Essays in Industrial Organization

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

Kailin R. Clarke

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

Thomas J. Holmes and Amil Petrin, Advisors

June, 2016

Acknowledgements

I am very grateful to my advisors, Thomas J. Holmes and Amil Petrin, and my committee members, Joel Waldfogel and Naoki Aizawa, for their continued advice and support. Their guidance has provided me with an invaluable blueprint for choosing and approaching research questions.

Additionally, my research has benefitted from conversations with Dirk van Duym, Joaquin Garcia-Cabo Herrero, Ellen McGrattan, Keaton Miller, Imke Reimers, Boyoung Seo, Matt Shapiro, Ethan Singer, Dominic Smith, Thomas Quan, Maria Ana Vitorino, Kevin Williams, Thomas Youle, members of the University of Minnesota Applied Microeconomics workshop, and participants of conferences and seminars at which I have presented while attending the University of Minnesota. Special thanks to Keaton Miller who co-authored the final chapter of my dissertation.

I acknowledge the financial support of the Graduate School of the University of Minnesota. My work was supported in part by the University of Minnesota Doctoral Dissertation Fellowship and by the University of Minnesota Graduate Research Program Partnership Fellowship. I also acknowledge support from the Minnesota Supercomputing Institute for providing computational resources. Finally, I would like to thank the Department of Economics staff, especially Caty Bach, Sara Shuford, and Wendy Williamson.

Dedication

To my family and friends. To Meghan for her unwavering patience and support.

Abstract

This dissertation is comprised of three essays. The first two study the welfare effects of dynamic ticket pricing, meaning when a seller adjusts the price of a given seat at a given event as that event approaches. The third studies the effectiveness of a particular merger policy.

In the first chapter, I analyze two motivations for dynamic ticket pricing: (1) price discrimination across consumer arrival time and (2) re-optimizing due to changes in the perceived product, in this case due to team performance. I estimate a flexible model of Major League Baseball ticket demand that takes both forces into account, then solve for optimal pricing over simulated team performance paths. I use an original data set of daily sales, prices, and product characteristics for over 400 games. I find that, on average, product changes affect price more strongly until the final week before a game, when the shift on consumer types plays a larger role. Prices therefore tend to oscillate and then increase, consistent with observed behavior. Resellers play a key role, dropping prices and dampening the franchise's final-week price increase. I find that dynamic pricing leads to substantial revenue gains, compared to a pricing policy which depends on date of purchase but not dynamic product characteristics. In aggregate consumers lose, with those low-willingness-to-pay consumers who happen to face higher prices being hit the hardest.

While the first chapter centers on an agent with high market power, the second chapter analyzes the dynamic pricing problem faced by agents with comparatively little market power: ticket resellers. Not only do they face steep demand curves but, unlike franchises, they have extremely limited inventory: most listings offer only two tickets. More than half of two-ticket resellers practice uniform pricing, suggesting a price adjustment cost. I assume that each of these resellers posted her price expecting never to change it, and by estimating demand and assuming optimal uniform pricing I recover each reseller's "scrap value," the value of still having tickets after markets have closed on the day of the game. The scrap value turns out to be negative for 41% of listings, suggesting either risk aversion or bounded rationality. Assuming that these scrap values

are actually zero, I then simulate each reseller's expected revenue under a counterfactual where that reseller uses full dynamic pricing. The increase in expected revenue is greatest for those with zero scrap value, followed far behind by those with scrap value above zero and less than face value, followed by those with scrap value above face value.

In the third chapter, Keaton Miller and I examine the effectiveness of U.S. pre-merger notification policy by studying the acquisition behaviors of cable telecommunication companies. We construct a novel dataset of acquisitions in the cable industry from 2000-2012, making use of the large sample and the ability to reasonably define the sets of actual potential mergers. We find that the Hart-Scott-Rodino disclosure threshold only affects firm behavior when acquiring firms whose geographic coverage overlaps with their own. In other words, The disclosure threshold appears to be successful in discouraging or preventing anticompetitive acquisitions.

Contents

List of Tables

List of Figures

Chapter 1

Dynamic Ticket Pricing with Changes in Consumers and Products

1.1 Introduction

Ticket-selling businesses often adjust prices in the days and weeks leading up to a given event. This practice, alternatively known as dynamic pricing (DP), yield management, or revenue management, was adopted by the airline industry decades ago but has more recently gained popularity in the arts and entertainment industry. Sport franchises, Broadway theaters, concert promoters, and even amusement parks are increasingly likely to adjust ticket prices on a regular basis, thanks to technological advancements in tracking inventory and forecasting demand.[1]

The growing empirical literature on dynamic ticket pricing has framed it mainly as intertemporal price discrimination. This framing is well-justified: in many industries, the consumers who shop for tickets earlier, relative to a given event, are likely to have different willingness to pay than those who shop later. For example, it is well known

[1]More broadly, any seller with market power over a perishable good can potentially use dynamic pricing to increase its revenue, but much of the empirical literature has focused on ticket prices, where market power tends to be especially large (relative to a grocery store selling poultry, for example) so that firms have more to gain by optimizing their prices over time.

that consumers shopping for airline tickets closer to a flight tend to be business travelers with low sensitivity to price, hence rising airfares as a flight date approaches [Williams, 2013, Lazarev, 2013].

In the arts and entertainment context, however, price adjustments are usually portrayed as a response to unforeseen changes in how consumers perceive the product. There are numerous intuitive examples. A promoter of an upcoming outdoor event may drop ticket prices if the weather forecast worsens: the product, or event, has changed in the eyes of consumers. Broadway theaters may increase ticket prices if a show receives rave reviews; again, the product has changed [Healy, 2014]. A flight may experience a "product change" when the location of a large event is announced, though the airline is less likely to know about such changes and immediately respond to the demand shift. The recent adoption of DP by Major League Baseball (MLB) franchises, the subject of this paper, provides a particularly interesting example: franchises often describe the practice as a means to respond to increased or decreased excitement around an upcoming game, which is closely tied to current team performance [Dunne, 2012, Rishe, 2012].

To sum up, there exists a discrepancy between the empirical DP literature's emphasis on intertemporal *rotations* of the demand curve and the emphasis by industry observers and participants, particularly in arts and entertainment, on intertemporal *shifts* of the demand curve across all consumers. Assuming that capacity constraints do not bind, the relative importance of these two factors will have clear implications for consumer welfare. If there are stark differences between those who tend to shop early and later then price discrimination will be paramount: early shoppers will face lower prices, while later shoppers will face higher prices, relative to a regime of fixed pricing. On the other hand, if product changes are the key motivator, then the allocation of "wins" and "losses" to consumers is more random: the winners are those who happen to buy when the event's value is *underestimated*, while the losers are those who buy when the value is *overestimated*, since they face lower and higher prices, respectively, relative to a fixed pricing regime.[2]

This paper uses the empirical setting of MLB ticket pricing to weigh the importance

[2] Indeed, if product changes were the only motivator, and product value followed a random walk, the only definite comparison one could make between early consumers and late consumers is that the late-arriving consumers face less variation in wins and losses.

of these two motivations in terms of the overall revenue gains from DP relative to fixed pricing, and to measure the weflare implications outlined above. To compare the importance of consumer changes and product changes, I simulate fixed and dynamic pricing under alternative scenarios where one of these channels is turned off. I also measure how close a baseball franchise can get to optimal dynamic pricing by choosing prices ahead of time, without being able to react to product changes. As far as I know, the only other work on DP to use data on dynamic product characteristics and to estimate their significance is Sweeting [2012]'s study of MLB ticket resellers. As I discuss in the related literature, that paper studies a different variety of DP that is almost exclusively based on opportunity/scarcity shadow costs, practiced by sellers in a highly competitive market.

This is an attractive industry setting for distinguishing between these different motivations for dynamic pricing. First, price changes cannot be explained by changing costs—the marginal cost of each ticket sold is effectively zero. Second, while the number of seats in a stadium is limited, suggesting the presence of a variable shadow cost of capacity, I show that this constraint is rarely binding in my data. Congestion is not a significant issue, even though *ex ante* it might have been. Third, there are objective measures of team performance that are shown in this paper to shift demand. Industry observers cite the *probability of entering the playoffs* as a key shifter of team popularity [Rishe, 2012], so I use this statistic.[3]

The data I collect includes a large original dataset of daily prices and purchases for 2014 regular season single-game tickets, in every section of six stadiums, across both the primary marketplace—tickets offered directly by the franchises—and the most popular secondary marketplace, StubHub. The six franchises practice DP to widely varying degrees.[4] The data also includes static event characteristics, such as whether the game takes place at night, and daily estimates of the probability that a team enters the playoffs.

[3]The probability is calculated by a reputable third-party website, FanGraphs.com, that I discuss in the data section. I do not currently use data on weather forecast and expected starting pitcher, but I have collected this data and plan to incorporate it in a future version of the paper. The challenge of incorporating these observables is that consumers are unlikely to care about them until the final 7-10 days before a game.

[4]For example, the Pirates, in keeping with their stated policy, kept all prices fixed except those of bleacher seats in the final week before a game, while on the other end of the spectrum the Giants changed the price of any given ticket once every five days on average.

Using this data, I estimate a flexible model of consumer demand using ticket purchase data for the four franchises that had sufficient sales for nonbiased estimation. Consumers arrive each day leading up to a game, and make a discrete choice among the ticket varieties on both the primary and secondary markets, or choose not to buy. Both game characteristics and the distribution of the two consumer types can change across time. I simultaneously estimate demand parameters and the parameters of a type-probability function using a general method-of-moments (GMM) framework. Here I develop a handy "trick": I must aggregate sales over multi-day periods in order to get an accurate estimation, but I still make use of daily inventories, prices and characteristics to increase the number of moments and improve accuracy. I also face the challenge of finding appropriate instruments, since marginal costs are zero. I instead use variables related to opportunity cost, as well as the instruments found to be optimal by Berry et al. [1995].

First, the estimation finds that the home team's probability of entering the playoffs is a significant determinant of the utility a consumer derives from a ticket to a future game. Second, it rationalizes the data with the interpretation that consumers with higher willingness to pay shop closer to a given game. It does so despite the fact that the number of shoppers, estimated using webpage visits, is assumed to increase dramatically in the final week before a game. While these two consumer types are universal over the four stadiums, the type-probability function is allowed to depend on the stadium. For example, high willingness-to-pay consumers tend to start shopping sooner for Giants games than for Cardinals games, which may explain why the Giants tend to raise prices consistently over the months before a game while the Cardinals usually wait until the final two weeks to raise price.

I then use the estimated parameters to conduct counterfactual simulations of fixed and dynamic pricing for a typical game, assuming that capacity constraints do not bind. Specifically, I run simulations of the team's "playoff probability path" in the weeks leading up to the game, using a process whose parameters are estimated from the playoff probability paths of all 30 MLB teams, using a maximum likelihood estimation. I then solve for optimal uniform pricing and optimal dynamic pricing by the franchise, taking the secondary market's response function into account so that the club optimizes over its residual demand. First I analyze a counterfactual in which the consumer distribution does not change but playoff probability can oscillate, and another in which playoff probability is constant but the consumers change. In the final four days, the change in consumer distribution increase price more than the average increase or decrease in price from movements in team performance. In both worlds, the gain from dynamic pricing relative to uniform pricing is significant, though the gains in the second world are nearly 50% higher. Typical shifts in the consumer base thus appear to "outweigh" typical variation in team performance, in terms of extracting revenue from consumers.

However, I then run a counterfactual scenario in which both channels operate—consumers and products both "change"—but I isolate the use of intertemporal price discrimination by considering a pricing strategy which depends only on the date of purchase and not on the playoff probability. I compare uniform pricing to this strategy, and this strategy to full dynamic pricing. The ability to change prices over time *and* in response to product changes leads to much larger revenue gains than those associated with pricing policies which depend only on the date of purchase—about 5.75% compared to 0.113%. In other words, mis-optimization is very costly: the gains from intertemporal price changes are nearly erased if these prices are not allowed to re-optimize around unforeseen demand shifts. If the franchise were constrained to the latter type of pricing policy, prices would only increase very close to a game, so that the effect on consumer surplus is negligible, while use of the more flexible pricing policy decreases consumer surplus by an average of 1.4%, with low-willingness-to-pay consumers who shop close to the game bearing the brunt of the welfare loss.

The 5.75% revenue gain from dynamic pricing would be larger if not for downward price movements in the secondary market. I estimate the secondary market equilibrium pricing response as a function of state variables and incorporate this function in the above counterfactuals.[5] As I will show, the franchise does retain residual market power, as consumers apparently view these primary market tickets as differentiated goods from secondary market tickets. The presence of the secondary market response function influences the counterfactuals in two ways: first, resellers' prices tend to move up or down as team performance improves or worsens, consistent with shifts in the resellers' expected values of attending a game and thus their opportunity costs. This fact has

[5]As I discuss further in the paper, this method has a problem: the response function might change if ticket resellers were aware that the team was pursuing a different pricing strategy. I outline a method to overcome this problem that can be used in future research.

an ambiguous effect on the franchise's revenue gains from DP: when team performance improves, the franchise benefits as resllers increase prices, while when team performance worsens, the franchise must compete with lower prices. Note that magnitudes of both the franchise's discounts and price hikes are increased. Second, their prices generally fall as a game approaches, due to falling opportunity costs, so that the franchise is less able to take advantage of the arrival of higher willingness-to-pay consumers.

Finally, I acknowledge that for some teams and some games, capacity constraints do play a role. The shadow price of capacity will change when the team's playoff probability changes, as tickets are now more valuable. however, simulating the value function in the dynamic choice of pricing is prohibited by the size of the state space. I therefore make an attempt to estimate the value of dynamic pricing, relative to uniform pricing, in the capacity-constrained case by searching over a constrained family of pricing policies with very few parameters. These results are preliminary and are included in the Appendix.

1.1.1 The Role of the Secondary Market

I will show that primary market tickets and secondary market tickets should be modeled as differentiated goods in the eyes of consumers. How differentiated they are, and whether consumers harbor a bias toward or against these tickets, is a key determinant of the firm's market power and ability to price discriminate.

A convenient feature of MLB ticket markets is that resale provides a source of competition but not an opportunity for arbitrage. Due to transaction costs, the Coase Theorem does not apply: single-game tickets are rarely purchased in order to be resold, so tickets on the resale market are almost always part of a season package purchased at a discount price at the start of the season [King and Fisher, 2011, Sweeting, 2012, Zhu, 2014].[6] Not having data on season ticket purchases, I treat season package purchases as exogenous when I simulate counterfactual pricing by the franchise. This exogeneity restriction may be forgivable as I do not allow prices to go below season ticket discount prices, consistent with the stated policy of all franchises that practice DP.

The determinants of resale ticket prices differ from the determinants of the franchise's

[6] King and Fisher [2011] find that "92 percent of baseball tickets sold on StubHub come from season accounts." Zhu [2014] uses season packages and single-game sales data, provided by an anonymous franchise, and comes to a similar conclusion. His analysis and mine make the same simplifying assumption of an exogenous initial stock of secondary market listings.

prices: resellers have no reason to price-discriminate, but should change prices when the product changes. It is instructive to consider Sweeting [2012]'s model of a ticket reseller with a single listing. A reseller's price is his opportunity cost, or the expected value of holding tickets the following day, plus a markup. Assuming other resellers offer nearly identical tickets, the markup is small, so the opportunity cost is the key factor. It is directly correlated with a seller's value of attendance and therefore with changes to the product. To take an extreme example of a product change affecting resellers' prices, when Serena Williams was unexpectedly knocked out of the U.S. Open Tennis Championship in the semifinal, resale ticket prices for the Women's final dropped by an average of about \$300 [Garcia, 2015].

On the other hand, traditional price discrimination has no place here, since the reseller has a very limited supply and no market power. It is not surprising, then, that resale prices tend to fall significantly leading up to a given game. They fall particularly steeply over the final week before a game, precisely when DP-practicing franchises tend to raise their own prices.

As I explain in the counterfactual results, I endogenize the behavior of resellers in my counterfactuals, as opportunity costs depend on the franchise's current and predicted behavior. For now, I estimate an "equilibrium pricing policy function" that depends only on the franchise's current prices for a given game. However, I outline a method to be used in future work that would allow reseller behavior to depend not only on current prices but expected future prices. This method would more rigorously simulate what happens when the team uses alternative pricing strategies, such as fixing prices over time.

1.1.2 Industry Background

DP has only recently become common in the North American sport industry. About half of the clubs in the National Basketball Association and the National Hockey League practice some form of it [Shapiro and Drayer, 2013], while I find that 25 of the 30 MLB franchises have adopted DP strategies between 2009 and the present. Much of this development can be attributed to software from external vendors which tracks ticket demand—both direct sales and StubHub sales—and makes black-box pricing recommendations [Xu et al., 2015].

The most popular software vendor in this niche industry, Qcue, was founded by former StubHub employees. Their relationship with StubHub provides them easy access to transaction data that they use to analyze demand shifts [Interview, 2014]. According to Qcue CEO Barry Kahn, "Baseball definitely presents the best opportunity for dynamic pricing. Compared to the NBA or NHL, MLB plays twice as many games, its venues are twice as large and the percentage of season ticket holders represents a smaller fraction of total sales. Other key factors include… a fairly cut-throat playoff system where only a limited number of teams advance" [Rishe, 2012].

In addition to the five franchises who do not practice DP, there are likely many franchises who rarely use it. Three of the six teams in my sample make price adjustments less than three times, on average, in the last 60 days before a game.[7] Apparently the cost of full DP adoption to these franchises is significant. In this paper I provide an estimate of the unconditional profit benefit of DP adoption, while in future research I plan to estimate this profit benefit for various teams and derive bounds on the cost of adoption.

1.1.3 Related Literature

The term "dynamic pricing" obviously refers to changing prices over time, but different strands of literature confer different emphases on its meaning. For example, theoretical work in the operations research literature thinks of "dynamic pricing" as optimization over stochastic demand and capacity limits [Gallego and Ryzin, 1994, Talluri and van Ryzin, 2004], with price discrimination between multiple consumer types playing a tertiary role if any.[8]

More relevant to my work is the theoretical literature on price discrimination. Clerides [2000] derives optimal pricing implications across various demand models and definitions of price discrimination. Stokey [1979] looked at intertemporal price discrimination in which prices fall as consumers tire of a good, while Rosen and Rosenfield [1997] specifically describe a model of ticket pricing that involves second-degree price discrimination. Relevant to the competition between primary and secondary markets in my setting is Holmes [1989], which analyzes the competitive effects of oligopoly when

[7] Many clubs only use Qcue's visualization tools but not its price recommendations [Interview, 2014].

[8] Elmaghraby and Keskinocak [2003] provide a comprehensive survey of this literature.

firms use third-degree price discrimination.[9]

Verboven [2008] and Lambrecht et al. [2012] provide useful surveys of empirical studies of price discrimination, which "has only recently become an area of rigorous empirical research." Intertemporal price discrimination has been studied across many markets including airlines [Lazarev, 2013, Williams, 2013],[10] Broadway theater [Leslie, 2004], and video games [Nair, 2007].[11] As stated above, these studies did not have data on product characteristics that change over time. Lazarev [2013] focuses on price discrimination but also simulates flight purchases under free resale. My work suggests, not surprisingly, that "free resale" is a theoretical extreme: as stated above, transaction costs make the resale of single-game tickets a rarity.

The problem of baseball ticket pricing in particular has received empirical attention from a few works which, interestingly enough, are rare in that they do not look directly at price discrimination. Xu et al. [2015] and Zhu [2014] both use sales data from anonymous MLB clubs to estimate a demand model and simulate revenue from optimal pricing, but do not allow consumer tastes to evolve systematically over the days leading up to a game.[12] Price discrimination of consumer types is implicit in Zhu [2014]'s time dummies but not explicitly analyzed, while in Xu et al. [2015] dynamic pricing takes place only because prices are allowed to fall below season ticket prices so that the capacity constraint frequently *does* bind.[13] Sweeting [2012] analyzes MLB resale ticket prices on StubHub and eBay in 2007, before any franchises used DP. His is the only paper I know of that uses dynamic event characteristics—various measures of team performance—and he does find them to be significant determinants of demand and

[9] Tirole [1988] and Stole [2007] provide a thorough discussion of the different kinds of price discrimination.

[10] Williams [2013] has a similar flavor to this paper, as it disentangles two motivations for DP. In his case, he studies the importance of capacity limits when demand is stochastic, relative to the importance of price discrimination. Capacity limits play a limited role in my setting.

[11] Another work analyzing price discrimination in the arts and entertainment is Courty and Pagliero [2012], which estimates revenue gains from price discrimination in the concert industry. However, in their setting the firm prices differentially across concerts, rather than across time leading up to a concert.

[12] Xu et al. [2015] excludes the secondary market from its analysis, while Zhu [2014] takes a similar approach to mine by collecting simultaneous StubHub sales, combining the primary and secondary markets into one, and differentiating the two classes of tickets.

[13] Xu et al. [2015] provide a caveat that the anonymous franchise's "strategic long-term rationale related to consumer behavior and pricing integrity" led it not to follow the paper's recommended strategy. Still, this paper is interesting in its own right: if the revenue gains are large enough, and a large fraction of season tickets are resold and cannibalize primary single-game ticket sales, one might question the wisdom of offering season packages at all. This question is beyond the scope of this paper.

pricing, as I do. The analysis finds reseller pricing to be consistent with a model in which consumers are not dynamic, i.e. they do not shop for tickets over multiple days. This finding, together with additional evidence, will justify my use of a model with static consumers.[14]

1.1.4 Outline

The rest of the paper proceeds as follows. Section 2 describes the data collected for this study. Section 3 presents the model. Section 4 discusses the econometric specification and identification of the model parameters. Section 5 presents the results of demand estimation. Section 6 analyzes the implications of these results on the effects of dynamic pricing. The conclusion follows.

1.2 Data

The empirical analysis uses data for single-game tickets to the regular season home games of six MLB franchises in 2014. Data was collected both from Tickets.com, where tickets are purchased directly from the teams, and from the dominant secondary market, StubHub. I also collected relevant game characteristics, some of which change over time leading up to a game, as detailed below.

1.2.1 Primary market: Tickets.com

Tickets.com has a partnership with the six franchises in my sample, among others. They are the Oakland Athletics, the St. Louis Cardinals, the San Francisco Giants, the Baltimore Orioles, the Pittsburgh Pirates, and the Minnesota Twins. Consumers can purchase tickets directly from the team for any given future game by viewing a seatmap that is updated at least hourly.[15] This seatmap provides the current price and number of seats remaining in each section of the stadium.[16]

[14] The same simplifying assumption is made in Williams [2013] and Leslie [2004], though Zhu [2014] estimates one model with static consumers and one with dynamic consumers.

[15] This statement is backed by an experiment in which I collected data for 30 random games, five for each team, every half hour over a one-week period.

[16] To give an idea of scale: the six teams' stadiums vary in size from 35,067 seats (Athletics) to 46,861 seats (Cardinals). Sections usually contain 150 to 200 seats, so that a stadium is divided into 200 to 300 sections.

Data collection began between April 15 and April 25 for all teams with the exception of the Pirates, for whom data collection started on June 6. I collected this data each morning at 1 AM, for all future games, through the last day of the season on September 29th. An observation is a game, search date, and stadium section. For each observation I observe price and inventory. For a subset of games I also collected inventories of each section one hour in, to get an idea of the magnitude of day-of purchases. The seatmap shows inventory available across all channels including internet, phone, and box office [Interview, 2014].

Data collection and cleaning details are provided in the appendix.[17] I went to great lengths to determine the fee schedules of each franchise in the sample, so that the prices I use below are final prices including fees.

1.2.2 Secondary market: StubHub

In order to model the consumer's choice set more accurately, I also collected data from StubHub and eBay using their web APIs.[18] Data downloads took place each day at 3 AM and included seat section, price, and inventory for every listing in the 2014 MLB regular season.

Currently I do not use the eBay data, as it does not include changes in inventory except when the listing sells out completely. However, the number of transactions per day on eBay is less than a tenth of the number observed on StubHub, likely due to StubHub's official partnership with MLB and the ease of printing out a digital ticket after purchase on StubHub.[19]

Data collection and cleaning details are provided in the appendix. Fortunately, earlier in 2014 StubHub began using "what you see is what you get" prices for consumers,

[17] Collecting data from this website was non-trivial. On a given day leading up to a given game, one can look at a seat map for that game and, using a mouse, hover the cursor over a particular section to see a popup stating the price and number of seats remaining. The seat maps are Flash-animated, so the data was not available via HTML source. This obstacle was overcome using a Python script involving windows automation, where a simulated mouse literally moved around a screen, pausing to wait for these dialog boxes to pop up. Ultimately, data collection was accomplished using a series of automated mouse movements, screenshots, and sophisticated text recognition tools. See the Appendix for more details.

[18] As discussed in section 1.4, secondary market data also allowed me to better account for price endogeneity in the demand estimation, since the instruments for primary firm price are relatively weak.

[19] In future work that is more focused on competition by resellers I may exploit this data, as the choice to create a listing can be used to estimate, or at least bound, a reseller's value of attendance.

as all fees were charged to the seller, so no price conversion was necessary.

1.2.3 Game characteristics

I use team fixed effects and several game-level attributes, such as time of day, in the demand model. Details on these attributes are provided in subsection 1.2.8. I also collect all 30 MLB teams' probabilities of entering the playoffs on each day of the regular season from a reputable third-party website called FanGraphs.com.

Two other dynamic event characteristics were collected but not yet incorporated into the model and estimation. I mention them here for the curious reader: (1) 10-day weather forecasts for the six sample stadiums, and (2) expected starting pitchers for each day leading up to each game as reported by ESPN.com.[20]

1.2.4 Stadium zones and areas

Each franchise divides its stadium into pricing zones. For example, Figure 1.1 shows the Twins' stadium with zones differentiated by color.

For the demand estimation I aggregate zones up to seating "areas." Areas were chosen to combine adjacent pricing zones that share similar primary market prices; intuitively seats in such zones are likely to be highly similar products. The number of areas per stadium therefore depended on the amount of price variation on the primary market.[21]

1.2.5 Example of primary market data

Figure 1.2 is an example of the data: it shows price paths and inventory levels for four seating zones in the St. Louis Cardinals' stadium for their September 21st game against the Cincinnati Reds. These graphs are intended merely as an example and are not meant to illustrate any "typical" pricing or sale pattern. Price paths do, however, tend to move more often as a game approaches, particularly 14 and 7 days out, and to slope upward. I elaborate on these facts in subsection 1.2.7.

[20] These names do frequently change in the two weeks before a game.
[21] At most, the difference between the lowest and highest primary market prices in a stadum area is \$10.

FIGURE 1.1 TWINS STADIUM SEATING ZONES

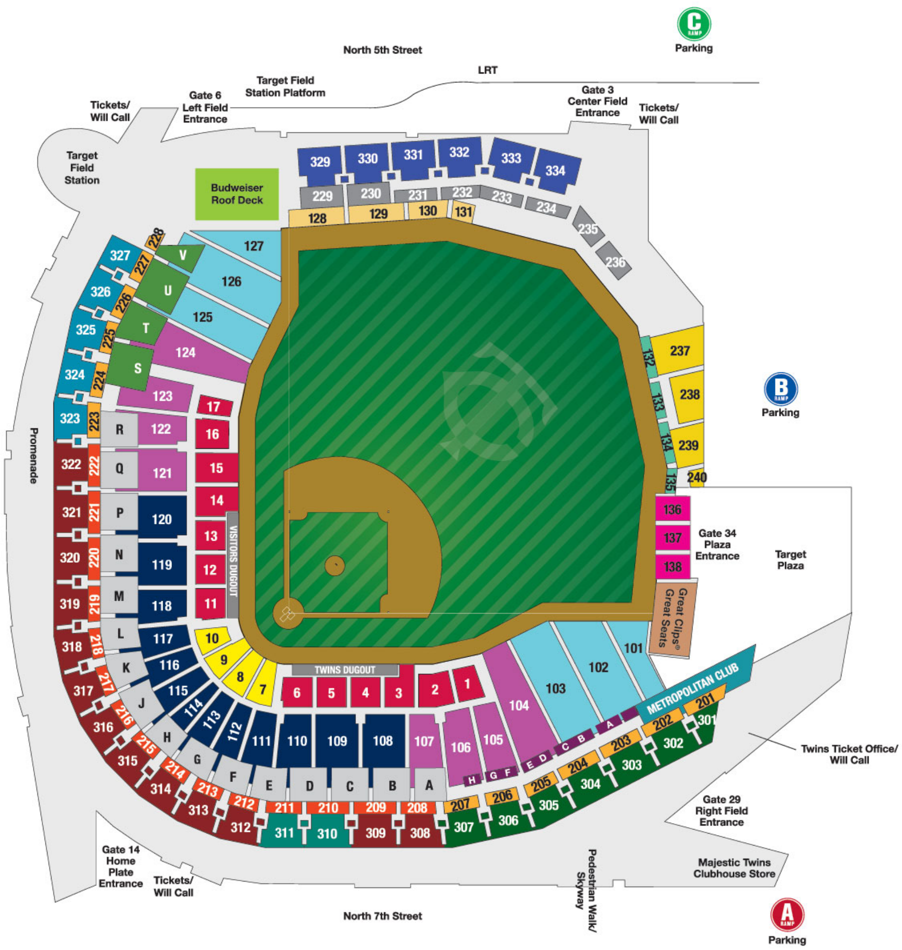

1.2.6 Summary of ticket sales and characteristics

In order to measure the relative importance of the two motivations for dynamic pricing, and to calculate welfare effects relative to fixed pricing, I will estimate a demand model and compute counterfactual demand under fixed pricing. In this subsection and Subsection 1.2.7, I describe the data and motivate the model that will be presented in Section (1.3). Key points about the data are numbered and italicized.

For each game, my data contains a number of "search days" leading up to it. I index

Figure 1.2 Example of Price and Inventory Paths: Cardinals-Reds, 9/18/2014

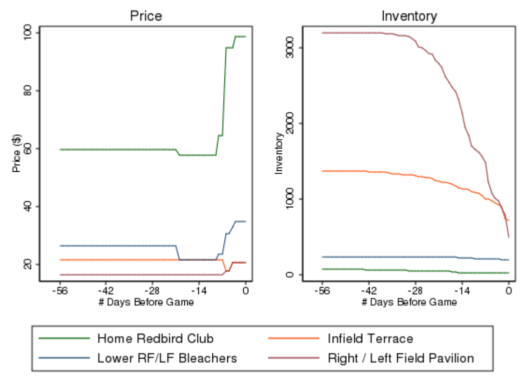

a "search day" or "day prior to game" by

$$t = (\text{date of observation}) - (\text{date of game})$$

so that, leading up to a given game, t moves from some negative number up to zero.

(i) In the demand estimation I use a subset of the games and search days in my data: four of the six teams, only games whose observations start on or before $t = -28$, and only days $t = -28, \ldots, -1$. First, I exclude $t = 0$ because I was not able to observe secondary market sales on this day.[22] As I discuss in Section 1.4, the demand estimation relies on having sufficient quantities of observed sales. At a reasonable level of aggregation by area and time periods, there were insufficient sales for the Oakland Athletics or the Baltimore Orioles.

Table 1.1 shows, by franchise, the number of games for which I have data, the number of stadium pricing zones defined by the teams, the number of stadium areas up to which I aggregated. It also provides the minimum, mean, and maximum number of search days observed for the various games in my data. For example, the minimum number of search days equals 1 for the Giants because they hosted a game on April 15th, the first day of data collection.

(ii) Total sales, as well as the fraction of sales that take place on the secondary

[22]The reason is twofold: StubHub does not provide day-of-game sales data, and there are likely to be a large number of secondary market sales that occur outside the stadium rather than through StubHub.

Table 1.1 Summary of data

	Athletics	Cardinals	Giants	Orioles	Pirates	Twins
Full Sample						
# of games	71	75	69	73	52	71
# of stadium zones	13	31	23	22	16	23
# of stadium areas	6	8	6	8	5	7
Min. # of search days	4	8	1	2	3	2
Mean # of search days	58	81	65	88	54	45
Max. # of search days	126	157	144	160	110	101
Estimation Sample						
# of games	0	50	34	0	25	19
# of search days	.	28	28	.	28	28

market, increase as a game approaches, as shown in Table 1.2. For example, the average number of tickets sold between $t = -7$ and $t = -1$, over all 389 games for which I fully observe sales, is 1,804 on the primary market and 1,208 on the secondary market, for a total of 3,012. This total number, and the ratio of secondary sales to primary sales, are both much larger than any previous week leading up to the game.

59% of total mean ticket sales on the primary market take place during the final 28 days prior to a game. While I cannot calculate the analogous figure for the secondary market without day-of-game sales, it is likely to be even higher given that the fraction of sales that take place on StubHub is rising as the game approaches. The estimation sample accounts for 44% of all mean ticket sales on the primary market. See figure **??** in Appendix **??** for a more complete picture of what data I collected, and for mean sales by search day.

(iii) Primary and secondary tickets in the same stadium pricing zone are different goods in the eyes of consumers. 29.6% of tickets purchased on the primary market were bought when at least four tickets in the same pricing zone were available for $10 less on the secondary market. The fraction of tickets purchased on the secondary market when at least four were available on the primary market for $10 less was 58.2%. Of course, I do not observe the row of primary market tickets, but it is hard to imagine that taking this into account would bring these percentages close to zero.

Table 1.2 Mean Number of Tickets Sold Over Time, by Game

Days Prior to Game	Obs.	Primary	Secondary	Total
Full Sample				
0	0	1589	.	.
1-7	389	1804	1208	3012
8-14	363	1077	506	1582
15-28	310	1644	682	2323
29-56	216	2092	854	2942
57-112	71	2091	964	3056
Estimation Sample				
1-7	128	1523	1567	3090
8-14	128	1079	676	1755
15-28	128	1723	889	2612

Obs. = number of games containing data over the entire period specified.
Full sample is Athletics, Cardinals, Giants, Orioles, Pirates, and Twins, with data collection starting on 6/6/2014 for the Pirates and between 4/15/2014 and 4/25/2014 for the other franchises. Estimation sample is limited to Cardinals, Giants, Pirates, and Twins, last 28 search days.

(iv) Even at the most popular games, teams never run out of tickets to sell. Franchises use a loose definition of capacity to define a "sellout."[23] Even the Giants, who claim a multi-year "sellout streak," usually have 200 to 500 tickets remaining unpurchased throughout any given game [Bulwa, 2012]. Table **??** in the Appendix, which shows remaining ticket inventory at $t = -56, -28, 0$, and one hour into a game, supports this fact: only 5% of games in my sample have less than 60 primary market tickets remaining, one hour into the game. Furthermore, consumers who shop last minute are unlikely to be turned off by lack of seating options: only in 3% of the 411 games observed do more than two stadium areas (out of 5 to 8 total areas, depending on the franchise) sell out by the day of the game.

(v) The quality and price of secondary market tickets is higher on average, as shown in Table **??** in the Appendix. This table provides statistics on the characteristics of remaining tickets over all ticket, day-prior-to-game observations. The average ticket

[23]Specifically, in the industry a "sellout" means that the distributed number of tickets exceeds the number of seats, which can easily happen because there are a very large number of "standing room only" tickets available.

on StubHub is 17 feet (or 5%) closer to home plate, 26% more likely to be on the first floor, and 36% higher in face value.[24] A couple reasons explain these facts. First, a large portion of remaining primary market tickets tend to be in the upper deck, especially in the final weeks before a game. Second, most secondary market tickets were originally purchased as part of a season package, which are usually higher-quality seats.

The difference in mean prices between the two markets is also shown in Table **??** and is also striking. Even excluding secondary market tickets listed at prices more than triple their face value,[25] the mean price of available secondary market ticket observations is double that on the primary market.

(vi) The Cardinals and Giants changed ticket prices with greater frequency than other franchises, also illustrated by Table **??**. Cardinals and Giants tickets each have an 8% and a 20% probability, respectively, of changing price on any given day, compared to 1% for a Twins ticket (and for a Pirates ticket, but this fact is less surprising given that franchise's fixed-prices-except-bleachers policy). Meanwhile, secondary market tickets exhibit roughly the same frequency of price changes (8-10%) regardless of team.

There are several possible explanations. None of these teams, including the Cardinals and the Giants, employ workers full time to worry about pricing, so it comes down to how much the franchise trusts the DP software vendors and how worried teams are about "undermining ticket value" [Interview, 2014]. Experience may play a role: the Cardinals and the Giants signed on with software vendor Qcue in 2011 and 2009, respectively, while the other teams did so in 2012 or later.[26] It is entirely possible that the Athletics, Orioles, Pirates, and Twins are not carrying out any kind of systematic dynamic pricing policy, so I focus on the Giants and Cardinals below.

[24] "Face value" refers to the price of the ticket on the primary market 56 days before the game, as I do not directly observe face values.

[25] Posting a listing on StubHub takes little time and no payment, and a particular pair of seats for multiple games is almost as easy as listing them for one game. Many sellers may find it convenient to post all of their season tickets, even the ones they strongly desire to use, just in case an exuberant buyer is willing to pay.

[26] To be specific, the Twins signed a contract with Digonex, Qcue's main competitor, which also makes pricing recommendations for Six Flags.

1.2.7 Pricing by primary and secondary market suppliers

Figure 1.4 looks more closely at these franchises' price changes. It looks at the entire pool of available tickets in my data, summarizing over day prior to game for each team. The probability of a price increase or decrease by either team increases as the game approaches, particularly in the final two weeks prior to the game. The Giants increase price more frequently but, but when they increase price they do so in smaller magnitudes: the median increase is roughly \$1, compared to the Cardinals' median increase of \$5. Giants discounts are much less frequent, but with similar magnitudes to the Cardinals'.

(vii) Consumers do not seem to strategically delay purchase. As shown in Figure 1.4, the average fraction of tickets that see a price increase goes from to from $t = -5$ to $t = -4$. If consumers were strategic, they would likely observe that the Cardinals have a particularly high propensity to raise prices on day $t = -4$. We would expect to observe a spike in sales on day $t = -5$, but we do not: mean Cardinals ticket sales increase smoothly over time leading up to games.

Figure 1.4 is complicated by compositional effects and looks only at median magnitudes; it mainly shows that different teams have different pricing styles on a superficial level. We now turn to the mean price paths in Figure 1.3 for a more substantive comparison.

(viii) The Giants and Cardinals both have upward sloping average price paths, but with different shapes. To generate Figure 1.3, I aggregated the data first by taking price-to-face-value ratio, then by averaging across seating zones using the same zone weights for all games and days prior to game, so that the graph shows only true price changes rather than compositional effects.[27] The Cardinals and Giants reach the same average price-to-face-value ratio by the day of a game, but the Cardinals only begin increasing prices in the last two weeks while the Giants steadily raise prices over at least the last 8 weeks.

(ix) Secondary market prices are more likely to fall, especially in the final week.

[27] On the primary market, all tickets in a given zone have the same price, by definition. On the secondary market, "price in each stadium zone" refers to the 10th percentile of available tickets in that zone. The mean transacted price tends to oscillate around this statistic, so I take it to be the typical price a consumer might expect to pay. The reason I included the secondary market price path will become clear momentarily.

Figure 1.3 Mean Price-to-Face Ratio Across Time

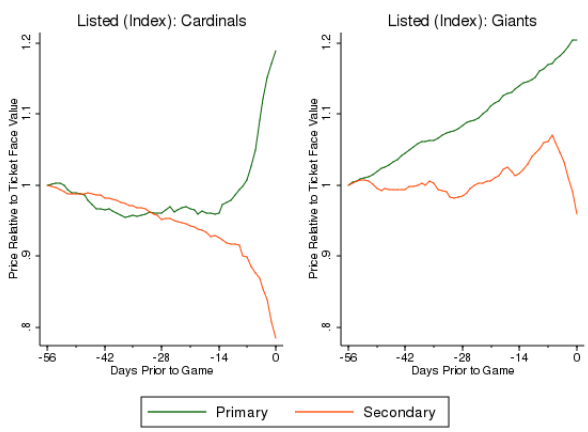

Figure 1.4 Frequency and Magnitude of Franchise Price Changes Across Time

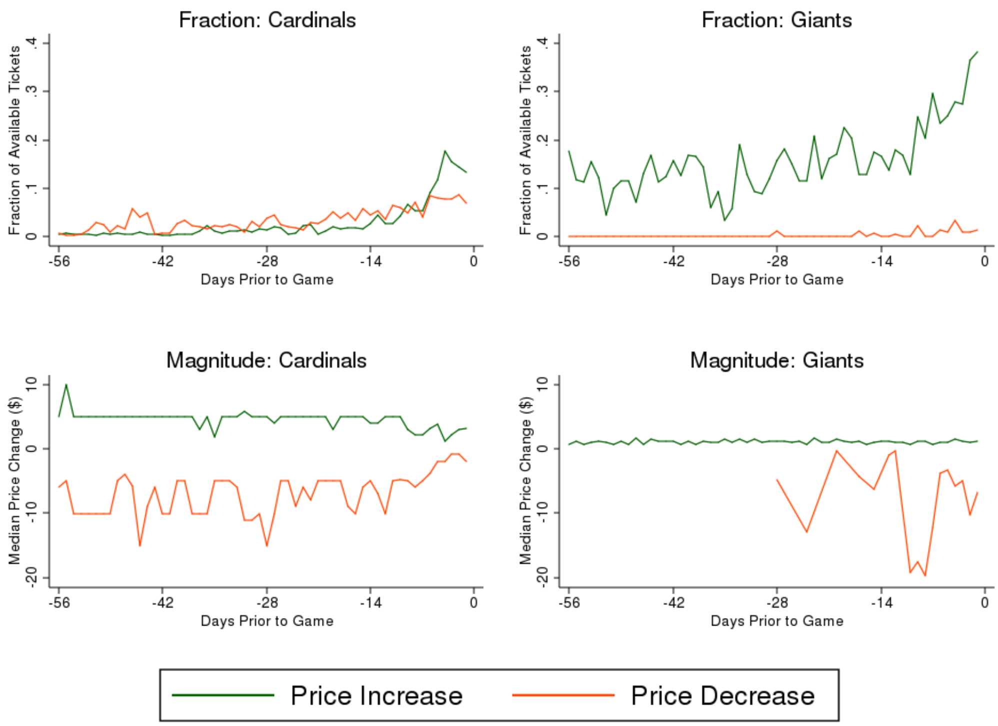

These sellers lack market power: they underbid each other and price near opportunity cost. By "opportunity cost" I mean the expected value of still holding a pair of tickets tomorrow.[28] This expected value decreases as a game approaches: with fewer opportunities to sell the tickets, either the probability of sale or the price must fall. This dramatic difference between primary and secondary market price paths was hinted at back in Table **??**: when price changes occur, the mean price change is roughly zero (-$0.64) on the primary market but -$9.49 on the secondary market.

[28] I am somewhat generalizing; not all sellers are trying to sell a pair of tickets. About 75% of listings on StubHub offer a single pair of tickets. Multiple listings might be managed by the same seller.

1.2.8 Changing products

(x) Team performance affects pricing in both markets, according to preliminary evidence. As measures of team performance, I use the home and away team's probability of entering the playoffs, matching these values to my observation days. The probabilities are calculated by Fangraphs, a popular baseball statistics website.[29] Figure 1.5 shows the 2014 probability paths for each home team in the estimation sample, though I also use away team playoff probabilities in the estimation.

FIGURE 1.5 TEAM PERFORMANCE IN THE 2014 REGULAR SEASON

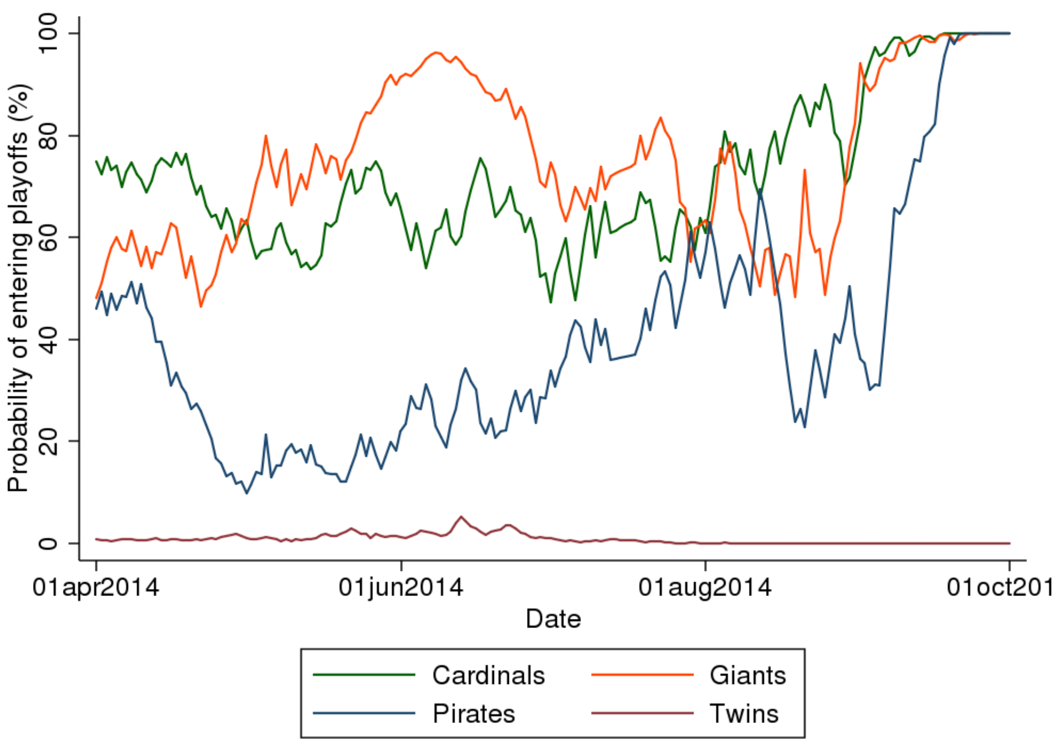

In the first two columns of Table 1.3 I regress log price on log characteristics, including home and away team playoff probabilities, as well as game and seat zone fixed effects. Consistent with the finding above, this reduced-form pricing function is increasing in time while secondary prices are decreasing in time. Home and away team probability

[29]See http://www.fangraphs.com/coolstandings.aspx?type=4 for information about their proprietary algorithm.

of entering the playoffs are positive and significant in both markets. A 100% increase in, or doubling of, home team playoff probability is associated with an 8% increase in primary market price and an 11% increase in secondary market price, while a doubling of away team playoff probability corresponds to a 1% increase in primary market price and a 1.6% increase in secondary market price. These results are consistent with the idea that fewer fans care about the opposing team than the home team.

Table 1.3 Primary and Secondary Price Determinants

	(1) Primary Price	(2) Secondary Price
playoff probability	0.126***	0.116***
	(0.00352)	(0.0152)
away's playoff probability	0.0198***	0.0169***
	(0.000475)	(0.00257)
time	0.0106***	-0.0443***
	(0.000613)	(0.00474)
game FE	Yes	Yes
seat FE	Yes	Yes
N	240946	34737
R^2	0.918	0.753

Standard errors in parentheses

All dependent and independent variables are in logs.

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

When I estimate demand I also use static game characteristics. These are summarized, along with the playoff probability variables, in table **??** in the appendix.

1.3 Model

In this section I build a structural model of demand and supply for seats at an event. On each search day leading up to each game, I assume events take place in the following order:

1. Time-varying characteristics (home and away-team playoff probabilities) are realized.
2. The primary supplier chooses prices for each venue area.
3. Secondary suppliers see these prices and choose prices for their own listings.
4. An exogenous number of consumers arrive, interested in attending the event. Each consumer chooses to purchase a ticket variety (defined in 1.3.1) or to exit the market forever.

No consumer purchases a single-game ticket with intent to resell. Non-purchase changes to the stock of secondary market tickets are taken as exogenous.

1.3.1 Consumer demand

As defined in section 1.2.6, a day prior to a game t equals the shopping date minus the game date. t will therfore start at some negative integer $\bar{t}$ and count up to 0, the day of hte game. In the estimation and counterfactual I let $\bar{t} = -28$, so that $t = -28, -27, \ldots, -1, 0$. Baseball games are indexed by $g = 1, \ldots, G$. A market is a pair (g, t), a single shopping day leading up to a given game.[30] For the remainder of this section I focus on a single game and omit the g subscript, so that t will denote markets.

A ticket variety j is defined by its seating area and whether it is offered on the primary or secondary marketplace. Let the number of seating areas be A. Inside goods are indexed by $j = 1, \ldots, 2A$, where $j = 1, \ldots, A$ are primary products and $j = A+1, \ldots, 2A$ are secondary products. Let n_{jt} denote inventory, or the number of remaining tickets belonging to product j, on day t. The set of available inside goods

[30] Tickets to all other events are included in the outside good (defined below), as assumed in Leslie [2004], Williams [2013], and Xu et al. [2015], for example.

is $\bar{J}_t = \{j \in \{1, \ldots, 2A\} : n_{jt} > 0\}$ and the set of *all* available goods is $J_t = \bar{J}_t \cup \{0\}$. $j = 0$ denotes the outside good, which is not purchasing a ticket.

On day t an exogenous number of consumers m_t arrive. Each consumer $i = 1, \ldots, m_t$ makes a discrete choice over the products $j \in J_t$.[31] Each product j is characterized by a price p_{jt}, a secondary market dummy $S_j = 1_{j \in \{A+1, \ldots, 2A\}}$, and a vector of area-specific characteristics $\mathbf{x}_j$ (such as distance to home plate).[32] There is also a vector of game-specific characteristics $\mathbf{x}_t$, both static (such as a night-game indicator) and time-varying (including home and away team playoff probabilities). Finally, consumers see a product attribute ξ_{jt} that is unobserved to the econometrician, likely a composite of unobserved event characteristics and (potentially product-specific) promotional effort.

Each consumer belongs to a type $r = 1, \ldots, R$. The conditional indirect utility that a consumer i of type r derives from choosing product j in market t is

$$u_{ijt} = \delta_{rjt} + \varepsilon_{ijt} \tag{1.1}$$

where ε_{ijt} is an idiosyncratic demand shock and δ_{rjt} is the average utility a type-r consumer derives from product j. δ_{rjt} is assumed to be a linear combination of the price and attributes, or

$$\delta_{rjt} = \alpha_r p_{jt} + \beta_{0r} + S_j \beta_{1r} + \mathbf{x}_j \boldsymbol{\beta}_{2r} + \mathbf{x}_t \boldsymbol{\beta}_3 + \xi_{jt}, \quad \text{with} \quad \delta_{r0t} = 0. \tag{1.2}$$

The α and β terms represent consumers' tastes for price and other product characteristics. I assume that consumer types agree in their tastes for game characteristics, so β_3 has no r subscript. Let the utility parameters be summarized by the vectors

$$\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_R)' \quad \text{and} \quad \boldsymbol{\beta} = (\beta_{01}, \ldots, \beta_{0R}, \beta_{11}, \ldots, \beta_{1R}, \boldsymbol{\beta}_{21}', \ldots, \boldsymbol{\beta}_{2R}', \boldsymbol{\beta}_3')'.$$

Following the discrete choice demand literature, consumer i chooses product j in market t if and only if

$$u_{ijt} \geq u_{ij't}, \forall j' \in J_t.$$

[31] For expositional simplicity I assume for now that all consumers in a period face the same menu J_t. If one wanted to be rigorous one could have consumers choose products sequentially, with the choice set evolving as purchases are made. The sellout of one of these products is a rare event in the data.

[32] In a future revision I plan to estimate area fixed effects by letting $\mathbf{x}_j$ consist of area dummy variables. Also, note that a primary product j has the same stadium area as secondary product $j+A$, so $x_j = x_{j+A}$.

This model provides an intuitive, flexible representation of market demand for differentiated products. I make a standard assumption on the distribution of the idiosyncratic demand shock:

Assumption 1. *Consumer idiosyncratic preferences are distributed* $\varepsilon_{ijt} \overset{iid}{\sim}$ *Type 1 Extreme Value (T1EV).*

The fraction of consumers in market t who are type r is γ_{rt}, which is assumed to be given by some (possibly game-specific) function with parameter vector $\boldsymbol{\phi}$ that will be specified in Subsection 1.4.4. We then have the following closed-form solutions for the choice probabilities:

$$s_{rjt}(\boldsymbol{\delta}_{rt}) = \frac{\exp\left(\delta_{rjt}\right)}{\sum_{k \in J_t} \exp\left(\delta_{rkt}\right)} \quad \text{and} \quad s_{jt}(\boldsymbol{\delta}_t; \boldsymbol{\phi}) = \sum_{r=1}^{R} \gamma_{rt} s_{rjt}(\boldsymbol{\delta}_{rt}). \tag{1.3}$$

where $\boldsymbol{\delta}_{rt} = (\delta_{r,1,t}, \ldots, \delta_{r,2A,t})'$ and $\boldsymbol{\delta}_t = (\boldsymbol{\delta}'_{1t}, \ldots, \boldsymbol{\delta}'_{Rt})'$. The expanded share function is $s_{jt}(\mathbf{p}_t, \mathbf{x}_t, \boldsymbol{\xi}_t, \mathbf{n}_t; \boldsymbol{\theta})$; this way of writing it emphasizes the underlying components of the δ_{rjt}'s and the dependence of the choice set J_t on the vector of inventories $\mathbf{n}_t$.[33]The vector of parameters is $\boldsymbol{\theta} = (\boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\phi}')'$, containing the type-probability parameters $\boldsymbol{\phi}$ and the taste parameters that determine δ_{rjt}.

Location normalization of utility is given by $\delta_{r0t} = 0$ for all r, t, while scale normalization is given by the standardized variance of ε_{ijt}.

1.3.2 Primary firm supply

Let J_t^P and J_t^S denote the set of available products on the primary and secondary markets, respectively, so that $J_t^P = J_t \cap \{1, \ldots, A\}$ and $J_t^S = J_t \cap \{A+1, \ldots, 2A\}$. Vectors of prices and characteristics can be defined accordingly: let $\mathbf{p}_t^P = (p_{jt})_{j \in J_t^P}$ and $\mathbf{p}_t^S = (p_{jt})_{j \in J_t^S}$ be the prices of the goods available on these two markets. Let the state space include inventories and game characteristics, $\boldsymbol{\Omega}_t = (\mathbf{n}_t, \mathbf{x}_t, \boldsymbol{\xi}_t)$.

The baseball franchise is a Stackelberg leader. It knows that if it does dynamic pricing (DP) then secondary market prices will be $\mathbf{p}_t^S = \boldsymbol{\sigma}_t^{DP}(\mathbf{p}_t^P, \boldsymbol{\Omega}_t)$, where $\boldsymbol{\sigma}_t^{DP}(\cdot)$ is

[33] I have omitted S_j and x_j from this function: these do not change over time, and can therefore be considered primitives of $s_{jt}(\cdot)$.

the secondary market's equilibrium pricing response function. If it is in a fixed pricing (FP) regime then secondary market prices will be $\mathbf{p}_t^S = \boldsymbol{\sigma}_t^{FP}(\mathbf{p}^P, \boldsymbol{\Omega}_t)$.

If the firm does dynamic pricing, I allow it to take capacity constraints into account,[34] so that the firm solves the following recursive problem:

$$V_t^P(\boldsymbol{\Omega}_t) = \max_{\mathbf{p}_t^P} \sum_{j \in J_t^P} q_{jt} p_{jt} + E_t \left[V_{t+1}^P(\boldsymbol{\Omega}_{t+1}) | \boldsymbol{\Omega}_t, \mathbf{p}_t^P, \mathbf{p}_t^S \right] \tag{1.4}$$

$$\text{s.t.} \begin{cases} n_{j,\bar{t}} & \text{is given } \ \forall j & \text{(exogenous starting inventory)} \\ n_{jt+1} & = n_{jt} - q_{jt}\left(\mathbf{p}_t^P, \mathbf{p}_t^S, \boldsymbol{\Omega}_t\right) \ \forall j & \text{(inventory transition)} \\ q_{jt} & = \min\left[n_{jt}, m_t s_{jt}\left(\mathbf{p}_t^P, \mathbf{p}_t^S, \boldsymbol{\Omega}_t\right)\right] \ \forall j & \text{(capacity limits)} \\ \mathbf{p}_t^S & = \boldsymbol{\sigma}_t^{DP}(\mathbf{p}_t^P, \boldsymbol{\Omega}_t) & \text{(secondary market response)} \\ V_1^P & = 0 & \text{(value=0 on day after game)} \end{cases}$$

where $V_t^P(\boldsymbol{\Omega}_t)$ is the primary firm's value function on day t in state $\boldsymbol{\Omega}_t$. I write the demand function $q_{jt}\left(\mathbf{p}_t^P, \boldsymbol{\sigma}_t^{DP}(\mathbf{p}_t^P, \boldsymbol{\Omega}_t), \boldsymbol{\Omega}_t\right)$ as above for expositional convenience. A messier, but more realistic, model would allow for consumers to substitute to other products if one becomes sold out mid-period.

Because the demand model is deterministic—there is no uncertainty about today's demand, given $\mathbf{p}_t^P$—I can now derive a clean Lerner Index rule below. Letting $c_{jt}(\mathbf{n}_t)$ denote the shadow prices of capacity, the firm faces the following problem:

$$\begin{aligned} V_t^P(\boldsymbol{\Omega}_t) &= \max_{\mathbf{p}_t^P} \sum_{j \in J_t^P} m_t s_{jt}\left(\mathbf{p}_t^P, \mathbf{p}_t^S, \boldsymbol{\Omega}_t\right) p_{jt} - \sum_{j \in J_t^P} c_{jt}(\mathbf{n}_t)\left[m_t s_{jt}\left(\mathbf{p}_t^P, \mathbf{p}_t^S, \boldsymbol{\Omega}_t\right) - n_{jt}\right] \\ &= \max_{\mathbf{p}_t^P} \sum_{j \in J_t^P} m_t s_{jt}\left(\mathbf{p}_t^P, \boldsymbol{\sigma}_t^{DP}(\mathbf{p}_t^P, \boldsymbol{\Omega}_t), \boldsymbol{\Omega}_t\right)(p_{jt} - c_{jt}(\mathbf{n}_t)) + \sum_{j \in J_t^P} c_{jt}(\mathbf{n}_t) n_{jt} \end{aligned} \tag{1.5}$$

Note that value functions (1.4) and (1.5) are equivalent.

Suppose the franchise faces a constrained problem where it can only make proportional adjustments to base prices, and that shadow prices of capacity are proportional

[34] Capacity constraints rarely bind in the data, so it is likely that, for most teams and most games, the choice of today's price does not affect tomorrow's continuation value. That is, it is usually the case that

$$\frac{\partial E\left[V_{t+1}^P(\boldsymbol{\Omega}_{t+1}) | \boldsymbol{\Omega}_t, \mathbf{p}_t^P, \mathbf{p}_t^S\right]}{\partial \mathbf{p}_t^P} = 0.$$

This condition will be assumed in the majority of the counterfactual simulations in Section 1.6.

to base prices as well:

$$p_{jt} = a_t \bar{p}_j, \ c_{jt}(\mathbf{n}_t) = b_t(\mathbf{n}_t)\bar{p}_j, \quad j \in J_t^P. \tag{1.6}$$

Let the "base value of primary tickets sold" be $\bar{q}_t = \sum_{j \in J_t^P} m_t s_{jt} \bar{p}_j$. Since the base price weights $\bar{p}_j$ are constant, $\bar{q}_j$ can be thought of as a quantity index, a way of summarizing the total quantity purchased at time t. The optimization problem becomes

$$V_t^P(\mathbf{\Omega}_t) \quad = \quad \max_{a_t} \bar{q}_t(a_t - b_t(\mathbf{n}_t)) + \sum_{j \in J_t^P} b_t(\mathbf{n}_t)\bar{p}_j n_{jt}.$$

Letting ε_t^P be residual primary demand elasticity with respect to a_t. This is the elasticity of the primary share with respect to a simultaneous proportional shift of all prices. The first order condition above yields the following Lerner's index rule:

$$\frac{a_t - b_t}{a_t} \quad = \quad -\left(\frac{\partial \bar{q}_t}{\partial a_t}\frac{a_t}{\bar{q}_t}\right)^{-1} = -\left(\frac{1}{\varepsilon_t^P}\right).$$

Note that ε_t^P is the franchise's residual demand, taking into account the pricing response of the secondary market:

$$\begin{aligned}
\frac{\partial \bar{q}_t}{\partial a_t} = \sum_{j \in J_t^P} m_t \bar{p}_j \frac{\partial s_{jt}}{\partial a_t} \quad &= \quad \sum_{j \in J_t^P} m_t \bar{p}_j \left(\frac{\partial s_{jt}}{\partial \mathbf{p}_t^{P'}}\frac{\partial \mathbf{p}_t^P}{\partial a_t} + \frac{\partial s_{jt}}{\partial \mathbf{p}_t^{S'}}\frac{\partial \boldsymbol{\sigma}_t^{DP}}{\partial \mathbf{p}_t^{P'}}\frac{\partial \mathbf{p}_t^P}{\partial a_t}\right) \\
&= \quad \sum_{j \in J_t^P} m_t \bar{p}_j \left(\frac{\partial s_{jt}}{\partial \mathbf{p}_t^{P'}}\bar{\mathbf{p}}^P + \frac{\partial s_{jt}}{\partial \mathbf{p}_t^{S'}}\frac{\partial \boldsymbol{\sigma}_t^{DP}}{\partial \mathbf{p}_t^{P'}}\bar{\mathbf{p}}^P\right)
\end{aligned}$$

If capacity is large enough that the probability of the constraint binding is 0, then $b_t = 0$ and the team simply maximizes daily revenue, which gives $\varepsilon_t^P = 1$ and $a_t = -\bar{q}_t\left(\frac{\partial \bar{q}_t}{\partial a_t}\right)^{-1}$ at the optimal a_t.

When I simulate dynamic pricing I will assume that the team is constrained to the family of pricing described above.

Finally, I will briefly describe the team's problem under fixed pricing, though there is no need to explicitly render this problem with notation. If the franchise were constrained to a fixed pricing regime then it would solve over a single vector of prices for each game, taking the alternative response function $\mathbf{p}_t^S = \boldsymbol{\sigma}_t^{FP}(\mathbf{p}^P, \mathbf{\Omega}_t)$ into account. The franchise would maximize its total expected revenue over all possible evolutions of the state $\boldsymbol{\Omega}_t$, given the state at the beginning of the season. When I simulate fixed pricing I will make the same assumption as in equation (1.6), except that a_t will be constant over t.

1.3.3 Secondary market supply

I present the following model of a secondary supplier for two reasons. First, the intuition for opportunity costs will be useful in instrumenting for price in the demand estimation. Second, it illustrates a feasible way to endogenize secondary market pricing in a subsequent version of this research.

Seller behavior follows the theoretical model in Sweeting [2012]. Letting the prime symbol (e.g. V') denote the value of a variable next period, I omit the search day subscript t for the remainder of this section. On a given search day, an optimizing risk-neutral seller with a single listing l will choose a price p_l to maximize his value, defined by the Bellman equation

$$V_l = \max_{p_l} p_l q_l \left(p_l, p_{-l}\right) + \left[1 - q_l(p_l, p_{-l})\right] EV_l', \tag{1.7}$$

where $q_l(p_l, p_{-l})$ is the probability that l sells in the current period and EV_l' is the *opportunity cost* of selling the listing, or the current-period expected value of still having listing l next period. The value and expected value functions are both functions of the state, but I leave this out for now. If the current period is the day of the event then $EV_l' = E_0 V_l^1 = V_l^1$ is the suppliers known value of being able to attend the event using the tickets in l. Under standard regularity conditions, a first-order condition shows that the optimal price p_l^* is equal to a markup plus the opportunity cost of sale:

$$p_l^* = \frac{q_l\left(p_l^*, p_{-l}\right) + \left[1 - q_l(p_l^*, p_{-l})\right]\left[\partial EV_l'/\partial p_l\right]}{\left|\partial q_l(p_l^*, p_{-l})/\partial p_l\right|} + EV_l'. \tag{1.8}$$

With many secondary market sellers offering tickets for any game on any search day, one can follow Pang, Berman, and Hu [2015] in assuming that $\partial EV_l/\partial p_l = 0$, at least as perceived by a secondary supplier.[35] Equation (2.2) becomes a simple Lerner Index

[35] Sweeting [2012] finds that the most likely violation of this assumption happens when a higher p_l may cause an interested buyer to wait, increasing future demand. As I will argue in section 1.5 that consumers are higly unlikely to delay purchase to a later date. Even allowing for this possibility, though,

Rule:

$$p_l^* = \frac{q_l\left(p_l^*, p_{-l}\right)}{\left|\partial q_l(p_l^*, p_{-l})/\partial p_l\right|} + EV_l'. \tag{1.9}$$

This formula provides the natural story that secondary market price is an opportunity cost plus a markup. The markup is expected to be small as $|\partial q_l(p_l^*, p_{-l})/\partial p_l|$ is expected to be large: with so many similar listings being offered, consumers should be highly elastic with respect to price among listings *within a product*.

In theory, $|\partial q_l(p_l^*, p_{-l})/\partial p_l|$ could be estimated and the opportunity costs EV_l' could be backed out for each period and listing. This approach could be valuable, as I describe in the following paragraph, but estimating *listing* demand, rather than product demand, turns out to be a difficult maximum-likelihood estimation that is still in progress.

One benefit of that approach is that, for any listing l still available on the day of the game, I would know V_l^1. This is the value to the seller of attending the game, and includes some primitive value plus the utility effect of the realized playoff probability on the day of the game. One could estimate EV_l' as a function of the primitive and the state variables. Because I have data on the Pittsburgh Pirates secondary market, which is "responding" to a fixed-pricing franchise, I would then have two opportunity cost functions, one for dynamic pricing and one for fixed pricing, estimable in counterfactuals. Equation (2.3) would then yield the two pricing functions $\boldsymbol{\sigma}^{DP}(\mathbf{p}_t^P, \boldsymbol{\Omega}_t)$ and $\boldsymbol{\sigma}^{FP}(\mathbf{p}_t^P, \boldsymbol{\Omega}_t)$ used in the primary supplier model above (Subsection 1.3.2).

For now, however, due to the difficulty of estimating $|\partial q_l(p_l^*, p_{-l})/\partial p_l|$, I assume that $\boldsymbol{\sigma}^{DP}(\mathbf{p}_t^P, \boldsymbol{\Omega}_t) = \boldsymbol{\sigma}^{FP}(\mathbf{p}_t^P, \boldsymbol{\Omega}_t) = \boldsymbol{\sigma}(\mathbf{p}_t^P, \boldsymbol{\Omega}_t)$ and I estimate $\boldsymbol{\sigma}(\mathbf{p}_t^P, \boldsymbol{\Omega}_t)$ simply using a linear regression on the states and on game and seating area fixed effects. This estimation method is detailed in Section 1.4.

1.4 Estimation

In this section I describe my approach to estimating demand parameters $\boldsymbol{\theta}$. I then explain the need to aggregate sales over multi-day periods prior to each game. I then describe how I calculate the exogenous market size, and thus am able to calculate shares over these multi-day periods. Finally, I present two estimation strategies making use of

---

a secondary supplier is unlikely to consider the negligible impact of his price on future demand: in 95% of game, search day combinations there are at least X listings on the secondary market.

the multi-day aggregation.

The first-pass strategy, used as a benchmark, assumes only one consumer type and redefines markets as game, multi-day-period-prior-to-game combinations. The second strategy is richer and more flexible: it returns to the original model specified in Subsection 1.3.1 where markets are game, day-prior-to-game combinations, but still relies on using total sales over multi-day periods. It also allows for two types of consumers.

The first strategy is the simple logit inversion method presented in Berry [1994] and again in Berry et al. [1995] (hereafter referred to as BLP for Berry, Levinsohn and Pakes). The second strategy closely resembles the generalized method of moments (GMM) approach taken by BLP, except that I allow the distribution of types to vary systematically over time and I use a slightly different set of "share-matching" moments.

1.4.1 Aggregating sales over multi-day periods

The approach taken in both Berry [1994] and BLP starts with matching the model's choice probabilities to the shares observed in the data *exactly*. These conditions can be thought of as moments on which the econometrician places infinite weight. These moments will not be viable given my data. If I were to use these moments right out of the box, they would be

$$s_{jgt}(\mathbf{p}_{gt}, \mathbf{x}_{gt}, \mathbf{n}_{gt}, \boldsymbol{\xi}_{gt}(\boldsymbol{\theta}); \boldsymbol{\theta}) - s_{jgt} = 0, \quad g = 1, \ldots, G, \ t = \bar{t}, \ldots, -1, \ j \in J_{gt}, \quad \text{(not viable)} \tag{1.10}$$

meaning that shares match for all markets and products. I would be solving for the vector $\boldsymbol{\xi}(\boldsymbol{\theta})$ of unobserved product attributes, whose existence and uniqueness is proven by Berry [1994] under mild regularity conditions.

But the estimation strategy implicitly assumes that the market size is large enough such the law of large numbers applies and observed shares equal the population's average choice probabilities. This assumption does not hold for my data: even after aggregating up to stadium area and source to define a "good," I observe many markets (g, t) and available goods $j \in J_{gt}$ such that the sales quantities in my data are $q_{jgt} = 0$, which is equivalent to $s_{jgt} = 0$. The true choice probabilities may indeed be quite small, but they cannot possibly be matched to zero in a logit choice model. The problem is that the sample is not large enough to measure the true choice probabilities accurately. And unfortunately, dropping such observations from the estimation will bias the sample.[36]

Recall from the descriptive evidence (Subsection 1.2.6) that my estimation sample includes $t = -28, \ldots, -1$. I now designate multi-day periods $\tilde{t} = -3, -2, -1$, and let $\tau_{\tilde{t}}$ denote the subset of $\{-28, \ldots, -1\}$ corresponding to $\tilde{t}$:

$$\begin{aligned}\tau_{-3} &= \{-28, \ldots, -15\},\\ \tau_{-2} &= \{-14, \ldots, -8\}, \text{ and}\\ \tau_{-1} &= \{-7, \ldots, -1\},\end{aligned}$$

Aggregating sales q_{jgt} across these multi-day periods almost entirely eliminates the "zeros" problem. Table 1.4 shows ticket sales by game, "search period," and stadium area, where the meaning of "search period" depends on the row: "search period translates to "day" for the first four rows and "multi-day period" for the last two rows. The table shows that by focusing on the estimation sample, and then aggregating across multi-day periods, I can nearly eliminate observations with low or zero sales. Specifically, the table says that

$$\sum_{t \in \tau_{\tilde{t}}} q_{jgt} \geq 4 \text{ for 95\% of the primary obs } (g, \tilde{t}, j) \in \{g = 1, \ldots, G, \, , \tilde{t} = -3, -2, -1, \, j = 1, \ldots, A\}$$

and

$$\sum_{t \in \tau_{\tilde{t}}} q_{jgt} \geq 18 \text{ for 95\% of the secondary obs } (g, \tilde{t}, j) \in \{g = 1, \ldots, G, \, , \tilde{t} = -3, -2, -1, \, j = A+1, \ldots, 2A\}$$

1.4.2 Market size assumption

As is common in the literature, I specify an exogenous market size based on the best information I have about the number of consumers who may be looking for a ticket on a given day leading up to a given game.

I use the average number of visits to game webpages of an anonymous franchise for

[36] Suppose all these "zero goods" are unpopular because they share some characteristic that particularly repels consumers. Not including these in the sample means the econometrician will estimate a taste for this characteristic that is biased toward zero. Gandhi et al. [2013] provides a more thorough explanation of this problem.

Table 1.4 Summary Statistics for Game, Search Period, Available Product Observations

	Primary Market					Secondary Market				
	Obs.	Mean	S.D.	5th %	95th %	Obs.	Mean	S.D.	5th %	95th %
Full Sample, Daily										
# Tickets Remaining	184761	1265	1468	13	4724	185270	162	330	2	717
# Tickets Purchased	182819	15	49	0	62	182605	7	16	0	30
Estimation Sample, Daily										
# Tickets Remaining	23694	563	830	6	2214	24072	67	141	2	250
# Tickets Purchased	23694	23	56	0	95	24072	17	24	0	59
*Estimation Sample, Aggregate**										
# Tickets Remaining	2567	723	932	12	2615	2586	200	280	23	660
# Tickets Purchased	2567	216	295	4	805	2586	155	138	18	447

Note: I have limited sales data for the day of the game, hence the discrepancy between inventory and purchases in number of observations.

* As in table 1.2, the time periods are the following intervals of search days:
[-112, -57], [-56, -29], [-28, -15], [-14, -8], [-7,-1], and 0.

days $t = -14, \ldots, -1$, as provided by Zhu [2014].[37] His averages lie nearly exactly on the curve

$$m_{0t} = 750 + \frac{3000}{-t+1}. \tag{1.11}$$

I scale this formula according to pop_g, the population around the stadium where game $g \in G$ is played. Specifically, I use the population of the Designated Market Area (DMA) surrounding the stadium, which approximates the number of consumers who think of that stadium's team as their "local" team. I scale the asymptotic minimum from equation (1.11) to 0.1% of DMA population, or

$$m_{gt} = .001\,(pop_g)\left(\frac{m_{0t}}{750}\right) = pop_g\left(.001 + \frac{.004}{-t+1}\right), \tag{1.12}$$

yielding minimum market sizes ranging from 555 for the Orioles to 1,258 for the Giants and Athletics.

I then use the sales quantities q_{jgt} in my data to calculate the "observed" shares. Observed shares are

$$s_{jgt} = \frac{q_{jgt}}{m_{gt}}, \quad g = 1, \ldots, G,\ t = -28, \ldots, -1,\ j = 1, \ldots, 2A.$$

More importantly for my estimation, observed *multi-day period* shares are

$$\tilde{s}_{jg\tilde{t}} = \frac{\sum_{t \in \tau_{\tilde{t}}} q_{jgt}}{\sum_{t \in \tau_{\tilde{t}}} m_{gt}}, \quad g = 1, \ldots, G,\ t = -28, \ldots, -1,\ j = 1, \ldots, 2A.$$

1.4.3 Preliminary, aggregate-market demand estimation

As a first pass, I allow only one consumer type and I distort the demand model by letting a "market" consist of tickets offered/chosen over each multi-day period $\tilde{t}$. I will refer to this alternative model as the aggregate-market demand model. Specifically, I make the following assumptions:

Assumption 2. *The number of consumer types is $R = 1$.*

[37] This market size function doubles from $t = -7$ to $t = -1$, and increases by 67% from $t = -1$ to $t = 0$ (day of game). This appears to be a realistic assumption, given the typical acceleration in sales that I observe in that final week.

Assumption 3. *All consumers in period $\tilde{t}$ face the same choice set $J_{g\tilde{t}} = \cup_{t \in \tau_{\tilde{t}}} J_{gt}$, the same prices $\tilde{\mathbf{p}}_{\tilde{t}}$, and the same game characteristics $\tilde{\mathbf{x}}_{\tilde{t}}$. Let $\tilde{p}_{j\tilde{t}}$ be j's mean transacted price and $\tilde{\mathbf{x}}_{\tilde{t}}$ be j's mean transacted characteristics over $t \in \tau_{\tilde{t}}$.*

Assumption 4. *The unobserved attributes ξ_{jt} do not change over aggregate periods. Let $\tilde{\xi}_{j\tilde{t}g}$ denote its common value, so that $\xi_{jtg} \equiv \tilde{\xi}_{j\tilde{t}g}\ \forall t \in \tau_{\tilde{t}}$, $g = 1, \ldots, G$, $j \in J_{g\tilde{t}}$.*

I will omit the g subscript for the remainder of this subsection and in subsection 1.4.4. Predicted choice probabilities are simply

$$\tilde{s}_{j\tilde{t}}(\boldsymbol{\delta}_{\tilde{t}}, \mathbf{n}_{\tilde{t}}) = \frac{\exp\left(\delta_{j\tilde{t}}\right)}{\sum_{k \in J_{\tilde{t}}} \exp\left(\delta_{k\tilde{t}}\right)} \quad \text{where} \quad \delta_{j\tilde{t}} = \alpha \tilde{p}_{j\tilde{t}} + \beta_0 + S_j\beta_1 + \mathbf{x}_j\beta_2 + \tilde{\mathbf{x}}_{\tilde{t}}\beta_3 + \xi_{j\tilde{t}} \tag{1.13}$$

The share-matching illustrated in Equation (1.10) now becomes $\tilde{s}_{j\tilde{t}}(\boldsymbol{\delta}_{\tilde{t}}, \mathbf{n}_{\tilde{t}}) - \tilde{s}_{j\tilde{t}} = 0,\ t = \bar{t}, \ldots, -1,\ j \in J_t$, which amounts to solving for the unique vector of average utilities $\boldsymbol{\delta}$. In this particular case, it can be solved for using the standard logit inversion given in Berry [1994]:

$$\delta_{j\tilde{t}} = \ln(\tilde{s}_{j\tilde{t}}) - \ln(\tilde{s}_{0\tilde{t}}), \quad \tilde{t} = -3, -2, -1,\ j \in J_{\tilde{t}} \tag{1.14}$$

Once I have the $\delta_{j\tilde{t}}$ terms I regress on the aggregate characteristics to estimate the demand parameters in equation (1.13).

There are two unsatisfactory elements of this estimation strategy. First, Assumption 3 disagrees with reality: this assumption says that consumers shopping in period $\tilde{t}$ choose among $J_{\tilde{t}}$, the set of all products available on some day $t \in \tau_{\tilde{t}}$. In reality, different consumers face different choice sets over the aggregate period $\tilde{t}$—some goods do sell out mid-period. Moreover, the true price and game characteristics *are not* $\tilde{p}_{j\tilde{t}}$ and $\tilde{x}_{\tilde{t}}$. Even if no sell-outs ever occurred, so that $J_{gt} = J_{g\tilde{t}}$ over $t \in \tau_{\tilde{t}}$, the demand function is nonlinear: demand for mean products does not equal mean demand of actual products. In other words, this model does not produce the correct predictions to match to $s_{j\tilde{t}}$.[38]

[38] To be fair, using mean prices and mean characteristics is often the only option, and is therefore quite common in the literature. And I am still guilty of this offense in that I use mean-transacted secondary market prices as my daily product prices for $j = A + 1, \ldots, 2A$.

1.4.4 Two-type daily-market demand estimation

In this preferred specification, whose results I use to simulate pricing policies in Section 1.6, I keep Assumption 4 but drop Assumptions 2 and 3: I continue to assume that ξ_{jt} is the same for all $t \in \tau_{\tilde{t}}$, but now $R = 2$ and consumers face the actual, daily products in the data. By allowing for two types of consumers, the model can explain patterns in how product shares change over time, even when characteristics and choice sets stay the same. By using daily markets, I now take into account mid-period sell-outs and mid-period changes in prices and characteristics.

Probability on type

I will assume, without loss of generality, that Type 2 consumers tend to shop closer to a game than Type 1 consumers, so that the probability of being a type 2 consumer is monotone increasing in t. My justification for the monotonicity assumption is a hypothesis that, as in airlines, higher-willingness-to-pay consumers will tend to shop closer to a game. Two possible stories might support this hypothesis. One story is that, as in airlines, these type 2 consumers are business consumers: a company entertains a client from out of town and spares no expense to get good seats. This story may be dubious, however, as businesses tend to hold season tickets. Perhaps a more convincing story is that type 1 consumers are lower-income rural consumers who, due to their distance from the stadium, must plan ahead and shop earlier, while type 2 consumers are higher-income urban professionals for whom it makes more sense to shop closer to the game.

Recall that γ_{rgt}, the fraction of consumers shopping for game g tickets on day-prior-to-game t who are type $r = 1, 2$, was assumed to result from a function with parameters $\boldsymbol{\phi}$. I now specify γ_{2gt} to be the output of a logistic function of t:

$$\gamma_{2gt} = \frac{1}{1 + \exp(-\phi_1(t - \phi_{0l(g)}))}, \quad \gamma_{1gt} = 1 - \gamma_{2gt} \tag{1.15}$$

where $l(g)$ denotes the stadium location of g. Since there are four stadiums in my sample, the vector $\boldsymbol{\phi}_0$ has four elements. Each element is the logistic function's stadium-specific midpoint, the point prior to the game at which half of consumers are Type 1 and half are Type 2. The scalar ϕ_1 sets the function's steepness. The type-probability parameter vector is $\boldsymbol{\phi} = (\phi_0', \phi_1)'$. Note that all four stadiums are assumed to face the same two types of consumers, but the fraction of consumers who are of a particular type on a particular day prior to the game depends on the stadium location. Without loss of generality, let $\phi_1 \geq 0$ so that Type 2 consumers tend to arrive later than Type 1 consumers. This model is flexible, however: it allows for $\phi_1 = 0$, meaning a static type distribution, and it allows Type 2 consumers to be more or less price sensitive than Type 1 consumers.

BLP Moments

The model's predicted choice probabilities over multi-day periods are

$$\begin{aligned}\tilde{s}_{j\tilde{t}}\left((\mathbf{p}_t, \mathbf{x}_t)_{t\in\tau_{\tilde{t}}}, \boldsymbol{\xi}_{\tilde{t}}, \mathbf{n}_t; \boldsymbol{\theta}\right) &= \frac{1}{\sum_{t\in\tau_{\tilde{t}}} m_t} \sum_{t\in\tau_{\tilde{t}}} m_t s_{jt}\left(\mathbf{p}_t, \mathbf{x}_t, \xi_{\tilde{t}}, \mathbf{n}_t; \boldsymbol{\theta}\right) \\ &= \frac{1}{\sum_{t\in\tau_{\tilde{t}}} m_t} \sum_{t\in\tau_{\tilde{t}}} m_t \sum_{r=1}^{2} \gamma_{rt} s_{rjt}\left(\mathbf{x}_t, \mathbf{p}_t, \boldsymbol{\xi}_{\tilde{t}}, \mathbf{n}_t; \boldsymbol{\theta}\right). \end{aligned} \quad (1.16)$$

The share-matching moments illustrated in Equation (1.10) now become

$$\tilde{s}_{j\tilde{t}}\left((\mathbf{p}_t, \mathbf{x}_t)_{t\in\tau_{\tilde{t}}}, \boldsymbol{\xi}_{\tilde{t}}(\boldsymbol{\theta}), \mathbf{n}_t; \boldsymbol{\theta}\right) - \tilde{s}_{j\tilde{t}} = 0, \quad \tilde{t} = -3, -2, -1,\, j \in J_{\tilde{t}}. \quad (1.17)$$

This moment matching is equivalent to solving for the vector $\boldsymbol{\xi}(\boldsymbol{\theta})$ of unobserved attributes for a given parameter vector $\boldsymbol{\theta}$. Existence and uniqueness can be proven using a similar argument to that of Berry [1994].[39]

The second set of moments assume that at the true parameter vector $\boldsymbol{\theta}_0$, the unobserved attribute $\xi_{jt}(\boldsymbol{\theta}_0)$ for any product j on day-prior-to-game t (for a given game) is uncorrelated with a vector of instruments $\mathbf{z}_{jt}$, or

$$E[\xi_{jt}(\boldsymbol{\theta}_0)|\mathbf{z}_{jt}] = 0. \quad (1.18)$$

Product j's non-price characteristics are valid instruments for themselves. For prices, which are likely correlated with ξ_{jt}, I use two vectors of additional instruments: $\boldsymbol{\zeta}_{jt}^P$ applies exclusively to products on the primary market (i.e. $\boldsymbol{\zeta}_{jt}^P = \mathbf{0}$ if $S_j = 1$) and $\boldsymbol{\zeta}_{jt}^S$

[39] Details on the necessary regularity conditions will be in a subsequent version of this paper.

applies exclusively to products on the secondary market (i.e. $\boldsymbol{\zeta}_{jt}^{S} = \mathbf{0}$ if $S_j = 0$). These are described in Subsubsection 1.4.4. The vector of instruments is therefore

$$\mathbf{z}_{jt} = \left(1, S_j, \mathbf{x}_j', \mathbf{x}_t', \boldsymbol{\zeta}_{jt}^{P\prime}, \boldsymbol{\zeta}_{jt}^{S\prime}\right)'.$$

Objective function

Define the set of markets as $M = \{1, \ldots, G\} \times \{-28, \ldots, -1\}$ and the set of data observations as $D = \{(g, t, j) : (g, t) \in M, j \in J_{gt}\}$. Let $\mathbf{w}_{jgt} = (s_{jgt}, \mathbf{x}_{jgt}, \mathbf{z}_{jgt})$ be the data for some observation $(g, t, j) \in D$. Letting

$$\mathbf{h}(\mathbf{w}_{jgt}, \boldsymbol{\theta}) = \xi_{jtg}(\boldsymbol{\theta})\mathbf{z}_{jtg},$$

where $\xi_{jtg}(\boldsymbol{\theta})$ comes from Equation (1.17), the following population moment conditions hold at the true $\boldsymbol{\theta}_0$:

$$\mathbf{G}(\boldsymbol{\theta}_0) \equiv E[\mathbf{h}(\mathbf{w}_{jgt}, \boldsymbol{\theta}_0)] = 0, \;\; \forall\, (g, t, j) \in D.$$

Hansen [1982] shows that the optimal (two-step) GMM estimator takes the form

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta} \in \Theta} \mathbf{G}^*(\boldsymbol{\theta})'\mathbf{G}^*(\boldsymbol{\theta}), \tag{1.19}$$

where $\mathbf{G}^*(\boldsymbol{\theta}) = \mathbf{a}(\tilde{\boldsymbol{\theta}})\hat{\mathbf{G}}(\boldsymbol{\theta})$.[40] $\hat{\mathbf{G}}(\boldsymbol{\theta})$ is the sample analogue to $\mathbf{G}(\cdot)$,

$$\hat{\mathbf{G}}(\boldsymbol{\theta}) = \frac{1}{|D|} \sum_{(g,t,j)\in D} \mathbf{h}(\mathbf{w}_{jgt}, \boldsymbol{\theta}_0),$$

and $\mathbf{a}(\tilde{\boldsymbol{\theta}})$ is a consistent estimate of the "square root" of the inverse of the asymptotic variance-covariance matrix of the moments (obtained using a preliminary consistent estimate $\tilde{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}_0$).[41] Let $\mathbf{V} = E[\mathbf{h}(\mathbf{w}_{jgt}, \boldsymbol{\theta}_0)\mathbf{h}(\mathbf{w}_{jgt}, \boldsymbol{\theta}_0)']$, the variance-covariance matrix of $\mathbf{h}(\mathbf{w}_{jgt}, \boldsymbol{\theta}_0)$ with respect to the true parameter values. Let $\boldsymbol{\Gamma} = E\left[\frac{\partial \mathbf{h}(\mathbf{w}_{jgt}, \boldsymbol{\theta}_0)}{\partial \boldsymbol{\theta}'}\right]$, the gradient of the moments with respect to the true parameter values. Asymptotic variance of $\sqrt{|D|}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)$ is given by

$$(\boldsymbol{\Gamma}'\boldsymbol{\Gamma})^{-1}\boldsymbol{\Gamma}'\mathbf{V}\boldsymbol{\Gamma}(\boldsymbol{\Gamma}'\boldsymbol{\Gamma})^{-1}.$$

In the following section I report standard errors by estimating $\boldsymbol{\Gamma}$ and $\mathbf{V}$ using consistent estimates

$$\hat{\boldsymbol{\Gamma}}(\hat{\boldsymbol{\theta}}) = \frac{1}{|D|} \sum_{(g,t,j)\in D} \frac{\partial \mathbf{h}(\mathbf{w}_{jgt}, \hat{\boldsymbol{\theta}})}{\partial \boldsymbol{\theta}} \quad \text{and} \quad \hat{\mathbf{V}}(\hat{\boldsymbol{\theta}}) = \frac{1}{|D|} \sum_{(g,t,j)\in D} \mathbf{h}(\mathbf{w}_{jgt}, \hat{\boldsymbol{\theta}})\mathbf{h}(\mathbf{w}_{jgt}, \hat{\boldsymbol{\theta}})'.$$

[40] In practice, I solve the GMM minimization using mathematical program with equilibrium constraints (MPEC), shown to be consistent in Dubé et al. [2012]. This method turned out to be faster than using the Berry [1994] inversion (though at one point, I did run the estimation using the inversion once in C++ and got very similar parameter estimates). Define

$$\mathbf{g}_D(\boldsymbol{\xi}) = \frac{1}{|D|} \sum_{(g,t,j)\in D} \xi_{jtg}\mathbf{z}_{jtg}.$$

For a given weighting matrix $\mathbf{W}$, the estimation procedure equates to solving the minimization problem

$$\begin{aligned} &\min_{\theta} \boldsymbol{\eta}'\mathbf{W}\boldsymbol{\eta} \\ &\text{subject to } \mathbf{g}_D(\boldsymbol{\xi}) = \boldsymbol{\eta}, \\ &\qquad\qquad \tilde{\mathbf{s}}(\boldsymbol{\xi}; \boldsymbol{\theta}) = \tilde{\mathbf{s}}^{actual}. \end{aligned}$$

[41] Specifically,

$$\tilde{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}\in\Theta} \hat{\mathbf{G}}(\boldsymbol{\theta})'\mathbf{S}_{\mathbf{zz}}^{-1}\hat{\mathbf{G}}(\boldsymbol{\theta})$$

where $\mathbf{S}_{\mathbf{zz}}$ is the sample analogue of $\boldsymbol{\Sigma}_{\mathbf{zz}}$, the variance-covariance matrix for the instrumental variables. This yields a two-stage least-squares estimate of $\boldsymbol{\theta}$, $\tilde{\boldsymbol{\theta}}$. The resulting estimate of the unobserved attributes $\hat{\boldsymbol{\xi}}$ is then used to calculate $\hat{\mathbf{S}} = \frac{1}{|D|}\sum_{(g,t,j)\in D} \mathbf{h}(\mathbf{w}_{jgt}, \tilde{\boldsymbol{\theta}})\mathbf{h}(\mathbf{w}_{jgt}, \tilde{\boldsymbol{\theta}})'$, the sample analogue of $\boldsymbol{\Sigma}_{\mathbf{hh}}$. $\mathbf{a}(\tilde{\boldsymbol{\theta}})$ is the "square root" of $\hat{\mathbf{S}}^{-1}$.

Characteristics and instruments

Area-specific characteristics $\mathbf{x}_{jg}$ are

- $1\{$area is on first floor of stadium$\}$ and
- distance from home plate to seating area, in hundreds of feet;

while game characteristics $\mathbf{x}_{tg}$ are

- $1\{$night game$\}$, $1\{$weekend game$\}$,
- $1\{$division rival$\}$, $1\{$geographic rival$\}$, $1\{$interleague rival$\}$,
- game date relative to start of season, and
- home and away teams' probabilities of entering playoffs (only these vary across t)

where $1\{\cdot\}$ is the indicator function.

For the secondary market instruments $\boldsymbol{\zeta}_{jt}^{S}$ I use three instruments that correlate with the seller's opportunity cost of sale, and hence his optimal price, and one instrument that proxies for the level of within-product competition:

- *The search date t.* All else equal, a secondary seller's opportunity cost of selling a ticket should decrease as the event approaches.
- *The average listing start date of j-type tickets.* Sellers who post earlier likely have less interest in attending the game, or lower opportunity costs.
- *Indicators of night game and weekend game*, as these sellers are more likely to be able to attend the game.
- *Number of listings within product j.* The more listings, the more likely a seller's product is not well-differentiated, the smaller the markup above opportunity cost.

Finding strong instruments for primary market prices, $\boldsymbol{\zeta}_{jt}^{P}$, presented a greater challenge. Because areas of the stadium so rarely sell out before the last day (which is not included in my sample), the team's opportunity costs are unlikely to vary much as the event approaches. Most price movement is strictly about revenue maximization, where unobserved attributes play a role. However, the team should price lower when there is more competition from the secondary market, so I use the optimal instruments derived by BLP, applied to the constant term and the product-specific characteristics $\mathbf{x}_j$ only. Specifically, the vector of BLP instruments for j in market t are[42]

$$\boldsymbol{\zeta}_{jt}^{P} = \left(\left| J_t^P - 1 \right|, \left| J_t^S \right|, \sum_{j' \in J_t^P, j' \neq j} \mathbf{x}_{j'}', \sum_{j' \in J_t^S} \mathbf{x}_{j'}' \right)'$$

I do not use analogous instruments for the game-level characteristics, nor do I use them for secondary market sellers, as these cause collinearity issues.[43]

1.5 Results

Identification of these type-probability parameters arises from changes in shares over time in the data. Even with the exogenous market size that more than doubles from $t = -14$ to $t = -1$, purchases tend to increase even faster during this period, even when prices are flat or increasing. The estimation rationalizes this observation by allowing the later-arriving type to have a lower price-sensitivity.

Table 1.5 presents the results of simple logit demand estimations. Again, these estimations use "aggregate periods" where product characteristics that vary day to day (price, home/away playoff probability) are simply summarized by their transaction-weighted average.

Table 1.6 follows, and presents the results of the GMM two-type consumer model estimation. Again, here I use daily product characteristics but match aggregate period predicted shares to actual shares to get aggregate-period product-level unobservables that go into the moment conditions of the GMM estimation.

[42] The instruments used in Berry et al. [1995] assume a Bertrand-Nash pricing game. The qualities of other within-firm products and the qualities of other products $j' \neq j$ are assumed to be uncorrelated with the unobserved attribute ξ_{jt} but affect the degree of competition and therefore correlate with price p_{jt}.

[43] The first instrument obviously cannot be used for secondary market products: Sellers each act independently, so there are no other within-seller products. The sum of x_{rkt} over other sellers' available products is simply $\sum_{r \in J_t^S} x_{krt}$, since the seller's product is almost always offered by another seller as well. This value is nearly collinear with the constant term if only one stadium location is used in estimation, since the choice set rarely varies and these are static seat characteristics.

Table 1.5 Demand Estimation Results

	Primary/Secondary			Secondary Only	
	OLS logit	IV logit	IV logit	OLS logit	IV logit
Seat-Level					
price	-0.000255	-0.00387**	-0.00489**	-0.000295	-0.0176***
	(0.000170)	(0.00130)	(0.00151)	(0.000151)	(0.00392)
secondary market	-0.0105	0.133*	0.171**		
	(0.0469)	(0.0591)	(0.0617)		
Game-Level					
playoff probability	0.0138***	0.0125***	0.0129***	0.0195***	0.00552
	(0.00292)	(0.00289)	(0.00307)	(0.00348)	(0.00487)
away's playoff probability	0.00177	0.00233*	-0.00121	0.00000801	0.00343
	(0.00114)	(0.00115)	(0.00419)	(0.00134)	(0.00259)
division rival	-0.0177	0.0192		0.0290	0.367
	(0.0797)	(0.0830)		(0.0903)	(0.241)
geographic rival	-0.207	-0.0395		-0.130	1.093*
	(0.127)	(0.140)		(0.152)	(0.437)
interleague rival	0.250*	0.303*		0.429**	0.731**
	(0.125)	(0.125)		(0.144)	(0.266)
other†	Yes	Yes	Yes	Yes	Yes
away team FE	No	No	Yes	No	No
median elasticity	-0.0149	-0.226	-0.286	-0.0172	-1.031
N	7066	7066	7066	3553	3553
R^2	0.147			0.179	

Standard errors in parentheses

†Includes home team dummy variables, date of game, and indicators for night and weekend games.

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

Table 1.6 Demand Estimation Results

	Primary/Secondary, 2 Types‡	
	Type 1	Type 2
Seat-Level		
price	-0.0826***	-0.0114
	(0.0106)	(0.0148)
secondary market	-2.493***	2.847***
	(0.270)	(0.451)
Game-Level		
playoff probability	0.0106***	
	(0.000962)	
away's playoff probability	0.00489	
	(0.0134)	
division rival	-0.211	
	(0.113)	
geographic rival	0.463***	
	(0.107)	
interleague rival	0.377***	
	(0.0219)	
other†	Yes	Yes
away team FE	No	No
median elasticity	-6.482	-0.893
N	5153	

Standard errors in parentheses

†Includes home team dummy variables, date of game, and indicators for night and weekend games.

‡These columns are one GMM estimation, solved as an MPEC problem using the KNITRO solver (see Subsection 1.4.4 for further details). It uses daily markets, unlike the other estimations which use aggregate-period markets. Estimated arrival parameters:

$$\phi_1 = 0.179,\ \phi_{0,Cardinals} = -8.434,\ \phi_{0,Giants} = -22.291,\ \phi_{0,Pirates} = -19.998,\ \phi_{0,Twins} = -8.219.$$

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

Referring back to equation (1.1), coefficients in this estimation represent tastes for characteristics, specifically the marginal effects of characteristics on average utility δ_{rjt}. The signs are as expected for nearly all specifications. "Home team playoff probability" is significant as positive, but I cannot reject the null hypothesis that the coefficient on "away team playoff probability" is zero. The first two columns show that using instruments to control for price endogeneity significantly increases the magnitude of the price coefficient. The third column shows that using away team fixed effects, rather than the division rival, geographic rival, and interleague rival indicator variables also appears to reduce the endogeneity problem but to a lesser degree. Columns four and five show that selecting only secondary market data and applying instruments results in strong price coefficients. This is not surprising as the instruments do more to control for secondary market price endogeneity than they do for primary market price endogeneity. Of course, using only this market could present a selection bias. The final two columns show the results of my preferred specification, and in this estimation I include both primary and secondary market data. Given the size of the median elasticities, the reduced bias of share predictions associated with the 2-type specification with daily markets seems to compensate for the limited quality of primary market instruments for price.

The demand estimation shown Table ? is estimated by solving estimated by solving the minimization problem in (1.19). It differs from the specifications in Table 1.6 in that (1) there are two consumer types with different arrival patterns and (2) for a given θ I predict daily sales and aggregate up according to equation (1.16). (1) means this model can be used to think about intertemporal price discrimination, while (2) means, among other things, that I treat sellouts in a more correct manner: if product j sells out midway through aggregate period $\tilde{t}$, the model calculates predicted quantity using only those periods when j was available.[44] It is simply not possible to accurately deal with sellouts in simple logit specifications where products are the "mean product" over an aggregate period.

The first two-type estimation finds the type 2 consumers, those who tend to arrive later, do indeed have lower price sensitivity. Given that the largest (in magnitude)

[44] Expected sales on the day of the sellout may remain slightly "over-predicted" for a given $\boldsymbol{\theta}$. One could, in principle, eliminate this problem by excluding the exact day of the sellout from estimation. This is a task for a future revision.

median elasticity of the previous models was -1.031, a median elasticity of -6.482 for the Type 1 and -.893 for Type 2 appears reasonable. I also find that later-arriving consumers harbor a bias toward StubHub while earlier-arriving consumers are biased against it. In fact, these coefficients are so large (one can see this with a quick-and-dirty "conversion" from utils to dollars using the coefficients on price) that the reality is likely that consumers completely ignore one market when they make consumption decisions. This behavior could be rationalized by the presence of a search cost: consumers look at whichever market they think is likely to have lower prices; consumers who usually shop early would know that Tickets.com is usually the better option while consumers who usually shop late would know that StubHub is usually the better option.[45]

The second two-type estimation, where the seating-area-specific characteristics are seating area dummy variables rather than distance to home plate and a first floor dummy, yielded very similar results. The only significant difference was in the StubHub coefficient, now roughly 1.5 in both specifications. Even in an alternative specificatoin where both types are constrained to have equal states for these seating area dummies (I was concerned about overfitting), this result still held: the two consumer types value the "StubHub label" equally.

Finally, observe that the home team's playoff probability is a significant determinant of utility.

Recall that in the preferred specification I also estimate type-probability parameters $\boldsymbol{\phi}$, and that I allow for a separate function for each franchise. These estimates are in the footer of Table (1.6). Figure (1.6) illustrates the probability functions associated with these parameters, while Figure (1.7) shows the market size, broken down by type.

Observe that Type 2 consumers begin arriving much sooner for the Giants than for the Cardinals. This difference could largely explain the different price paths shown in Figure 1.3. The next section seeks to more rigorously understand the value of price discrimination, relative to product changes, in DP in this industry.

[45]In future work I plan to estimate an alternative demand model in which some consumers look at only one of the two markets.

Figure 1.6 Type-probability Functions

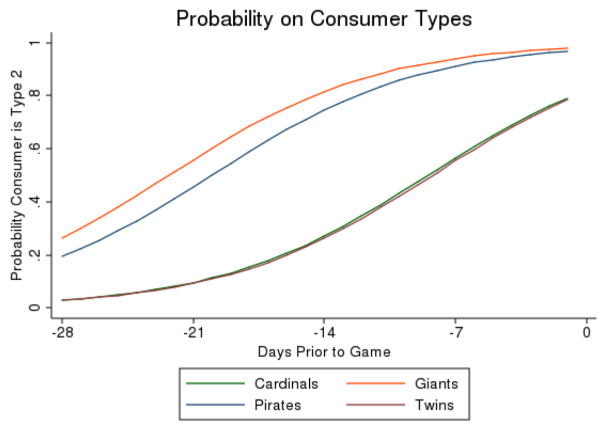

1.6 Analysis of the estimated model

The first goal of this section is to conduct a welfare analysis: to measure the revenue gains, and consumer welfare effects, from using dynamic pricing, relative to uniform pricing across time. The second is to explore how important the change in consumer mix over time is relative to the importance of product changes in generating the revenue gains. In this section, "product changes" refers to changes in the home team's playoff probability: the away team's playoff probability is not found to be significant in the demand estimation (Table 1.6), so I do not focus on it here.

Capacity constraints are unlikely to play a role for most baseball games, which do not even approach a sell-out. Allowing for such constraints to play a role introduces many computational challenges: rigorously simulating dynamic pricing requires computing the value function in Equation (1.4) many times, and size of the state space (which would include inventories on the primary and secondary markets, as well as the home team's playoff probability) yields a "curse of dimensionality." In Subsection **??** of the appendix I present second-best solution to this problem: I approximate optimal dynamic pricing, when capacity constraints matter, by searching over a family of pricing functions with a small number of parameters.

Instead, I assume that capacity is unlimited, both on the primary and secondary markets. Even if capacity did often play a role in baseball games, there is nothing wrong

Figure 1.7 Populations of consumers by type

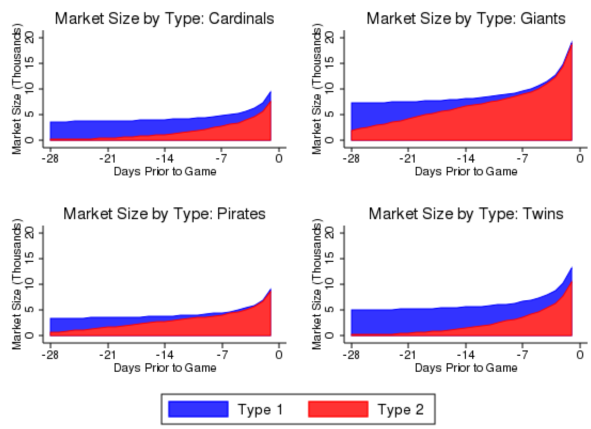

with abstracting away from it in order to understand the roles of price discrimination and product changes in isolation.

The secondary market pricing response function $\sigma(\cdot)$ is essentially the same as the one estimated in Table 1.3, except that I do not include the away team's playoff probability. I estimate the variance of the error term ε_σ for use in the simulations below.

I first construct a "typical game." Arbitrarily, I choose to use the Giants stadium areas, with their first floor dummies and distances to home plate, as my products; however all other characteristics are averages. For example, the "night game" variable is 0.47 rather than being 0 or 1. The game therefore has the average utility, relative to the outside good, over all games I observe. The secondary market pricing function is "average" in its area fixed effects. Finally, I calculate the "average face value," meaning average earliest-observed price, for each section of the Giants stadium, scaled by the overall average over all teams. Let these face values be denoted by $\bar{p}_j$ for products on the primary market $j \in J_t^P$.

I assume that playoff probability is a martingale sequence and estimate the transition probability $f(x_t|x_{t-1})$ by parametrizing it and doing a maximum likelihood estimation (MLE) using observed playoff probability paths of all 30 teams over the 2014 season. Details are provided in Subsection **??** of the Appendix. I then simulate 1000 playoff probability paths, where all paths start at $x_{\bar{t}} = 50$. Four of the simulated transition paths are shown in Figure 1.8. These paths are quite similar to the behavior of the

Figure 1.8 Simulated Playoff Probability Paths

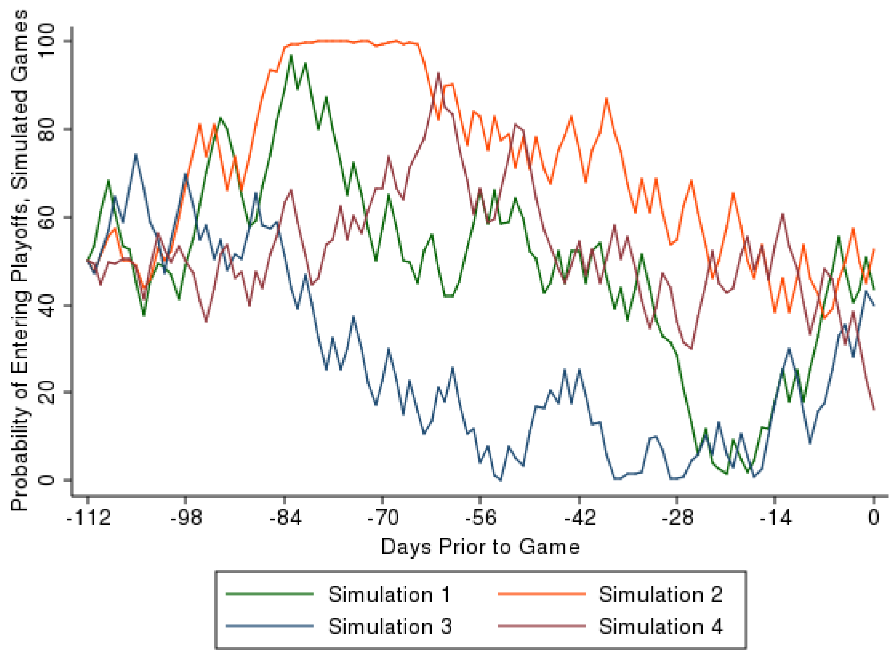

actual, observed paths (see Figure 1.5).

I also simulate 1000 draws of the secondary market policy function error term $\varepsilon^{\sigma}_{jtg}$ and of the unobserved attributes ξ_{jt}, based on their estimated standard deviations.

I assume that the franchise keeps the ratios between the prices of the various stadium areas the same, so that prices for a given simulated game g are

$$p_{jt} = a_{tg}\bar{p}_j, \quad j \in J^P. \tag{1.20}$$

Note that available primary products are denoted by J^P rather than J^P_t, to reflect the fact that capacity is unchanging in these simulations.

Because current pricing does not affect the potential revenue in future periods, "dynamic pricing" simply means solving for optimal price multipliers a_{tg} each market (g,t) separately. "Fixed pricing" means that $a_{tg} = a$ for all (g,t): prices are scaled by the same factor in all periods, for all games. By maximizing the average revenue over all games, I find the fixed price that the team would choose to maximize its expected revenue. Finally, "date-of-purchase pricing" means that $a_{tg} = a_t$ for all (g, t): the scale factors can differ across time, but not across games. Effectively, the maximization finds the price path, before knowing how its playoff probability will evolve, that maximizes its expected revenue.

First I simulate dynamic pricing under two simple counterfactuals, one where the consumer distribution is fixed over time and one where playoff probability is fixed over time, in order to see which one tends to produce larger price changes over time. In order to be sure that I'm only picking up price changes due to playoff probability or due to the consumer mix, I set $\xi_{jtg} = 0$ and $\varepsilon^{\sigma}_{jtg} = 0$, for all j, t, g.

The comparison between price movements depends on the time horizon. Comparing prices 7 days from the game going forward, the consumer mix influences price, relative to the $t = -7$ price, much more strongly on the day of the game, but comparing prices 112 days from the game going forward the two factors have nearly equal influence on the absolute percentage in price change by the day of the game. Figure 1.9 shows the average absolute factor change in price from the starting day, for four different starting days. For the counterfactual where the consumer distribution is fixed, price is just as likely to move down as up, while for the counterfactual where products are fixed, price only moves up. Because playoff probability can only move as high as 100 or as low as 0 and the taste parameter on playoff probability is small, in the final 3-4 days it takes a backseat to consumer distribution changes a motivator for moving price. Notice that playoff probability "matters" all the time, while consumer shifts only influence price in the final two weeks, due to the type probability function of this hypothetical "average" team (recall Figures 1.6 and 1.7).

I then take the final two weeks before the game and solve for optimal dynamic pricing and uniform pricing over time under these counterfactuals, except that I use the simulated draws of unobserved attributes ξ_{jtg} and idiosyncratic resale pricing $\varepsilon^{\sigma}_{jtg}$. These two factors, taken separately in this way, also have nearly equal impact on the revenue gain achieved by dynamic pricing, as compared to uniform pricing over time, when I look at the final two weeks. In the world where the consumer mix stays constant over time, the use of dynamic pricing leads, on average, to a 5.14% increase in revenue relative to uniform pricing over time. In the world where all product characteristics

Figure 1.9 Absolute price change factor from starting day

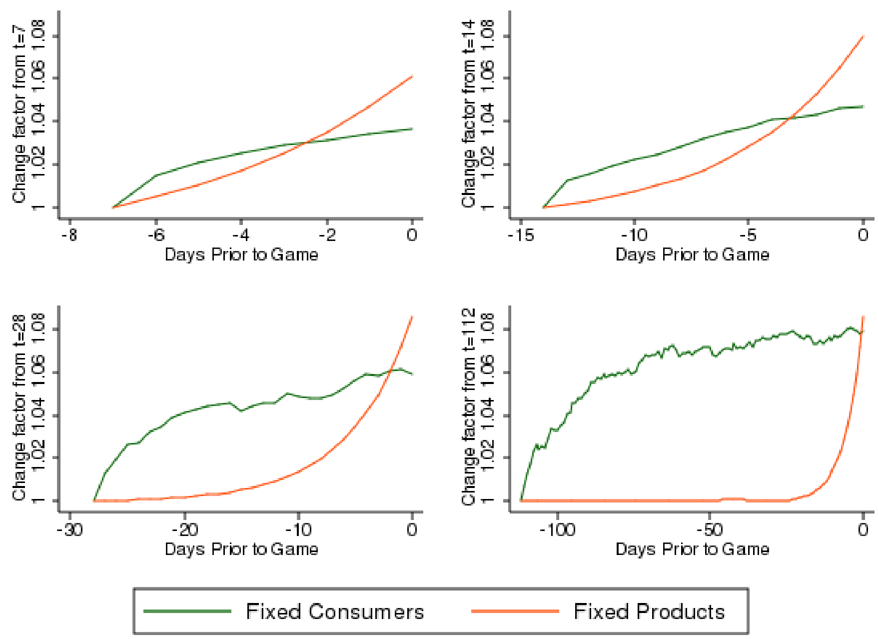

remain constant over time, the use of dynamic pricing leads, on average to a 5.50% increase in revenue relative to uniform pricing over time.

I then do the same exercise for the "true" setting in which both consumers and products are changing (again, where the degree to which they change is an average across the four teams of focus). I solve not only for optimal unform pricing over time and optimal dynamic pricing but also optimal day-of-purchase pricing, i.e. the optimal time path of prices if one is constrained to set them ahead of time rather than each day. Results are shown below:[46]

By considering a pricing strategy which depends only on the day of purchase and not on the playoff probability, I isolate the use of intertemporal price discrimination. I compare uniform pricing to this strategy, and this strategy to full dynamic pricing. The ability to change prices over time *and* in response to product changes leads to much larger revenue gains than those associated with pricing policies which depend only on

[46]For curious readers, Table **??** in the Appendix also includes the results "fixed consumers" and "fixed products" results.

Table 1.7 Average Welfare Effects of Dynamic Pricing (and Day-of-Purchase Pricing) Relative to Uniform Pricing Over Time

	actual setting	actual setting, preset path
Revenue	+5.75%	+0.113%
CS_1	-3.07%	+0.217%
CS_2	-1.27%	-0.469%
CS	-1.4%	-4.186%

the date of purchase— 5.75% compared to 0.113%. The inability to optimize is costly: if prices are not *dynamic*, i.e. not allowed to re-optimize around unforeseen demand shifts, revenue is only slightly higher than under uniform pricing.

Looking now at only the third column, the intuitive idea that high willingness-to-pay consumers would lose more under price discrimination turns out not to be the case: price is not adjusted so much as to make up for their relatively small disutility from spending. Hence the decrease in consumers surplus for Type 1 consumers is actually larger. This is observed because there are still a significant number of Type 1 consumers shopping in the final week, when the optimal price increases most dramatically.

1.7 Concluding Remarks

Academics and industry observers studying the arts and entertainment industry have provided two broad rationales for changes in ticket prices across time: changes in the consumer mix and unexpected changes to the product. In this paper, I study the effects of these two changes on changes in demand for baseball tickets over time leading up to a game. This setting is convenient because there are observable product characteristics that change over time which shifts demand, and because capacity constraints play only a minor role. I assemble a novel dataset of daily listings and transactions, both on the primary and secondary markets. I estimate a structural model of demand in which the distribution of consumers changes over time. The estimation finds a similar story to that of the airline industry: as the game approaches, shoppers are more likely to be consumers with high willingness to pay. However, unlike in the airline industry, there are demand shifts that have to do with changes in the event: a future event becomes more or less popular depending on current team performance.

Optimal pricing, relative to fixed pricing, increases revenue by 5.75% in my simulations. One reason these increases are not larger is that the firm chooses its optimal price given prices on the secondary market, which tend to plunge in the final week before a game. This is precisely the time when the team would like to do price discrimination. The secondary market acts a restraining force, though only to a limited extent because secondary products are well-differentiated from primary products, and because transaction costs are too high for arbitrage of single-game tickets. Another reason these increases are not larger is that the firm gets to choose uniform prices only 14 days before the game, whereas in actuality, franchises using uniform pricing must set prices pre-season when their trajectory of team performance is less certain.

I show that changes in the consumer mix tend to influence optimal price most dramatically in the final week before a game, while oscillating product characteristics play a larger role more up to that point. Because playoff probability moves in an unpredictable fashion, having to choose a preset price path 14 days out constrains the franchise so much that revenue is barely higher than under uniform pricing. On the other hand, if I were to repeat the exercise for 3 days out then a preset price path might accomplish most of the revenue gains.

It may be that other changes, such as a change in the starting pitcher, could shift demand as well. I am constrained by limited data, though in future work I hope to incorporate data I have collected on expected starting pitcher. One could also allow the pricing strategy to depend on the unobserved attributes estimated from the demand specification, as these attributes presumably are known by the franchise.

Another improvement on this research would be to allow for capacity constraints to play a role, by computing value functions under fixed pricing and dynamic pricing. I would have to find a way to narrow the state space. This approach would have to address what information the baseball franchise pays attention to in the secondary market. It may be that franchises only look at minimum or mean prices, and are not sophisticated enough to predict secondary market behavior.

Chapter 2

The Gains from Dynamic Pricing for Ticket Resellers

2.1 Introduction

Chapter 1 asked about the gains from dynamic pricing to a baseball franchise, which offers a product that at least some consumers apparently perceive as different from any substitute, including resold season tickets. Here I focus on the problem of a ticket reseller who is generally reselling season tickets. These sellers have barely any market power and capacity is quite limited: the majority of listings offer only two tickets. This chapter provides a simplified method to estimate the gains from daily price adjustments to StubHub sellers and applies it to the same secondary market transaction data used in Chapter 1.

In future work this method shown here may be used to improve the work done in Chapter 1: this method could extend and allow one to simulate the counterfactual season ticket purchases that take place when a different franchise pricing policy is expected.

The problem of simulating alternative pricing policies by a ticket reseller is more complicated than it might at first appear. First, aside from limited inventory and market power, a ticket reseller also differs from a franchise in that he may have a positive scrap value associated with the ticket. In other words, if he does not sell the ticket, he may enjoy some utility from either attending the game himself or giving the ticket to friends. In the final market period, his opportunity cost of sale is this scrap value. In earlier periods, his opportunity cost of sale is a combination of that scrap value and future opportunities to sell before time runs out.

Second, many sellers never change price, suggesting the presence of a price adjustment cost. Indeed, one question this paper aims to answer is: how large would the price adjustment cost have to be to rationalize observed behavior, and how reasonable does that lower bound appear to be?

The complication of counterfactual pricing policies arises from the fact that different sellers may have very different scrap values. I fully address this issue here. It also arises from the price adjustment cost. In this paper I get around this issue by focusing only on sellers who never change price and assuming the price adjustment cost was so large that they posted the price never intending to change it. This is only one possible assumption to allow identification of the scrap value, but it is a good place to start.

When I describe this chapter as providing a "simplified method" I mean that there are some elements missing from the model that one would take issue with. The main one is the effect of the franchise's pricing. Here I assume that the behavior is expected and regular enough that a time variable, as well as higher order terms of this variable and cross terms with price, account for the team's behavior. But to predict how optimal reseller pricing would change under an alternative franchise pricing policy, one would have to model these effects more explicitly. Another missing element is dynamic product characteristics like team performance. In theory, StubHub sellers might want to think take the distribution of possible product characteristics in the future into account when they price today.

To give a preview of my findings: Roughly 40% of sellers have a scrap value of zero or "below", another 40% value the tickets at less than the franchise's face value, and 20% value it above the franchise's face value. There are two interesting facts here. First, it seems that the way many sellers price cannot be rationalized, because scrap value cannot be negative, and maybe these sellers are victim to the sunk cost fallacy. However, a more likely explanation is that some of the elasticities are poorly estimated; it seems highly unlikely that elasticity for one little StubHub listing would be inelastic. Second, 20% of sellers actually don't need to post their tickets on StubHub, since no one is likely to purchase from them, but they do so anyway because it's convenient to post all of one's season tickets at once and just set prohibitively high prices for some of them.

This is good news; it means I have a better chance of recovering the true distribution of scrap values and that the counterfactuals in Chapter 1 can be improved upon.

2.1.1 Outline

The rest of the chapter proceeds as follows. Section 2 describes the data collected for this study. Section 3 presents the model. Section 4 discusses the econometric specification and identification of the model parameters. Section 5 presents the results of demand estimation and the recovery of sellers' scrap values. Section 6 uses these results to find the distribution of gains from dynamic pricing. The conclusion follows.

2.2 Data

The StubHub dataset is described in Chapter 1. I have daily snapshots of inventory and price for each individual listing for each of over 400 baseball games. I know each listing's stadium section and row number. Unlike in Chapter 1, where I did not have the corresponding data for the primary market, here I make use of the row numbers as the primary market is ignored.

Caveats are in order. First, I do not have a seller ID, so I assume that each listing has its own seller and thus is a separate optimization problem. Second, I only see changes in inventory, and therefore assume that a fall in inventory is a sale unless it seems especially unlikely (details provided in the appendix). Third, I do not have any means of inferring sales on the day of the game as inventories are not available post-game.

2.2.1 Description of the data

At least 73% of all StubHub ticket sales take place in the final 28 days before a game, on average. This does not include day-of-game sales, for which I do not have data, but judging from industry reports, StubHub sales continue to increase from the day before to the day of the game.[1] Because most sales take place in the final 28 days, I limit the analysis to these days for all games in my data. I also limit the analysis to Giants

[1]One might think that consumers would be worried about not getting their tickets in time, but an increasing number of consumers know that StubHub the official partner of Major League Baseball and therefore provides digital tickets immediately.

ticket resellers, as differences in income, fan loyalty, and stadium capacity may lead to radically different demand functions across stadiums.

A common issue with demand estimation is price endogeneity: the seller knows more about the product than the econometrician does and price is therefore correlated with unobservable product characteristics, positively biasing the price coefficient. To eliminate this concern I focus only on only on one zone of the stadium. Since the franchise prices all sections in the zone identically, I assume that quality is homogenous (except for row number, which will be included in the demand function).

Finally, demand may differ for differently-sized listings. A listing of four tickets is likely to be more popular than a listing of two tickets as one can bring more friends with her to the game. Again I take advantage of the luxury of having so much data and I simply drop any observations of listings that offer more than 2 tickets.[2]

Figure shows the face values of the "View Reserve Outfield" zone for the games in the sample, where face value is defined as the franchise's price 56 days prior to the game. Interestingly, by $t = 56$ the Giants priced nearly every game differently. We see large variation in price, as characteristics like rival, time of day, and day of week, among other things lead to very differently-valued games.

Figure 2.2 shows the distribution of the ratio of price to face value. As t counts down, meaning the game is approaching, prices tend to go lower, though many prices do not move.

Figure 2.3 shows the same distributions for the subset of observations in the View Reserve Outfield stadium zone. In the estimation and counterfactuals I focus on this section. The figure shows that this section is broadly representative of pricing behavior across the stadium.

Nearly a majority of listings never exhibit a price change, as seen in Figure 2.4. Figure 2.5 shows that a price change becomes far more likely as a game approaches. But this figure is misleading, since there are also simply more listings as the game approaches. Figure shows the same histogram but looking only at those listings posted on or before Day 28 prior to the game.

[2] 54% of all Giants listings offer 2 tickets. 4 tickets and then 8 tickets are the next most common intial stocks, and anything else is extremely rare.

Figure 2.1 Histogram of Face Value

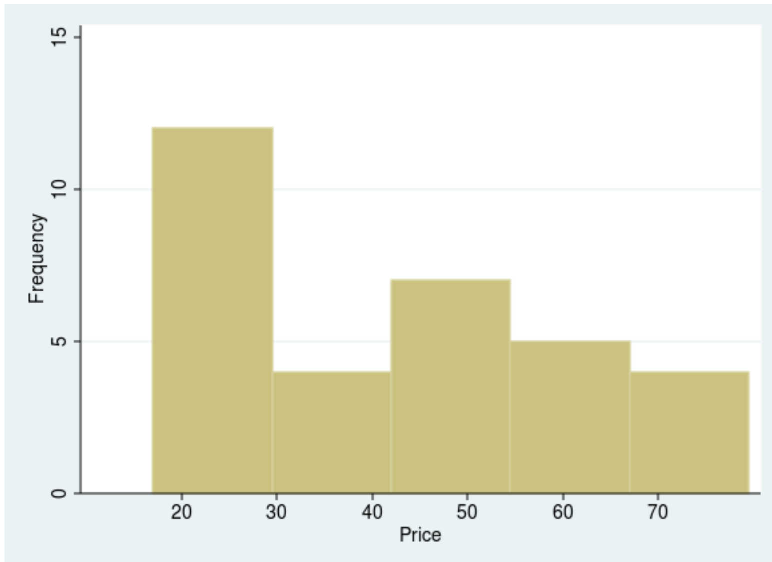

2.3 Model

A listing is defined as an offer for a pair of tickets, posted on StubHub. The seller may change the per-ticket price each market period $t$leading up to a game. The market periods count down, i.e. $t = T, T-1, \ldots, 2, 1$ where T is the earliest date on which a seller may post tickets.

Demand for listing l is a probability function $q_{lt}(p_l, p_{-l})$ (there may be product characteristics as well but these are ignored for ease of notation).

Seller behavior follows the theoretical model in Sweeting [2012]. Letting the prime symbol (e.g. V') denote the value of a variable next period, I omit the search day subscript t for the remainder of this section. On a given search day, an optimizing risk-neutral seller with a single listing l will choose a price p_l to maximize his value, defined by the Bellman equation

$$V_l = \max_{p_l} p_l q_l \left(p_l, p_{-l}\right) + \left[1 - q_l(p_l, p_{-l})\right] EV_l', \tag{2.1}$$

where $q_l(p_l, p_{-l})$ is the probability that l sells in the current period and EV_l' is the *opportunity cost* of selling the listing, or the current-period expected value of still having

FIGURE 2.2 HISTOGRAM OF PRIVE-FACE VALUE RATIO

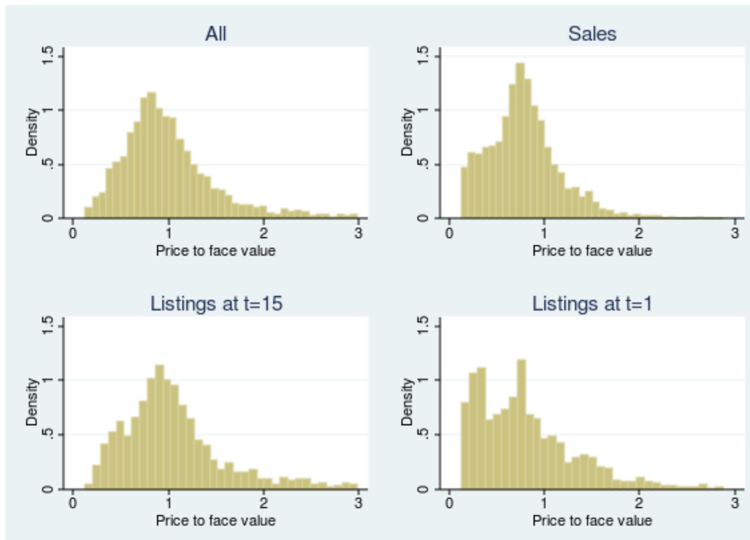

listing l next period. The value and expected value functions are both functions of the state, but I leave this out for now. If the current period is the day of the event then $EV_l' = E_0 V_l^1 = V_l^1$ is the suppliers known value of being able to attend the event using the tickets in l. Under standard regularity conditions, a first-order condition shows that the optimal price p_l^* is equal to a markup plus the opportunity cost of sale:

$$p_l^* = \frac{q_l\left(p_l^*, p_{-l}\right) + \left[1 - q_l(p_l^*, p_{-l})\right]\left[\partial EV_l' / \partial p_l\right]}{\left|\partial q_l(p_l^*, p_{-l}) / \partial p_l\right|} + EV_l'. \tag{2.2}$$

Because there are so many secondary market sellers offering tickets for a given game on any given search day, I follow Pang et al. [2015] in assuming that $\partial EV_l / \partial p_l = 0$, at least as perceived by a secondary supplier.[3] Equation (2.2) becomes a simple Lerner Index Rule:

$$p_l^* = \frac{q_l\left(p_l^*, p_{-l}\right)}{\left|\partial q_l(p_l^*, p_{-l}) / \partial p_l\right|} + EV_l'. \tag{2.3}$$

[3]Sweeting [2012] finds that the most likely violation of this assumption happens when a higher p_l may cause an interested buyer to wait, increasing future demand. As I will argue in section 1.5 that consumers are higly unlikely to delay purchase to a later date. Even allowing for this possibility, though, a secondary supplier is unlikely to consider the negligible impact of his price on future demand: in 95% of game, search day combinations there are at least X listings on the secondary market.

Figure 2.3 Histogram of Prive-Face Value Ratio: View Reserve Outfield

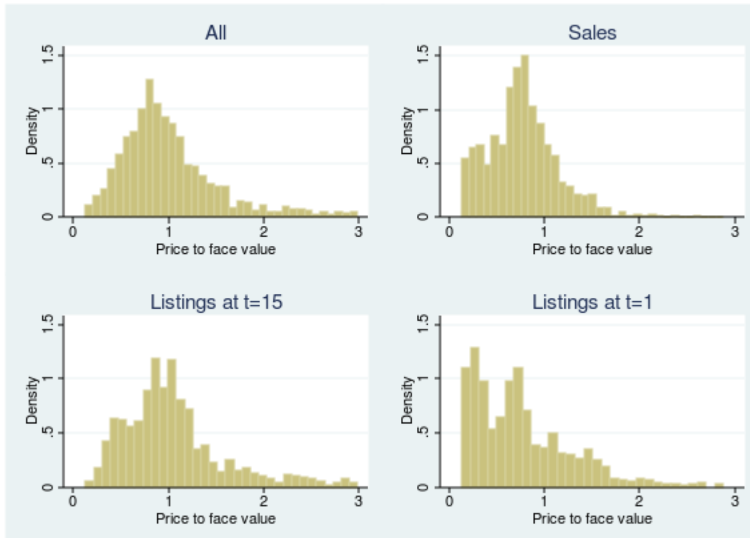

This formula provides the natural story that secondary market price is an opportunity cost plus a markup. The markup is expected to be small, as $|\partial q_l(p_l^*, p_{-l})/\partial p_l|$ is expected to be large: with so many similar listings being offered, consumers should be highly elastic with respect to price among listings *within a product*.

2.4 Estimation

2.4.1 Daily listing demand

I assume that a given pair of tickets is purchased at time t with probability $q_{lt}(p_{lt}, p_{-lt}, X_{lt})$. I confine my data to a single pricing zone, so that the only observable characteristics that differ across products are price and row number, and each listing relates to the distribution of prices in the same way. Therefore the function need not depend on the distribution of prices within the vector p_{-lt}, as this distribution is simply a function of t.

I flexibly estimate the demand function as a probit function of log price and days before game. Letting $\tilde{p}_{lt} = \ln(p_{lt})$, and letting r denote the row number, the probability

Figure 2.4 Histogram of Number of Price Changes

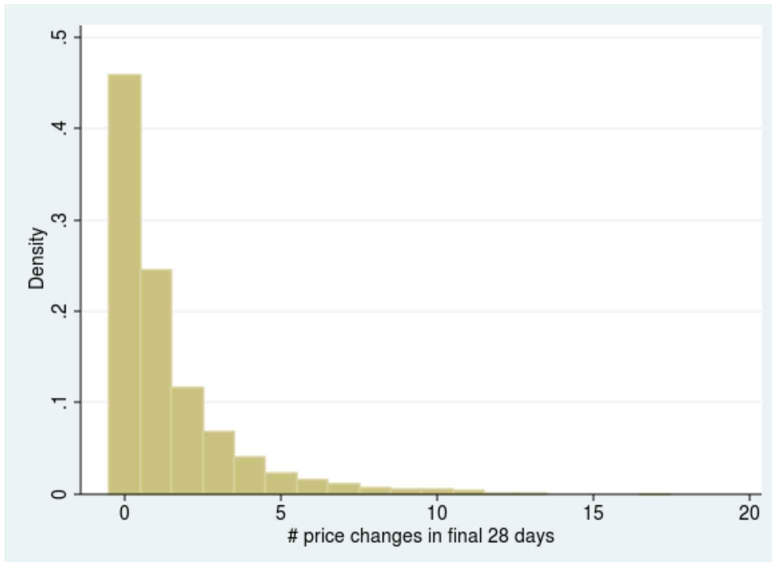

function for game $g \in G$ is

$$q_{glt} = \Phi\left(h\left(\tilde{p}_{lt}, t, r, g\right)\right) \quad \text{where} \quad h(\tilde{p}, t, r, g) = \sum_{i=0}^{3} \sum_{j=0}^{3-i} \beta_{ij} \tilde{p}^i t^j + \gamma_1 r + \gamma_2 r^2 + +F_g \tag{2.4}$$

where F_g is a game fixed effect. I allow for clustered sandwiched errors within a listing, across different days leading up to a game.

I do not have day-of-game (the final market period) sales, but one can make an out-of-sample prediction to get the probability that a listing sells at $t = 1$, the final market period.

2.4.2 Listing demand across time leading up to game

A seller with constant price is, as stated earlier, assumed to have set that price believing that she would never subsequently change it. When she set this price I assume she was simply concerned with the probability that her listing sells at some point before the game. One could, of course, estimate this probability by using the daily sale probability function described above, but a more accurate estimation comes from using the data more directly. I estimate a function which looks very similar to (2.4) but in which the

Figure 2.5 Histogram of Day Prior to Game of Price Changes

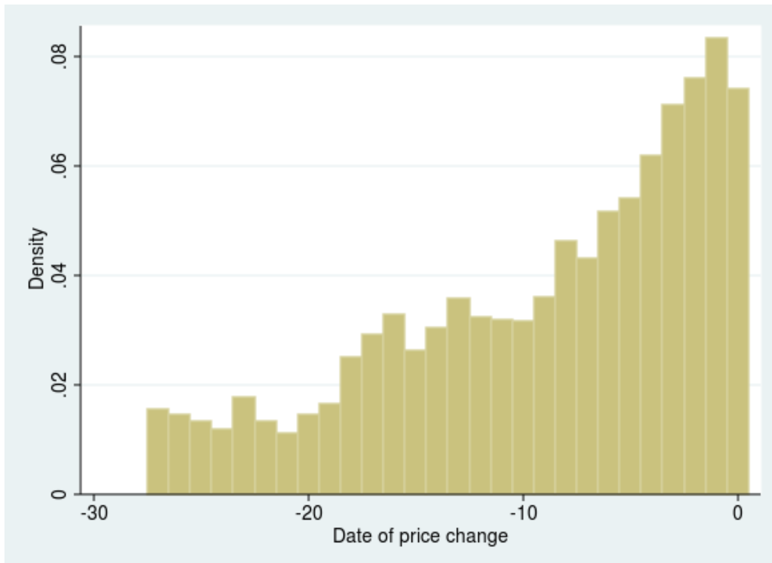

dependent variable is the probability $\mathring{q}_{glt}$ that game g's listing l will sell at some t' where $1 < t' \leq t$.

In order to get the probability that l will sell at some t' where $1 \leq t' \leq t$ (note the change from strict to weak inequality), I use the out-of-sample $t = 1$ probability described above:

$$\bar{q}_{glt} = \mathring{q}_{glt} + (1 - \mathring{q}_{glt})\, q_{gl0} \tag{2.5}$$

2.4.3 Solving for scrap value

Alternatively from Equation (2.4), a seller who does not expect to change price sets her price according to the Lerner Index Rule

$$\bar{p}^*_{lt} = \frac{\bar{q}_{lt}\,(p^*_{lt}, p_{-lt})}{\left|\partial q_{lt}(p^*_{lt}, p_{-lt})/\partial p_{lt}\right|} + EV_{l0}. \tag{2.6}$$

The elasticity (or, if you prefer, simply the partial derivative) of "eventual sale probability" with respect to price can be calculated using estimates of $\boldsymbol{\beta}, \boldsymbol{\gamma}$, and $\boldsymbol{F}$, and knowing the price set at the time of listing one can back out EV_{l0} for any listing.

Figure 2.6 Histogram of Day Prior to Game of Price Changes, Old Listings

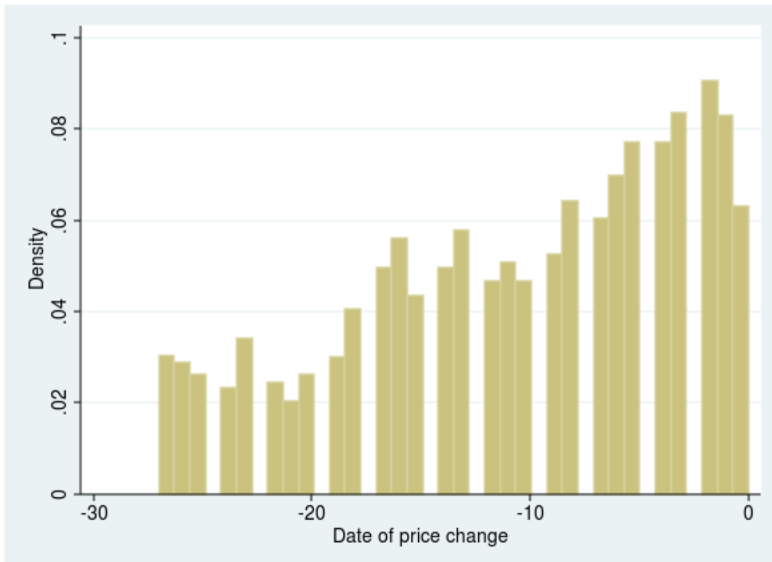

2.4.4 Expected revenue from optimal daily pricing

Calculating the expected revenue from optimal daily pricing is simply a matter of recursively solving Equation (2.4) for p^*_{lt} for each listing l, from $t = 0$ back to the day on which the listing was posted. After each solution, one can calculate the value of having tickets today as $EV_{lt} = q_{lt}p^*_{lt} + (1 - q_{lt})EV_{l,t-1}$. In this way, I find the expected value of having two tickets to sell at the time of the listing.

I compare this expected value to the expected value of having two tickets to sell at the time of listing and only being able to set price once and never altering it. To be consistent with the previous calaculation, I use the daily sale probabilities to recursively calculate expected values, but of course without recalculating optimal prices this time.

In this way I obtain counterfactual expected revenue (where "revenue" includes scrap value) for each listing. By keeping the demand function the same in the counterfactual, I implicitly assume that this alternative behavior does not change the other listings' behaviors.

Figure 2.7 Graphs of demand functions

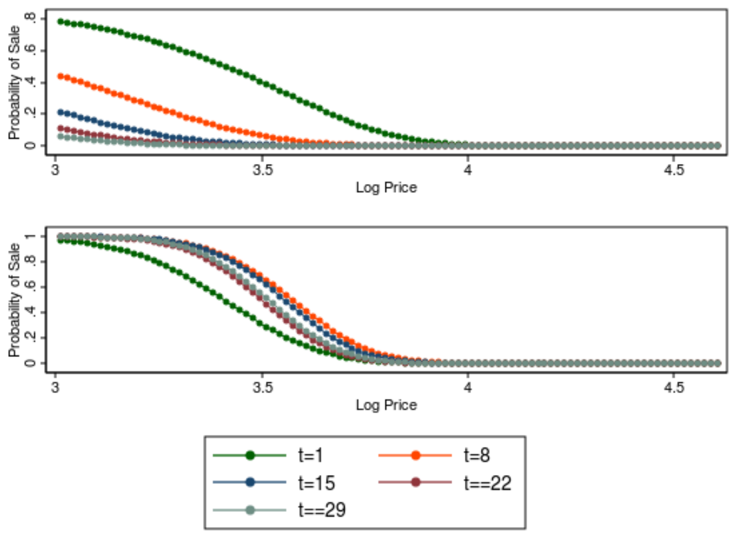

2.5 Results

The daily demand estimation and the "Before $t = 1$" demand estimations are given in Table 2.1, though they are difficult to intuit. What one really wants to know is how price and time change probability of sale, holding other variables equal. This question is answered by Figure 2.7. The graphs in this figure are fairly intuitive: sales probability (demand) is always decreasing in price; the probability of selling sometime before $t = 1$ increases the more days left for it to happen, while the probability of selling on a particular day t increases as t decreases toward the day of the game. Note that the estimation is fitting the fact that listings with high outlier prices almost never sell.

Using the Lerner Index Rule, I find a wide distribution of "scrap values" as shown in Figure 2.8. Note that it looks truncated because it is: I did not include observations with price greater than three times face value in the analysis. Intuitively, those with higher scrap value should gain very little from dynamic pricing, since they are so unlikely to be able to sell their tickets for a higher price than the value they already place on them.

I find that 41% of listings show scrap values below \$0, 40% are positive and below face value, and 19% of listings are associated with a scrap value above face value.

Figure 2.8 Value to Face Value

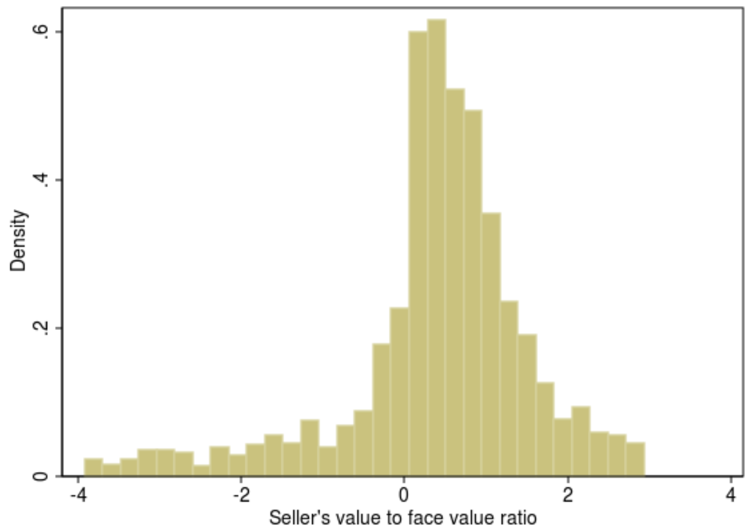

The finding that 41% of listings evidently have scrap values below zero is a puzzle. There are a number of potential explanations. One is that I have assumed risk-neutrality when, in fact, many sellers only post a few listings and are likely to be risk-averse, leading to a lower price. A future avenue of this research would be to redo this analysis assuming a utility function featuring risk-aversion, and to see how many listings still turn out to have a scrap value below zero. Another is that sellers simply do not accurately estimate the ex-ante probability of sale when they post their tickets.

2.6 Analysis of the Gains to Dynamic Pricing

As explained in Subsection 2.4.4, calculating the gains to dynamic pricing under this simple model amounts first to solving for optimal price at $t = 0$, where opportunity cost is simply the seller's scrap value. Next, I solve for optimal price at $t = 1$ where opportunity cost is the value evaluated at $t = 0$, and follow this process recursively until I reach the period t at which the listing was posted. Figures 2.9 and 2.9 show histograms of gains to dynamic pricing for various groups of listings. The former uses level change while the latter shows percent change.

For the 41% of sellers whose scrap values lie below \$0, I find that the average increase

Figure 2.9 % Change in Expected Revenue

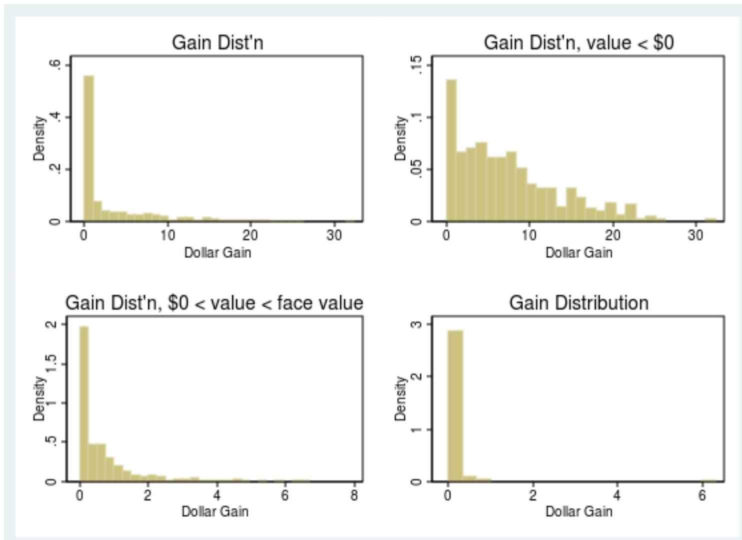

in revenue is 23.23% (with a standard deviation of 18.33 percentage points). For the 40% of sellers whose scrap values lie between \$0 and face value, the average increase is 2.69% (with a standard deviation of 6.16 percentage points). For the 19% of sellers whose values lie above face value, the average increase in revenue is only 0.078% (with a standard deviation of 0.214 percentage points).

As expected, sellers who have a scrap value greater than price (i.e. sellers who greatly enjoy attending games, or greatly enjoy giving gifts to friends) do not gain very much by using dynamic pricing as opposed to uniform pricing: in either case they set prices so high that the probability of purchase is negligible.

For sellers who were previously behaving as if they had a negative scrap value, I assume their actual scrap value is \$0. These sellers gain the most from dynamic pricing.

2.7 Conclusion

This chapter investigated the pricing behavior of StubHub sellers an measured the gains that would result from using dynamic pricing rather than uniform pricing. The discussion from the introduction bears repeating, however: there may be a price adjustment cost. It is quite possible that what I have recovered is a distribution of lower bounds on

Figure 2.10 % Change in Expected Revenue

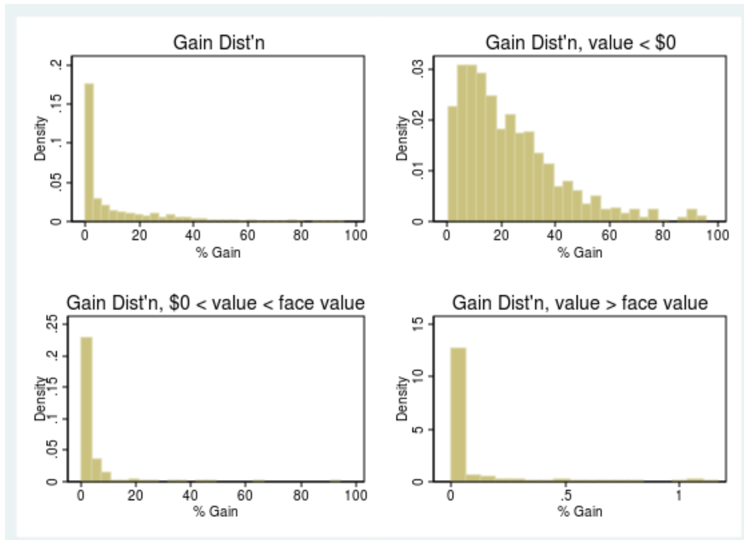

price adjustment costs.

Regardless, as long as we can assume that uniform pricers *planned* on being uniform pricers, we can back out their scrap values. This information would be very useful to baseball franchises looking to price their season tickets optimally. It would also be informative for policymakers and taxpayers involved in the ongoing debate over stadium subsidies: if season ticket holders derive a large amount of surplus from the secondary market, then perhaps stadiums are worthwhile projects to fund. Finally, if this study were repeated over multiple years one could ask about the effect on scrap value of of promotion, advertising, or improvements made to the stadium. One could juxtapose this with the variable's effect on single-game consumers' valuations using a demand-based estimation of the values.

Table 2.1 Daily and "Before $t = 1$" Demand Estimations

	Prob of sale today		Prob of sale before t=1	
	(1)	(2)	(3)	(4)
main				
t	0.0709*	0.236***	0.833***	0.914***
	(0.0324)	(0.0490)	(0.0741)	(0.0887)
$\tilde{p}$	-6.964***	-17.37***	-12.82***	6.222
	(0.684)	(1.523)	(2.209)	(5.192)
t^2	0.00718***	-0.00259	-0.0225***	-0.0298***
	(0.00136)	(0.00171)	(0.00166)	(0.00238)
$t\tilde{p}$	-0.136***	-0.146***	-0.237***	-0.219***
	(0.0148)	(0.0271)	(0.0381)	(0.0484)
$\tilde{p}^2$	2.552***	6.498***	4.650***	-2.125
	(0.218)	(0.514)	(0.716)	(1.809)
t^3	-0.000161***	-0.0000429	0.000259***	0.000383***
	(0.0000256)	(0.0000322)	(0.0000243)	(0.0000359)
$t^2\tilde{p}$	0.000840***	0.00219***	0.00222***	0.00278***
	(0.000248)	(0.000317)	(0.000337)	(0.000486)
$t\tilde{p}^2$	0.0144***	0.00426	0.0202***	0.0117
	(0.00211)	(0.00432)	(0.00529)	(0.00728)
$\tilde{p}^3$	-0.285***	-0.846***	-0.514***	0.0876
	(0.0223)	(0.0559)	(0.0740)	(0.195)
rowno	-0.0199*	-0.0356***	0.00566	0.0968**
	(0.00834)	(0.00974)	(0.0292)	(0.0360)
r^2	0.00205***	0.00215***	0.000739	-0.00418*
	(0.000415)	(0.000481)	(0.00149)	(0.00176)
night game	0.0629**		0.303***	
	(0.0225)		(0.0732)	
Game FE	No	Yes	No	Yes
N				
R^2	60764	60764	20511	20511
r2				

Standard errors in parentheses

I do not report elasticity estimates, as the most important estimate is the price elasticity of all-time demand, meaning any day on or after the listing. Summary statistics are reported in the text. Also note that weekend was not included as it is perfectly collinear with night-game for Giants Games.

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

Chapter 3

Does Premerger Notification Matter? Evidence from Cable Television

3.1 Introduction

In 1976 Congress passed the Hart-Scott-Rodino (HSR) Act as a response to several criticisms of anti-trust policy. While the Clayton Antitrust Act of 1914 had given broad powers to the Federal Trade Commission (FTC) and the Department of Justice (DOJ), these powers were largely reactive. Enforcement agencies had difficulties challenging anticompetitive actions after they had occured and often found restoring a market to competitive status a costly endeavor. The HSR Act sought to address these concerns by forcing individuals and firms to report certain asset transfers or purchases and obtain pre-clearance before completing the transaction Fed [2009].

While the FTC and DOJ have reported the number of disclosures they've received on an annual basis (though individual filings are generally private) and used the powers granted under the HSR Act to challenge several large proposed mergers, it has been difficult to cleanly test the effectiveness of the policy.[1] A fully specified model must detail the information and choice sets, as well as choice-specific payoff functions for all

[1] See discussions in, for example, Allen [1984], Eckbo [1992], Nelson and Sun [2002], Crandall and Winston [2003].

parties involved in merger activities as well as regulators. The state space of such a model easily exceeds modern computational capacity. Reduced-form approaches also encounter empirical problems. In many markets where large mergers are observed, it is difficult to obtain data on the holdings of the parties – and thus the value of the merger activity. It is also difficult to estimate the consumer level impacts of cross-industry mergers, and thus to get a sense of the degree of regulatory scrutiny transactions face. Additionally, in order to test the effectiveness of the policy we need to have a sense of the mergers that did not happen – not just those that did. Finally, for statistical power, econometricians must observe a large number of potential and realized mergers.

The cable television industry provides a solution to these concerns. Firms providing cable television service must register their 'cable communities' with the Federal Communications Commission (FCC), which makes these registrations, as well as ownership changes, public. We combine this data with data from the Census Bureau to obtain a complete picture of the cable industry in the United States from 2000 to 2012. Not only does this data allow us to identify actual acquisitions, we can also construct the universe of possible acquisitions.[2] Additionally, this data allows us to understand the degree of horizontal competition - referred to in the industry as 'overbuild' - these firms face in their individual communities. We assume any merger which includes overbuilt communities - in other words, any merger that results in a local market shift from duopoly to monopoly - would face increased scrutiny from regulators. However, even in the absence of overbuild, regulators may wish to apply scrutiny to proposed cable mergers if they may have an effect on the market for content provision.

To test the effect of HSR, we develop a model of firm valuation which includes terms representing the cost of regulator scrutiny, particularly for acquisitions which contain a horizontal component. We then take this model to the actual and potential acquisition sets for the top four firms in the market: Comcast, Time Warner, Charter, and Cox. We concentrate on the top firms largely to reduce the assumptions we must make about their choice sets. First, given their size and the (general) strength of their balance sheets throughout the time period we study, it is reasonable to believe their marginal decision to acquire a particular small regional competitor is not based on financial constraints.

[2]We implicitly assume that cable television firms are primarily interested in acquiring other cable television firms, an assumption that is largely satisfied with some large exceptions, such as Comcast's bid for Disney in 2004 (and successive acquisitions).

This allows us to consider each decision independently, instead of considering a bundle of multiple acquisitions. Second, given their geographic spread, it is reasonable to believe they consider acquisitions across the entire extent of the United States. This allows us to remain agnostic about which small firms entire the decision set of the large firms.

In some sense, our exercise can be thought of as estimating a series of acquisition hazard rates for these small firms by their larger competitors as a function of observables, somewhat in the vein of Rust [1987], though we do not take dynamic incentives into account explicitly. As a robustness check, we therefore combine the observed acquisitions across our four large firms into a single specification that estimates the overall hazard rate of acquisition for small cable companies as a function of observables.

Our results indicate the HSR filing requirement discourages firms from pursuing mergers involving a horizontal component. However, the results also suggest acquiring firms do not change their behavior for firms which are over the HSR size threshold and are under the size threshold. We attribute this later result to a lack of statistical power.[3] Our results are robust to several variations of our empirical specification.

The literature has struggled to understand the effects of anti-trust policies. Clougherty and Seldeslachts [2013] use a conditional probabilities methodology from the economics of crime literature to study understand the deterrence effects of various policy instruments across a number of US industries. They find the an increase in the rate at which the government challenges mergers deters future horizontal merger activity (defined as firms with identical 4-digit SIC codes), but they do not test the filing requirement and size thresholds. Clougherty et al. [2015] details a similar exercise on EU data that includes more accurate information on the regulators' views of potential competitive effects.

The remainder of this paper proceeds as follows: In Section 2 we provide brief background on both the cable industry and merger policy. Section 3 introduces a valuation model for potential acquisitions which leads to our simple test. In Section 4 we describe our novel dataset constructed from FCC, Census, and FTC data, highlighting the difficulty posed by limited information on former cable providers. Details on our data cleaning methods are left for an appendix. In Section 5 we detail our empirical strategy

[3] In the limit as the number of observations increases, we should be able to detect the effect of the HSR filing fee, if nothing else.

and provide the results of our test. Section 6 concludes.

3.2 Background

3.2.1 Hart-Scott-Rodino Anti-Trust Improvements Act

The primary effect of HSR was the creation of the FTC and DOJ?s Merger Prenotification Program. Under the program, parties considering a sizeable transaction must file a ?Notification and Report Form? and pay a substantial fee based upon the size of the proposed transaction. The parties must wait 30 days during which regulatory agencies may request additional information or time to review the transaction. If the reviewing agencies believe a proposed transaction violates antitrust laws, they may attempt to prohibit completion of the transaction by filing for an injunction in federal district court. Information provided to regulators during this process, including the original filing, is not subject to public disclosure, though court filings are generally available.

If the parties are conducting routine transactions or have experience with the system, they may file a request for Early Termination of the waiting period. If the Early Termination is approved, the transaction is made public as part of the Federal Register. While this data can be used to give a flavor of the types of transactions generally seen by regulators as unlikely to have anticompetitive effects, it cannot be used to identify the entire universe of attempted or successful purchases, since not all transacting parties request Early Terminations and not all Early Terminations requests are approved.

Transaction reports are necessary when either the value of the assets or the size of the parties reaches certain thresholds. These rules are designed to take effect cumulatively, so a firm which slowly acquires the assets of a competitor through multiple transactions will be forced to report even if each individual transaction is small. Thresholds are adjusted periodically by the FTC and DOJ to reflect inflation. Figure B.2 illustrates the various reporting thresholds based on the size of the parties and transaction denominated in dollars. As of 2016, reporting is required if the acquiring party will hold assets of \$312.6 million or more, or if one party is worth at least \$15.6 million, the other is worth at least \$156.3 million, and the assets transferred are worth at least \$78.2 million [Register, 2013]. An additional set of reporting requirements exist based on the percentage of assets transferred: filing is required if the transaction involves \$78.2 million in assets consisting of at least 50% of a company.

3.2.2 History of cable

Cable television began in the early 1950s as a way to improve the reception of over-the-air broadcast channels in remote communities. High demand for broadcast television coupled with the Federal Communications Commission?s 1948 “freeze” on licenses to construct new stations led to the creation of Community Antenna Television (CATV) systems [of Broadcast Communications, 2013]. Instead of a separate antenna required for each household who wanted to receive broadcasts, a single, more sensitive antenna could be placed in a centralized location and connected to households through wiring.

Demand for cable systems spread rapidly, and by the 1970s even large metropolitan areas were wired for cable. Local governments executed ad-hoc franchise agreements with cable operators; in exchange for the (sometimes exclusive) right to provide cable services to the area, cable operators would guarantee certain benefits such as educational and governmental channels or special rates for particular segments of the population [Commission, 2012].

Exclusive channels began appearing on cable systems starting with Home Box Office in 1972 and quickly became a large draw for subscribers. With the increased bandwidth available through wired technology, cable operators were able to offer a much wider variety to consumers than the broadcast alternative [Association, 2013b, Commission, 2012, Eisenmann, 2000].

Today, over 90% of households have access to cable television and roughly 60% of households are active subscribers [Association, 2013a, Nielsen, 2011]. Cable operators, empowered by the Telecommunications Act of 1996 discussed below, have also used the two-way properties of the communication technology to offer internet and phone services.

Competition in the video space comes mainly from Direct Broadcast Satellite technology, a subject previously studied in detail by Goolsbee and Petrin [2004]. Competition in the market for data provision comes from Digital Subscriber Line and fiber-to-the-home technologies.

3.2.3 Telecommunications Act of 1996

The Telecommunications Act of 1996, which amended the Communications Act of 1934, is the primary law regulating cable operators (as well as the rest of the telecommunications industry) today. The law?s main goal was to promote competition by removing entry restrictions in telecommunications markets. In essence, the law was designed ?to let any communications business compete in any market against any other.? [Commission, 2011] Additionally, the law sought to update the FCC?s regulatory authority and framework to encompass the Internet.

The 1996 Act removed most price controls from the market and encouraged local franchise authorities to allow additional firms to construct physical capital and enter local service markets. It was believed these so-called ?overbuilders,? along with entry from telephone service providers, would provide effective competition in major markets [Padilla, 2001]. These overbuilders are the source of the true horizontal purchase opportunities available to cable incumbents such as Comcast. Emmons and Prager [1997] finds empirical evidence that this change in market structure created increased incentives for monopoly power in the cable industry while Kelly and Ying [2003] examined the feasibility of overbuild and concluded profitable opportunities were rare.

3.3 Model

This section introduces a flexible model of acquisitions, focusing on the valuation the large firm makes. In our model, a large firm values individual communities individually and aggregate these valuations to value the entire set of communities served by a small firm. An acquisition is more likely to take place if the acquirer values the communities in the acquisition more than the company being acquired does, since there exists a merger price (in between the two valuations) such that both firms benefit. With this model in hand, we describe a simple test of the effectiveness of the HSR threshold which we can take to our data.

3.3.1 Environment

At the beginning of each discrete time period t there is a set J_t of large firms and a set K_t of small firms. Each firm i (where $i \in J_t \cup K_t$) operates systems which serve a set of

CUIDs L_{it}. There is a set of communities M and a function $m_t(\cdot)$ which maps CUIDs to communities for a given period t. A community $c \in M$ has overbuild if two firms each have at least one CUID in the community; that is, there exist firms i, $i' \in J_t \cup K_t$, $i \neq i'$, and CUIDs $l \in L_i$, $l' \in L_{i'}$, such that $m_t(l) = m_t(l') = c$. Let $O_t(c)$ indicate whether community c is part of a community with overbuild at time t.

For a given community $c \in M$, let $R_{mt}(c)$ be the revenue that a monopolist can make in period t. Let N_c be the number of households who choose to subscribe in times of monopoly. We do not make any assumptions on how $R_{mt}(c)$ relates to N_c. However, we assume the decrease in revenue from a shift to duopoly is proportional to N_c; that is revenue decreases by $\theta_o N_c$ if the market goes from monopoly to duopoly. Total revenue from any CUID l at time t is then[4]

$$R_t(l) = R_{mt}(m_t(l)) - O_t(m_t(l))\theta_o N_c$$

Operational costs are modeled as a per-subscriber marginal cost m_i and a fixed per-period cost $f_i(d)$ that varies by firm and depends on some measure of distance d from the CUID to the nearest other CUID that the firm owns.

In period t large firms observe which firms have operations in which communities, and information about the small firms that is unobservable to the econometrician. This information comes in two varieties: community-level information on the net per-subscriber savings in marginal costs ϵ_{jlt} and small-firm information on the fixed costs of acquisition ν_{jkt}. ϵ_{jlt} is drawn from a normal distribution with variance $\sigma_j^2, j \in J_t$ to reflect the intuition that some large firms (e.g. Comcast) produce higher gains in efficiency.

3.3.2 Valuation

Suppose a large firm j and a small firm k meet to determine whether an acquisition of k by j would be mutually beneficial. This amounts to determining whether j values k's operating systems more than k does. The difference between j's present value of k's

[4] For expositional and notational simplicity, we write this as if there is no difference in per-subscriber revenue between small firms and large firms. In practice, large firms may be able to obtain more revenue than small firms either through expansions of service or greater negotiating power with advertisers. In our estimation, we cannot separately identify subscriber-level changes on the cost side and the revenue side.

systems and k's present value of k's systems, which we will call $DV(j,k)$, is partially made up of the sum of the differences between present values $v_j(l)$ and $v_k(l)$ of each CUID $l \in L_k$. It is also made up of the differences in j's valuations of each of its existing CUIDs $r \in L_j$, written as the difference between its pre-merger valuation $v_j(r)$ and its post-merger valuation $v'_j(r)$:

$$DV(j,k) = \left(\sum_{l \in L_k} (v_j(l) - v_k(l)) \right) + \left(\sum_{r \in L_j} (v'_j(r) - v_j(r)) \right)$$

It can be shown, using both the formulation for revenue and costs above as well as the restrictions imposed by our data (detailed below), that the change in value simplifies to

$$DV(j,k) = \beta_0 s_k + \beta_1 \{\text{horiz}\} s_k + F(d) + \{\text{HSR}\}(\beta_2 + \beta_3 \{\text{horiz}\}) + \epsilon_{jlt} s_k + \nu_{jkt}$$

where

$$s_k = \sum_{l \in L_k} N_{m_t(l)}$$

is the number of subscribers acquired, $\{\text{horiz}\}$ is a flag for a horizontal or duopoloy-to-monopoly transition, $F(d)$ is the fixed cost of merging as a function of the distance d between j and k[5] and $\{\text{HSR}\}$ is a flag for the Hart-Scott-Rodino threshold. In practice, we estimate $F(d)$ according to

$$F(d) = \alpha_0 + \alpha_1 d + \alpha_2 d^2 + \xi$$

A large firm j will execute an acquisition if $DV(j,k) > 0$. This implies the probability of observing an acquisition follows a known distribution.

[5]We set d to the minimum distance between the CUIDs operated by j and the CUIDs operated by k.

3.4 Data

In order to capture an accurate picture of the cable industry through time and understand the effect of merger policy on consolidation in the cable industry, we combine data on cable television systems from the FCC with market-level data on household counts from the Census and geographic location data from the United States Board on Geographic Names to create a novel dataset. We supplement this data with Annual Reports submitted to the FCC by cable providers, Early Termination Notices from the FTC, a series of letters Comcast wrote to the FCC informing the Commission of completed acquisitions, and a small number of public transaction size disclosures.

Our data on cable television systems was collected from FCC?s internet-based Cable Operations and Licensing System (COALS) using an automated process. For a given Community Unit (known as a CUID in FCC parlance), COALS lists the current and previous service providers. COALS also provides access to administrative or regulatory filings made by the system operator that relate to the cable system, including ownership change forms and annual reports.

Table B.2 presents a summary of the CUID ownership file. Just under half of CUIDs undergo legal-entity changes at some point throughout the study period, and the average number of unique parent companies responsible for a CUID was 1.85.

We identified individual acquisition events by looking at groups of CUIDs that switched from (say) Owner A to Owner B within a short time period. We verified our purchase identification process using data collected from a series of public disclosures Comcast made to the FCC about its acquisitions from 2003 to 2008. We distinguish between horizontal and conglomerate purchases with a simple process: For each CUID involved in the event, we examined the list of the acquiring company's existing properties at the time of the event for an exact community name match. If a match is found, the CUID is flagged as a horizontal acquisition. The remaining purchases are considered conglomerate.

To understand the value of controlling any particular cable system, we obtained population and household count data from the U.S. Census at the Census Place level. To understand the value of geographic clustering, we collected data on the location of the various systems (i.e. latitude and longitude) from the Gazetteer created by the Board on Geographic Names. We matched these data to our FCC community information by community name and county.

Finally, to understand the effect of the HSR disclosure requirement, we needed to map the financial value requirement to the context of our community-level data. We used limited public disclosures on acquisition prices to estimate a value of $4000 per subscriber and use annual report and industry data to estimate subscription rates across years. On average, the estimated acquisition value per household was $2600. We then applied the monetary threshold of $71 million to arrive at a threshold value of 27,000 households. While the monetary thresholds change throughout the study period, they are tied to the rate of inflation, which should roughly track the rate of growth in the value of a single subscriber.

3.4.1 Acquisitions

Table B.5 provides summary statistics for our final set of 712 purchases made by top firms, covering 15,357 communities (or CUIDs, in FCC parlance) during our study period. Most purchases covered a relatively small area; the median number of communities involved in a single transaction was 3 and the median population affected was 31,123.[6] Figure B.3 shows the distribution of merger size as measured in households.

The existence of clustered purchases is immediately apparent: the average mean distance between CUIDs involved in an acquisition and the set of CUIDs already owned by the acquiring firm was 4.7 miles. Since distances are calculated using centroids, this suggests many purchases involved systems essentially adjacent to the acquiring firm's pre-existing properties.

Table B.6 provides the same summary statistics for each large firm we study. Comcast had 43% of the acquisitions covering 45% of the total acquired CUIDs and 44% of the population transferred during the period. As such, the summary statistics for Comcast largely drive the overall numbers reported in table B.5. That being said, the acquisition strategies for the other firms implied by the summary statistics are remarkably similar. The average number of CUIDs involved in a single event are almost identical, except for Adelphia which was impacted by its bankruptcy during the period.

The median number of households involved in purchases was below the threshold

[6] Compare to the median population of all cities and towns of the U.S. of 41,994.

value of 27,000 for all firms except AT&T, suggesting a large amount of the consolidation in this industry was done without regulator scrutiny. Time Warner's significantly larger average purchase size was driven mostly by a few very large purchases in the New York and New England region.

Additionally, the average minimum distance between the acquired CUIDs and the firm's pre-existing CUIDs was also similar for all companies besides AT&T. Even AT&T's relatively large distance, 42.7 miles, equates to most acquisitions taking place within a space similar in size to the average US county.[7]

This clustering is apparent visually. Figure B.4 shows Comcast's holdings by county in 2001. By 2003, shown in Figure B.5, Comcast had not just consolidated its holdings in places such as Florida, it had also bought clustered operations in the Mountain West. Finally, by 2010 (Figure B.6) Comcast had expanded to the market leadership position largely through additional regional purchases. In this way, as shown in Figure B.7, Comcast has expanded its reach from roughly 10 million households to over 60 million by 2013. This implies that today, over 50% of households are in Comcast's territory (Figure B.8).

Horizontal acquisitions

Of the 15,357 CUIDs that were acquired by one of the large firms during the study period, 190 were considered horizontal purchases. These 190 switches were part of 74 distinct acquisition events – 10% of the total number of events seen.

Within the 74 events that included a horizontal component, the median percentage of CUIDs involved in the purchase that were considered horizontal was 12.5%. The mean was 33.2%. Several small purchases that consisted of a completely horizontal takeover contributed significantly to this mean – these tended to be municipality-run networks that were sold.

Of the 23 acquisitions with more than 50% of the CUIDs considered horizontal, the median number of households involved was 24,504, implying that many of these purchases required disclosure and scrutiny under HSR.

[7]In fact, this large distance is largely driven by a single acquisition 560 miles from the nearest AT&T-owned CUID.

3.5 Testing Hart-Scott-Rodino

The data presented in the previous section lend themselves to two clear hypotheses:

1. Outside of true horizontal purchases, Hart-Scott-Rodino has little effect on merger strategy.

2. Firms place a high value on "near-horizontal" or highly-clustered acquisitions.

Both of these hypotheses are testable. First, if HSR filing rules place a major burden on transactions over a certain size, large firms should be less willing to pursue those transactions, relative to the opportunities available in the marketplace. Second, if firms value clustered systems, they should be more willing to pursue those transactions relative to the available opportunities.

To test these hypotheses, we ran a simple exercise. For each year in our study period, we created a list of cable systems the large firms could have acquired based on the ownership records.

We then used a simple logistic regression to estimate the probability of a successful acquisition event based on the size of the acquisition and the percentage of the potential purchase's horizontality based upon the acquiring firm's presence in the communities involved at the time of the purchase. We added a dummy variable representing the necessity of Hart-Scott-Rodino disclosure, as well as year dummy variables to reflect changing macroeconomic conditions.

The most important decision in the execution of this exercise is the selection of the decision set available to the firms. A number of questions emerge:

1. Can large firms partially acquire smaller firms? How do we determine the possible subsets?

2. What level of horizontality is allowed?

3. Can firms acquire other large firms?

4. Should potential targets acquired by other large firms be included?

The first question essentially defines the cardinality of the set, and we look to the data to provide guidance. Though partial acquisitions do occur, they are relatively rare. Additionally, many partial acquisitions lead to further transactions with the same target later in the study period – meaning the "cumulative size" portion of the HSR rules applies. For this reason, we opt to model acquisitions as absolute: you either buy the whole company, or you buy nothing.[8]

Since we observe several truly horizontal purchases in the data, we allow any level of horizontality in our potential purchase set. Additionally, since the only "purchase" of a large firm (Adelphia purchased by Comcast and Time Warner) was the result of a bankruptcy process, we do not allow the large firms to acquire each other.

The final question is also the most vexing. Unfortunately, we have no data covering behind-the-scenes overtures and negotiations, so we are unable to observe (for instance) targets of mutual interest, bidding wars, and other types of strategic activity. Therefore, we estimate the model with several variants of the data representing alternative answers to this question.

The first variant considers the probability of a smaller firm being purchased by any of the larger firms – the overall hazard rate of acquisition. In this variant, we calculate the distance variables according to the nearest distance to any cable system owned by any of the large firms. In the second variant, we estimate separate models for acquisition by the individual large firms but exclude any small firm acquired by other firms from the set of potential purchases available to the firm in question. This assumes any negotiation process acts as a truth-telling device and large firms with the highest internal valuation always have the first option to purchase small concerns. In essence, if Firm B acquires Firm C, that event is viewed as evidence that Firm C was never truly an option for Firm A. In the third variant, we estimate separate models for each of the major firms and allow them the possibility of acquiring any other firm in the market. This assumes the negotiation process may break down and firms may end up acquiring a target despite a different firm's higher valuation. Furthermore, if Firm B acquires only some CUID operations of Firm C, then we reason that this subset of Firm C's CUID's was also a potential purchase by Firm A.

The results for the overall hazard specification are shown in Table B.7. Parameter estimates for the second and third variations are shown in Table B.8 and Table B.9

[8] An alternative interpretation of this assumption is: you either execute a transaction with the firm or not.

respectively. For clarity, we discuss the results related to each of the hypotheses in separate subsections.

3.5.1 Does Hart-Scott-Rodino have an effect?

Across our specifications, a couple of patterns emerge. First, the HSR disclosure flag on its own has a positive coefficient and is highly significant. This implies that firms aren't dissuaded from pursuing large acquisitions by the HSR rules alone. However, when interacted with the horizontal flag, HSR disclosure has a negative effect, though the effect is much less significant. While we refuse to believe regulators do not scrutinize large mergers with a strong horizontal component, this suggests such scrutiny is not particularly burdensome, particularly compared with the benefits of horizontality as measured by the horizontal flag on its own.

3.5.2 How important is clustering?

Though the minimum distance parameter is not significant in any estimation apart from for AT&T Broadband, the parameter is negative in every specification estimated. This suggests that while firms pursue purchases that are located close to their current holdings, it is not an overwhelming factor in their decision. Alternatively, given the relative crudity of our distance measure, it is possible our model is insufficiently nuanced to capture the true value. An ideal measure of distance would combine a concept of adjacency and the amount of right-of-way needed to combine physical systems.

3.6 Conclusion and future work

Although many have tried to measure the effectiveness of U.S. merger policy in an empirical way, these attempts have largely been stymied by the problem of sample size [Clougherty and Seldeslachts, 2011, Carlton, 2009]. This project has attempted to cast the problem into the context of a specific industry, cable television service, in order to achieve enough variation to provide an empirically robust answer.

The results of our simple 'potential acquisition' exercise suggest policy may be too focused on particular types of acquisitions without considering the industry at large. In particular, it is not difficult to imagine that regulators in 1999 may have rejected a proposal to combine the cable television access of 50% of U.S. households into a single company.[9] Yet this is precisely what has occurred.[10]

While this paper lays out the acquisition history and strategy of the largest players in the cable provider market, questions remain. In particular, the individual acquisition strategies of major firms may have varied considerably by geography or for other competitive reasons. Further work is necessary to understand these dynamic incentives.

[9]If this thought experiment does not convince you, consider a proposal to combine cable television, internet, and voice services for 50% of American households into a single company that *also* controls a quarter of the broadcast television market.

[10]To be clear, we are not making any claims about consumer or firm welfare through this period of consolidation. Rather, we believe regulators may have opted for additional scrutiny.

3.7 Appendix: Data details

Our main sources of data are the U.S. Census Bureau and the Cable Operations and Licensing System (COALS), operated by the FCC. We also obtained information on Early Terminations from the Federal Trade Commission and supplemented our procedures with several additional sources. This appendix gives details of our various data collection and processing procedures.

3.7.1 Early Terminations

The Federal Trade Commission maintains lists of all early terminations granted each week under the Hart-Scott-Rodino Act.[11] We manually searched these lists for events that included the large firms we were concerned with.

3.7.2 Comcast Letters

As part of a public comment period on proposed ownership rules in the cable industry, Comcast voluntarily submitted quarterly letters detailing their acquisition activity to the FCC, which subsequently published them on their website. We collected all of the letters available.

3.7.3 Geographic Data

We obtained population data at the Census Place[12] level from the National Historical Geographic Information System for Census 2000 and 2010 and augmented this data with 2010 data directly from the Census Bureau. For a city that crosses county lines, population counts are available for each "county-part" of the city while household counts are only available for the city as a whole. We imputed 2010 household counts for multi-county cities by taking the city-wide ratio of households to population and multiplying it by the population of each "county-part." Population and household counts were also available for the balance of counties (or other civil divisions) that are unincorporated – similar to the FCC community classifications described below.

[11] Available at http://www.ftc.gov/bc/earlyterm

[12] This includes Census Designated Places

We estimated the 2010 household counts for unincorporated communities by using a simple linear regression of household count on total population interacted with state dummies for all communities for which household data was available. We then used the growth rates of household counts by county from 2000-2010 to impute CUID household counts from 2000-2010.

Finally, we incorporated latitude and longitude data from the State Gazetteer prepared by the United States Board on Geographic Names,[13] matching by place name and the Census' internal unique identifiers. Where exact matches weren't available, we used the geographical centroid of the containing county or township.[14] Additionally, several manual links were made to account for changes in the definitions of certain political units (i.e. changes in county and city boundaries) throughout the country during our study period.

3.7.4 COALS

Overview of COALS and FCC identifiers

COALS consists of a database of cable system information, with a publicly accessible front end, as well as secured-access options for cable systems owners and administrators.[15]

Cable systems regulated by the FCC (and collected in COALS) are identified through Physical System Identification numbers (PSIDs) and the communities they service are identified through Community Unit Identification numbers (CUIDs). In towns where more than one physical system operates, multiple CUIDs are created. Additional CUIDs may also be created when towns cross county lines. For example, the city of Minneapolis, Minnesota, which is currently served by Comcast, is assigned a single CUID, MN0180. That CUID is ?owned? by PSID 011339, which serves the greater Twin Cities area. On the other hand, Kansas City, Missouri, which spans four separate counties, is host to five separate CUIDs serviced by three PSIDs representing Comcast, Time Warner, and Surewest. The presence of two CUIDs with identical community names does not necessarily imply true overbuild; many of these cases occur in large geographic areas,

[13]Available at http://geonames.usgs.gov/domestic/fips55codedef.html
[14]This ensures every CUID can be included in distance calculations.
[15]COALS is available at https://apps.fcc.gov/coals/

such as the non-incorporated portions of counties.

CUIDs may also represent unincorporated areas and communities at a variety of scales. At the low end of the spectrum, a single CUID may represent a single 'private' settlement such as an apartment complex or hotel. A CUID may be created for an unincorporated community regardless of Census status. A single CUID may also be used to represent the 'balance' of a county: the total area of that county not included in any incorporated city contained within that county. Table B.1 shows the distribution of CUIDs by FCC community type classification.

Data Collection

Our data collection process begins with a exhaustive list of every CUID in the United States, taken from an FCC-provided current-status digest.[16] This CUID list is used as the input to a Python script which opens the public COALS page, parses the source HTML, and saves relevant information on providers and filings.[17] The primary output of this script is a dataset of every CUID/provider combination in the COALS system.

Merging COALS and Census Data

With our geographic data and CUID data collected at the finest levels possible, we use a "specific-to-general" process to combine the data. We map the Census Place classifications to the FCC CUID classifications according to table B.4. We then match the community type and the community, county, and state names as closely as possible. An overview of the match quality is tabulated in Table B.3. Of the 45,146 CUIDs in the FCC file, we match 31,598 to Census locations. Of those 31,598 matches, 5,517 are unincorporated communities and therefore use imputed household data. Though all major cities match successfully, the CUID file contains many unmatched entries. While some of the unmatched CUIDs consist of individual housing developments or government facilities, most are unincorporated communities or areas which do not qualify as a Census Designated Place.

---

[16] Available at http://www.fcc.gov/mb/vax/registeredcuid.xls

[17] See figure B.9 for an example CUID shown in COALS.

3.7.5 Data cleaning

The first step in our analysis is a manual cleaning process focusing on the 9,506 unique legal entities that control CUIDs at various points in time throughout our raw dataset. The vast majority of the changes come from either missing address information or typographical errors in the legal name or address.[18] Many additional changes are made through the identification of franchised or otherwise split legal entities which are in fact owned by a single corporation. These entities were identified either through analysis of their names or publicly available business databases maintained by Business Week and Funding Universe.[19] We also used SEC filings to identify lists of subsidiaries in 2000.[20]The result of this process is a mapping that links each of the 9,506 "raw" legal entities to one of "cleaned" 3,889 entities. These cleaned entities are then merged back into the original providers dataset.

With the legal entities cleaned, it is now the case that several "switches" in a single CUID may now actually be multiple entries of the same parent company. We perform a sifting procedure on the dataset to identify the earliest date a CUID was controlled by each of the legal entities which ever controlled the community during the period covered by COALS data. The result is a pared-down list of unique legal entities controlling CUIDs at different points in time.

We refactor this list into a set of switches, by combining multiple observations in our source data into a single observation for each switch containing information on the prior owner, the new owner, and the date of the switch. We group these switches by the two owners in question and the calendar quarter of the switch to identify mergers. These so-called "switch groups" represent the universe of possible merger events in our data.

These groups require additional manual cleaning. Although FCC rules require cable providers to inform the FCC of changes in the legal status of a CUID or cable system within 30 days of such a change,[21] we find several instances where the bulk of a change is consummated (according to the COALS providers data) on one day, and a few additional

[18]See figure B.10 for examples of these two cases.

[19]Figure B.11 has examples of this sort of cleaning.

[20]Comcast: http://www.cmcsa.com/secfiling.cfm?filingID=950159-00-66

[21]47 C.F.R 76.1610, available at http://www.gpo.gov/fdsys/pkg/CFR-2010-title47-vol4/pdf/CFR-2010-title47-vol4-sec76-1610.pdf

changes are made some days or months later. An example of this phenomenon is shown in table B.10. This process reduces the number of observed switch groups (and thus the number of mergers we report) from 896 to 713.

As a check on our data cleaning procedures, we compare our final Comcast merger list (including dates) to the data we collected from the Comcast letters. We successfully match nearly all of the 119 reported Comcast acquisitions.[22]

To understand the geographic layout of the merger, we compare the distance of each CUID within a switch group with all of the CUIDs owned by the acquiring company at the time of the switch (excluding other CUIDs acquired within the same group). Distances are calculated from latitude/longitude data with the Equirectangular Approximation which has high accuracy over the relatively short distances we observe.

3.7.6 Horizontal purchases

We distinguish between horizontal and conglomerate purchases with a simple process: For each CUID involved in the acquisition event, we examined the list of the acquiring company's existing properties at the time of the merger for an exact community name match. If a match is found, the CUID is flagged as a potential horizontal merger. Since we cannot confirm overbuild directly, we excluded those CUIDS which referred to townships or unincorporated areas of counties and parishes. It is unlikely that companies would pursue an overbuild strategy in these rural areas.

3.7.7 FCC Annual Report Data

To ground our subscription rate assumptions, we acquired all annual report data from 2002-2009 from the FCC. The FCC requires all cable systems with greater than 20,000 subscribers, as well as a random sample of smaller systems, to submit an annual report with details of their coverage, subscription rates, and offerings. These reports are filed at the Physical System level and are integrated into COALS upon submission. While this data is considered public, the FCC has agreed to an industry request to hold the report data for three years before release.

[22]We believe our mismatches are due to differences in the names of entities as reported by Comcast and recorded in COALS.

Unfortunately, due to the design of COALS, the annual report data does not contain any point-in-time geographic linkage information. In other words, we cannot identify which historical annual report corresponds to which CUIDs. Whenever a CUID is attached to a new PSID, it is immediately linked to all filings for that PSID and all previous linkages are destroyed. For example, Verizon registered a CUID for Medford, MA (CUID MA0484) in 2012 and attached it to their existing regional PSID, 020666. COALS lists a 2008 annual report as a relevant filing for this CUID, despite the CUID's failure to exist in that year. Unfortunately, there does not seem to be a solution to this obstacle at this time.[23]

While we cannot precisely identify which physical systems controlled which CUIDs, we regress the number of subscribers on the number of households covered by the system interacted with year dummies. This regression captures the overall decline in cable subscription rates and is used to ground the value assumptions made in our potential merger exercise.

[23] We asked the FCC to release any geographic link data (beyond the "present-time view" available in COALS) they possess under the Freedom of Information Act. Mike Perko, the Chief of the FCC's Office of Communications and Industry Information, asserted no such information existed, and that storing such information was "not in the public interest." Since the lack of such information significantly reduces the usefulness of annual report data and hampers the FCC's ability to make informed decisions, we must disagree.

References

Druce T Allen. Merger statistics and merger policy. *Review of Industrial Organization*, 1(2):78–92, 1984.

National Cable & Telecommunications Association. Industry data, May 2013a. URL `http://www.ncta.com/industry-data`.

National Cable & Telecommunications Association. Our story, May 2013b. URL `http://www.ncta.com/who-we-are/our-story`.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995. URL `http://www.jstor.org/stable/10.2307/2171802`.

Steven T Berry. Estimating Discrete-Choice Models of Product Differentiation. *The RAND Journal of Economics*, 25(2):242–262, 1994. ISSN 0741-6261. doi: 10.2307/2555829. URL `http://www.jstor.org/stable/2555829`.

Demian Bulwa. SF Giants' sellout streak a numbers game. *SF Gate*, sep 2012.

Dennis W. Carlton. How to measure the effectiveness of U.S. merger policy, March 2009. URL `http://www.voxeu.org/article/how-measure-effectiveness-us-merger-policy`.

Sofronis K. Clerides. Price Discrimination with Differentiated Products: Definition, Theoretical Foundation, Identification. *Department of Economics, University of Cyprus, Discussion Paper*, 09, 2000.

Joseph A. Clougherty and Jo Seldeslachts. The deterrence effects of U.S. merger policy instruments. CEPR Discussion Papers 8482, C.E.P.R. Discussion Papers, 2011. URL `http://EconPapers.repec.org/RePEc:cpr:ceprdp:8482`.

Joseph A. Clougherty and Jo Seldeslachts. The deterrence effects of us merger policy instruments. *Journal of Law, Economics, and Organization*, 29(5):1114–1144, 2013. doi: 10.1093/jleo/ews024. URL `http://jleo.oxfordjournals.org/content/29/5/1114.abstract`.

Joseph A Clougherty, Tomaso Duso, Miyu Lee, and Jo Seldeslachts. Effective european antitrust: Does ec merger policy generate deterrence? *Available at SSRN 2697464*, 2015.

Federal Communications Commission. Telecommunications act of 1996, May 2011. URL `http://transition.fcc.gov/telecom.html`.

Federal Communications Commission. Evolution of cable television, March 2012. URL `http://www.fcc.gov/encyclopedia/evolution-cable-television`.

Pascal Courty and Mario Pagliero. The Impact of Price Discrimination on Revenue: Evidence from the Concert Industry. *Review of Economics and Statistics*, 94(February): 359–369, 2012. ISSN 0034-6535. doi: 10.1162/REST{_}a{_}00179.

Robert W Crandall and Clifford Winston. Does antitrust policy improve consumer welfare? assessing the evidence. *The Journal of Economic Perspectives*, 17(4):3–26, 2003.

Jean-Pierre Dubé, Jeremy T. Fox, and Che-Lin Su. Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation. *Econometrica*, 80(5):2231–2267, 2012. ISSN 0012-9682. doi: 10.3982/ECTA8585.

Patrick Dunne. Dynamic pricing trend sweeps across Major League Baseball, feb 2012. URL `http://www.ticketnews.com/news/Dynamic-pricing-trend-sweeps-across-Major-League-Baseball021222303`.

B Espen Eckbo. Mergers and the value of antitrust deterrence. *The Journal of Finance*, 47(3):1005–1029, 1992.

Thomas R. Eisenmann. Cable TV: From community antennas to wired cities, July 2000. URL `http://hbswk.hbs.edu/item/1591.html`.

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

William M. Emmons and Robin A. Prager. The effects of market structure and ownership on prices and service offerings in the U.S. cable television industry. *RAND Journal of Economics*, 28(4):732–750, Winter 1997.

Hart-Scott-Rodino Premerger Notification Program Introductory Guide. Federal Trade Commission, March 2009.

Guillermo Gallego and Garrett Van Ryzin. Optimal Dynamic Demand Pricing of Inventories with Stochastic Demand over Finite Horizons. *Management Science*, 40(8): 999–1020, 1994.

Amit Gandhi, Zhentong Lu, and Xiaoxia Shi. Estimating Demand for Differentiated Products with Error in Market Shares. pages 1–76, 2013.

Ahiza Garcia. Serena's upset sends U.S. Open final tickets plunging, sep 2015. URL `http://money.cnn.com/2015/09/11/news/serena-williams-us-open-ticket-prices/`.

Austan Goolsbee and Amil Petrin. The consumer gains from direct broadcast satellites and the competition with cable TV. *Econometrica*, 72(2):351–381, March 2004.

Lars Peter Hansen. Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica*, 50(4):1029–1054, 1982.

Patrick Healy. Ticket Pricing Puts 'Lion King' Atop Broadway's Circle of Life. *New York Times*, mar 2014.

Tj Holmes. The effects of third-degree price discrimination in oligopoly. *American Economic Review*, 79(1):244–250, 1989. ISSN 0002-8282. URL `http://www.tcd.ie/Economics/staff/ppwalsh/papers/holmes.pdf`.

Interview. Interview with Sales VP of an anonymous MLB Franchise, 2014.

Mary T Kelly and John S Ying. On measuring competitive viability and monopoly power in cable: An empirical cost approach. *Review of Economics and Statistics*, 85 (4):962–970, 2003.

Bill King and Eric Fisher. A second look at secondary ticketing. *Sports Business Journal*, oct 2011. URL `http://www.sportsbusinessdaily.com/Journal/Issues/2011/10/24/In-Depth/Main.aspx`.

Anja Lambrecht, Katja Seim, Naufel Vilcassim, Amar Cheema, Yuxin Chen, Gregory S. Crawford, Kartik Hosanagar, Raghuram Iyengar, Oded Koenigsberg, Robin Lee, Eugenio J. Miravete, and Ozge Sahin. Price discrimination in service industries. *Marketing Letters*, 23(2):423–438, 2012. ISSN 09230645. doi: 10.1007/s11002-012-9187-0.

John Lazarev. The Welfare Effects of Intertemporal Price Discrimination: An Empirical Analysis of Airline Pricing in U.S. Monopoly Markets. *American Economic Review (Revised and Resubmitted)*, 2013.

Phillip Leslie. Price Discrimination in Broadway Theater. *RAND Journal of Economics*, 35(3):520–541, 2004. ISSN 07416261. doi: 10.2307/1593706. URL `http://ideas.repec.org/a/rje/randje/v35y20043p520-541.html`.

Harikesh Nair. Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games. *Quantitative Marketing and Economics*, 5(3):239–292, 2007. ISSN 1570-7156. doi: 10.1007/s11129-007-9026-4. URL `http://link.springer.com/10.1007/s11129-007-9026-4`.

Philip Nelson and Su Sun. Consumer savings from merger enforcement: A review of the antitrust agencies' estimates. *Antitrust Law Journal*, 69(3):921–960, 2002.

Nielsen. Nielsen estimates number of U.S. television homes to be 144.7 million, May 2011. URL `http://www.nielsen.com/us/en/newswire/2011/`

nielsen-estimates-number-of-u-s-television-homes-to-be-114-7-million.html.

The Museum of Broadcast Communications. Freeze of 1948, May 2013. URL http://www.museum.tv/eotvsection.php?entrycode=freezeof1.

Alex Padilla. Cable 'overbuild' puts choice in hands of consumers. *L.A. Times*, January 21 2001.

Zhan Pang, Oded Berman, and Ming Hu. Up Then Down: Bid-Price Trends in Revenue Management. *Production and Operations Management*, 24(7):n/a–n/a, 2015. ISSN 10591478. doi: 10.1111/poms.12324. URL http://doi.wiley.com/10.1111/poms.12324.

Federal Register. *Federal Register*, 2013. 78(No. 8; January 11), 2406.

Patrick Rishe. Dynamic Pricing : The Future of Ticket Pricing in Sports, jan 2012. URL http://www.forbes.com/sites/prishe/2012/01/06/dynamic-pricing-the-future-of-ticket-pricing-in-sports/.

Sherwin Rosen and Andrew M. Rosenfield. Ticket Pricing. *The Journal of Law and Economics*, 40(October):351–376, 1997. ISSN 0022-2186. doi: 10.1086/467376.

John Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033, 1987.

Stephen L. Shapiro and Joris Drayer. An examination of dynamic ticket pricing and secondary market price determinants in Major League Baseball. *Sport Management Review*, jun 2013. ISSN 14413523. doi: 10.1016/j.smr.2013.05.002. URL http://linkinghub.elsevier.com/retrieve/pii/S1441352313000326.

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

LA Stole. Price Discrimination and Competition. In *Handbook of industrial organization*, volume 3, pages 2224–2299. 2007. doi: 10.1016/S1573-448X(06)03034-2. URL http://www.sciencedirect.com/science/article/pii/S1573448X06030342.

Andrew Sweeting. Dynamic Pricing Behavior in Perishable Goods Markets : Evidence from Secondary Markets for Major League Baseball Tickets. *Journal of Political Economy*, 120(6):1133–1172, 2012.

Kalyan Talluri and Garrett van Ryzin. Revenue Management Under a General Discrete Choice Model of Consumer Behavior. *Management Science*, 50(1):15–33, jan 2004. ISSN 0025-1909. doi: 10.1287/mnsc.1030.0147. URL `http://pubsonline.informs.org/doi/abs/10.1287/mnsc.1030.0147`.

Jean Tirole. *The Theory of Industrial Organization*, volume 56. 1988. ISBN 0262200716. doi: 10.2307/2554286. URL `http://mitpress.mit.edu/catalog/item/default.asp?ttype=2{&}tid=8224`.

Frank Verboven. Price Discrimination (empirical studies). In Steven N Durlauf and Lawrence E Blume, editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, Basingstoke, 2008.

Kevin R Williams. Dynamic Airline Pricing and Seat Availability. *working paper*, pages 1–53, 2013. URL `kevinrwilliams.com/kwilliamsJMP.pdf`.

Joseph Xu, Peter Fader, and Senthil Veeraraghavan. Evaluating the Effectiveness of Dynamic Pricing Strategies on MLB Single-Game Ticket Revenue. In *MIT Sloan Sports Analytics Conference*, 2015.

Jian-da Zhu. Effect of Resale on Optimal Ticket Pricing : Evidence from Major League Baseball Tickets. *working paper*, 2014.

Appendix A

Appendix to Chapter 1

A.1 Data Collection and Cleaning

Collection of ticket sales on Tickets.com and StubHub.com took place each day between 2 AM and 5 AM every day from April 15th to the end of the 2014 season, September 29th.

A.2 Asymptotic GMM Standard Errors

As explained in Subsubsection **??**, I report standard errors by estimating $\mathbf{\Gamma}$ and $\mathbf{V}$ using consistent estimates

$$\hat{\mathbf{\Gamma}}(\hat{\boldsymbol{\theta}}) = \frac{1}{|D|} \sum_{(g,t,j)\in D} \frac{\partial \mathbf{h}(\mathbf{w}_{jgt}, \hat{\boldsymbol{\theta}})}{\partial \boldsymbol{\theta}} \quad \text{and} \quad \hat{\mathbf{V}}(\hat{\boldsymbol{\theta}}) = \frac{1}{|D|} \sum_{(g,t,j)\in D} \mathbf{h}(\mathbf{w}_{jgt}, \hat{\boldsymbol{\theta}})\mathbf{h}(\mathbf{w}_{jgt}, \hat{\boldsymbol{\theta}})'.$$

The gradients involve some special derivations due to the use of the type-probability logistic function. Below I derive the analytical solution for $\hat{\mathbf{\Gamma}}(\hat{\boldsymbol{\theta}})$:

1. Observe that
$$\hat{\mathbf{\Gamma}}(\hat{\boldsymbol{\theta}}) = \frac{1}{|D|} \left(\sum_{(g,t,j)\in D} \frac{\partial \xi_{jtg}(\hat{\theta})}{\partial \boldsymbol{\theta}'} \mathbf{z}_{jtg} \right)$$
since $\mathbf{z}_{jtg}$ does not depend on $\hat{\boldsymbol{\theta}}$.

2. Use implicit function theorem to derive $\frac{\partial \xi(\hat{\theta})}{\partial \theta}$, and use quantities rather than shares for convenience:

$$\begin{aligned}\frac{\partial \xi(\hat{\theta})}{\partial \theta'} &= -\left[\frac{\partial \tilde{s}(\xi;\hat{\theta})}{\partial \xi'}\right]^{-1} \frac{\partial \tilde{s}(\xi;\hat{\theta})}{\partial \theta'} \\ &= -\left[\frac{\partial \tilde{q}(\xi;\hat{\theta})}{\partial \xi'}\right]^{-1} \frac{\partial \tilde{q}(\xi;\hat{\theta})}{\partial \theta'}\end{aligned}$$

3. First matrix on RHS, inside the brackets, has the following elements:

$$\begin{aligned}\frac{\partial \tilde{q}_{j\tilde{t}}(\xi;\hat{\theta})}{\xi_{j'\tilde{t}'}} &= \sum_{t\in\tau_{\tilde{t}}} m_t \sum_{r\in\{L,H\}} \phi_{rt} \frac{\partial s_{rjt}}{\partial \xi_{j'\tilde{t}'}} \\ &= \begin{cases} \sum_{t\in\tau_{\tilde{t}}} m_t \sum_{r\in\{L,H\}} \phi_{rt} s_{rjt}(1-s_{rjt}) & \text{if } \tilde{t}=\tilde{t}', j=j' \\ -\sum_{t\in\tau_{\tilde{t}}} m_t \sum_{r\in\{L,H\}} \phi_{rt} s_{rjt} s_{rj't} & \text{if } \tilde{t}=\tilde{t}', j\neq j' \\ 0 & \text{otherwise} \end{cases}\end{aligned}$$

so $\left[\frac{\partial \tilde{q}(\xi;\hat{\theta})}{\xi'}\right]^{-1}$ is a matrix of block diagonals.

4. Second element on RHS has the following elements:

   (a) If parameter is a taste coefficient:

$$\begin{aligned}\frac{\partial \tilde{q}_{j\tilde{t}}(\xi;\hat{\theta})}{\partial \beta_k} &= \sum_{t\in\tau_{\tilde{t}}} m_t \sum_{r\in\{L,H\}} \phi_{rt} \frac{\partial s_{rjt}}{\partial \beta_k} \\ &= \sum_{t\in\tau_{\tilde{t}}} m_t \sum_{r\in\{L,H\}} \phi_{rt} s_{rjt}(x_{jk} - \sum_{j'\in J_t} x_{j'k} s_{rj't})\end{aligned}$$

   (b) If parameter is ϕ_{0f}, then for all g with $l(g)=f$:

$$\begin{aligned}\frac{\partial \tilde{q}_{j\tilde{t}}(\xi;\hat{\theta})}{\partial \phi_{0f}} &= \sum_{t\in\tau_{\tilde{t}}} m_t \sum_{r\in\{L,H\}} \frac{\partial \phi_{rt}}{\partial \phi_{0f}} s_{rjt} \\ &= \sum_{t\in\tau_{\tilde{t}}} m_t \phi_1 \phi_{Lt} \phi_{Ht} (s_{Ljt} - s_{Hjt})\end{aligned}$$

(otherwise partial derivative is zero)

(c) If parameter is ϕ_1, then

$$\begin{aligned}\frac{\partial \tilde{q}_{j\tilde{t}}(\xi;\hat{\theta})}{\partial \phi_1} &= \sum_{t\in\tau_{\tilde{t}}} m_t \sum_{r\in\{L,H\}} \frac{\partial \phi_{rt}}{\partial \phi_1} s_{rjt} \\ &= \sum_{t\in\tau_{\tilde{t}}} m_t \phi_{Lt}\phi_{Ht}(t-\phi_0)(-s_{Ljt}+s_{Hjt})\end{aligned}$$

A.3 Estimating the transition probability of team performance

In order to do realistic simulations of playoff probability paths, I estimate the distribution $f(x_t|x_{t-1}, d_t)$, where x_{t-1} is the team's old playoff probability and d_t is the date (I use the number of days into the season, starting on April 1st 2014). Fortunately, I have a large amount of data: I collect playoff probability paths for all 30 teams over the entire 2014 regular season.

Of the parametrizations I tried, the highest maximum likelihood was achieved with x_{t+1} distributed log-normally around x_t, *toward the middle* ($x = 50$), where the standard deviation of the normally-distributed variable is a linear function of x_t and and d_t. By *toward the middle*, I mean that the estimated distribution $f(x_t|x_{t-1} = 25)$ is the mirror image of $f(x_t|x_{t-1} = 75)$. Playoff probability never jumps by more than 7.5 in the data, so I truncate the distribution. Figure A.1 shows the two distributions previously mentioned, assuming the date is July 10th, 2014 ($d_t = 100$), before and after truncating:

I estimate the standard deviation to be $1.42424802 + 8.20585730e{-}04(d_t) - 2.80592858e{-}02(x_t)$.

FIGURE A.1 PLAYOFF PROBABILITY PATHS

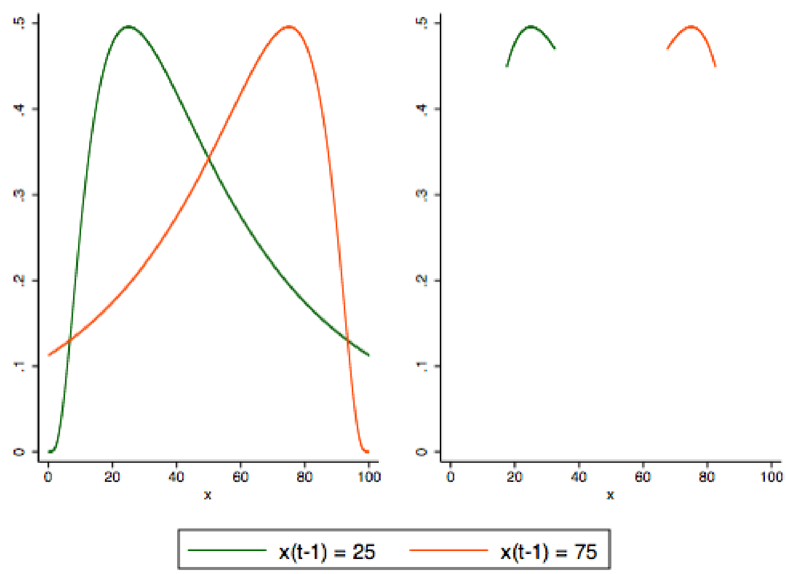

Appendix B

Appendix to Chapter 3

B.1 Data Collection and Cleaning

Our main sources of data are the U.S. Census Bureau and the Cable Operations and Licensing System (COALS), operated by the FCC. We also obtained information on Early Terminations from the Federal Trade Commission and supplemented our procedures with several additional sources. This appendix gives details of our various data collection and processing procedures.

B.1.1 Early Terminations

The Federal Trade Commission maintains lists of all early terminations granted each week under the Hart-Scott-Rodino Act.[1] We manually searched these lists for events that included the large firms we were concerned with.

B.1.2 Comcast Letters

As part of a public comment period on proposed ownership rules in the cable industry, Comcast voluntarily submitted quarterly letters detailing their acquisition activity to the FCC, which subsequently published them on their website. We collected all of the letters available.

[1] Available at http://www.ftc.gov/bc/earlyterm

B.1.3 Geographic Data

We obtained population data at the Census Place[2] level from the National Historical Geographic Information System for Census 2000 and 2010 and augmented this data with 2010 data directly from the Census Bureau. For a city that crosses county lines, population counts are available for each "county-part" of the city while household counts are only available for the city as a whole. We imputed 2010 household counts for multi-county cities by taking the city-wide ratio of households to population and multiplying it by the population of each "county-part." Population and household counts were also available for the balance of counties (or other civil divisions) that are unincorporated – similar to the FCC community classifications described below.

We estimated the 2010 household counts for unincorporated communities by using a simple linear regression of household count on total population interacted with state dummies for all communities for which household data was available. We then used the growth rates of household counts by county from 2000-2010 to impute CUID household counts from 2000-2010.

Finally, we incorporated latitude and longitude data from the State Gazetteer prepared by the United States Board on Geographic Names,[3] matching by place name and the Census' internal unique identifiers. Where exact matches weren't available, we used the geographical centroid of the containing county or township.[4] Additionally, several manual links were made to account for changes in the definitions of certain political units (i.e. changes in county and city boundaries) throughout the country during our study period.

B.1.4 COALS

Overview of COALS and FCC identifiers

COALS consists of a database of cable system information, with a publicly accessible front end, as well as secured-access options for cable systems owners and administrators.[5]

Cable systems regulated by the FCC (and collected in COALS) are identified through

[2] This includes Census Designated Places

[3] Available at http://geonames.usgs.gov/domestic/fips55codedef.html

[4] This ensures every CUID can be included in distance calculations.

[5] COALS is available at https://apps.fcc.gov/coals/

Physical System Identification numbers (PSIDs) and the communities they service are identified through Community Unit Identification numbers (CUIDs). In towns where more than one physical system operates, multiple CUIDs are created. Additional CUIDs may also be created when towns cross county lines. For example, the city of Minneapolis, Minnesota, which is currently served by Comcast, is assigned a single CUID, MN0180. That CUID is ?owned? by PSID 011339, which serves the greater Twin Cities area. On the other hand, Kansas City, Missouri, which spans four separate counties, is host to five separate CUIDs serviced by three PSIDs representing Comcast, Time Warner, and Surewest. The presence of two CUIDs with identical community names does not necessarily imply true overbuild; many of these cases occur in large geographic areas, such as the non-incorporated portions of counties.

CUIDs may also represent unincorporated areas and communities at a variety of scales. At the low end of the spectrum, a single CUID may represent a single 'private' settlement such as an apartment complex or hotel. A CUID may be created for an unincorporated community regardless of Census status. A single CUID may also be used to represent the 'balance' of a county: the total area of that county not included in any incorporated city contained within that county. Table B.1 shows the distribution of CUIDs by FCC community type classification.

Data Collection

Our data collection process begins with a exhaustive list of every CUID in the United States, taken from an FCC-provided current-status digest.[6] This CUID list is used as the input to a Python script which opens the public COALS page, parses the source HTML, and saves relevant information on providers and filings.[7] The primary output of this script is a dataset of every CUID/provider combination in the COALS system.

Merging COALS and Census Data

With our geographic data and CUID data collected at the finest levels possible, we use a "specific-to-general" process to combine the data. We map the Census Place classifications to the FCC CUID classifications according to table B.4. We then match

[6] Available at http://www.fcc.gov/mb/vax/registeredcuid.xls

[7] See figure B.9 for an example CUID shown in COALS.

the community type and the community, county, and state names as closely as possible. An overview of the match quality is tabulated in Table B.3. Of the 45,146 CUIDs in the FCC file, we match 31,598 to Census locations. Of those 31,598 matches, 5,517 are unincorporated communities and therefore use imputed household data. Though all major cities match successfully, the CUID file contains many unmatched entries. While some of the unmatched CUIDs consist of individual housing developments or government facilities, most are unincorporated communities or areas which do not qualify as a Census Designated Place.

B.1.5 Data cleaning

The first step in our analysis is a manual cleaning process focusing on the 9,506 unique legal entities that control CUIDs at various points in time throughout our raw dataset. The vast majority of the changes come from either missing address information or typographical errors in the legal name or address.[8] Many additional changes are made through the identification of franchised or otherwise split legal entities which are in fact owned by a single corporation. These entities were identified either through analysis of their names or publicly available business databases maintained by Business Week and Funding Universe.[9] We also used SEC filings to identify lists of subsidiaries in 2000.[10]The result of this process is a mapping that links each of the 9,506 "raw" legal entities to one of "cleaned" 3,889 entities. These cleaned entities are then merged back into the original providers dataset.

With the legal entities cleaned, it is now the case that several "switches" in a single CUID may now actually be multiple entries of the same parent company. We perform a sifting procedure on the dataset to identify the earliest date a CUID was controlled by each of the legal entities which ever controlled the community during the period covered by COALS data. The result is a pared-down list of unique legal entities controlling CUIDs at different points in time.

We refactor this list into a set of switches, by combining multiple observations in our source data into a single observation for each switch containing information on the

[8]See figure B.10 for examples of these two cases.

[9]Figure B.11 has examples of this sort of cleaning.

[10]Comcast: http://www.cmcsa.com/secfiling.cfm?filingID=950159-00-66

prior owner, the new owner, and the date of the switch. We group these switches by the two owners in question and the calendar quarter of the switch to identify mergers. These so-called "switch groups" represent the universe of possible merger events in our data.

These groups require additional manual cleaning. Although FCC rules require cable providers to inform the FCC of changes in the legal status of a CUID or cable system within 30 days of such a change,[11] we find several instances where the bulk of a change is consummated (according to the COALS providers data) on one day, and a few additional changes are made some days or months later. An example of this phenomenon is shown in table B.10. This process reduces the number of observed switch groups (and thus the number of mergers we report) from 896 to 713.

As a check on our data cleaning procedures, we compare our final Comcast merger list (including dates) to the data we collected from the Comcast letters. We successfully match nearly all of the 119 reported Comcast acquisitions.[12]

To understand the geographic layout of the merger, we compare the distance of each CUID within a switch group with all of the CUIDs owned by the acquiring company at the time of the switch (excluding other CUIDs acquired within the same group). Distances are calculated from latitude/longitude data with the Equirectangular Approximation which has high accuracy over the relatively short distances we observe.

B.1.6 Horizontal purchases

We distinguish between horizontal and conglomerate purchases with a simple process: For each CUID involved in the acquisition event, we examined the list of the acquiring company's existing properties at the time of the merger for an exact community name match. If a match is found, the CUID is flagged as a potential horizontal merger. Since we cannot confirm overbuild directly, we excluded those CUIDS which referred to townships or unincorporated areas of counties and parishes. It is unlikely that companies would pursue an overbuild strategy in these rural areas.

[11] 47 C.F.R 76.1610, available at http://www.gpo.gov/fdsys/pkg/CFR-2010-title47-vol4/pdf/CFR-2010-title47-vol4-sec76-1610.pdf

[12] We believe our mismatches are due to differences in the names of entities as reported by Comcast and recorded in COALS.

B.1.7 FCC Annual Report Data

To ground our subscription rate assumptions, we acquired all annual report data from 2002-2009 from the FCC. The FCC requires all cable systems with greater than 20,000 subscribers, as well as a random sample of smaller systems, to submit an annual report with details of their coverage, subscription rates, and offerings. These reports are filed at the Physical System level and are integrated into COALS upon submission. While this data is considered public, the FCC has agreed to an industry request to hold the report data for three years before release.

Unfortunately, due to the design of COALS, the annual report data does not contain any point-in-time geographic linkage information. In other words, we cannot identify which historical annual report corresponds to which CUIDs. Whenever a CUID is attached to a new PSID, it is immediately linked to all filings for that PSID and all previous linkages are destroyed. For example, Verizon registered a CUID for Medford, MA (CUID MA0484) in 2012 and attached it to their existing regional PSID, 020666. COALS lists a 2008 annual report as a relevant filing for this CUID, despite the CUID's failure to exist in that year. Unfortunately, there does not seem to be a solution to this obstacle at this time.[13]

While we cannot precisely identify which physical systems controlled which CUIDs, we regress the number of subscribers on the number of households covered by the system interacted with year dummies. This regression captures the overall decline in cable subscription rates and is used to ground the value assumptions made in our potential merger exercise.

B.2 Tables and Figures

[13] We asked the FCC to release any geographic link data (beyond the "present-time view" available in COALS) they possess under the Freedom of Information Act. Mike Perko, the Chief of the FCC's Office of Communications and Industry Information, asserted no such information existed.

Municipality Type	**CUIDs**
Incorporated Borough	1,733
Incorporated City	10,873
Incorporated Town	8,878
Incorporated Village	4,211
Privately owned settlement	1,072
State or Federal Reservation	440
Unincorporated area adjacent to incorporated community	1,478
Unincorporated area commonly known as	5,809
Unincorporated unnamed area within a County or Parish	4,211
Grand Total	45,146

Table B.1 CUID types identified by the FCC

Total number of CUIDs	45,146
Average number of providers per CUID	1.85
Std. dev.	1.14
CUIDs with single provider	22,986
CUIDs with more than 5 providers	690

Table B.2 Summary of cleaned provider data.

Match Type	**CUIDs**
Full (County, community type and name)	21,158
County and name	5,610
Type and name	2,210
Name only	2,620
Unmatched	13,548
Total	45,146

Table B.3 Breakdown of CUID/Census match quality

CUID classification	**CDP classification**
Incorporated Borough	City
Incorporated City	City
Incorporated Town	Town
Incorporated Village	Town
Privately owned settlement	Private
State or Federal Reservation	Reservation
Unincorporated area adjacent to incorporated community	Balance
Unincorporated area commonly known as	CDP
Unincorporated unnamed area within a County or Parish	Balance

Table B.4 Mapping CUID classifications to CDP classifications

Number of mergers	712
Median CUIDs per merger	3
Average CUIDs per merger	21.6
Std. Dev. CUIDs per merger	82.6
Median mean distance to nearest owned CUID	.267
Average mean distance to nearest owned CUID	4.66
Std. Dev. of minimum distance to nearest owned CUID	29.98
Total CUIDs acquired	15,357
CUIDs missing population data	3,894
Median population per merger	31,123
Average population per merger	540,510.4
Std. Dev population per merger	2,861,357
Median households per merger	12,637
Average households per merger	201,099.4
Std. Dev households per merger	1,029,455

Table B.5 Acquisition summary. Note: Household statistics include missing data for some rural CUIDs.

	Comcast	**Time Warner**	**Charter**	**Cox**	**AT&T**	**Adelphia**
Num. of mergers	307	157	139	34	26	49
Median CUIDs	3	3	3	3	4	4
Average CUIDs	22	23	22	19	22	13
Std. dev CUIDs	95	79	84	39	52	26
Median mean dist	0.20	0.17	0.35	1.27	3.13	0.79
Average mean dist	1.12	2.9	1.33	10.2	42.7	18.2
Std. dev mean dist	7.70	13.73	3.35	31.73	116.05	59.86
Total CUIDs	6,849	3,612	3,066	628	566	636
Missing pop data	1,830	814	900	69	108	173
Median pop	53,284	24,541	13,680	30,980	99,166	65,201
Average pop	555,237	940,561	219,075	329,859	450,493	272,204
Std dev pop	2,426,169	4,953,978	877,393	611,282	787,075	512,692
Median HH	20,898	8,759	5,633	12,446	37,691	26,114
Average HH	213,551	333,154	83,246	127,692	172,818	100,223
Std dev HH	934,400	1,718,539	323,470	236,221	302,860	183,336

TABLE B.6 ACQUSITION SUMMARY BY ACQUIRING FIRM.

	Acquired
Mean Distance	-0.673
(per 100 miles)	(0.826)
Horizontal Flag	0.388
	(0.259)
HSR Flag	1.270***
	(0.0919)
HSR * Horizontal	-0.641*
	(0.317)
Num. Households	1.33e-07*
	(6.47e-08)
Year Dummies	Yes
Constant	-4.665***
	(0.167)
N	39813

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.7 Parameter estimates for the "megafirm" specification of our lives exercise.

	(1) Comcast	(2) AT&T Broadband	(3) Cox	(4) Time Warner	(5) Charter
Mean Distance	-0.00219	-0.00832**	-0.0119	-0.00686	-0.0617
	(0.00206)	(0.00295)	(0.00839)	(0.00545)	(0.0379)
Horizontal Flag	0.663*	-9.431	1.880*	1.661***	1.802***
	(0.277)	(666.0)	(0.758)	(0.287)	(0.274)
HSR Flag	1.545***	1.340**	1.304**	0.870***	0.194
	(0.146)	(0.455)	(0.402)	(0.222)	(0.297)
HSR * Horizontal	-0.604	10.77	-1.026	-0.692	-0.427
	(0.330)	(666.0)	(1.021)	(0.395)	(0.484)
Num. Households	0.000000111	3.29e-08	-0.000000114	3.55e-08	-0.000000101
	(8.66e-08)	(0.000000369)	(0.000000443)	(0.000000132)	(0.000000308)
Year Dummies	Yes	Yes	Yes	Yes	Yes
Constant	-6.542***	-6.623***	-6.483***	-6.224***	-5.206***
	(0.386)	(0.532)	(0.457)	(0.364)	(0.265)
N	33443	5926	27391	39813	36636

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.8 Parameter estimates for our 'potential merger' exercise assuming large firms were able to buy any small firm.

	(1) Comcast	(2) AT&T Broadband	(3) Cox	(4) Time Warner	(5) Charter
Mean Distance	-0.00227	-0.00829**	-0.0118	-0.00685	-0.0622
	(0.00209)	(0.00292)	(0.00836)	(0.00544)	(0.0380)
Horizontal Flag	0.654*	-9.435	1.868*	1.653***	1.795***
	(0.277)	(669.9)	(0.758)	(0.287)	(0.274)
HSR Flag	1.574***	1.398**	1.332***	0.919***	0.240
	(0.146)	(0.456)	(0.402)	(0.222)	(0.297)
HSR * Horizontal	-0.619	10.78	-1.026	-0.721	-0.435
	(0.330)	(669.9)	(1.021)	(0.395)	(0.484)
Num. Households	0.000000114	5.65e-08	-0.000000111	4.27e-08	-8.55e-08
	(8.61e-08)	(0.000000348)	(0.000000441)	(0.000000131)	(0.000000306)
Year Dummies	Yes	Yes	Yes	Yes	Yes
Constant	-6.534***	-6.582***	-6.477***	-6.224***	-5.203***
	(0.386)	(0.531)	(0.457)	(0.364)	(0.265)
N	33095	5723	26904	39320	36121

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

TABLE B.9 PARAMETER ESTIMATES FOR OUR 'POTENTIAL MERGER' EXERCISE ASSUMING SMALL FIRMS BOUGHT BY OTHER LARGE FIRMS WERE UNAVAILABLE.

Date	**CUIDs**
January 15, 2008	364
April 25, 2008	2
August 1, 2008	2
Total	368

Table B.10 An example of different dates within a "switch group." The event shown took place between Comcast and Insight Communications Co.

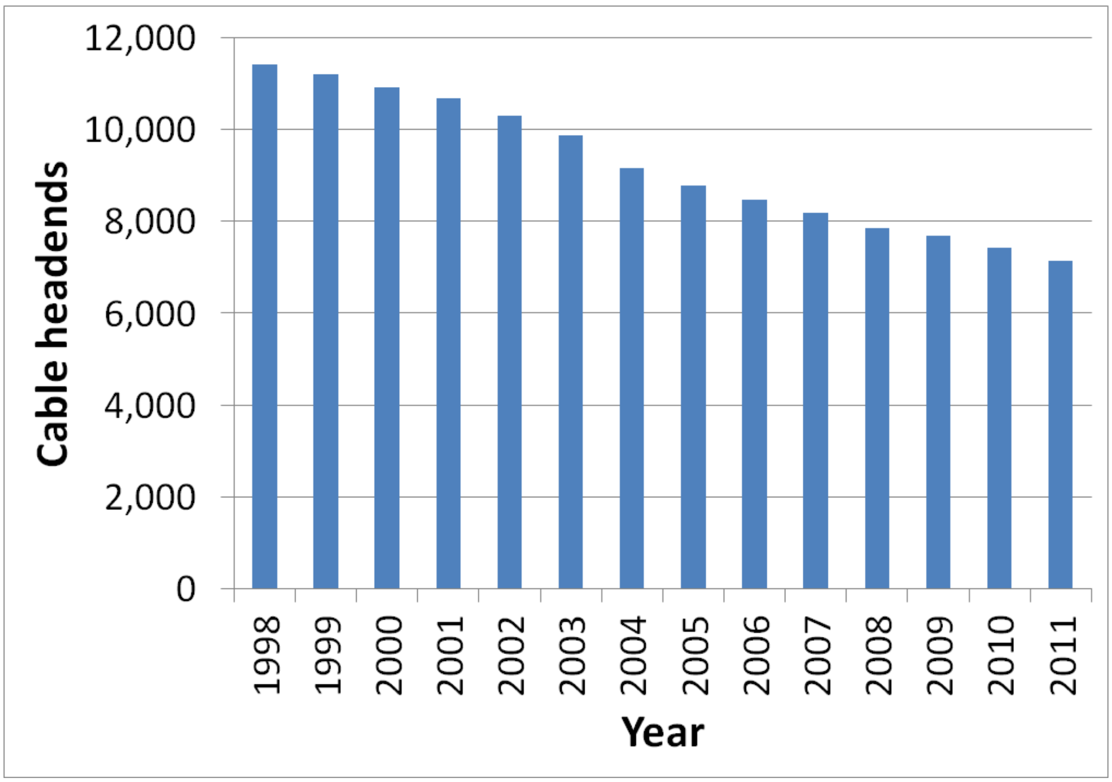

FIGURE B.1 THE NUMBER OF CABLE HEADENDS (PHYSICAL LOCATIONS USED TO RECIEVE AND DISTRIBUTE PROGRAMMING) HAS DECREASED EVERY YEAR SINCE 1998. SOURCE: ASSOCIATION [2013A]

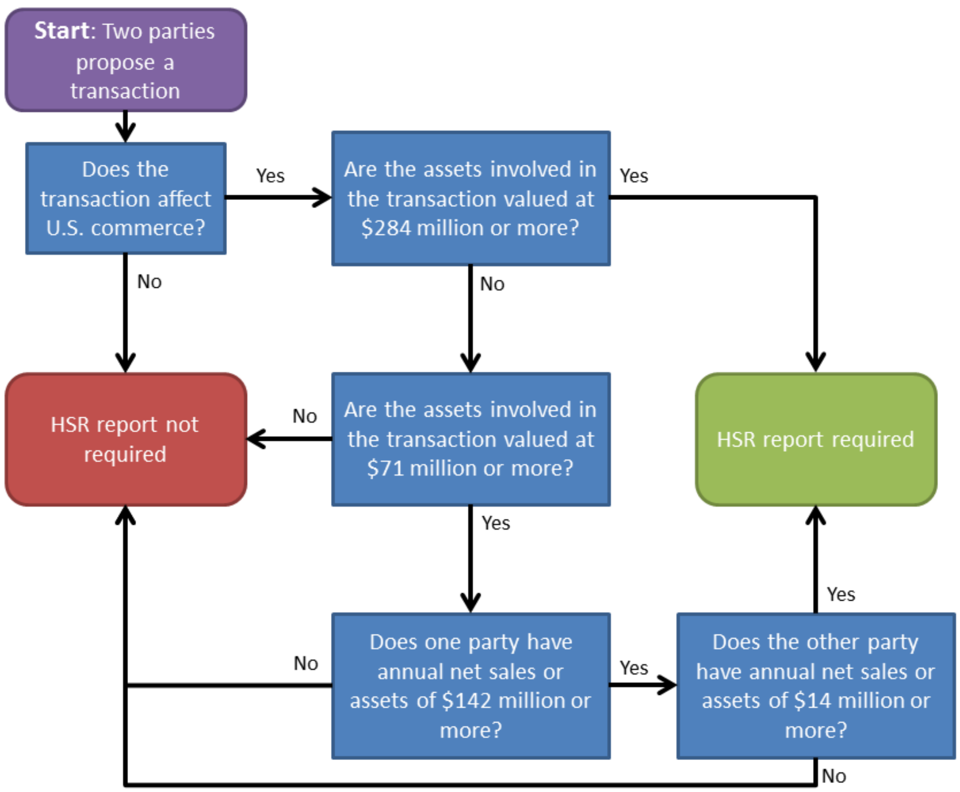

Figure B.2 Flowchart of the 2013 Hart-Scott-Rodino reporting thresholds

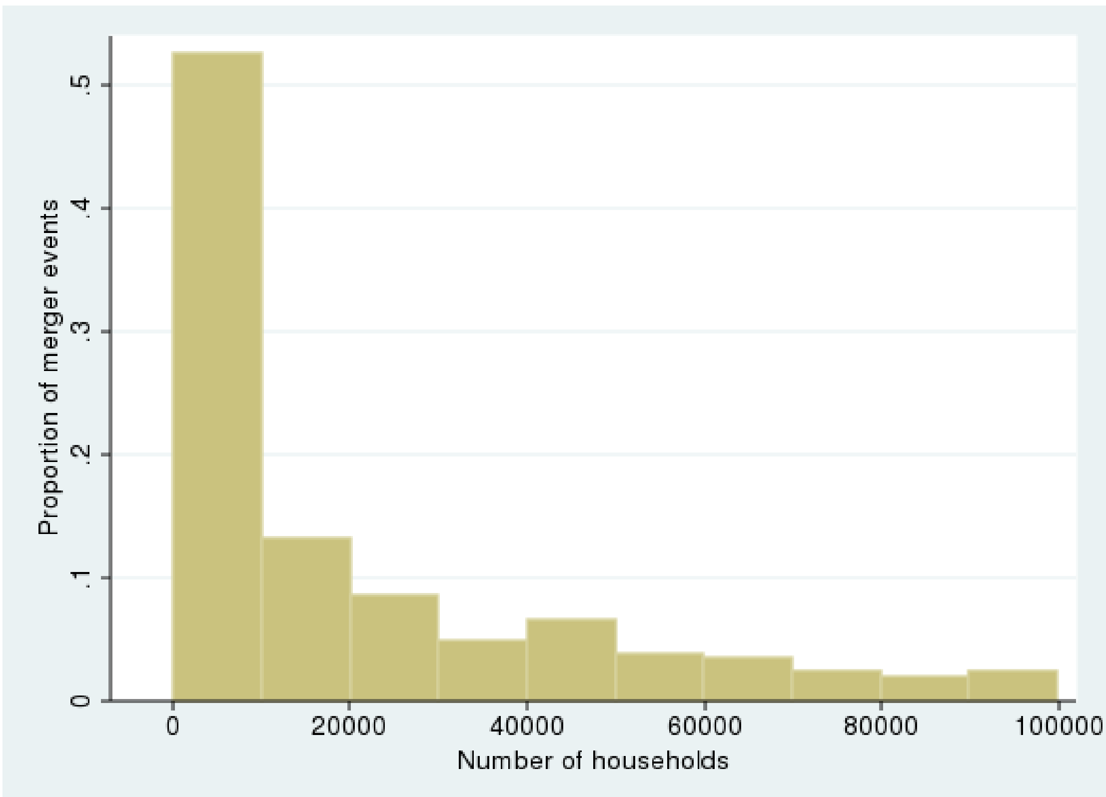

Figure B.3 Histogram of the size of the 712 acquisitions we study. This chart removes a small number of extremely large transactions for clarity.

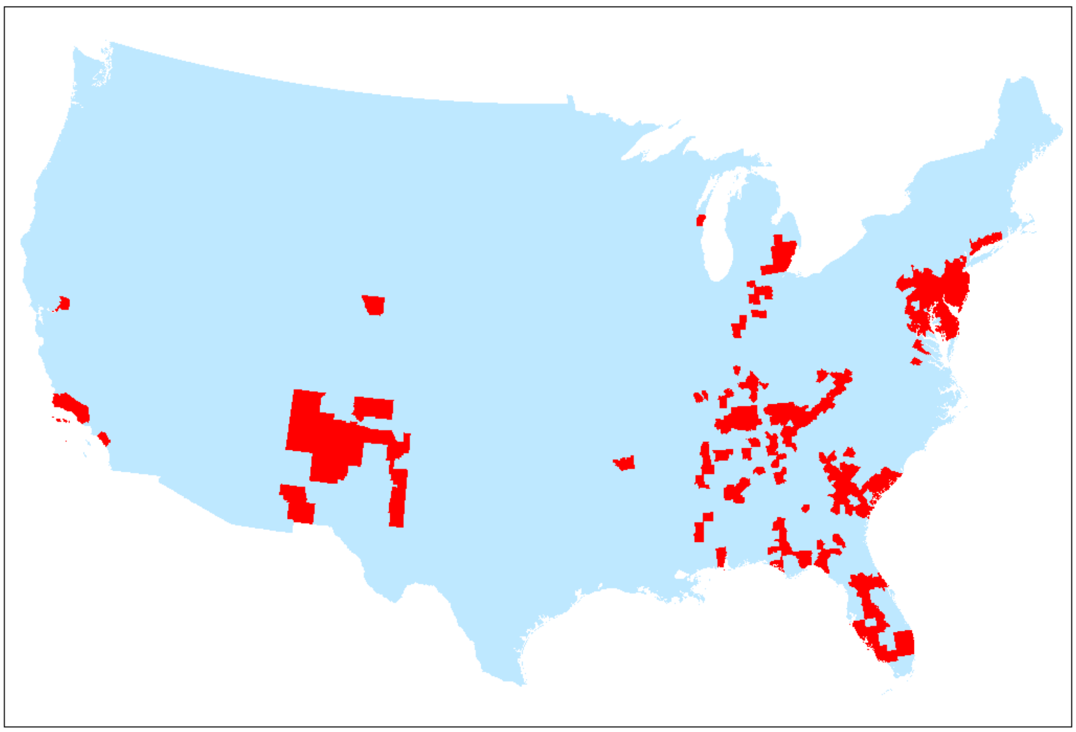

Figure B.4 Map of Comcast's holdings by county in 2001. Counties are red if Comcast serves at least one community in the county.

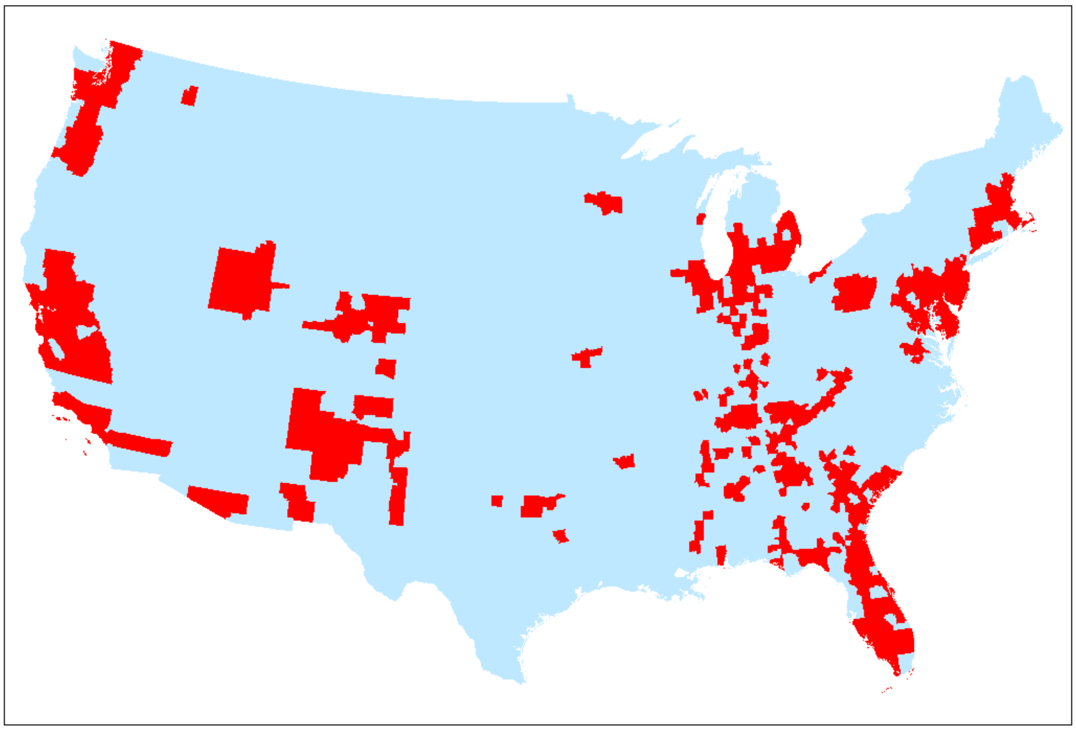

Figure B.5 Map of Comcast's holdings by county in 2003. Counties are red if Comcast serves at least one community in the county.

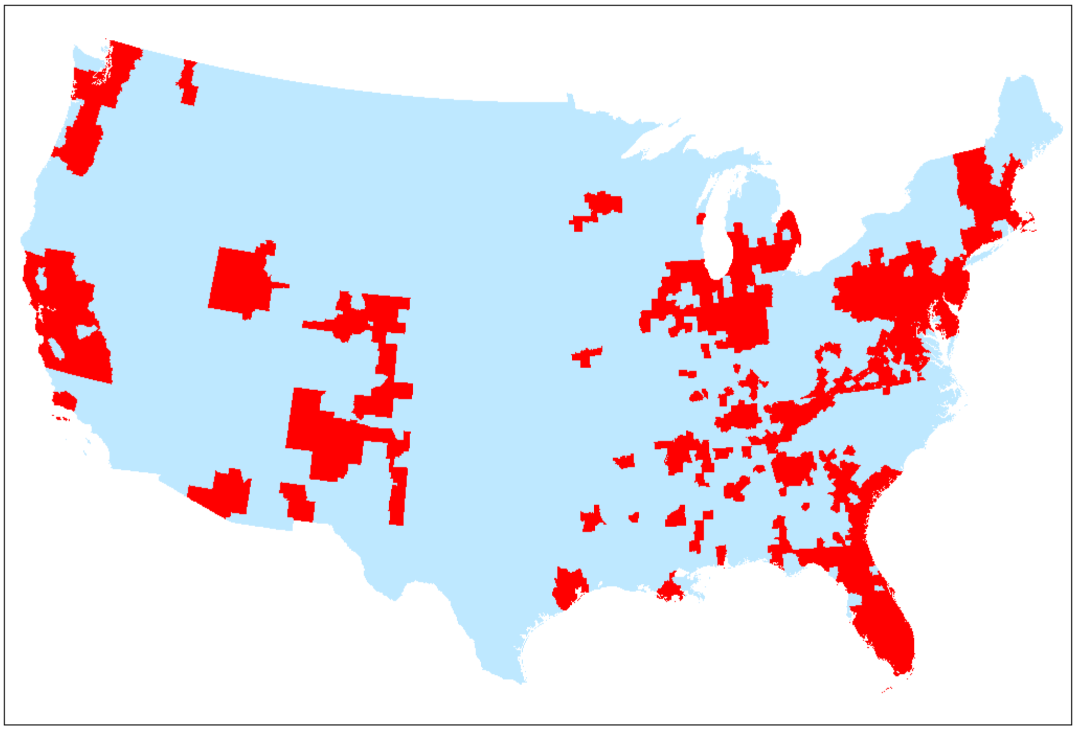

Figure B.6 Map of Comcast's holdings by county in 2010. Counties are red if Comcast serves at least one community in the county.

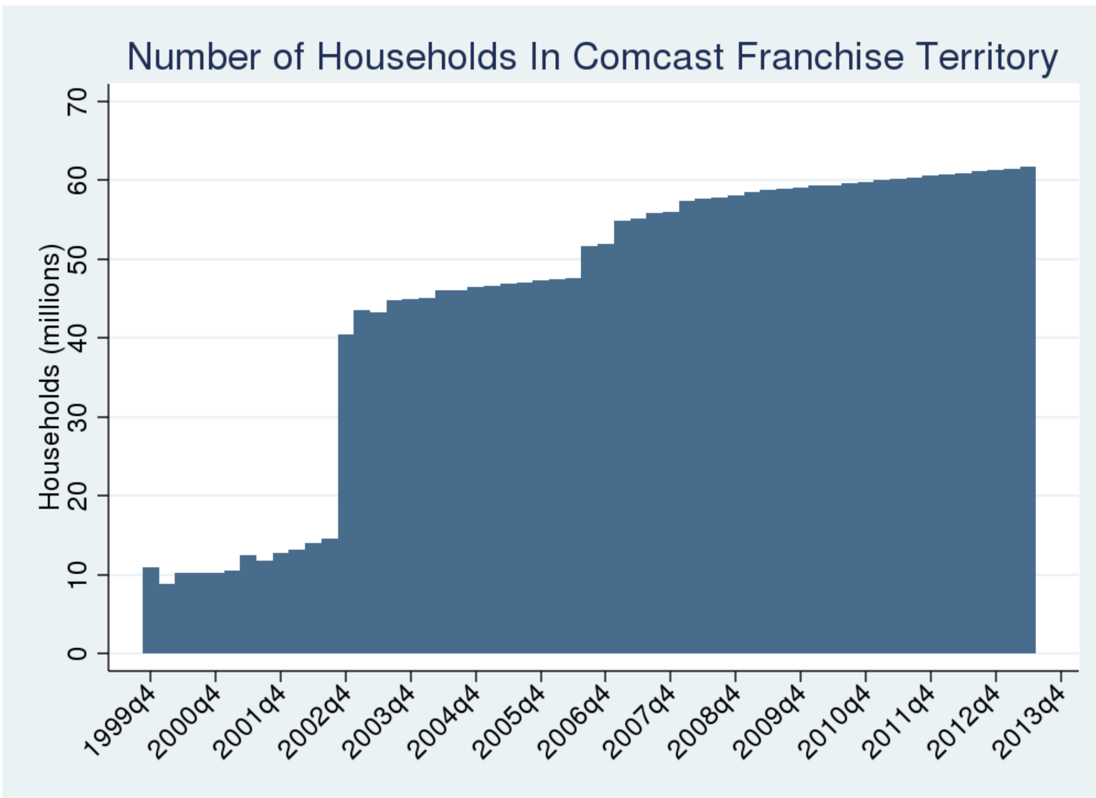

Figure B.7 The number of households within Comcast's franchise territory (as identified through our PSID/Census match process) has increased steadily throughout our study period. The large jumps in 2002 and 2006 are the result of the AT&T Broadband and Adelphia acquisitions, respectively. Quarterly household counts are imputed using 2010 Census levels and 2000-2010 growth rates by county.

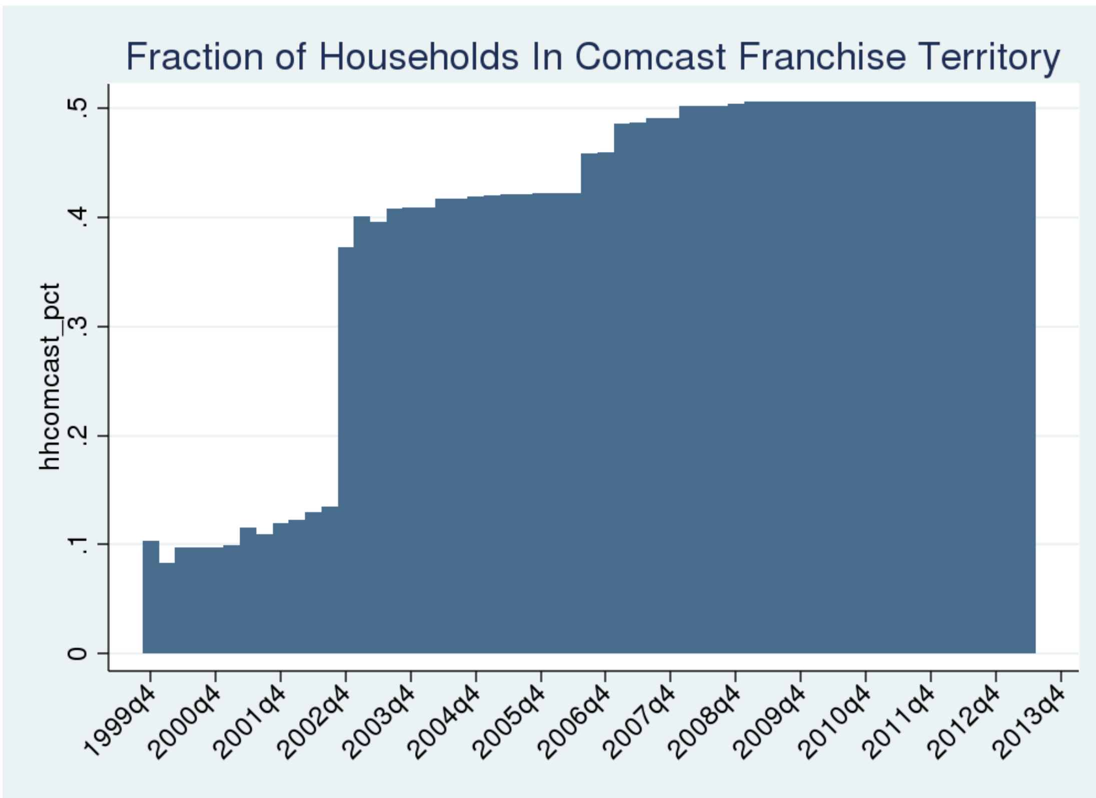

Figure B.8 The percentage of households within Comcast's franchise territory (as identified through our PSID/Census match process) has increased steadily throughout our study period. The large jumps in 2002 and 2006 are the result of the Adelphia and Susquehanna acquisitions, respectively. Quarterly household counts and percentages are imputed using 2010 Census levels and 2000-2010 growth rates by county.

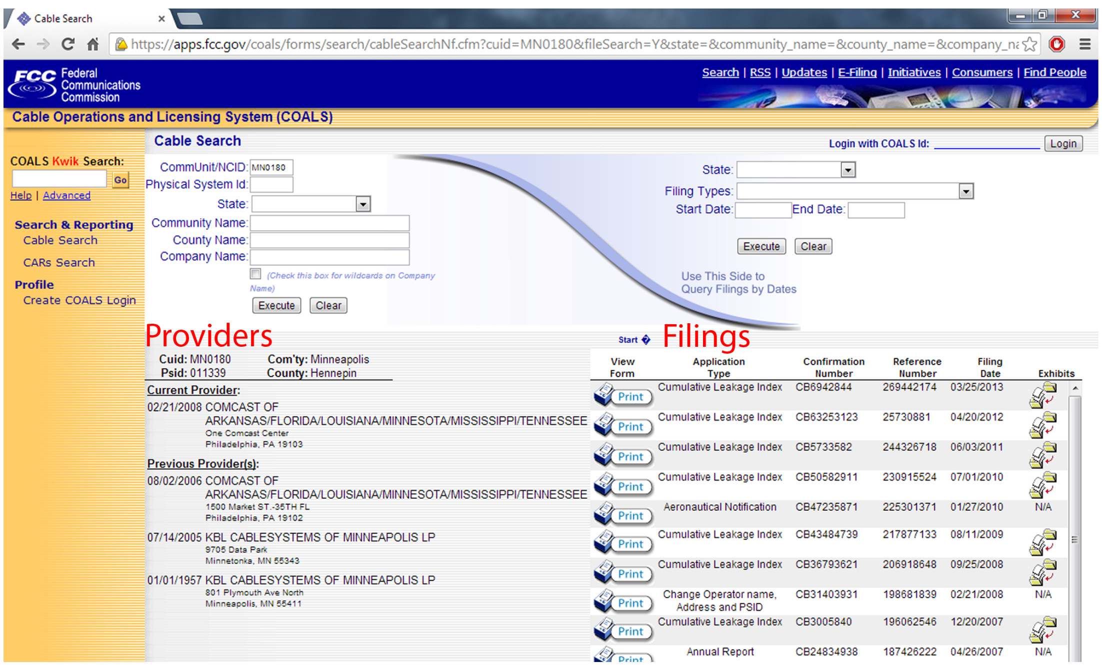

Figure B.9 A screenshot of the COALS page for the cable system in Minneapolis Minnesota, with emphasis on the providers and filings information we scraped.

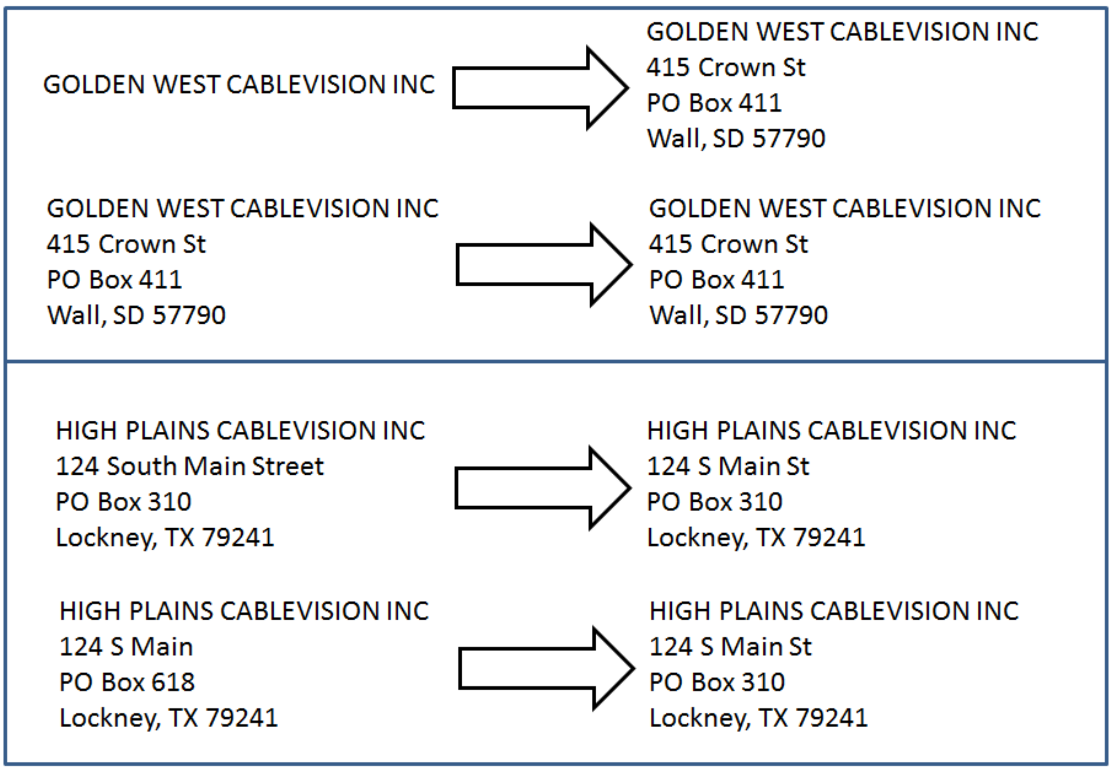

Figure B.10 Top: Some legal entity entries were missing address data. We filled in missing addresses using entries with identical names where available. Bottom: When multiple addresses were found (or when addresses had typos), we used the most-common entry for all identically named entities.

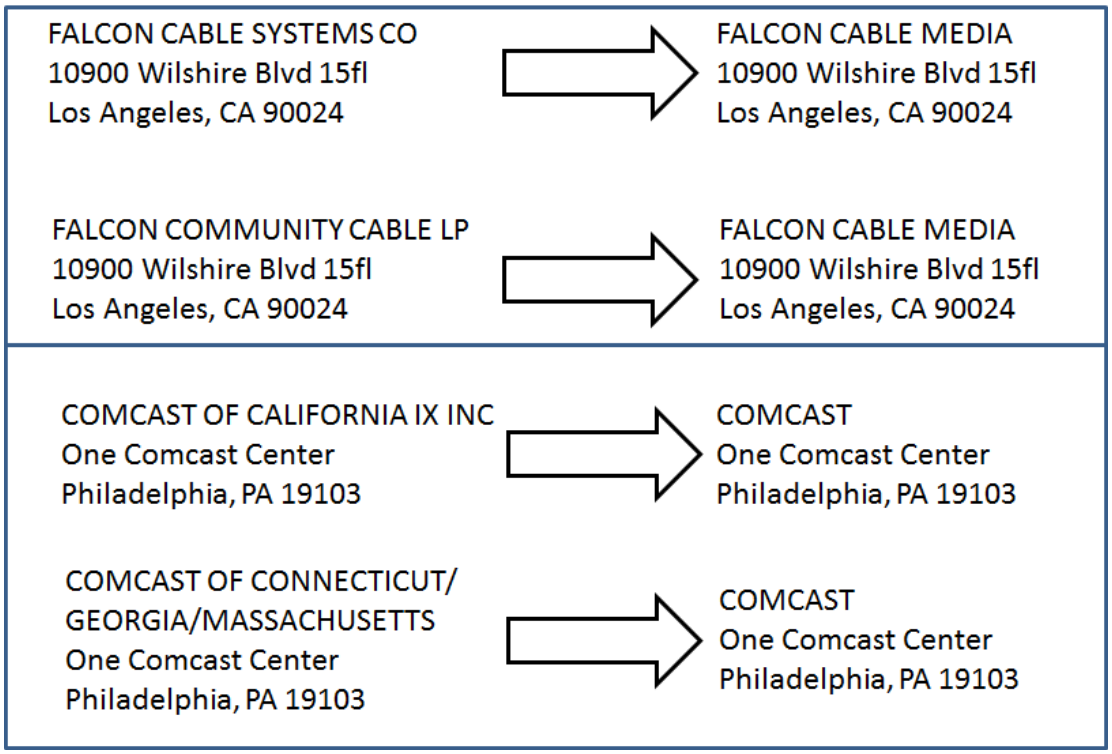

FIGURE B.11 TOP: SOME LEGAL ENTITY DIFFERENCES CAME FROM SUBSIDIARIES WITH SLIGHTLY DIFFERENT NAMES. BOTTOM: MANY CABLE OPERATORS OPERATE THROUGH FRANCHISED OR REGIONALLY-BASED SUBSIDIARIES.