Platform Pricing at Sportscard Conventions *

Ginger Zhe Jin
University of Maryland
NBER

Marc Rysman
Boston University

December 28, 2009

Abstract

Using a new data set of US sportscard conventions, this paper extends and tests the pricing theory of two-sided markets. Since most sportscard conventions are local and convention organizers must set fees to attract both consumers and dealers, we have detailed information on consumer price, dealer price and platform competition. Consistent with the theory, we present two findings: first, consumer pricing decreases with competition at any reasonable distance, but pricing to dealers is insensitive to competition and in longer distance even increases with competition. One consistent explanation is that dealers multi-home within a short distance but single-home across longer distance. Second, when consumer price is zero (and presumably cannot go beyond zero), dealer pricing is more likely to decrease with competition. These two findings confirm a link of pricing between the two sides of the market.

*We thank Mark Armstrong for guidance at an early stage, and to John List, Glen Weyl, Julian Wright, and seminar participants at the conference on the Future of Academic Communication at the University of Michigan for advice and comments. David Rapson, Haizhen Lin, Supatcha Mahathaleng and Lauren Moon provided excellent research assistance. All errors are our own.

1 Introduction

The theory of two-sided markets is an important recent development in industrial organization (see Armstrong, 2006; Rochet & Tirole, 2006; Rysman, 2009; Weyl, 2009), but few empirical papers have tested it directly. To fill in this gap, we first present a theoretical model that verifies and extends previous work on the relationship between platform pricing and platform competition; and then test it in the context of US sportscard conventions. As shown below, sportscard convention is a classical example of two-sided markets. Lessons learned from these conventions could be relevant for a broad class of two-sided markets.

By definition, two-sided markets involve two (or more) groups of agents that (1) interact through an intermediary and (2) participation or usage of each group affects the utility of the other groups. The intermediating firm is often referred to as a "platform" and each group of agent is referred to as one "side" of the market. The economics of two-sided markets focus on agent choice of platform and the price decision of competing platforms. [1]

For example, consumers value video game consoles that are served by many game developers and developers value consoles that attract many consumers. In this case, the console producer is the platform firm, accounting for interactions between game players and game developers. Similarly, a local sportscard convention provides a platform for dealers of sportscards (most

[1] This focus differentiates the literature on two-sided markets from the literature on network effects. The definition of a two-sided market is very similar to the definition of a market with indirect network effects. However, the literature on network effects tended to focus on technology adoption and network size, although this distinction is not perfect. For more on network effects, and definitional issues in two-sided markets, see Farrell & Klemperer (2007), Rochet & Tirole (2006) and Rysman (2009).

time acting as sellers) to interact with consumers that visit the convention. Sportscard dealers prefer conventions with many consumers (holding competition constant) and consumers prefer conventions with many dealers.

For several reasons, sportscard conventions provide an excellent environment to test the theory of two-sided markets, particularly for the relationship between platform pricing and platform competition. First, conventions are two-sided markets. A successful convention requires both buyers and dealers to appear at the convention, and a convention organizer must take this into account in setting prices. Second, pricing is very simple and observable. Dealers and consumers pay separate entrance fees only. There are no transaction fees or other complicating issues. We observe these fixed fees in a set of uniformly formatted classified advertisements in a trade magazine. Third, there are a huge number of conventions in the United States, more than two thousand per month at the height of their popularity, which gives us tremendous leverage for econometric estimation, as well as important panel variation in market structure. In comparison, many two-sided markets (e.g. yellow pages directories, radio, internet search engine) have zero price on one side and therefore restricte platform pricing to the other side only. Even if a platform (say video game console) charges prices on both sides, the number of important platforms and game developers are less than 25, contracts are complex and secret, and technological change makes time series variation difficult to interpret.

It is important for us to develop a theoretical model that allows lower platform prices to bring in new customers or new dealers. The previous literature typically ignores the market expansion effect, which is clearly unrealistic. Moreover, the nature of platform competition could vary greatly depending on the degree of market expansion. When the market expansion effect is present, sportscard conventions compete for not only the hard-core fans but also the casual collectors who are on the margin of attending or not attending any sportscard show. The latter could be essential for dealer attraction, pricing decision on both sides, and the final profitability of the show.

To address this shortcoming, we present a new model of platform competition based on Armstrong (2006), allowing for a market expansion effect from lower platform prices. This model yields several important results. The first is that, if dealers multi-home (i.e. attend multiple conventions) and consumers single-home (i.e. attend a single convention), increased competition between conventions leads to lower prices for consumers but not for dealers. This result is standard in the literature, and we confirm it in our model. Our second result is new to the literature: if both sides single home but competition affects consumers more than dealers, increased competition between conventions could lead to lower prices for consumer but *higher* prices to dealers. In section 3, we explain how this result is driven by our inclusion of a market expansion affect. Finally, as foreshadowed in Armstrong & Wright (2007), if competition between platforms increases, platforms reduce price more on the dealer side if their prices on the consumer side are constrained not to move. We argue below that this set of results would be difficult to rationalize without the theory of two-sided markets.

We find support for these theoretical results in our empirical work. We show that consumer pricing responds to increases in competition at any reasonable distance. However, pricing to dealers is more complex: it does not respond to competition within relatively long distances, up to 100 miles. That is consistent with a model in which dealers multi-home and consumers single-home. As we consider even longer distances, up to 150 miles, dealer prices actually increase in competition. That result is consistent with a model in which dealers single-home across longer distances, and consumers are more sensitive to distant competition than dealers.

The final theoretical result depends on conventions whose prices are constrained on one side of the market. To capture this issue, we use conventions that charge zero admission fee to consumers, about half the conventions in our data set. Presumably, these conventions cannot lower price beyond zero in response to a small increase in the number of competitors. Consistent with the theory, we show that conventions with free consumer admission reduce the price to dealers in response to competition, whereas conventions with positive consumer admission do not change dealer prices in response to competition.

Several papers seek to empirically evaluate two-sided markets. Rysman (2004) estimates the positive feedback loop between advertising and entry in the Yellow Pages market and evaluates the welfare effects of entry. Kaiser & Wright (2006) study pricing in the German magazine market and Argentesi & Filistrucchi (2007) study market power in the Italian newspaper market. Lee (2008) evaluates exclusive contracting in the video-game market. Genakos & Valletti (2007) show that lower call termination revenue for cellular phone providers leads to higher fees to subscribers. Chandra (2009) study newspaper pricing in the context of mergers.

Our paper differs in that we directly test the pricing strategy of competing platforms, arguably the major results of the two-sided market literature so far. Taking a "reduced-form" approach, we seek to determine whether correlations in the data are consistent with the proposed theories. In contrast, most other papers[2] estimate an explicit theoretical model using structural

[2]Genakos & Valletti and Chandra are exceptions

techniques, which makes it difficult to detect if the model does not hold. A direct test of the theory seems natural given that both the theoretical and empirical literatures are at such an early stage. [3]

2 Industry and Data

Collecting sportscards and sports memorabilia is a popular pastime in the United States. Sportscards are small cards with a picture of a professional player and the player's statistics. Baseball cards are the most popular. Collectors value cards of top players in top or rookie years, as well as complete sets and well produced cards. Collectors are often interested in other types of memorabilia, such as game balls, jerseys or player signatures. The popularity of collecting cards can vary a great deal, including seasonally with whether a sport is in season, and regionally with the success of the local team. A major event in our data set is the labor strike in Major League Baseball in 1994, which hurt the popularity of the league and of collecting baseball cards.

Sportscard conventions provide short events for dealers and consumers to come together. While a number of dealers establish retail shops, many dealers trade entirely at conventions. A small convention may last one day and consist of 10 tables set up at a mall. The largest conventions have at least 250 tables, last a week and take over a large convention center in a major city. Convention organizers rent the location, advertise the convention and charge fees to dealers and consumers. Conventions sometimes contract

[3] A related empirical literature focuses on indirect network effects, such as Gandal, Kende & Rob (2000), Saloner & Shepard (1995) and Ackerberg & Gowrisankaran (2006). Consistent with the theoretical literature on network effects, these papers focus on technology adoption rather than pricing by an intermediary.

for the appearance of professional athletes who will provide signatures for free. Organizers primarily profit from the entrance fees, although some organizers are also dealers who will trade cards at the convention. Both the organizer and dealer markets are extremely unconcentrated and are characterized by many small participants, many of whom have separate full-time jobs unrelated to sportscards.

Pricing at conventions is very simple. Consumers and dealers pay a fixed fee to the convention organizer. Typically, consumers pay less than $2, with about half of the conventions in our data set offering free admission. Dealers pay the "table fee", typically $25 to $100. The table fee allows the dealer to set up a table at the convention. Prices at multi-day events may be more complicated, with prices varying by day (for instance, weekend prices are typically higher) or with lower per-day fees for admission over multiple days. Also, we observe some discounts from the table fee for purchasing multiple tables. [4]

Our data set is based on the trade magazine *Becket Baseball Card Monthly*. This magazine provides articles on baseball and collecting, market prices for a huge number of cards, and most importantly for our purpose, listings for sportscard conventions (the "Convention Calendar"). Listing is free and, as we understand it, every convention would be sure to place a listing in this magazine. The magazine requires that listed conventions have at least 10 ta-

[4] In practice, dealers can buy from consumers and dealers can trade with other dealers, as can consumers with other consumers. Hence, the important distinction is not who buys cards and who sells cards but who pays the table fee and who pays the consumer fee. There could be substitution between entering with a table and not doing so. Such substitution could be problematic for our theoretical predictions since we do not include it in our model, but should bias our results away from finding differences between the two sides.

bles, although this does not appear to be binding (see below). Each calendar covers the month of the issue, so the October 1997 issue has listings for all conventions in that October.

Our data set consists of the convention listings from a selection of issues of this magazine. Convention organizers fill out a standard form and listings follow a uniform pattern, which is amenable to computer interpretation. To create our data set, we scanned all of the listings and used an Optimal Character Recognition (OCR) program (in particular ABBYY PDF Transformer 1.0) to convert these scans to text files. Then we wrote computer programs to parse the results into a usable data set. To ensure data quality, we compared the original copy with each parsed listing and corrected errors by hand.

Each convention lists the city or town in which it occurs. We match these towns to a list of towns from the U.S. Census and assign the longitude and latitude of the town to the convention. Hence, we assume that each convention is located in the population center of the town in which it occurs. [5] We drop conventions that do not occur in the continental United States. We dropped some listings that did not provide town names that we could reliably match to a location in the census. Altogether, we have data on 50,450 conventions in 36 months over 9 years.

For each listing, we use the dates of the convention, the town and state, the number of tables, the admission fee for consumers and the table fee for dealers. [6] For prices, we always took the price for a single day of admission

[5]The listings provide addresses which in principal could be used to more accurately identify locations. However, many addresses are descriptive ("VFW Hall" or "Westgate Mall") and therefore are difficult to geocode. Even for the entries that provide a street address, cleaning them would be an enormous task.

[6]We discarded some information: the exact location, the times of day of each convention and the contact name and number. The contact names are potentially very interesting

if there were discounts for multi-day admission. We used the simple average of prices if there were different prices for different days. We took the price for a single table if there were discounts for multiple tables.

Our selection of magazines range from April, 1989 to December, 1997, for a total of 36 issues. As the magazines are drawn from a personal collection, it is not a continuous set of magazines. [7] We purposely stopped collecting data after 1997, which coincides with the popularity of the World Wide Web. There is a significant decline in the number of conventions during the late 1990's which makes our approach difficult since we rely on the presence of competition to create our tests. [8] Table 1 lists the issues of the magazine in the data set, along with the number of conventions in each issue. Figure 1 graphs this series. There is a peak in activity in the Summer and Fall of 1992 when there are regularly more than 2000 conventions in a month. There is a steady decline afterwards, presumably due to the baseball strike in 1994 and the popularity of the Internet. In 1997, there are less than 1000 conventions per month.

We are interested in oligopoly interactions, so it is useful to get a sense

but difficult to clean reliably.

[7]Our collection of magazines is drawn from those we found for sale at several conventions, and some contributions from John List (for which we are very grateful). We made a number of attempts to find missing issues, for instance at public libraries. We believe that our selection of magazines is random.

[8]The impact of the Web on the convention market represents an interesting topic in its own right, as in Emre, Hortascu & Syverson (2005) for booksellers and travel agents. Jin & Kato (2007) present a detailed study on the online and offline trading of sportscards. Here we shy away from the post-Internet months because it would be difficult to determine the channel by which the Web affects conventions. Not only does the Web represent an alternative method for trading cards, the Web represents an alternative leisure activity which substitutes for card collecting altogether.

Date		mean	count	Date		mean	count
1989	Apr.	1.90	497	1994	Oct.	3.30	1701
	Aug.	1.80	386	1995	Feb.	2.78	1250
1990	Nov.	2.80	1276		Apr.	2.99	1457
	Dec.	2.84	1278		May.	2.71	1196
1991	Nov.	3.97	2206		Aug.	2.72	1169
1992	Jan.	3.90	1805		Nov.	2.97	1216
	Apr.	4.23	2477	1996	Feb.	2.56	1051
	Jul.	4.35	2294		Jun.	2.69	1217
	Oct.	3.99	2233		Sep.	2.61	1098
	Nov.	4.16	2294		Oct.	2.56	918
1993	Feb.	3.57	1797	1997	Feb.	2.39	933
	Mar.	3.75	1950		Apr.	2.48	989
	Apr.	3.70	2084		May.	2.41	957
	Jul.	3.64	1840		Jun.	2.40	942
1994	Feb.	3.41	1646		Jul.	2.31	746
	May.	3.63	1827		Aug.	2.41	862
	Jul.	3.38	1563		Oct.	2.37	946
	Aug.	3.32	1516		Dec.	2.43	832

Table 1: Number of conventions and average by 3 digit zip code for each month in data set.

Figure 1: Number of conventions by date.

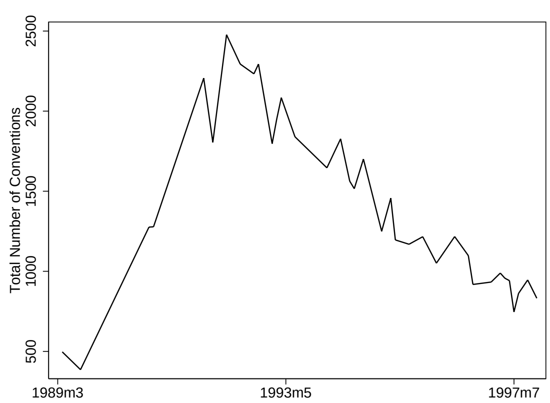

Table 2: Number of conventions per 3-digit zip code.

Count	Number	Perc.	Cum. Perc.
1	6,886	43.04	43.04
2	3,285	20.53	63.57
3	1,745	10.91	74.48
4	1,154	7.21	81.69
5	717	4.48	86.17
6	458	2.86	89.04
7	376	2.35	91.39
8	243	1.52	92.91
9	207	1.29	94.2
10	169	1.06	95.26

Count	Number	Perc.	Cum. Perc.
11	135	0.84	96.1
12	88	0.55	96.65
13	91	0.57	97.22
14	65	0.41	97.62
15	52	0.33	97.95
16	49	0.31	98.26
17	46	0.29	98.54
18	23	0.14	98.69
19	25	0.16	98.84
>19	185	1.2	100

of the number of conventions in any given region. Table 1 provides the mean number of conventions per 3-digit zip code by month for zip codes that have at least one convention. The overall average is 3.15, and this ranges from 1.90 to 4.35 in months with low and high activity. Not surprisingly, the distribution underlying these means is highly skewed. Table 2 displays the number of 3-digit zip code-months with each count of the number of conventions. For instance, there are 6,886 zip code-months in which we observe only 1 convention in a month, which represents 43.04 percent of the data. Zip code-months with three or less conventions represent almost 75% of the data, and 10 or less represents 95% of the data. There is a tail of observations with a large number of conventions, the maximum being 49 conventions in a 3 digit zip code in a single month.

The number of tables at a convention is an important explanatory variable for price. We treat the number of tables as an exogenous measure of the quality of the convention. Clearly, the quantity of dealers that purchase a table at the convention may be endogenous to the price of a table. However, the "number of tables" listed in the calendar is determined well before a final count of how many dealers will appear is available. We regard the posted number as "cheap talk" that serves to inform readers of the expected size of the event. Consider that the "number of tables" variable falls disproportionately on multiples of 5 (like 10, 15, 20 etc.), unlike a true measure of quantity. The variable is also highly correlated with other measures of quality, such as the number and quality of athletes that will be available to sign autographs. [9] As the variable is not verified, organizers could choose it in a misleading way. Our approach assumes the extent of misrepresentation does not vary systematically with competition.

Table 2 describes the distribution of the number of tables. The mean is 41.6 and the median is 35. The distribution is approximately log normal. The 99th percentile is 160. The magazine states that conventions have at least 10 tables to be listed but this does not appear to be binding. A number of conventions list less than 10 tables and the number of conventions listing 10 is not large compared to surrounding numbers. For instance, 589 conventions list 10 tables and 1,502 list 15, and 4,212 list 20. We find missing listings or listing of 0 number of tables at 1,853 observations and drop these in our statistical work.

Most conventions, 77.1%, last only one day. Almost all (98.8%) last three days or less. Most take place on weekends. In our data, 49% cover a Saturday,

[9] We do not observe other measures of quality. However, some conventions take out display advertisements in the calendar section, and from these it is clear that larger conventions offer extra features such as autograph sessions with athletes, door prizes and free raffles.

Perc.	Tables		Perc.	Tables
1	10		75	50
5	15		90	70
10	20		95	90
25	25		99	160
Median	35		Mean	41.6

Table 3: Distribution of the number of tables.

Perc.	Table Fee		Perc.	Table Fee
1	10		75	50
5	15		90	80
10	20		95	102.5
25	25		99	165
Median	35		Mean	43.7

Table 4: Distribution of table fee

53.7% cover a Sunday and 81% cover Saturday or Sunday.

The dependent variables in our empirical work are the prices. Table 4 displays the distribution of the table fee. The mean is \$43.7 and the median is \$35. The distribution is approximately log normal, with a long tail of expensive conventions. The 99th percentile is \$165 but we observe a few with table fees greater than \$1000.

A striking feature of the distribution of the admission fee is that 52.9% of conventions feature free admission. A further 29.6% charge a fee of \$1. There is little further variation, with much of it falling on multiples of 50 cents. The 95th percentile is \$2. These features lead us to model the admission fee as a binary variable so we simply predict whether admission is free or not. With more than 80% of the observations choosing 0 or 1, this seems like a decent approximation. Furthermore, our identification strategy relies on fixed effects, which are difficult to incorporate into most non-linear models

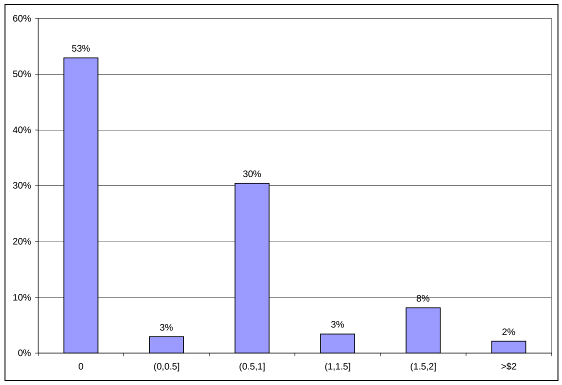

Figure 2: Admission fee distribution

besides the binary logit case. Figure 2 graphs the distribution of admission fees in a histogram.

We compute the number of competitors that a convention faces by counting the number of conventions within a given range of time and geographic distance. For example, we calculate the number of conventions on the same day within 25 miles. To do so, we use any competing convention that has at least one day that overlaps with the convention in question. As stated above, we calculate distance based on the latitude and longitude of the relevant towns in the U.S. Census. Table 5 provides the average number of competitors by different distances. Note that when computing the "within three days" variable, we treat the observation as missing for any convention for which we do not have data on conventions within three days. So for instance, a convention on April 30, 1989 would be problematic since we do not have the May, 1989 issue so we cannot count all conventions within three days. Hence, Table 5 displays a lower number of observations for the "within three days" row than the "same day" row. Given that the great majority of conventions appear on the weekends, we interpret "within three days" as

	Distance in Miles								
	5	10	25	50	75	100	150	200	obs.
Same Day	1.24	1.44	2.5	4.27	5.93	7.77	11.87	16.84	50450
	(0.57)	(0.88)	(2.54)	(4.74)	(6.40)	(7.94)	(11.39)	(15.56)	
Within	1.68	2.39	5.80	11.03	15.45	20.06	30.02	42.10	38801
3 Days	(1.12)	(2.11)	(6.84)	(12.30)	(15.70)	(18.62)	(24.94)	(32.90)	

Table 5: Average number of competitors by distance

essentially meaning "same weekend". Given these small numbers and the very local nature of the organizing market, we expect organizers to know with some accuracy in advance the number of conventions they will face on any given day. Hence, prices respond to competition although prices and competition are announced publicly at the same time.

3 Model

In reality, competition among sportscard conventions can be summarized as the number of competitors that a convention faces in the same time frame within a specific distance. However solving a theoretical model that allows for variation in both the number of platforms and distance between platforms is a challenge. Like the existing theoretical literature on two-sided markets, we restrict our model to have two platforms and and we use the distance between platforms to proxy for the level competition. In Section 5 , we explain how we interpret predictions from this model in the context of sportscard conventions.

We base our work closely on the model of Armstrong (2006) and the extensions in Armstrong & Wright (2007). These models are useful for representing the sportscard convention market because they address competition between two platforms that charge only a fixed fee to each consumer and do not charge based on the number of trades a consumer makes through the platform. We extend their models to allow for a market expansion effect from lowering prices. While doing so complicates the model such that we analyze some results numerically, it also provides more realistic results to take to data.

Consider two sides of a market, with one set of agents on each side. Side C is made up of consumers and side D is make up of dealers. We index sides by m. Consumers and dealers are distributed along separate real number lines with density one. They are distributed across the entire line. [10] Agent i in market m is located at l_i (we do not index l_i by market for notational convenience). Consumer i bears no relationship with dealer i in side D. That is, they make their choices independently. There are two conventions, or "platforms", 1 and 2, indexed by j. Throughout, we assume that platforms have no costs. The platforms sell to both sets of consumers simultaneously. The location of platform j in side m is l_j^m. In this set-up, we can consider comparative static in l_j^C but not l_j^D, as if a platform could change its location with respect to consumers but not dealers. Clearly, this is an abstraction since an actual change in geographic location would affect all types of agents, but we think of this is an approximation to a situation in which one set of agents cares about location much more than the other. Alternatively, one may assume conventions have the same location on each side of the market (that is, $l_j^C = l_j^D$) but with consumers having higher travel costs than dealers ($t^C > t^D$). We solve this alternative model and find the same qualitative results, although the equations are much more complicated and we must

[10] The important assumption is not that agents extend forever across each line but that they extend past whatever location would generate sales for zero prices, so there is always a demand expansion effect to lower prices.

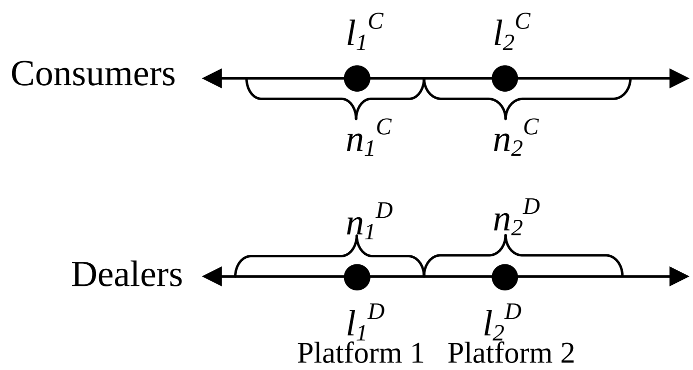

Figure 3: Consumers, dealers, platform locations and quantities.

resort to numerical analysis to a greater degree. For this reason, we only present the model that assumes equal travel cost but allows platforms to choose locations on the two sides asymmetrically.

Agents value a platform based on how many agents the platform serves in the other side. Suppose platform j sells to n_j^C consumers and n_j^D dealers. The utility to agent i in side m is u_{ij}^m, defined to be:

$$u_{ij}^C = v^C + \alpha n_j^D - p_j^C - t|l_i - l_j^C|$$
$$u_{ij}^D = v^D + \alpha n_j^C - p_j^D - t|l_i - l_j^D|$$

where v^C and v^D represent the stand-alone utility of purchase to consumers and dealers, α determines the value conferred by sales in the other side of the market, p_j^m is the price of platform j in side m, and t parameterizes the travel cost. We assume that agents could instead use some outside good with utility of zero. Figure 3 displays our model visually.

Note that we are treating dealers symmetrically with consumers. We do not model price-setting, competition or other externalities between dealers. Doing so would complicate our analysis unnecessarily. Wright (2005) presents a model with within-group competition that turns out to be just a change of variables from the model we consider. Also, we assume there is agent heterogeneity in l_i but not α. Weyl (2009) explores a model with heterogeneity in both terms but we have not done so as we believe that geographic heterogeneity is the most important issue in our application.

Throughout, we parameterize the model as follows, which we regard as without loss of generality:

Assumption 1 $\quad l_1^C = -1 \qquad l_1^D = -1 \qquad t = 1.$

If we want to consider changes in competitiveness on only one side of the market, we will assume $l_2^D = 1$ and consider comparative statics of prices in l_2^C. If we want to consider changes on both sides, we assume $l_2^C = l_2^D = l_2$ and consider comparative statics of prices in l_2. We assume that the two platforms choose prices simultaneously and we solve for a Nash equilibrium of the game. Platforms are symmetric and we find only symmetric equilibria. Hence, we always find that $P_1^m = P_2^m$ for $m = \{C, D\}$.

We are interested in cases in which there is a strategic interaction between the two platforms, so we consider cases in which all consumers located between the two platforms prefer purchasing from either platform to no purchase. Armstrong (2006) obtains simple analytic solutions to a similar model in which agents are assumed to be on finite lines and platforms are at the ends of these lines. However, the lack of a demand expansion effect leads to some unrealistic implications. The principal innovation of our model is to incorporate a demand expansion effect by allowing agents to be located on both sides of the platforms. While our solutions are less elegant, they are more relevant for our empirical work.

3.1 Single-homing

To begin, we consider the case in which consumers and dealers both single-home. That is, agents pick the single platform that gives the highest utility, or choose no platform if the value of purchase does not exceed their reservation value of zero. We are interested in how competition on one side affects pricing, so we assume that $l_2^D = 1$ and consider comparative statics in l_2^C. For simplicity, we assume that $v^C = v^D = v$.

For each platform, profit is $\pi_j = p_j^C n_j^C + p_j^D n_j^D$. Demand for each platform on each side of the market is:

$$n_j^m = \frac{\left(\alpha n_j^{-m} - p_j^m\right) - \left(\alpha n_{-j}^{-m} - p_{-j}^m\right)}{2} + \frac{l_2^m - l_1^m}{2} + v + \alpha n_j^{-m} - p_j^m$$

where n_{-j}^{-m} refers to the number of agents purchasing the other platform's product in the other side. The first two terms refer to the profit drawn from agents between the two platforms. The first term will be zero in a symmetric equilibrium and the second term increases as the platforms become farther apart. The last term captures the profits drawn from agents on the other sides of the platforms, and generates the demand-expansion effect.

We take first-order conditions from the profit functions for each price and we solve for prices by solving the four first-order conditions simultaneously. [11] Algebraic manipulation shows that:

$$\frac{dp_1^D}{dl_2^C} < 0 \quad \texttt{if} \quad 0 < \alpha < \frac{5}{8}$$

This result implies that p_1^D decreases as platform 2 becomes farther away on side C, if the network effect is not too large.[12] That is, higher competition on one side increases prices in the other. Finding that prices increase

[11] All derivations in this paper are available in Mathematica files and PDF output on our web site, at xxxx

[12] It turns out that this condition also implies that all agents between the two platforms

in competition would be a very surprising result that would be difficult to replicate in a model with only single-sided interaction.

The condition that the network effect be not too large turns out to be non-binding. As is well-known in the network effects literature, large network effects lead to intense price competition. There exists a critical value of α such that prices become zero. It takes on a particularly simple form in our model: when $l_2^C = 1$, prices are greater than zero if $\alpha < 1/2$. Obviously, this lies below 5/8. Hence, we would never consider values of α such that $dp_1^D/dl_2^C > 0$.

To get a sense of the whole set of prices, consider the specification with $v = 1$ and $\alpha = 0.3$. These parameters imply that for $l_2^C = 1$, all agents on both sides between the platforms purchase. Then, prices are:

$$\begin{aligned} p_1^C &= 0.27 + 0.19 l_2^C \qquad & p_1^D &= 0.49 - 0.03 l_2^C \\ p_2^C &= 0.27 + 0.19 l_2^C \qquad & p_2^D &= 0.49 - 0.03 l_2^C \end{aligned}$$

Not surprisingly, platform 1 increases its price to consumers as platform 2 becomes farther away on that side. However, we also see that platform 1 *decreases* its price to dealers as platform 2 becomes *farther away* on the consumer side.

What is the intuition for this surprising result? As platforms become closer together, they serve fewer consumers. Thus dealers, who are attractive in part because they allow the platform to raise price on consumers, are less attractive. Hence, platforms raise price to dealers. Note that this result is unique to our setting because we have a demand expansion effect. Armstrong

are served as long as v is high enough, in particular $v > 7/6$. For lower values of v, we require α to be above some low value, but this requirement could be eliminated by lowering t.

(2006) finds that changing the distance between platforms on one side does not affect prices on the other. This follows from Armstrong's assumption that consumers are fully served.

It is difficult to construct alternative stories for why a price might increase in competition. Without specifying an explicit model, it is surely true that most multi-product oligopoly models predict that prices fall in competition. One exception has been developed for prescription drugs, in which we have seen brand-name drugs raise prices after entry by a generic drug (for instance, see Ching, 2009). These models rely on the entrant successfully capturing the high elasticity customers, leaving the incumbent to exploit the remaining low elasticity customers, where elasticity is driven by the type of insurance that a consumer has. That kind of differentiation between platforms and heterogeneity in elasticity seems unlikely to be relevant for the card convention market.

3.2 Multi-homing

It is possible for agents to attend multiple shows. In practice, it seems more likely that dealers would do so than consumers as we will be focusing on shows that are more than 50 miles apart. Where traveling so much in a day to attend multiple shows seems difficult for consumers, dealers could split their collections and have different employees attend different shows.

We can capture this situation in our model by allowing dealers to purchase from both platforms if they would like, and then obtain the sum of utilities from each show $u^D_{i1} + u^D_{i2}$. We continue to assume that consumers single-home. [13] Our result matches that found in the previous literature, and

[13] Armstrong (2006) argues that markets often arrive at a situation where one side single-homes and one side multi-homes. For instance, consumers may read one newspaper

hence a brief discussion of the results suffices: None of the prices are affected by locations on the dealer side. This follows because there is no strategic interaction between the platforms on side D. In this model, utility for the multi-homing agents is additively separable across the two platforms and so they make their decision at each platform independently of their decision at the other platform. Thus, there is no competition on the dealer side and so locations do not matter.

Suppose we exogenously move platform 2's location on both sides simultaneously. Let $l_2^C = l_2^D = l_2$ and consider $\partial P_j^m / \partial l^B$. We find the result similar to Armstrong (2006): as platform 2 becomes closer to 1, consumer prices drop but dealer prices do not change. As in previous work, the single-homing side benefits from competition whereas the multi-homing side does not. Intuitively, dealers can only reach the consumers at a convention by going to that convention, so the presence of consumers gives conventions market power over dealers. Hence, competition for consumers is fierce whereas competition for dealers is not.

3.3 Constrained prices

Our last theoretical point is that competition affects price-constrained platforms differently. Armstrong & Wright (2007) generate this result by allowing the network parameter α to differ across sides of the market, but we find it more natural in our setting to consider stand-alone demand differing across sides: we assume $v^D > v^C$. However, we can generate similar results if we allow for side heterogeneity in other parameters.

We consider the case in which consumers and dealers single-home. As v^D

whereas advertisers appear in multiple newspapers. Consumers use a single payment card whereas merchants can accept cards from multiple networks.

rises, the price to dealers rises and sales to dealers increase. The presence of consumers allows platforms to raise price to dealer and hence, more dealers makes consumer more valuable. As a result, consumer prices drop. We assume that the platforms cannot set prices below zero, and we are interested in values of v^C and v^D such that this constraint is binding.

First, we note that there always exists a value of v^D such that $P_j^C = 0$. Conditional on α, it is a simple linear function of v^C. For instance, for $\alpha = 0.3$ and $l_2 = 1$, $v^D > 2.47 + 5.93v^C$ implies $P_j^C = 0$. Our approach is to compute the equilibrium prices as a function of l_2 on both sides of this boundary, and show that the effect of l_2 on P_j^D is larger when the zero-pricing constraint is binding.

More formally, let $v^{D*}(v^C, \alpha, l_2)$ be the value of v^D above which $P_j^C = 0$. Then consider $P_j^D(v^C, v^D, \alpha, l_2)$, the equilibrium dealer price. We show numerically that:

$$\left.\frac{\partial P_j^D(v^C, v^D, \alpha, l_2)}{\partial l_2}\right|_{v^D > v^{D*}(v^C,\alpha,l_2)} > \left.\frac{\partial P_j^D(v^C, v^D, \alpha, l_2)}{\partial l_2}\right|_{v^D < v^{D*}(v^C,\alpha,l_2)} \quad \forall v^C, \alpha, l_2 \tag{1}$$

Naturally, we only consider combinations of v^C, α, l_2 such that agents between the two platforms are fully served. The calculation in (1) is straightforward since the pricing functions are linear in l_2 with no interaction between l_2 and v^C, so only α appears. Numerical calculations show 1 holds for all values of $\alpha < 0.5$. Recall that $\alpha > 0.5$ implies that prices on both sides become zero, and is out of our range of interest.

Intuitively, as platforms become closer together, they would like to lower price to consumers. When they cannot do so any more, they turn to an alternative strategy for attracting consumers: attracting dealers. Hence, the effect of competition on dealer prices becomes particularly intense when

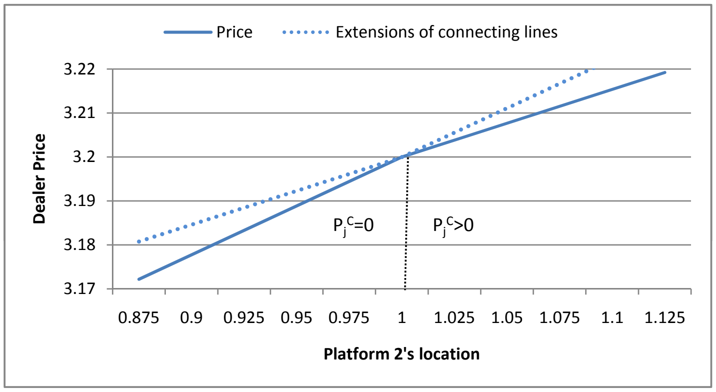

Figure 4: Dealer prices with and without the zero-pricing constraint.

platforms are constrained on the consumer side.

As an example, consider Figure 4. It displays the equilibrium dealer price in the solid line for different values of l_2 around $l_2 = 1$. For this picture, we have imposed $\alpha = 0.3$ and $v^D = 8.4$. For these values, the zero-pricing constraint on consumer prices binds for $l_2 < 1$ and does not bind for $l_2 \geq 1$. Note the change in slope around $l_2 = 1$. But for the kink, the lines are linear. The dotted lines are extensions of these lines and accentuate the change in slope.

4 Empirical implementation

This section discusses empirical implications of the preceding theoretical model and our strategy for implementation. We group the single-homing and multi-homing results in a single discussion, and discuss the results from the model of the constrained prices separately.

4.1 Single- and multi-homing

Considering the single- and multi-homing results together in Sections 3.1 and 3.2, we find the following:

Empirical implication: Competition that affects one side of the market more than the other leads to lower prices on that side and unchanged or higher prices in the other.

How do we identify competition that affects one side of the market and not the other? We focus on competition that is relatively far away, for instance, between 25 and 100 miles away, or between 50 and 150. For instance, it is possible that only dealers are willing to multi-home at these distances, because dealers could split their collections and attend multiple conventions simultaneously, whereas a consumer would have to drive from one to another. If dealers multi-home but consumers do not, that is consistent with our model in Section 3.2.

Even if dealers attend only one convention, it can still be possible that far competition affects consumer prices more than dealer prices. Recall that although our theoretical model is based on moving the location of a platform on one side and not the other, that is equivalent to the platform moving location on both sides when travel costs are very different on the two sides. Hence, if dealers are willing to travel much further than consumers to reach a convention, we have our result. For instance, consider two conventions, one that faces a competitor 200 miles away (out of the range we consider in our empirical work) and one that faces a competitor 100 miles away. If dealers have low travel costs, they may view those situations very similarly, whereas consumers may view the second case as a greater increase in competition.

Alternatively, it could be that dealers that do travel long distances are not price-sensitive. These stories are consistent with our model in Section 3.1.

When analyzing the effect of far competition, unchanged dealer prices would be consistent with dealers that multi-home whereas increased dealer prices is consistent with our model of single-homing dealers. In practice, there are some dealers of each type. Many dealers work independently and have relatively small collections, so that splitting their collection and hiring someone to administer a table at a convention would be costly. A larger dealer that had regular employees would find this less problematic. We do not observe dealer behavior and we do not take a position *a priori* on which type of behavior would be most prevalent empirically. Rather, we check our results to see with which story they are consistent. As we will see, we find that the results change from being consistent with multi-homing to single-homing dealers as we consider competition that is farther apart. [14]

Clearly, we have not established that the theory of two-sided markets is the *only* theory that could generate these results, but we believe that it is difficult to find an alternative, particularly for the result about increased prices.

4.2 Constrained prices

Section 3.3 generates our second empirical prediction.

[14] We do not model the endogenous change from multi-homing to single-homing as distance increases, but this seems straightforward to do address. For instance, if multi-homing requires hiring an employee to work extra hours and travel to the convention, the cost increases as teh convention gets farther away in a way that might make multi-homing unattractive at long distances.

Empirical implication:
Increased competition has a larger negative effect on the price of one side of the market if the convention faces a constraint on the price of the other side.

We interpret conventions that allow free admission as being price-constrained. Presumably these conventions would not reduce their fee in response to small changes in exogenous variables. There is a sense in which every convention is constrained as they tend to use round numbers for the admission fee so it does not appear to be a truly continuous choice variable. However, we view the conventions with free admission as the "most constrained": they would be least likely to adjust their admission fee in response to changes in market variables. Our theoretical point does not depend on competition affecting one side more than the other so we test this issue in "near" competition, which affects both sides of the market.

Our theoretical model also provides an explanation of why some conventions choose free admission and some do not. We should observe free admission when dealer demand for conventions is particularly large. Variation in relative demand arises from the volatile nature of consumer demand. As stated above, consumer interest can vary based on the season, the relative success of the local team and opportunities for alternative leisure activities, even within a month. Presumably the dealer population is more stable.

5 Empirics

We present a series of regressions that explore the questions raised in the Section 4. In addition to the issues raised already, we must also address econometric issues of unobserved heterogeneity and omitted variable bias. These issues are important when we consider the relationship between pricing and competition. Demand factors are not entirely observable, and we expect that high unobservable demand will lead to both high prices and more entry, which creates bias in our estimates. We address this issue by including location fixed effects, where locations are indicated by 3-digit zip codes.

To see how our strategy works, we start with regressions of price on nearby competition. These regressions do not address any issues raised by two-sided markets but rather serve to verify that our fixed effects strategy works appropriately. Results appear in Table 6. In the first two columns, we have results from a binary logit model predicting when admission fees are not zero, as a function of the number of nearby competitors.

We do not include location fixed effects in the first two columns. We do however include time fixed effects, at the level of the month, and controls for the days of the week. [15] We include time fixed effects and day-of-the-week controls in all specifications in this paper, although we do not report their effects. We include one control variable in all of the regressions: the log of the number of tables. We interpret this variable as a control for the quality of the convention and, not surprisingly, it is positive and both economically

[15] Our controls for the days of the week consist of dummy variables for each combination of days that appear in our data set more than 500 times. For instance, there is a dummy variable for Saturday-only conventions, one for Sunday-only conventions, and a separate dummy for conventions on both Saturday and Sunday. The full set of dummy variables capture 97% of the data set. We also include the duration of the convention (1, 2 or 3 days), which serves to better match the remaining 3% of conventions. We do not report any of these results in the paper. In all regressions, we reduce heterogeneity by dropping conventions that last more than three days, although we still use them for purposes of computing the number of competitors.

and statistically significant in every specification.

In column 1 of Table 6, we define competition to be the log of the number of competitors within 25 miles and in column 2, we use 50 miles. We add one to the number of competitors to address log zero issues. In both cases, we see that competition is positive and statistically significant, as if competition caused higher prices. However, this result appears to be driven by unobserved geographic heterogeneity. In columns 3 and 4, we use Chamberlain's conditional logit model (Chamberlain, 1980) to include location fixed effects. [16] In this specification, we see that competition has a negative effect on the likelihood of setting a non-zero admission fee and that this effect is significantly different from zero at a confidence level of 1%. [17] We conclude from this regression that our measure of competition and our fixed effects strategy addresses the omitted variable problem, at least in part.

We see similar results when we use the log of the table fee as the dependent variable. We use a linear model for the log of table fee and estimate with linear panel data techniques. In the first two columns of the second panel in Table 6, we present results without location fixed effects. In these cases, the effect of of competition on price is estimated to be statistically insignificant, and is very close to zero. However, the next two columns include location fixed effects, now we see that competition has a negative and statistically significant effect for distances of both 25 and 50 miles. The effect is not large – the elasticity is around -0.02. But we believe this to be an upper bound due to the positive correlation induced by remaining unobserved heterogeneity.

[16]We use conditioning to address the location fixed effects and capture the rest of the fixed effects as dummy variable regressors.

[17]Unfortunately, we cannot compute economic magnitudes for the conditional logit since marginal effects for non-linear models depend on the level of the explanatory variables and the fixed effects are not recoverable.

Dependent Variable	Admission Fee > 0					Log Table Fee			
Mileage	25	50	25	50		25	50	25	50
ln(Competition)	0.086 *	0.094 *	-0.135 *	-0.113 *		-0.001	-0.001	-0.021 *	-0.015 *
	(0.015)	(0.012)	(0.026)	(0.024)		(0.003)	(0.002)	(0.004)	(0.003)
ln(# of tables)	0.647 *	0.638 *	0.825 *	0.816 *		0.263 *	0.263 *	0.281 *	0.281 *
	(0.021)	(0.021)	(0.028)	(0.028)		(0.004)	(0.004)	(0.004)	(0.004)
Location FE's	No		Yes			No		Yes	
Observations	48,123		47,606			45,965		45,965	
Standard errors in parenthesis, * indicates signficant at 1% confidence level									

Notes: Admission fee is estimated by a logit or conditional logit model. Log table fee is estimated in a linear model. Location fixed effects are 3 digit zip codes. All models include fixed effects for time (monthly) and days-of-the-week the convention covers. Competition is the number of conventions within the "Mileage" number that overlap in the calendar (plus 1 to address log zero).

Table 6: Regressions with and without location fixed effects

Note that the number of observations is slightly lower when we use the table fee because the table fee is missing for a number of observations. Restricting the admissions fee regressions to observations in which the table fee is present does not change results. For both admission fee and table fee regressions with fixed effects, some observations are dropped because the dependent variable does not vary within a group. This happens more often for the binary admission fee variable than for the continuous table fee.

The key implications of our theoretical model relate to the asymmetric effect of competition on the two sides of the market. To the extent that distance affects consumers and dealers differently, we capture the asymmetry by distinguishing between "near competition" and "far competition". We define near to be those conventions within 25 or 50 miles, and far to be those within 100 or 150 miles, not including near competitors. Again, we require all competitors to overlap in calendar time by at least one day.

Recall from the model that we focus on a market with two platforms and the degree of platform competition only varies by the distance between platforms. In reality, we could observe more than two platforms in near or far competition. To address this discrepancy, we make two assumptions. First, one unit increase in the number of nearby competing platforms is equivalent to moving a platform from a really far distance that has a minimal competitive effect (say more than 200 miles away) to a distance that counts for near or far competition. Second, the effect of having competing platform(s) in near or far competition is similar no matter whether there are two or more than two platforms in the market.

Under these assumptions, we report the results of near and far competition in Table 7. In the first panel, we see that for all definitions of distance, an increase in both near and far competition makes free admission more likely. All effects are significant at 1% confidence levels. Not surprisingly, the coefficient on far competition is smaller, although this may in part be due to the different sizes of the near and far competition variables.

In the second panel, we see that near competition also drives down table fees, and does so with effects similar to what we found in Table 6. However, far competition measured within a 100 mile radius has no significant effect on table fees, and far competition measured within 150 miles has a *positive* effect on prices. This result is consistent with our theoretical prediction in Section 4.1. The key feature of our empirical result is the pairing of the negative coefficient on one side of the market with the zero or positive coefficient on the other side.

Dependent Var.	Admission Fee>0				Log Table Fee				Log Table Fee			
Near Mileage	25	25	50	50	25	25	50	50	25	25	50	50
Far Mileage	100	150	100	150	100	150	100	150	100	150	100	150
ln(Near Comp)	-0.121 *	-0.124 *	-0.097 *	-0.099 *	-0.021 *	-0.022 *	-0.015 *	-0.016 *	-0.002	-0.003	-0.005	-0.007
	(0.026)	(0.026)	(0.025)	(0.024)	(0.004) *	(0.004)	(0.003) *	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)
ln(Far Comp)	-0.094 *	-0.083 *	-0.09 *	-0.08 *	-3.7E-06	0.009 *	-0.001	0.009 *	3.0E-04	0.009 *	-1.9E-04	0.010 *
	(0.025)	(0.026)	(0.025)	(0.025)	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)
Admission Fee = 0									-0.032 *	-0.033 *	-0.039 *	-0.039 *
									(0.005)	(0.005)	(0.006)	(0.006)
ln(Near Comp) X(Ad Fee = 0)									-0.039 *	-0.038 *	-0.018 *	-0.017 *
									(0.005)	(0.005)	(0.004)	(0.004)
ln(# of tables)	0.826 *	0.827 *	0.826 *	0.822 *	0.281 *	0.281 *	0.281 *	0.281 *	0.273 *	0.273 *	0.273 *	0.273 *
	(0.028)	(0.028)	(0.028)	(0.028)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Observations	47,606				45,965				45,965			
Standard errors in parenthesis, * indicates signficant at 1% confidence level												

Notes: Admission fee is estimated by a conditional logit model. Log table fee is estimated in a linear model. All models include fixed effects for location (3 digit zip codes), time (monthly) and days-of-the-week that the convention covers). Near competition is the number of conventions with the "Near Mileage" number that overlap in the calendar. Far competition uses the "Far Mileage" number, and is not inclusive fo near competition. We add 1 to each number to address log zeros.

Table 7: Main results

Recall that the zero coefficient on competition within 100 miles is consistent with a model in which dealers multi-home. In contrast, the positive coefficient on competition within 150 miles is consistent with a model in which dealers single-home. Certainly, it is plausible that single-homing is more prevalent at longer distances given the small scale at which most dealers operate.

An alternative explanation for the positive coefficient on far competition is that it is due to endogenous entry, so that unobserved temporal-geographic heterogeneity causes this result. However, this idea is rather hard to formulate because we include time and location fixed effects and because we find a negative coefficient on nearby competition. While it is plausible that there is heterogeneity that varies over time and space jointly, it must also somehow operate in a wider area more strongly than a local area, and affect the two sides of the market in an asymmetric way.

The next theoretical point that we wish to include is that conventions constrained on one side of the market respond more strongly to competition than those that are not constrained. We address this issue by including the a dummy variable for whether the convention offers free admission, and importantly, the interaction of the dummy with the log of nearby competition. Results appear in the third panel of Table 7. These results are our preferred specification for table fees.

In these columns, the coefficient on nearby competition becomes insignificant and precisely estimated near zero. Instead, we see a negative and significant coefficient on the interaction term, with a substantially higher magnitude than we was on nearby competition in panel 2. That is, the negative effect of nearby competition that we observed earlier appears to be coming from the constrained firms. As theory predicts, constrained firms respond more strongly to competition than unconstrained firms. [18] As before, we see no effect of far competition when measured at 100 miles and a positive and significant effect at 150 miles.

These regressions paint a consistent story of how two-sided market issues affect conventions at different distances. Dealers multi-home and so competition between conventions has no effect on dealer prices over some middle range. When conventions are very close, admission fees are driven to zero and then conventions compete by lowering dealer fees. When conventions are very far from each other, dealers switch to single-homing, in which case dealer prices increase in competition (since competition for consumers is more intense then competition for consumers).

6 Conclusion

We present a theoretical model that verifies and extends results in previous work on two-sided markets. We consider a set of reduced-form empirical implications from this literature in a new data set on sportscard conventions in the United States in the early to mid 1990's. These conventions are a two-sided market since convention organizers must set admission and table fees to attract consumers and dealers. In our theoretical work, we find that when consumers single-home and dealers multi-home, competition between conventions affects consumer prices but not dealer prices. When dealers single-home but consumers respond more strongly to competition, prices to

[18] Note that in our model, both admission and table fees are determined simultaneously so there is a potential endogeneity problem with including admission fees on the right-hand side. We do not address this issue, so these results should be interpreted with some caution.

dealers can increase in competition. When the price to consumers reaches zero and is thus constrained, dealer prices decrease in competition between conventions.

These predictions are consistent with our empirical results, where we show that consumer prices respond to the number of conventions both nearby and far. In contrast, dealer prices do not respond to competition between conventions, except for conventions with free consumer admission, for whom competition lowers dealer prices, and for conventions with competition very far away (when dealer single-homing is likely), for whom competition raises price. Thus, we find that empirical evidence from the sports cards market supports the recent theoretical literature on two-sided markets.

References

Ackerberg, D. A. & Gowrisankaran, G. (2006). Quantifying equilibrium network externalities in the ach banking industry. *RAND Journal of Economics*, *37*, 738–761.

Argentesi, E. & Filistrucchi, L. (2007). Estimating market power in a two-sided market: The case of newspapers. *Journal of Applied Econometrics*, forthcoming.

Armstrong, M. (2006). Competition in two-sided markets. *RAND Journal of Economics*, *37*, 668–691.

Armstrong, M. & Wright, J. (2007). Two-sided markets, competitive bottlenecks and exclusive contracts. *Economic Theory*, *32*, 353–380.

Chamberlain, G. (1980). Analysis of covariance with qualitative data. *Review of Economic Studies*, *47*, 225–238.

Chandra, Ambarish; Collard-Wexler, A. (2009). Mergers in two-sided markets: An application to the canadian newspaper industry. *Journal of Economics and Management Science*, *18*, 1045–1070.

Ching, A. (2009). A dynamic oligopoly structural model for the prescription drug market after patent expiration. *International Economic Review*, *forthcoming*.

Emre, O., Hortascu, A., & Syverson, C. (2005). E-commerce and the market structure of retail industries. *Net Institute Working Paper*, *05-24*.

Farrell, J. & Klemperer, P. (2007). *Handbook of Industrial Organization, Vol 3*, chapter Coordination and Lock-In: Competition with Switching Costs and Network Effects. Elsevier.

Gandal, N., Kende, M., & Rob, R. (2000). The dynamics of technological adoption in hardware/software systems: The case of compact disc players. *RAND Journal of Economics*, *31*, 43–61.

Genakos, C. & Valletti, T. (2007). Testing the waterbed effect in mobile telephony. Unpublished Manuscript, University of Cambridge.

Jin, G. Z. & Kato, A. (2007). Dividing online and offline: A case study. *Review of Economic Studies*, *74*, 981–1004.

Kaiser, U. & Wright, J. (2006). Price structure in two-sided markets: Evidence from the magazine industry. *International Journal of Industrial Organization*, *24*, 1–28.

Lee, R. (2008). Vertical integration and exclusivity in platform and two-sided markets. Unpublished Manuscript, Stern School of Business.

Rochet, J.-C. & Tirole, J. (2006). Two-sided markets: A progress report. *RAND Journal of Economics*, *37*, 645–667.

Rysman, M. (2004). Competition between networks: A study of the market for yellow pages. *Review of Economic Studies*, *71*(2), 483–512.

Rysman, M. (2009). The economics of two-sided markets. *Journal of Economic Perspectives*, *?*, ?

Saloner, G. & Shepard, A. (1995). Adoption of technologies with network effects: An empirical examination of the adoption of Automated Teller Machines. *RAND Journal of Economics*, *26*, 479–501.

Weyl, E. G. (2009). Heterogeneity in two-sided markets. Unpublished Manuscript, Harvard University.

Wright, J. (2005). Why do firms accept credit cards? Unpublished Manuscript, National University of Singapore.