism by comparing the observed zone equilibrium with the calculated zone equilibria. The lowest number of observations with marginal sales is 306. The median difference between equilibrium and observed zone prices is \$0.007 and the mean difference is \$0.11. Lowe's observed annual profit is \$8,330,954 for the 53 stores in the sample. With our selection mechanism, we calculate equilibrium zone profits for Lowe's at \$8,301,573. For Home Depot, we obtain observed and equilibrium zone price profits of \$20,363,560 and \$20,736,293, respectively. Other selection mechanisms, such as the sum of total profits, yields unrealistic equilibrium profits given observed sales.

If both firms choose zone pricing in all markets, in all periods, we must solve for nearly 7,000 prices. The other three possible outcomes have even more prices to solve for. The optimality conditions on firms' problems are highly nonlinear, so we solve for equilibria using state of the art solvers. To search for equilibria, we set 1,000 random starts and solve for the fixed point. On average, around one-third of the starts converge to a fixed point – a BNE for the regime choice.

Since both firms utilize zone pricing, we solve for two parameters of the meta game: a lower point on the menu costs associated with adopted market level pricing. Table 2.9 provides the payoff matrix associated with the metagame. In the absence of menu costs, both Lowe's and Home Depot would adopt market level pricing; however, in this case, Lowe's obtains lower profits with market level pricing than with zone pricing. This is due to both higher demand for Home Depot products in general, as well as a cost advantage in the stores located around Salt Lake City, where Home Depot has a distribution center, but Lowe's does not. This allows Home Depot to undercut Lowe's and gain market share. This analysis shows that a finer degree of pricing – in this case the ability to discriminate at the market level – does not lead to larger profits, a possibility noted in Holmes

(1989). Indeed, additional competition hurts Lowe's, but provides a nearly 14% increase in profits for Home Depot.

**Table 2.9:** Metagame of market or current zone pricing

| Lowe's / HD | Zone Pricing | Market Pricing |
|---|---|---|
| Zone Pricing | \$8,301,573, \$20,736,293 | \$7,681,509 , \$21,667,885$-\mu^{HD}$ |
| Market Pricing | \$8,705,745$-\mu^{L}$, \$20,937,713 | \$7,818,741$-\mu_{L}$, \$23,632,457$-\mu^{HD}$ |

The zone numbers are equilibrium zone profits instead of observed profits.

Also in the absence of menu costs, we find that moving to market level pricing for a single firm increases profits for that firm. Lowe's sees a 4.8% increase in pre-menu cost profits by moving to market level pricing with Home Depot keeping its zone structure. Home Depot also sees modest gains with this regime at 0.97%. On the other hand, if Home Depot moves to market level pricing but Lowe's keeps its zone structure, Home Depot sees a 4.8% increase in profits, largely due to the ability to discriminate in monopoly markets. However, with this regime, Home Depot's competitive advantages, both in costs and demand, results in a 5.9% decrease in profits for Lowe's.

For zone pricing to be the solution of the metagame in pure strategies, it must be the case that $\mu^{L} \geq \$404,172$ and $\mu^{HD} \geq 2,694,744$. These numbers represent lower bounds on the menu costs associated with adopting market level pricing. These equate to 4.8% and 13.0% of profits for Lowe's and Home Depot, respectively. Together, this yields a menu cost of 10.3% of industry profits, about half the figure calculated using single product deviations. Other market level pricing equilibria exist, some giving much higher profits that match or exceed the menu costs of Section 5. The lower profits in the selected equilibrium could reflect the substitutability between products. Price decreases on one product could prompt the rival firm to discount other products. The increased competition on all products (instead of on only one product) may reduce market share, prompting further rounds of discounting, and lower profits.

Table 2.10 provides summary statistics at the market level across the various pricing regimes of the metagame. The table also provides a summary of the equilibrium in which firms use uniform pricing. We find Lowe's would earn higher profits under uniform pricing than zone pricing (utilizing our selection mechanism). This is consistent with Lowe's earning lower profits under market level pricing than zone pricing; that is, Lowe's benefits when Home Depot has limited ability to price discriminate. Under zone pricing, Home Depot balances the benefits of discriminating in monopoly markets with its desire to undercut Lowe's in duopoly markets. This allows Lowe's to capture market share in duopoly markets that it would not if Home Depot priced at the zone or market level. With uniform pricing, Home Depot obtains approximately \$18.8 million in profits, \$1.8 million less than when both firms use zone pricing and nearly \$5.0 million less than when both firms use market level pricing.

Since Home Depot operates several monopoly stores as part of larger zones, finer pricing results in monopoly prices in these markets, whereas with uniform and zone level pricing, Home Depot balances discriminating in these markets with competing with Lowe's in other markets. The relationship between zone structure and profits is opposite for Lowe's. With a cost disadvantage in the large Salt Lake City market, lower mean utility overall for products, and few monopoly stores, Lowe's does not capture additional profits from finer pricing. Instead, the chain benefits when Home Depot has reduced ability to discriminate in competitive markets.

### Consumer Surplus

Finally, we investigate how consumer surplus changes across pricing regimes. Given the nested logit demand system, the change in consumer surplus across pricing regimes for consumer $i$ at location $\ell$ can be written as

$$\mathbb{E}[\Delta CS_{i,\ell}] := \frac{1}{\alpha}\left[\ln\left(1+\sum_g D_g^{1-\sigma}(Z')\right) - \ln\left(1+\sum_g D_g^{1-\sigma}(Z)\right)\right],$$

where $D_g = \sum_{j\in J_g}\exp(\delta_j/1-\sigma)$, $Z'$ represents prices under counterfactual zones, and $Z$ represents prices under observed zones. Multiplying $\mathbb{E}[\Delta CS_{i,\ell}]$ by $M_\ell$ yields the total change in consumer surplus for a single market. We then aggregate across markets to calculate the total change in consumer surplus.

Table 2.11 quantifies the change in consumer surplus across four counterfactual scenarios. In the aggregate, we find that equilibrium uniform pricing results in the highest consumer surplus, increasing consumer surplus by nearly \$500,000 annually for the 128 stores in the sample. Market-level pricing results in higher prices for consumers and decreases consumer surplus by close to \$1 million. The move for one chain to adopt uniform pricing while the competitor sets optimal zone pricing under observed zones has varying impact. If Lowe's moves to uniform pricing, we find consumer surplus goes down by nearly \$100,000 annually; however, if Home Depot moves to uniform pricing, consumer surplus goes up by nearly \$300,000. These figures reflect both the estimated Home Depot costs advantage and that Home Depot operates many more monopoly stores in the sample.

**Table 2.10:** Chain Performance Across Pricing Regimes, summarized across markets

| | | (1) Zone BNE | (2) Market BNE | (3) Uniform BNE | (4) HD Dev. BNE | (5) Lowe's Dev. BNE |
|---|---|---|---|---|---|---|
| **Lowe's** | | | | | | |
| | mean | 276,719 | 260,625 | 290,048 | 256,050 | 290,192 |
| | median | 135,203 | 120,247 | 132,386 | 120,634 | 124,887 |
| | lower quartile | 51,632 | 57,659 | 61,037 | 63,012 | 64,053 |
| | upper quartile | 333,312 | 283,439 | 349,276 | 277,780 | 270,362 |
| | min | 39,470 | 26,298 | 50,449 | 42,432 | 25,423 |
| | max | 1,650,275 | 1,496,118 | 1,777,272 | 1,500,762 | 1,714,356 |
| | total profits | 8,301,573 | 7,818,741 | 8,701,427 | 7,681,509 | 8,705,745 |
| **HD** | | | | | | |
| | mean | 441,198 | 508,721 | 401,253 | 461,019 | 445,483 |
| | median | 243,045 | 274,062 | 176,794 | 262,462 | 241,781 |
| | lower quartile | 142,205 | 196,623 | 125,986 | 191,048 | 149,604 |
| | upper quartile | 435,670 | 496,624 | 382,818 | 418,650 | 438,875 |
| | min | 50,292 | 23,443 | 55,720 | 49,148 | 71,889 |
| | max | 2,979,718 | 2,906,893 | 2,868,559 | 2,842,305 | 2,916,245 |
| | total profits | 20,736,293 | 23,909,881 | 18,858,886 | 21,667,885 | 20,937,713 |

Results are annualized and aggregated to the market level.

**Table 2.11:** Consumer Surplus Across Pricing Regimes

| Regime Change | Agg. $\mathbb{E}[\Delta CS]$ | $\min_\ell (\mathbb{E}[\Delta CS_\ell])$ | $\max_\ell (\mathbb{E}[\Delta CS_\ell])$ |
|---|---|---|---|
| Zone → Uniform | \$473,762 | -\$135,721 | \$181,456 |
| Zone → Market | -\$965,073 | -\$196,841 | \$12,354 |
| Zone → Lowe's Unif. | -\$94,488 | -\$95,082 | \$39,646 |
| Zone → HD Uniform | \$285,748 | -\$46,164 | \$69,625 |

The direction of consumer surplus changes under uniform and market-level pricing are opposite of what they are for firms; however, the magnitudes differ substantially. After accounting for the change in profits under different regimes and zero menu costs, we find uniform pricing lowers total surplus by \$1,003,791 annually, whereas market level pricing increases total welfare by \$1,725,683 annually.

## 2.6 Conclusion

In this article, we document the prevalence of zone pricing in home improvement retail stores. Although product categories such as drywall and lumber have sizable price variation nationally, within regions there can be no price variation within firm. The size of zones, including the number of stores and the number of markets per zone, varies from product category to product category. We find some product categories with hundreds of zones, and for other categories, a single firm, or uniform pricing, is pursued by a firm. Having different zone structures by product category is not surprising given that retailers have separate marketing managers for different product categories. The choice of the zone structure reveals how firms balance discrimination and competition across markets. We postulate that the use of zone pricing, instead of a finer grade of pricing, such as by market or by store, is the result of firms facing a friction – “spatial menu costs”. These spatial menu costs have induced firms to set a constant price over multiple markets.

To provide a measure of the spatial menu costs needed to rationalize the use of zone pricing, we estimate a structural model of consumer demand on a detailed data set of retail drywall. We find that consumers consider the products of competing chains to be close substitutes, but the industry elasticity for drywall is inelasitic. Assuming firms are engaged in Bertrand price competition, we back out marginal costs to find that transportation costs are a small, but significant, component of costs.

Given our estimates on supply and demand, our menu costs are calculated by comparing the observed profits in zone level competition to the equilibrium profits in which firms adopt market level pricing. Since firms are offering multiple products, priced uniformly across several markets, multiple equilibria exist. To obtain menu cost estimates, we investigate small deviations in the current zone structure which result in unique counterfactual equilibria. Finding the equilibria for these alternative pricing regimes yields a lower bound on the spatial menu costs at 2.2% of current revenues or 22% of observed profits. While the previous literature on zone pricing has determined menu costs may be large, these articles have not taken into account the competitive interaction of firms. We find that ignoring competitive effects by fixing opponents prices implies much larger gains from market level pricing which overstates the spatial menu costs by upwards of 32%, at 3.3% of observed revenues.

Spatial menu costs force firms into a price zone system that prevents them from abusing their market power. In an industry like drywall with high transportation costs and inelastic demand, menu costs and zone pricing protect consumers in monopoly markets. The elimination of the menu costs would prompt a new level of strategic competition in the duopoly markets, but according to our estimates would still leave the retailer more profitable.

# References

AGUIRRE, I., S. COWAN, AND J. VICKERS (2010): "Monopoly Price Discrimination and Demand Curvature," *American Economic Review*, 100(4), 1601–15.

ALDERIGHI, M., M. NICOLINI, AND C. PIGA (2012): "Combined Effects of Load Factors and Booking Time on Fares: Insights from the Yield Management of a Low-Cost Airline," *working paper*.

BAJARI, P., C. L. BENKARD, AND J. LEVIN (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75(5), 1331–1370.

BERRY, S., M. CARNALL, AND P. T. SPILLER (1996): "Airline Hubs: Costs, Markups and the Implications of Customer Heterogeneity," Discussion paper, National Bureau of Economic Research.

BERRY, S., AND P. JIA (2010): "Tracing the Woes: An Empirical Analysis of the Airline Industry," *American Economic Journal: Microeconomics*, 2(3), 1–43.

BERRY, S. T. (1994): "Estimating Discrete-Choice Models of Product Differentiation.," *RAND Journal of Economics*, 25(2), 242 – 262.

BOARD, S. (2008): "Durable-Goods Monopoly with Varying Demand," *The Review of Economic Studies*, 75(2), 391–413.

CAPLIN, A., AND B. NALEBUFF (1991): "Aggregation and Imperfect Competition: On the Existence of Equilibrium," *Econometrica*, 59(1), pp. 25–59.

CHEN, Y., AND M. SCHWARTZ (2013): "Differential Pricing When Costs Differ: A Welfare Analysis," Discussion paper, Georgetown University, Department of Economics.

CHINTAGUNTA, P., J.-P. DUBÉ, AND V. SINGH (2003): "Balancing Profitability and Customer Welfare in a Supermarket Chain," *Quantitative Marketing and Economics*, 1(1), 111–147.

Clark, R., and N. Vincent (2012): "Capacity-Contingent Pricing and Competition in the Airline Industry," *Journal of Air Transport Management*, 24, 7–11.

Conlisk, J., E. Gerstner, and J. Sobel (1984): "Cyclic Pricing by a Durable Goods Monopolist," *The Quarterly Journal of Economics*, 99(3), 489–505.

Corts, K. S. (1998): "Third-Degree Price Discrimination in Oligopoly: all-out competition and strategic commitment.," *RAND Journal of Economics*, 29(2), 306 – 323.

Courty, P., and M. Pagliero (2012): "The Impact of Price Discrimination on Revenue: Evidence from the Concert Industry," *Review of Economics and Statistics*, 94(1), 359–369.

Cowan, S. (2012): "Third-Degree Price Discrimination and Consumer Surplus," *Journal of Industrial Economics*, 60(2), 333–345.

Dana, J. D. (1999): "Using Yield Management to Shift Demand When the Peak Time is Unknown," *The Rand Journal of Economics*, pp. 456–474.

Elmaghraby, W., and P. Keskinocak (2003): "Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices, and Future Directions," *Management Science*, 49(10), 1287–1309.

Escobari, D. (2012): "Dynamic Pricing, Advance Sales, and Aggregate Demand Learning in Airlines," *Journal of Industrial Economics*, 60(4), 697–724.

Gale, I. L., and T. J. Holmes (1993): "Advance-Purchase Discounts and Monopoly Allocation of Capacity," *The American Economic Review*, pp. 135–146.

Gallego, G., and G. Van Ryzin (1994): "Optimal Dynamic Pricing of Inventories with Stochastic Demand Over Finite Horizons," *Management science*, 40(8), 999–1020.

Gandhi, A., Z. Lu, and X. Shi (2013): "Estimating Demand for Differentiated Products with Rrror in Market Shares," Discussion paper, University of Wisconsin - Madison.

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.

Goolsbee, A., and C. Syverson (2008): "How do Incumbents Respond to the Threat of Entry? Evidence from the Major Airlines," *The Quarterly Journal of Economics*, 123(4), 1611–1633.

Holmes, T. J. (1989): "The Effects of Third-Degree Price Discrimination in Oligopoly.," *American Economic Review*, 79(1), 244.

Hotz, V. J., and R. A. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *The Review of Economic Studies*, 60(3), 497–529.

IATA (2013): "Profitability and the Air Transport Profitability and the Air Transport Value Chain," Discussion paper, International Air Transport Association.

Jones, C. (2012): "How theater ticket prices are changing like airline fares," *Chicago Tribune*.

Khan, R. J., and D. C. Jain (2005): "An Empirical Analysis of Price Discrimination Mechanisms and Retailer Profitability," *Journal of Marketing Research*, 42(4), pp. 516–524.

Lambrecht, A., K. Seim, N. Vilcassim, A. Cheema, Y. Chen, G. S. Crawford, K. Hosanagar, R. Iyengar, O. Koenigsberg, R. Lee, et al. (2012): "Price Discrimination in Service Industries," *Marketing Letters*, 23(2), 423–438.

Lazarev, J. (2013): "The Welfare Effects of Intertemporal Price Discrimination: An Empirical Analysis of Airline Pricing in U.S. Monopoly Markets," *working paper*.

Leslie, P. (2004): "Price Discrimination in Broadway Theater," *RAND Journal of Economics*, pp. 520–541.

Levy, D., M. Bergen, S. Dutta, and R. Venable (1997): "The Magnitude of Menu Costs: Direct Evidence From Large U. S. Supermarket Chains," *The Quarterly Journal of Economics*, 112(3), pp. 791–825.

Li, J., N. Granados, and S. Netessine (2013): "Are Consumers Strategic? Structural Estimation from the Air-Travel Industry," *INSEAD Working Paper*.

McAfee, P., and V. Te Velde (2006): "Dynamic Pricing in the Airline Industry," *Handbook on Economics and Information Systems*.

Miller, N. J., and M. Osborune (2011): "Competition Among Spatially Differentiated Firms: An Estimator with an Application to Cement," *Working Paper*.

Montgomery, A. L. (1997): "Creating Micro-Marketing Pricing Strategies Using Supermarket Scanner Data.," *Marketing Science*, 16(4), 315.

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

Nevo, A. (2001): "Measuring Market Power in the Ready-To-Eat Cereal Industry," *Econometrica*, 69(2), 307–342.

Pashigian, B. P., and B. Bowen (1991): "Why are Products Sold on Sale?: Explanations of Pricing Regularities," *The Quarterly Journal of Economics*, 106(4), 1015–1038.

Petrin, A. (2002): "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110(4).

Puller, S. L., A. Sengupta, and S. N. Wiggins (2012): "Does Scarcity Drive Intra-Route Price Dispersion in Airlines?," NBER Working Paper 15555, National Bureau of Economic Research, Inc.

Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, pp. 999–1033.

Soysal, G. P., and L. Krishnamurthi (2012): "Demand Dynamics in the Seasonal Goods Industry: An Empirical Analysis," *Marketing Science*, 31(2), 293–316.

Stokey, N. L. (1979): "Intertemporal Price Discrimination," *The Quarterly Journal of Economics*, pp. 355–371.

Su, C.-L., and K. L. Judd (2012): "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*, 80(5), 2213–2230.

Su, X. (2007): "Intertemporal Pricing with Strategic Customer Behavior," *Management Science*, 53(5), 726–741.

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.

Talluri, K., and G. Van Ryzin (2004): "Revenue Management Under a General Discrete Choice Model of Consumer Behavior," *Management Science*, 50(1), 15–33.

Talluri, K. T., and G. J. Van Ryzin (2005): *The Theory and Practice of Revenue Management*, vol. 68. Springer.

Vulcano, G., G. van Ryzin, and W. Chaar (2010): "Choice-Based Revenue Management: An Empirical Study of Estimation and Optimization," *Manufacturing & Service Operations Management*, 12(3), 371–392.

Zbaracki, M. J., M. Ritson, D. Levy, S. Dutta, and M. Bergen (2004): "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets," *The Review of Economics and Statistics*, 86(2), 514–533.

Zhao, W., and Y.-S. Zheng (2000): "Optimal Dynamic Pricing for Perishable Assets with Nonhomogeneous Demand," *Management Science*, 46(3), 375–388.

# Appendix A

# Appendix to Chapter 1

## A.1 Accuracy of Seat Maps

I perform two analyses to address the potential measurement error with using seat maps to proxy bookings. First, I match statistics of my data with reported carrier data found in the T100 US Segment tables. The T100 Segment tables report actual monthly enplanements for an origin, destination, air carrier. I aggregate the number of occupied seats the day flights leave to the monthly level for each route-month and compare with the T100. A scatter plot comparing monthly enplanements can be seen in Figure A.1 Most points lie very close to the 45-degree line (zero measurement error at the monthly level) indicating a close match between actual monthly enplanements and the totals calculated from seat maps. The points highlighted in blue correspond to the city pair Kansas City, Boston, operated by Delta Air Lines. Delta switched regional carriers in this market during August 2012. My programming scripts failed to pick up the changes in flight numbers, which resulted in more than half of the month's flights from being tracked. However, the other means are still close to the 45-degree line. By comparing these aggregate measures, I find my seat maps understate true enplanements by 0.81% of load factor at the monthly level.

**Figure A.1:** Comparison of monthly enplanements of T100 US Segment and sample seat maps

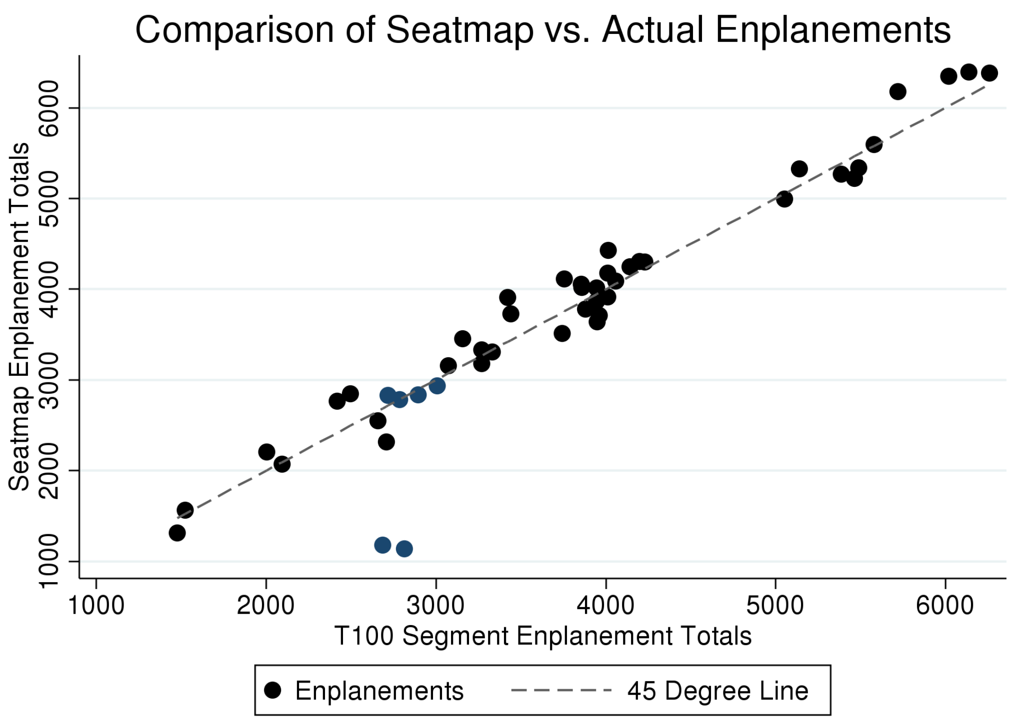


Notes: Each points corresponds to an airline, month, origin, destination traffic count. Points on the dotted line have zero monthly measurement error.

However, comparing enplanements with the T100 does not address the issue of consumers purchasing tickets but not selecting seats at the time of booking. To this end, I create an additional data set from an airline that provides flight loads and seat maps on its website. I collect 18,107 seat maps and reported flight load by randomly selecting routes and days until departure over a five month period. Unfortunately, the airline does does not operate in the markets I study, so this constitutes an out of sample analysis. Further, this data is only for business class

**Figure A.2:** Estimated measurement error by day before departure

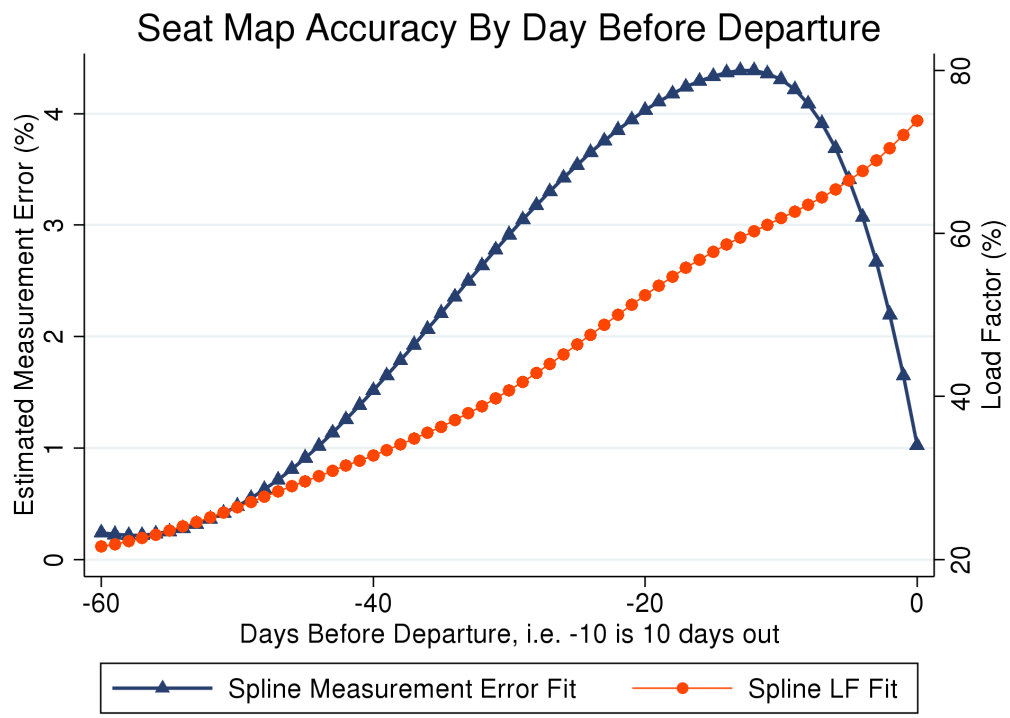


Notes: Spline fit comparing seat map total with reported flight load by day before departure. Also, load factor using reported flight load by day before departure.

cabins with a range of 25 to over 50 seats. Of course, a one seat difference is a much higher percentage of load factor compared to the capacities in my data. This may exaggerate the measurement error percentage. At the same time time, it may be that more economy class consumers do not select tickets at the time of booking. With this data set, I find that the seat maps understate actual loads by 2.3% on average. Few seat maps overstate the number of occupied seats, so the direction is the bias is mostly downward.

I use this secondary data set to estimate the measurement error by day before departure. Figure A.2 plots a nonparametric fit of the the measurement error along with average load factor by day before departure. This plot shows that the measurement error remains small across time, at between 0.0% and 4.3% of load factor. At the same time, the measurement error is not flat indicating some days are more accurate than others. The worst measurement error between two and three weeks prior to departure. Seat maps are most accurate either well in advance or close to the departure date.

# Appendix B

# Appendix to Chapter 2

## B.1 Monte Carlo Experiments

In this section we conduct Monte Carlo experiments of our cost estimator. We show the key parameter in the cost functions of firms is estimated with little bias and a low mean-squared error.

In the Monte Carlos, the true model specifies a profit maximizing monopolist using zone pricing and facing costs of the form

$$c_{jszt} = a_j + \kappa d_{sz} + \nu_{jzt}.$$

With this notation, store $s$ belongs to a zone $z$. Hence, in a model specifying $Z = 3$ and $S = 2$, this implies the firm has 3 zones, each with 2 stores. In all experiments, we use a discrete choice logit demand system where the market is set equal to the zone; that is, consumers select amongst all $J \times S$ products in a given purchasing period. The pricing decision can still be thought of as zone pricing because the firm operates several stores within the market, and these stores have different costs; however, the price for a product must be identical across stores. The Monte Carlos specify a monopolist only to ease the computational burden of solving for profit-maximizing prices. The procedure extends to models with competition by changing how optimal prices are found. The cost estimation procedure remains unchanged.

For the simulations, we specify the utility of the consumers as equal to $\mathcal{U} = \beta - \alpha p_{jzt} + \xi_{jszt}$. We set $\beta = 3, \alpha = -0.3$ and $\xi \sim \mathcal{N}(0,3)$. For costs, we assign

| | |
|---|---|
| Distance Coefficient | $\kappa = -0.01$ |
| Distance Data | $d_{sz} \sim \lvert N(300,40) \rvert$ |
| Product FE | $a_j \sim \lvert \mathcal{N}(8,3) \rvert$ |
| Cost Unobserved | $\nu \sim \mathcal{N}(0,9)$ |

With the data generating process defined, our Monte Carlo procedure first finds optimal prices $p^*$ by maximizing the profit function of the firm:

$$p^* \in \text{argmax}_p \sum_{j,s,z,t} (p_{jzt} - c_{jszt})\, M_{zt} s_{jszt}(p).$$

We use a state of the art solver to calculate profit-maximizing prices given costs.

The next step in the Monte Carlo procedure is to then back out costs given optimal prices and distance data. Our key parameter of interest is $\kappa$ – the parameter on distance; however, the estimates of the product-specific cost intercepts ($a_j$) could also be of interest. We estimate costs using the cost estimation approach in Section 3. As we discuss there, the traditional approach in the industrial organization literature is to back out costs using either Lerner's Index or the multiproduct analog. In a second step, the cost parameters are estimated using GMM. Our novel approach is to estimate the costs in a single step using mathematical programming with equilibrium constraints (MPEC). We proceed with this method because the first order conditions (FOCs) for the firm's problem contain all costs for a product within a zone. This means we cannot use Lerner's index to recover costs directly. One resitrction of our approach is that we assume that the unobserved cost term does not vary by store. If we allowed this flexibility, our approach would lead to an underindentification problem as discussed in Section 3.

With our proposed estimator, costs and the parameters associated with costs are recovered by solving

$$\begin{aligned} &\min_{\mathbf{c},\theta^S} \; m(\theta^S; \mathbf{W}) \\ \text{s.t.} \quad &\text{FOC}(\mathbf{c}; \theta^S) = 0, \end{aligned}$$

where $\mathbf{W} := [\mathbf{d}, \mathbf{a}_j]$ is data that enters the cost function, and $m(\cdot)$ is a loss function, such as a GMM criterion function. In the formulation above, identification comes from the $\dim(\theta^S)$ moment conditions in the criterion function and that the constraints are such that $\dim(\text{FOC}) = \dim(\nu)$.

Simulation results appear in Table B.1 where we report the bias of the estimator along with the mean squared error for the distance parameter for different values of $S, J$ and $T$. For all experiments, we assign $Z = 10$. The left set of experiments shows changes in the number of stores operated by the firm, and on the right, we report experiments with changes in the number of products offered. The table shows that the estimator performs well. For all experiments, the parameter of interest (distance) is estimated with a small bias and small mean-squared error. Our cost estimation procedure estimates the cost parameters precisely for all experiments, ranging from 4,500 cost variables to recover, to 54,000 cost variables.

**Table B.1:** Monte Carlo Experiments of Cost Estimator

| T | S | | | MPEC using FOCs | | T | J | | | MPEC using FOCs | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | Bias | MSE | | | | | Bias | MSE |
| 50 | 3 | | | 0.001 | 8.102E-6 | 50 | 3 | | | 0.001 | 1.016E-5 |
| | | 6 | | 0.001 | 8.593E-6 | | | 6 | | 0.001 | 7.261E-6 |
| | | | 9 | 0.002 | 1.784E-5 | | | | 9 | 0.001 | 5.813E-6 |
| 100 | 3 | | | 0.002 | 1.634E-5 | 100 | 3 | | | 0.001 | 9.882E-6 |
| | | 6 | | 0.001 | 1.051E-5 | | | 6 | | 0.001 | 5.921E-6 |
| | | | 9 | 0.002 | 1.019E-5 | | | | 9 | 0.001 | 3.828E-6 |

Notes: The data generating process specifies a profit maximizing monopolist, with a distance parameter of $\kappa = .01$. In the first set of experiments (left), we set $J = 6$. For the second set (right), we set $S = 3$. To estimate the cost parameters, we first calculate optimal prices given true costs, and then using optimal prices, back out costs. We run 200 iterations per experiment, recovering costs using the interior point - conjugate gradient algorithm in Knitro 9.0. Computations are done on on a 4-core Linux machine with a 4.1 Ghz OC Intel CPU.