Vertical Restraints in the Movie Exhibition Industry

Krzysztof Wozniak

March 15, 2012

Abstract

I conduct an empirical study of vertical restraints the movie exhibition industry, where retailers face uncertain product demand. In this paper I aim to answer two questions: what is the welfare costs of these restraints, and do they facilitate anticompetitive behavior by large upstream firms. I collect a detailed dataset of movie theater attendance for a group of independent exhibitors, which I use to estimate a model of consumer demand. I also construct a dynamic model of retailer behavior with product demand uncertainty, where exhibitors sequentially learn about the true movie quality from observed attendance figures. Counterfactual results indicate substantial consumer welfare is lost as a results of the prevailing restraints, as removing them allows movie theaters to increase attendance between 3 and 37%. The results also suggest these restraints allow major movie distributors to "squeeze out" movies sold by their smaller and independent competitors, though the effect is small.

1 Introduction

The size of the consumer choice set, the set of products available to consumers, is a crucial determinant of welfare in markets with differentiated goods and heterogeneous consumer tastes. As illustrated by the simple model developed by Hotelling (1929), when the number of firms/products increases consumers are at least as well off, absent price collusion between retailers. Unlike in the simple Hotelling setup, when every good produced is available for consumers to buy, in vertically-separated markets the exact composition of the consumer choice set is determined by the downstream retailers rather than directly by upstream producers. Producers' incentives are not necessarily aligned with those of the retailers, and they may try to influence the retailers' decision of which goods to offer using non-linear pricing or vertical restraints such as retail price maintenance, exclusive dealing, exclusive territories or tying.

Vertical restraints play a crucial role in determining the nature of competition in vertically-separated markets, however their exact impact on the market equilibrium is often ambiguous - theoretical models have been constructed to show both negative or positive effect they may have on consumer welfare. This is reflected in antitrust treatment of vertical restraints that's been inconsistent over time[1], leaving plenty of scope for abuse of market power. In light of such theoretical ambiguity empirical analysis is necessary to determine what kind of impact a combination of vertical restraints has on the market equilibrium. However, to date there have been few papers which seek to conduct structural empirical analysis of vertical restraints: Asker (2005), Brenkers & Verboven (2006) and Ho et al. (2010).

The aim of the paper is to conduct an empirical investigation of how vertical restraints can be used by upstream firms to impact the consumer choice set offered by downstream retailers, within the context of the movie exhibition industry. Exhibitors (cinemas) rent movies from upstream distributors and then sell consumers tickets to movie screenings, keeping a fixed percentage of the sales proceeds. While they cannot impose schedules directly on exhibitors, distributors use vertical restraints to restrict the flexibility exhibitors have when deciding movie rental and scheduling. The effect of this, combined with limited number of screens exhibitors have at their disposal, is that movie theaters only ever show a small fraction of all movies released; this is especially true for medium and small movie theaters. Although this is just one specific industry, findings from this paper readily translate into other vertically-separated, capacity-constrained industries such as radio and TV program scheduling, offline and online advertising or retail sales.

I pose two research questions in this paper. First, do vertical restraints imposed by movie distributors reduce consumer welfare by shrinking their choice set and, if so, by how much? Second, do vertical restraints help major distributors keep smaller distributors movies from being shown in the movie theaters? To answer these questions I construct a detailed model of industry demand and supply and estimate it using a detailed panel dataset of moviegoer attendance. I model consumer demand for movies in cinemas using a flexible random-coefficient logit framework that's best suited to discrete choice events such as going to the movies. The demand system accounts for competition across titles and screening times, allowing for a changing choice set depending on the exhibitor's scheduling decisions.

[1] Up until 1940s vertical restraints were considered lawful, however by the 1960s the attitude had shifted towards considering them all *per se* illegal. Since then it's been recognized that in most cases there is a need for empirical assessment of the welfare impact of vertical restraints, which lead to the current rule-of-reason approach (Lafontaine & Slade (2005)).

I supplement the demand system with a supply-side model of how exhibitors make rental and scheduling decisions given the set of movies available to them from distributors. In order to abstract from complicated competitive considerations I focus in this paper on movie theaters which are local monopolies. The model incorporates vertical restraints imposed by the distributors, and takes account of how exhibitors learn the *ex ante* uncertain movie quality. To answer the research questions, I utilize the demand and supply system estimates to construct counterfactuals in which vertical constraints are lifted, allowing exhibitors full flexibility when deciding their schedule.

The plan for the paper is as follows. In Section 2 I describe the movie exhibition industry and the vertical restraints imposed by distributors. Section 3 provides an overview of the demand and supply models, while Section 4 describes the dataset. Section 5 talks about estimation and identification, and Section 6 presents results of the estimation and counterfactual analysis. Section 7 concludes.

1.1 Vertical restraints in the movie industry

Movie distributors employ two types of vertical restraints to influence the movies being shown by exhibitors: *no screen-sharing* and *minimum run length.* The best way to understand their impact is in context of the theoretical economic models for *exclusive dealing* and *tying*, respectively. For the sake of the argument it helps to think of one week as the minimum length of time the distributor and exhibitor contract over.

No screen-sharing is a clause incorporated in the movie-rental contract between distributors and exhibitors which stipulates that the contracted-upon movie has the exclusive use of a given screen for the duration of the contract. This means that a movie theater with N screens can only screen N movies over the course of the week, and cannot put on a late-night show of a horror movie on a screen that during the day shows a kids movie, or reassign a screen on Friday and Saturday to provide additional seats when releasing a movie for which it expects the first screen to sell out[2]. One can view this practice as a form of *exclusive dealing*, a vertical restraint wherein a contract with an upstream firm prevents the retailer from carrying products of other firms[3]. If one considers each screen in the movie theater as a separate retailer than the *no screen-sharing* clause effectively means movie distributors engage in exclusive dealing for the duration of the contract.

Contract duration is another area over which the distributors exert control via vertical restraints. When renting a movie the exhibitor has to agree to keep the movie on the screens for at least a couple of weeks, the exact number of which is called the *minimum run length.* This can be thought of as an intertemporal form of *tying*, a vertical restraint wherein if the retailer wants to buy one product from the upstream firm he also needs to buy other *tied* products from the same firm[4]. Under this contractual restraint if an exhibitor wants to show the movie he automatically has to commit to showing it for two weeks or longer[5]. While there are instances in which this restraint is

[2] In cases where the exhibitor wants to show the movie on more than one screen the standard practice is to sign a separate contract for and additional screen, usually for the duration of the whole week.

[3] Examples of *exclusive dealing* include car dealerships which only carry one brand, gas stations and fast food restaurants

[4] Examples of *tying* include printers and ink cartridges, razors and replacement blades, cars and servicing at certified dealerships

[5] Through conversation with exhibitors I ascertained that the usual duration of *minimum run length* is 2-4 weeks, however I do not have exact data on this since often times there is no written contract, just an understanding between the exhibitor and distributor.

not binding (e.g. *Avatar* often ran in movie theaters for more than 2 months), exhibitors say they often have to keep on movies longer then they would absent the *minimum run length* restraint.

No screen-sharing and *minimum run length*can impact exhibitors' decisions in a variety of ways. First, because exhibitors have to commit to a movie for a long time, *ex ante* they may chose to play it safe with movie rental choices and go for the bigger movies usually offered by major distributors because of their appeal to wider audiences and, for most part, consistently high box-office draw. They may not be willing to risk taking on riskier movies whose appeal to local audiences is more uncertain, such as the movies offered by independent distributors. Second, they may not want to take on movies with a narrower appeal because they know there is only a small group of moviegoers who may want to see them - enough to warrant a few screenings, but not enough for a period of two or more weeks. Third, if the exhibitor finds *ex post* that one of the movies he's screening is attracting many fewer customers than expected he will not be able to drop this movie as quickly as he would like, which prevents other movies from being taken on. Finally, when constructing their schedule exhibitors often have 1-week holes to plug before a big release, but because of *minimum run length*restrictions they often cannot take on a new movie and instead need to keep on a movie they were already screening, longer than they would otherwise want to.

1.2 Exhibitor decisions without vertical restraints: examples

To illustrate how different movie rental and scheduling might be in the absence of the aforementioned vertical restraints I provide two examples. The first one is based on the experience of one (anonymous) exhibitor I have spoken to. In 2009 the exhibitor owned the only movie theater in a town of around 50,000 people. Since the movie theater was a small with only 4 screens, under *no screen-sharing* and *minimum run length* restraints imposed by the distributors the exhibitor could play only around 40-50 movies each year. Given this limitation, the type of movies shown were primarily big crowd-pleasers since they had the largest appeal to the masses and biggest potential to attract audiences for 2+ weeks. Recognizing untapped local demand for less standard fare (e.g. "high-brow" or independent movies) the exhibitor chose to defy his contractual obligations, hoping his movie theater was too small for this to raise the ire of the distributors, and started putting up a limited screening of an additional movie each week. Each week a new additional movie was screened once on Saturday, twice on Sunday and all day Monday, replacing screenings of the worst-performing movie that week. As soon as the local consumers got used to the "limited" release schedule the movie theater started attracting more patrons then before, with the moviegoers seeing the extra movie being mainly older, with more "sophisticated" tastes in movies. The cumulative attendance for the movie whose screenings were reduced fell hardly at all, with people who would want to see the movie choosing one of the remaining screenings[6].

Other examples come from countries such as France or Poland, where distributors do not impose vertical restraints which would limit the exhibitors' flexibility in movie rental and scheduling. Free of such restraints, exhibitors schedule limited screenings of kids movies on weekend mornings, allowing parents to take their children to the movies on days off work; they keep one screening per day of movies which had been playing for a while to capture any remaining demand, while filling the remaining slots on those screens with mainstream blockbuster movies; even blockbusters which

[6] It is also worth mentioning the same exhibitor subsequently decided to open a second, one-screen movie theater that would only play such "high-brow", artistic and independent movies, but found that screening these movies all day, every day of the week did not increase attendance much over the "limited" release schedule, which seemed to have captured most of the local demand for these movies

play daily are shown at different times - kids movies play mainly in the mornings and afternoons, while the majority of R-rated movie screenings happen in the late evenings. Finally, most movies have more screenings in their first week of release than in the following weeks, and last longer on the screens than in comparable movie theaters in the US.

1.3 Relationship to the Literature

The bulk of the theoretical literature on exclusive dealing has focused on determining its impact on competition and allocative efficiency. In a setting where retailers are also consumers Posner (1976) and Bork (1978), heading the so-called "Chicago critique", ague that exclusive dealing provisions can only have procompetitive impact; Comanor & Frech (1985) argue that they can only have an anticompetitive impact; others like Mathewson & Winter (1987) argue that the net impact on welfare depends on the exact nature of competition between manufacturers. Fumagalli & Motta (2006) address the question in a setting where retailers are not consumers and compete with each other, and also come away with the finding that the net impact on welfare is ambiguous, depending in their setting on the intensity of competition between retailers. Since all these arguments focus on price competition which is largely absent in the movie exhibition industry their applicability to my setting is limited, however all agree that a reduction in the size of the consumer choice set is welfare-decreasing, *ceteris paribus*. Theoretical literature on tying is similarly focused on its welfare implications. Representatives of the Chicago school of thought argue agains the "leverage theory" that a monopolist in market A can benefit from tying its product with good B, which is competitively supplied, and argue that alternative motivations for tying are not anticompetitive. Defending the leverage theory Whinston (1990) constructs a model where the monopolist in market A can use tying to achieve foreclosure of competing firms in market B, increasing profits as a result of his anticompetitive behavior.

Empirical investigations into exclusive dealings have been limited. Asker (2005) examines whether, in the Chicago beer market, exclusive dealing arrangements between upstream firms (brewers) and downstream sellers (beer distributors) leads to foreclosure; he finds no evidence that it does. Brenkers & Verboven (2006) evaluate the welfare impact of enhanced competition between downstream firms (car dealers) due to the removal of exclusive territory and exclusive dealing arrangements in the EU car market; they find that these vertical restraints do not increase manufacturer profits significantly, leading to the conclusion that they are most likely a by-product of other considerations. Finally, Ho et al. (2010) look at types of contracts offered in the movie rental industry by upstream firms (distributors) and contract chosen by downstream firms (movie retailers), with special interest in "full-line forcing", a form of tying under which the retailer needs to carry the full product range of the upstream firm; they find that the choice of contracts offered by distributors is profit-maximizing, but that the retailers choice of contract often is not.

There are many papers which empirically estimate demand in movie theaters. Ainslie et al. (2005) construct a model which uses aggregate box-office revenue data to determine consumers' demand for movies at the movie theaters. Basuroy et al. (2003) look at how demand for movies is impacted by factors such as critical reviews, movie stars and budget size, while Moul (2007) measures how it's impacted by word of mouth. Davis (2006) examines the way geographic dispersion of consumers and sellers affects the properties of markets by looking at movie exhibition, where moviegoers have preferences both over geographic and product characteristics. Einav (2007) looks at the US movie exhibition industry, with the aim to separate seasonality in underlying demand for movies from seasonality in the number and quality of movies released. In the management field Jehoshua

Eliashberg has a series of papers (Swami et al. (1999), Elberse & Eliashberg (2003) and Eliashberg et al. (2009), among others) in which he and his co-authors look at the way movie theaters go about choosing which movies to rent and when to screen them. Finally, Filson et al. (2005) and Gil & Lafontaine (2009) look at why movie-rental contracts between distributors and exhibitors employ revenue-sharing.

2 The Movie Exhibition Industry

The motion picture industry is a prominent economic sector, employing over 700,000 people and generating over $10.6 Billion in North American box office receipts in 2010[7] from over 1.3 Billion admissions. Moreover, motion pictures are a key driver of the market for entertainment products, one of the largest export markets the U.S.

Focusing on the movie exhibition industry, the value chain consists of three main stages: production, distribution, and exhibition. The process of production, which I do not look at in this paper, encompasses everything from the beginning, when a movie is only a concept in a screenwriter's mind, until the moment a movie is finalized in the form it will eventually be shown to paying customers in the theater. The biggest players at this level are the "majors", big studios which integrate production and distribution, as do the slightly smaller "mini-majors". In addition there are many "independent" producers who use either the big studios' distribution networks or work with independent distributors.

The next stage is distribution. A distributor designs and implements the strategy for releasing the finished movie onto multiple platforms, only the first of which is showing it in movie theaters[8]. For the theatrical release the most important choice is when to release the movie - when deciding this the distributor needs to take account of not only inherent seasonality of demand, but also of movies being released by the competition[9]. The timing of the release is crucial as most exhibitors who are constrained by vertical restraints and the number of screens will only be able to show a small portion of the 530+ movies being released each year on average in the period 2001-2010[10]. Other strategic decisions distributors make include what kind of release to pursue (wide or limited), the accompanying advertising campaign and deciding on movie rental contractual terms.

Finally, exhibitors are movie theaters owners, controlling anywhere from a single-screen theater in a local community to a nationwide chain of multiplexes. In 2000 32% of movie theaters had only 1 screen, a further 43% were "mini-plexes" with 2-7 screens, 20% had 8-16 screens and the remaining 5% had more than 16 screens[11]. Although the number of multiplexes has risen since then and some smaller movie theaters have closed since then, it is clear that movie theaters with 1-7 screens still serve a large portion of the population. With few exceptions they are not vertically integrated with distributors and are fully independent to pursue their own profit-maximizing strategies[12].

[7] source: MPAA (2010)

[8] Other media in which movies are releases are, in chronological order: movie theaters, on-demand, DVD/BluRay disks and VHS cassettes, cable and network television and, increasingly, the Internet; additionally, to the producers chagrin, increasingly movies also find their way to consumers via pirate channels

[9] For a more detailed discussion of distributors' decisions see Einav (2007) and Wen (2011)

[10] Source: MPAA (2010)

[11] Source: Davis (2006)

[12] After a period of vertical integration which peaked in 1945, in *United States v. Paramount Pictures* (1948) the Supreme Court decreed that studios were not allowed to own or directly control movie theaters (other practices which were disallowed in this ruling include: organizing exhibitors into runs, selling movies through blind block-booking

Exhibitors seek to maximize their profits from selling movie tickets and concessions[13]. At the beginning of the week exhibitors decide which movies currently showing to drop and which to keep, which new movies to rent from distributors and whether to get them "on the break" (the week they are released nationwide) or wait, and when to screen them throughout the week. Since movies are often shown for many weeks, exhibitors need to be mindful of the distributors' release schedule (which they know in advance) to ensure screens are available when big blockbuster movies are released (blockbusters are almost always shown on the break).

The exhibitor's scope to use price as a strategic variable is limited due to the industry business model of charging the same for all movies showing at the same time, although most engage in modest price discrimination between age groups (child and senior discounts), time of day (matinee discounts) and 2D/3D screenings. However, pressure from distributors ensures that ticket prices change little over time (see Chart 1 in Section 4) and in fact many theaters simply increase them once every couple of years to keep up with inflation. As a result, in my analysis I will abstract from price choice and assume ticket prices to be exogenous to the model.

2.1 Movie exhibition contracts

Exhibitors enters two types of contract with distributors they will usually have a 'master' contract with each distributor, as well as a separate 'movie rental' contract for each movie shown. The latter stipulate the exact revenue split for the movie, as well as restrictions on how flexible the exhibitor can be when scheduling the movie, including the *no screen-sharing* and *minimum run length* clauses. The process of 'booking' movies can either be handled directly by the exhibitor (typically small independent exhibitors and large chains do their own booking) or by professional 'movie bookers' who act as intermediaries and handle booking for multiple exhibitors at the same time.

Movie rental contracts between distributors and exhibitors involve revenue sharing, but are unusual in that there are no fixed fees. Historically the contracts operated on a *sliding scale*, wherein the distributor's revenue share (known as the *split*) started off high in the first week of a movie's release, but then fell each week the movie was on the screen, as described in Einav (2007) for the U.S. and Gil & Lafontaine (2009) for Spain. Over the last ten years, however, the industry in the United States has moved away from this model to one with a fixed split over time, such that in 2010 only 14% of movies in the sample were rented on sliding scale contracts[14]. The exact revenue split differs between movies, with exhibitors 'paying' more for blockbusters and less for niche and independent movies, but crucially differs little between movie theaters (see Table **??**). While historically the exhibitors/bookers could hope to negotiate the exact revenue split, today all but the largest nationwide exhibitor chains are effectively price-takers, with distributors offering standard, take-it-or-leave-it revenue split terms that are non-negotiable to all movie theaters[15]. While exhibitors do sometimes renegotiate the terms of movie rental *ex post*, this occurs rarely and

and assigning exclusive territories to exhibitors). While these rules have been relaxed slightly since then, allowing studios to take small interest in exhibitors, the vast majority of exhibitors today are not owned or controlled by the distributors

[13]The exhibitors' revenues come almost exclusively from screening movies, despite efforts in recent years to diversify to showing sports and cultural events

[14]The shift was driven by distributors who changed the type of contract for all their movies at some point in time, thus the type of contract offered on a movie is unrelated to the type of movie.

[15]This is driven by two trends which put more bargaining power in the hands of the distributors: first, the distribution market has become highly concentrated, with C(8) = 95% (source: Davis (2006)); second, the theatrical release's contribution to overall box office revenues has declined significantly over the last couple of decades.

mostly affects the terms in later weeks of the release, which have a much smaller impact on the exhibitor's bottom line than the early week.

3 The Model

In this section I present a model of demand and supply which I will then use to evaluate the impact of removing vertical restraints in the movie exhibition industry. Estimating this model will allow me to simulate the equilibrium without these restraints, and determine whether exhibitor profits and consumer welfare increase in the process.

The primary decision maker in the model is the exhibitor, whose objective is to maximize profits by deciding which movies and at what times to play at the movie theater. He is faced, on one side, with consumer demand for movies, while on the other side he has to choose from movies available to him from distributors. In the following two subsections I discuss the consumer demand model and exhibitor (supply) model.

3.1 Product Demand

3.1.1 Market Definition

Before we proceed further, it is important to find an appropriate market definition for this industry. I focus on movie theaters which are local monopolies[16], thus the population of the town in which the movie theater is located is a good starting point for defining the market in the geographic sense.

The temporal dimension should reflect the timeframe over which consumers make decisions and over which they explicitly compare alternatives. Since in the exhibition industry it has become customary for movie theaters to announce their schedules one week at a time, this provides a natural bound for the temporal dimension - I thus consider a market to be a movie theater/week combination.

3.1.2 Consumer Choice

I employ a discrete choice logic model to model consumer choice. Consumers choose directly from all movie/screening combinations available in a given market (henceforth referred to as *combos*). At the same time, I assume consumers only ever see a particular movie once.[17] Their choice thus depends on utilities of all combos available to them in a given week, excluding those with movies they had already seen.[18]

3.1.3 Agent's Utility Function

I assume the conditional indirect utility for consumer i from watching combo (m, s) (movie m, screening time s) in market tc (week t, movie theater/cinema c) to be of the form

[16]The closest potential competitor to a movie theater in my dataset is 26 miles away, which I consider to be a distance too large for most people to cover in order to see a movie.

[17]While this may overestimate the quality of movies such as "Avatar" or "Titanic" which attracted lots of repeated viewers, such movies are the fringe of movies in our dataset and thus I believe this to be a reasonable simplification.

[18]An advantage of this approach over a nested-logit model is that it allows us to capture that a consumer might inherently prefer movie A to movie B, but because the former does not screen at a convenient time for her she instead chooses to see movie B at a more convenient time.

$$u_{imsct} = u^M_{imt} + u^S_{isct} + \xi_{msct} + \epsilon_{imsct} \tag{1}$$

where u^M_{imt} is movie m's attractiveness in week t to agent i, u^S_{isct} is the attractiveness of screening time s in week t at cinema c to agent i, ξ_{msct} is a product characteristic observable to the consumer but not to the econometrician, while ϵ_{imsct} is an individual-movie-screening-week specific, mean-zero stochastic term that captures idiosyncratic consumer heterogeneity that's observable to the consumer but not to the econometrician.

Modeling consumer's utility in this fashion has important implications for how we believe consumers go about choosing whether to go to the movie theater and which movie to see. Additive separability between u^M_{imt} and u^S_{ist} reflects that consumers choose freely between movies and screenings, as described in Section 3.1.2. u^{MT}_{it} captures how if a consumer has not been to the movie theater in a long time she has a stronger desire to do so and is more likely to justify the outlay.

3.1.4 Attractiveness of movie m: u^M_{imt}

I model movie m's attractiveness to agent i as follows:

$$u^M_{imt} = x^M_m \beta^M_i + \omega_{im} + f_m(w_{mt}) \tag{2}$$

where x^M_m is a vector of observable movie characteristics, including its genre, MPAA rating, the box-office draw of the director/actors and awards, ω_{im} is a consumer-specific, mean-zero signal of movie m's quality and $f_m(\cdot)$ captures the decay of a movie's attractiveness as a function of w_{mt}, the number of weeks since the movie's release[19]. By allowing the decay function parameters to differ between movies I capture the fact that while some movies (e.g. genre movies, summer blockbusters) lose a lot of audience in weeks 2 and later, others (e.g. indie movies, Oscar contenders) actually become more attractive in the initial release period as word-of-mouth builds up and remain relevant to moviegoers for much longer periods.

3.1.5 Attractiveness of screening time s: u^S_{isct}

I model the attractiveness of screening time s as follows:

$$u^S_{isct} = p^A_{isct}\beta^P_i + \chi^S_s(p)\beta^S_i + x^O_{sct}\beta^O_i \tag{3}$$

where p^A_{isct} is the price of admission for individual i to screening s[20], $\chi^S_p(s)$ are indicator variables whether screening time s is in one of the time periods p (e.g. Mon 5-8pm) and β^M_i is the utility individual i derives from going to see a movie at time s. By making utility from screening times depend on consumers' demographic characteristics, I am able to capture that e.g. children are able to see movies earlier in the day while working people can only see movies in the evenings or on weekends.

x^O_{sct} is a vector of characteristics capturing factors impacting the opportunity cost of going to see a movie at time s, cinema c and in week t, including outside temperature, precipitation, dummy

[19]This effect is part of what Einav (2007) captures in the "decay factor" λ - the fact that advertising and "buzz factor" when the movie is release give people higher utility from seeing it early on. However, unlike in that paper, λ^M_m does not capture the fact that audience numbers fall as the pool of people who have not yet seen the movie shrinks - this effect will be captured explicitly by excluding movies people have already seen from their choice set, as described in Section 3.1.1.

[20]Although price does not differ between individual movies *per se*, it is generally higher for movies shown in 3D

for cultural events and holidays. In effect $x^O_{sct}\beta^O_i$ is the observable opportunity cost of going to the movies at this specific time and location that's common to all consumers, while the unobservable component is captured by ξ_{msct} and ξ_{0ct}.

3.1.6 Attractiveness of the outside good: u_{i0ct}

u_{i0ct} captures the value of the outside alternative not already explained by $x^S_{sct}\beta^S_i + \xi_{msct} + \epsilon_{imsct}$ (see Section 3.1.5):

$$u_{i0ct} = \epsilon_{i0ct} \tag{4}$$

where ϵ_{i0ct} is the idiosyncratic, consumer-specific value of the outside option.

3.1.7 Heterogeneity in consumer tastes

β^M_i, β^S_i, β^O_i, β^{MT}_i are consumer-specific coefficients which reflect heterogeneity in movie tastes in the population. We model them as multivariate normal with the mean dependent on observable demographic variables and parameters to be estimated, and a variance-covariance matrix to be estimated:

$$\begin{pmatrix} \beta^M_i \\ \beta^S_i \\ \beta^P_i \\ \beta^O_i \end{pmatrix} = \begin{pmatrix} \beta^M \\ \beta^S \\ \beta^P \\ \beta^O \end{pmatrix} + \Pi D_i + \Sigma v_i, \qquad v_i \sim N(0, I) \tag{5}$$

where Π captures how consumer demographics D_i impact their preferences and Σ captures idiosyncratic parameter variance between individuals[21].

3.1.8 Specification of the unobserved product characteristic: ξ_{msct}

We model the unobserved component ξ_{msct} as

$$\xi_{msct} = \xi_m + \xi_s + \xi_c + \xi_w + \xi_y + \xi_h + \Delta\xi_{msct} \tag{6}$$

where ξ_c captures the relative attractiveness of each movie theater compared to alternative forms of local entertainment; ξ_t captures the attractiveness of seeing movies in week t[22] while ξ_{mt} captures movie m's attractiveness that's not captured by its observable characteristics and how it diminishes differently to that predicted by the model laid out in Section 3.1.4. In the empirical application ξ_c, ξ_t and ξ_{mt} will be captured by movie theater-, week- and interaction movie/week-fixed effects, respectively. Finally, the leftover econometric error term $\Delta\xi_{msct}$ captures any remaining unobserved preferences.

3.1.9 Mean product and idiosyncratic utility

Substituting the definitions of β_i from (5) and ξ_{msct} from (6) into the indirect utility function in (1) I get:

$$u_{imsct} = \delta_{msct} + (x^M_m + p^A_{isct} + \chi^S_s(p) + x^O_{sct})(\Pi D_i + \Sigma v_i) + \epsilon_{imsct} \tag{7}$$

[21] My empirical implementation uses a diagonal Σ, although correlations between coefficients can easily be added.
[22] This will also capture the seasonality in the industry, as described by Einav (2007)

where

$$\delta_{msct} = \underbrace{x_m^M \beta^M + \xi_m}_{\lambda_m} + f_m(w_{mt}) + p_{isct}^A \beta^P + \chi_s^S(p)\beta^S + \xi_s$$
$$+ x_{sct}^O \beta^O + \xi_c + \xi_w + \xi_y + \xi_h + \Delta\xi_{msct} \tag{8}$$

is the mean utility for product $msct$. The term λ_m captures the quality of movie m that's not fully known to the exhibitor, which will feed into the learning model described in Section 3.2.3.

3.1.10 Uncertainty and scope for learning in the model

The part of u_{imsct} that depends on the movie m is as follows:

$$u_{imsct} = ...x_m^M \beta^M + f_m(w_{mt}) + \xi_m + \omega_{im} + \Delta\xi_{msct} + \epsilon_{imsct}$$

The following is a breakdown of uncertainty of the side of the consumer and the exhibitor:

Consumer:

$$u_{imsct} = ... \underbrace{x_m^M \beta^M + f_m(w_{mt})}_{\text{knows}} + \underbrace{\xi_m + \omega_{im}}_{\text{noisy signal of } \xi_m} + \underbrace{\Delta\xi_{msct} + \epsilon_{imsct}}_{\text{knows}}$$

The consumer receives a noisy signal of the unobserved movie quality, $\xi_m + \omega_{im}$, which informs his choice of which movie to see. One can think of ω_{im} as the positive or negative impression the consumer gets from seeing advertising, trailers and previews for the movie. In the context of this model I will not attempt to capture the way in which consumer learn ξ_m as too few moviegoers see a movie more than once for this process to have a sizable impact on my results.

Exhibitor:

$$u_{imsct} = ... \underbrace{x_m^M \beta^M + f_m(w_{mt})}_{\text{knows}} + \underbrace{\xi_m + \Delta\xi_{msct}}_{\text{noisy signal of } \xi_m} + \underbrace{\omega_{im} + \epsilon_{imsct}}_{\text{simulates}}$$

Exhibitor uncertainty, on the other hand, will play a central role in my paper. Exhibitors receive a noisy signal of movie m's quality, $\xi_m + \Delta\xi_{msct}$, which helps them learn the true quality of the movie. The full way in which this learning will proceed is described in Section 3.2.3.

3.1.11 Market shares resulting from choice model

Given the choice model described above, the set of consumers who choose combo (m, s) in week t at movie theaters c is defined as

$$A_{msct}(x_{msct}, \xi_{msct}; \theta) = \{(D_i, \epsilon_{imsct}) | u_{imsct} > u_{im's'ct} \forall m', s' s.t. (m', s') \neq (m, s) \text{ and } m \notin \iota_{it}\} \tag{9}$$

where x_{msct} and ξ_{msct} are the observable and unobservable characteristics, respectively, of combo (m, s), D_i and ϵ_{imsct} are observable and unobservable consumer characteristics, respectively, θ are all the model parameters and ι_{it} is the list of all the movies seen by consumer i up to period t.

By keeping track of consumer moviegoing history, ι_i, I am able to ensure consumers see a given movie only once, as described in Section 3.1.2. This allows me to explicitly model consumer selection, wherein the first batch of consumers to see a movie may be different from those who choose to see it in later weeks.

The aggregate market share of combo (m, s) in week t at movie theater c is the sum of all consumers who choose that option

$$s_{msct}(x_{msct}, \xi_{msct}; \theta) = \int_{A_{msct}} dP^*(D, \epsilon) = \int_{A_{msct}} dP^*(\epsilon) dP^*(D) \tag{10}$$

where $P^*(.)$ denotes a population distribution function and the second equality follows from an assumption of independence of D and ϵ.

3.1.12 Impact of capacity constraints on demand estimation

The model does not explicitly take into account capacity constraints, however I believe the impact this will have on demand estimates will be negligible, as in my sample less than 1% of screenings get sold out. One could expect, however, that consumers may choose to avoid going to screenings they *expect* could be close to sold out, even though they prefer this screening to all others, *ceteris paribus*. Not modeling this explicitly could lead to underestimating the value people place on the most popular screening times (e.g. Friday evening) if there is a substantial number of screenings which are sold close to capacity. However, since in my sample less than 3% of screenings sell more than 75% of capacity I believe any potential bias on the estimates from this source to be negligible.

3.2 Exhibitor Decision

Given the demand model described in Section 3.1 the exhibitor decides which movies to show and when so as to maximize her profits.

3.2.1 Model overview and dynamics

At the beginning of each period t the exhibitor makes a decision about which movies to show in the coming week and when to screen them[23]. When making this decision the exhibitor takes into account:

1. Profits in week t from movie exhibition and concession sales
2. The impact reduced uncertainty as to the quality of movies screened is likely to have on profits in future periods

The exhibitor aims to maximize the PDV of the stream of profits $\mathbb{E}[\sum_{\tau=t}^{\infty} \beta^{\tau-t} \boldsymbol{\pi}_\tau | \boldsymbol{s_t}]$, leading to the Bellman equation

$$V(\boldsymbol{s_{ct}}) = \max_{\boldsymbol{x_{ct}} \in X_{ct}} \mathbb{E}\left[\boldsymbol{\pi}(\boldsymbol{s_{ct}}, \boldsymbol{x_{ct}}; \lambda) + \beta V(\boldsymbol{s_{ct+1}}) \middle| \boldsymbol{s_{ct}}, \boldsymbol{x_{ct}}\right] \tag{11}$$

where X_{ct} is the set of all possible actions for exhibitor c in period t, $\boldsymbol{\pi}$ is the per-period expected profit function, $\boldsymbol{s_{ct}}$ are the state variables, λ is the unknown movie quality and β is the discount factor.

[23] As mentioned in Section 2 the exhibitor's capacity to use price as a strategic variable is severely limited, thus I will abstract from the choice of ticket prices when considering the exhibitor's profit maximization problem.

3.2.2 Per-period profit function

The exhibitor's expected per-period profit function is

$$\boldsymbol{\pi}(\boldsymbol{s_{ct}}, \boldsymbol{x_{ct}}; \lambda) = \sum I_{msct}^{z_{ct}^*} \int_{A_{msct}(\boldsymbol{s_{ct}})} (p_{imsct}^A r_{mt} + p_{ct}^C) dP^*(\epsilon) dP^*(D) - C_c \tag{12}$$

where $I_{msct}^{z_{ct}^*}$ is a an indicator function whether movie/screening combo (m, s) is shown by exhibitor c in week t under optimal schedule $z_{ct}^* = z(s_{ct}, x_{ct})$, r_{mt} is the revenue split for movie m in period t, p_c^C is the average concession profits per moviegoer in movie theater c,[24] and C_c are weekly costs of operating movie theater c.[25] It is also important to note that $A_{msct}(\boldsymbol{s_{ct}})$, the set of consumers who choose combo (m, s) in week t at movie theaters c, depends on $\{\iota_{it}\}_{\forall i}$, the set of movies seen by potential moviegoers up to period t.

3.2.3 Beliefs and Learning

The prior distribution $p_{ct}(\lambda_m)$ describes exhibitor c's belief about unknown movie quality λ_m at the beginning of period t. The exhibitor can obtain new information about λ_m by screening movie m in periods $S_{mct} \subset S_{ct}$, where S_{ct} is the set of all possible screenings for exhibitor c in period t, and observing the market share it captures. This information is captured by vector $\zeta_{mct} = \{\zeta_{msct}\}_{s \in S_{mct}}$. At the end of period t the exhibitor updates his prior: $p_{ct+1}(\lambda_m) = p_{ct}(\lambda_m | \zeta_{mct})$.

After observing the market shares in week t the exhibitor uses the demand system inversion to calculate the mean product utility δ_{msct}, as described in (8). The exhibitor knows the values of all variables in (8) apart from the movie quality λ_m and the unobserved product characteristic $\Delta\xi_{msct}$, and thus for each product (m, s) he observes a noisy signal $\zeta_{msct} = \lambda_m + \Delta\xi_{msct}$ of movie quality.

We assume the exhibitor's initial prior on λ_m is a normal distribution. Since the distribution of both $\Delta\xi_{msct}$ is normal and that of ζ_{sct} is multivariate normal the posterior is also normally distributed:

$$\begin{aligned}
\Delta\xi_{msct} &\sim N(0, \sigma_\xi^2) \\
\zeta_{sct} &\sim N(\mu_{ct}, \sigma_\xi^2) \\
p_{ct}(\lambda) &\sim N(\mu_{ct}, \sigma_{ct}^2) \\
p_{ct+1}(\lambda) &\sim N(\mu_{ct+1}, \sigma_{ct+1}^2)
\end{aligned}$$

where μ_{ct+1} and σ_{ct+1}^2 are calculated using bayesian updating on the basis of observed signals[26]. Intuitively, the larger the number of screenings of movie m the more information the exhibitor will gather on λ_m, reducing σ_{ct+1}^2 and providing a more accurate estimator μ_{mct+1}.

[24] I assume concession sales do not vary in a meaningful way between movies and screenings, which is supported by Gil & Hartmann (2007) who find that concession sales are roughly proportional to total attendance in Spanish movie theaters.

[25] These costs are independent of the exhibitor's choice of $\boldsymbol{x_{ct}}$, which is a reasonable assumption as long as the opening hours/days and number of screens operating are kept constant throughout the model

[26] In order to calculate μ_{ct+1} and σ_{ct+1}^2 for a movie m that was screened N_m times in period t we use the formulae

$$\mu_{mc(2)} = \mu_{mc(1)} + \frac{\sigma_{mc(1)}^2(\zeta_{msct(1)} - \mu_{mc(1)})}{\sigma_{mc(1)}^2 + \sigma_\xi^2}$$

$$\sigma_{mc(2)}^2 = \sigma_{mc(1)}^2(1 - \frac{\sigma_{mc(1)}^2}{\sigma_{mc(1)}^2 + \sigma_\xi^2})$$

s.t., abusing notation slightly, $\mu_{oct} = \mu_{mc(1)}$ and $\mu_{ct+1} = \mu_{mc(N_m)}$ (correspondingly for σ_{mct+1}).

3.2.4 State variables

The state variables in the exhibitor problem include, for week t:

- $\{\iota_{it}\}_{\forall i}$: the set of movies seen by consumer i up to period t;
- $\{\mu_{mct}, \sigma^2_{mct}\}_{\forall m \in M_{ct}}$: parameters summarizing exhibitor c's belief about unknown movie quality λ for movies he's considering playing in week t;[27]

3.2.5 Actions and optimal schedules

As described in Section 3.2.3 what drives the learning of unknown movie quality λ_m is the number of screenings of movie m in period t, and not the exact timing of those screenings. As such, the decision vector x_{ct} consists of the number of times each movie under consideration $m \in M_{ct}$ would be played in week t: $x_{ct} = (x_{1ct}, x_{2ct}, ..., x_{mct}, ..., x_{M_{ct}ct})$ where $x_{mct} \in \{0, 1, 2, ..., |S_{ct}|\}$ and $\sum_{m \in M_{ct}} x_{mct} \leq |S_{ct}|$.

Given a decision vector x_{ct} and state vector s_{ct} the exhibitor then needs to decide what is the optimal schedule $z^*_{ct} = z(s_{ct}, x_{ct})$ for week t:

$$z(s_{ct}, x_{ct}) = \underset{z_{ct} \in Z_{ct}}{\arg\max} \left(\boldsymbol{\pi}(\boldsymbol{s_{ct}}, \boldsymbol{z_{ct}}; \lambda) \right) \qquad \text{s.t.} \sum_s I^{z_{ct}}_{msct} = x_{mct} \forall m \in M_{ct} \tag{13}$$

where $z_{ct} = \{I_{msct} \forall m, s\}$ and Z_{ct} is the set of all possible schedules for exhibitor c in week t.

Intuitively, what the $z(\cdot)$ function does is it determines how to schedule x_{mct} screenings of movie m, $\forall m$, so as to maximize exhibitor c's profits in period t given by $\pi(\cdot)$. For example, if $x_{mct} = 2$ for movie m that appeals to adults, $z(\cdot)$ will ensure that the two screenings of movie m are scheduled at times convenient to adults e.g. evenings and weekends.

3.2.6 Movie theater characteristics

In my model I take movie theaters' location and number of screens as given. While this allows me to focus my investigation on the direct impact of removing vertical restraints, there may be potential long-term equilibrium consequences which my model will not capture. Since the removal of vertical restraints allows exhibitors to fit more movies on the same number of screens, some movie theaters may find that their optimal number of screens is lower than what they have. While existing movie theaters are unlikely to reduce the number of screens due to sunk costs and relatively small costs of screening movies on existing screens, this effect might lower the number of screens in newly-constructed movie theaters. Similarly, in a world without vertical restraints different competitive equilibria might arise in markets with more than one movie theater, however such considerations are beyond the scope of this paper.

3.3 Distributor behavior

I take the distributors' strategic behavior as given - this implies that in the counterfactual experiments distributors (a) release the same movies as in the base case (b) release movies at the same time as in the base case (c) do not change the terms of movie rental (revenue split). These are

[27] The set M_{ct} includes movies played in week $t-1$ and those that opened in week t and two weeks prior - in order to keep the state space manageable I'm excluding movies the exhibitor dropped in the past to which he will not be coming back to.

strong assumptions, and I need to carefully consider whether they are realistic and what is their likely impact on the results of the counterfactual experiments.

First, the assumption that distributors release the same movies as they did in the base case. As mentioned in Section 2, movies are released over multiple channels and their prospective theatrical box office revenue is only one factor which impacts whether they are given the green light in the production process, thus any impact on the set of movies that are produced is likely to be limited. Nonetheless, it is possible that of the movies that lose out most from removal of vertical restraints in exhibition (bad big-budget movies) some might never get made, while in the group that gains most (independent movies, niche movies) more movies are produced. Additionally, more movies might reach the US market from abroad. While the net welfare impact of these changes is impossible to determine with certainty, the fact that movie production is likely to more closely reflect consumer tastes suggest any such impact should be positive.

The assumption that movies are released at the same time as in the base case is a reasonable one, to the extent that movies' release dates are primarily driven by when production finishes and the movie is ready to screen, as well as seasonality of movie demand. While strategic considerations in the timing of movie releases, as described by Wen (2011), may change the release schedule slightly (and thus $\{M_{ct}\}_{\forall t}$), its impact on consumer welfare is likely to be limited, while the directionality is unclear.

The assumption that distributors do not change revenue split terms on movie rental in the counterfactual universe is the strongest, and most contentious one I make with regards to the distributors. Strategic considerations between distributors are complicated, and I am not able to model them as this would require me to observe the whole US movie market. It is possible that distributors could lower revenue split terms on movies which lose in the counterfactual universe to help them regain market share - not accounting for this might overestimate the changes to the consumer choice set and reduce the incremental number of movies shown in movie theaters, leading to an overestimation of consumer welfare rise. However, there are reason I believe any such bias would be small in magnitude. First, not only do distributors have a limited range over which to change the revenue split terms[28], but also their impact on exhibitors' decision is limited by the fact movie theaters earn substantial profits on concession sales, which are independent of the revenue split. Second, given the nature of movie demand, revenue split changes are unlikely to affect whether movie theaters take on new movies *per se*, which is the primary source of consumer welfare gain, and rather could only affect marginal movie screenings later in these movies' runs.

4 Data

The data required to estimate the consumer demand model consists of attendance data broken down by movie and screening time for each movie theater, distribution of consumer demographics within each market, movie characteristics and additional drivers of consumer demand, such as ticket prices and advertising spending by distributors. Additionally, in order to solve the exhibitor problem information is required on the revenue split for each movie available from the distributors, as well as information on movie-specific vertical restraints faced by the exhibitors.

Attendance data comes directly from exhibitors, which collect this information for accounting purposes. In total 4 exhibitors representing 5 movie theaters agreed to take part in this study,

[28] In my sample revenue split varies between 35% and 70%, however the majority of movies fall in the 45% - 60% range

giving me access to their attendance data sets. Each movie theater in the sample can reasonably be thought of as having a local monopoly on movie exhibition. The exact period for which the data was made available varies by movie theater. For each movie/screening combo I observe the total number of attendees; additionally, some movie theaters provide the breakdown of the total figure by type of ticket sold, allowing me to observe the regular/child/senior ticket split.

Table 1: Data summary statistics

	min	max	mean
Screens per theater	3	6	4.3
Avg. annual attendance	33,015	174,131	84,197
Avg. annual # movies	63.6	86.4	78.5
Avg. # weeks on screens (by movie)	2	3.8	2.9
Local population	4,380	44,737	13,134.3
Data period (mths)	14	64	35.8

NOTE: for now only data in the table are those for the movie theater I've been working with to date

I obtained the majority of movie characteristics from The Internet Movie Database (IMDB)[29], an online database of information related to movies, television shows, actors, production crew personnel and other entities in visual entertainment media. These characteristics include the movie's distributor, genres[30], the MPAA rating and Academy Awards nominations and wins; it also includes a measure of box office draw of the movie's main stars and director[31]. I obtain the movie's critical ratings from two sources: consumer ratings from IMDB and professional critic ratings from Metacritic[32], an aggregator service for visual entertainment. Additionally, I collect the distributors' spending on movie advertising nationwide from *Ad $ Summary*. The movie theaters provide information on ticket prices, which did not change considerably throughout the sample period (see Chart 1). The distribution of consumer demographics was obtained by sampling individuals from the US Census of Population.

[29]www.IMDB.com

[30]IMDB allows for users to add more than one genre tag to a title, leading to a multitude of genre types, thus in my analysis I only include those tags that describe more than 5% of the movies released in the period 2005-2010.

[31]Proxied by the combined box office revenues of the previous 5 movies made by (a) the movie's three main starts (b) the movie's director.

[32]www.Metacritic.com

Chart 1: Ticket prices by movie theater

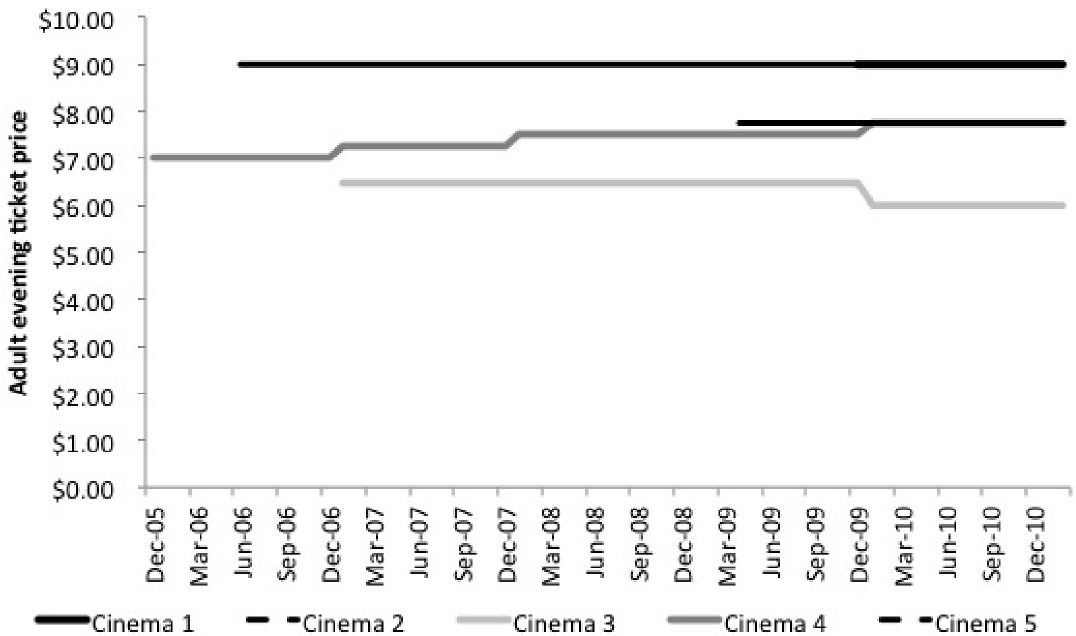

NOTE: Prices is for regular non-discount, non-3D, evening ticket; prices for other times/age groups generally move in tandem

Revenue split information, like attendance information, was obtained directly from exhibitors. Unfortunately, not all exhibitors in my sample were willing to divulge this information, and I was only able to collect this information for 3 of the movie theaters in my sample. I was also not able to obtain the information for movies which were not played by any of the movie theaters in my sample. The limitations of a thus constructed dataset are twofold. First, it does not cover all movie theaters in my sample. I rectify this by assuming that each exhibitor is offered the same revenue split terms by the distributor. Second, revenue split information is not available for all movies released over the period covered by the analysis, which could incorrectly limit M_{ct}, the set of movies exhibitors consider when making their rental/scheduling decision. I address this issue by constructing a reduced-form model which seeks to predict the revenue split offered by distributors using a movie's characteristics. For a full description of this approach, the model's fit and its impact on the analysis see Section 5.

4.1 Reduced Form Findings

Analyzing the data by how long exhibitors keep movies on their screens illustrates the choices they face under the restrictions currently imposed on them by the distributors. If an exhibitor wants to release the movie within the first 3 weeks of when it is released nationwide it will almost certainly have to commit to a 2+ weeks minimum run length - this is illustrated in Chart 2 by the fact most movies released early on run for 2-4 weeks, with almost no movies having a 1-week run. If, on the other hand, an exhibitor does not believe a movie can attract enough audience to justify keeping it on for the minimum run length required of an early release, he may choose to wait until the distributor is willing to reduce the minimum run length, at the expense of losing a part of the audience who will only see the movie if it is released early. For movies which were released in weeks 4 and beyond the vast majority only runs for one week. This finding is all the stronger for the fact that if movies are released late they often come with a more favorable revenue split for the movie theaters, which *ceteris paribus* encourages them to keep movies on for longer.

Chart 2: Number of weeks on the screen, by week of release

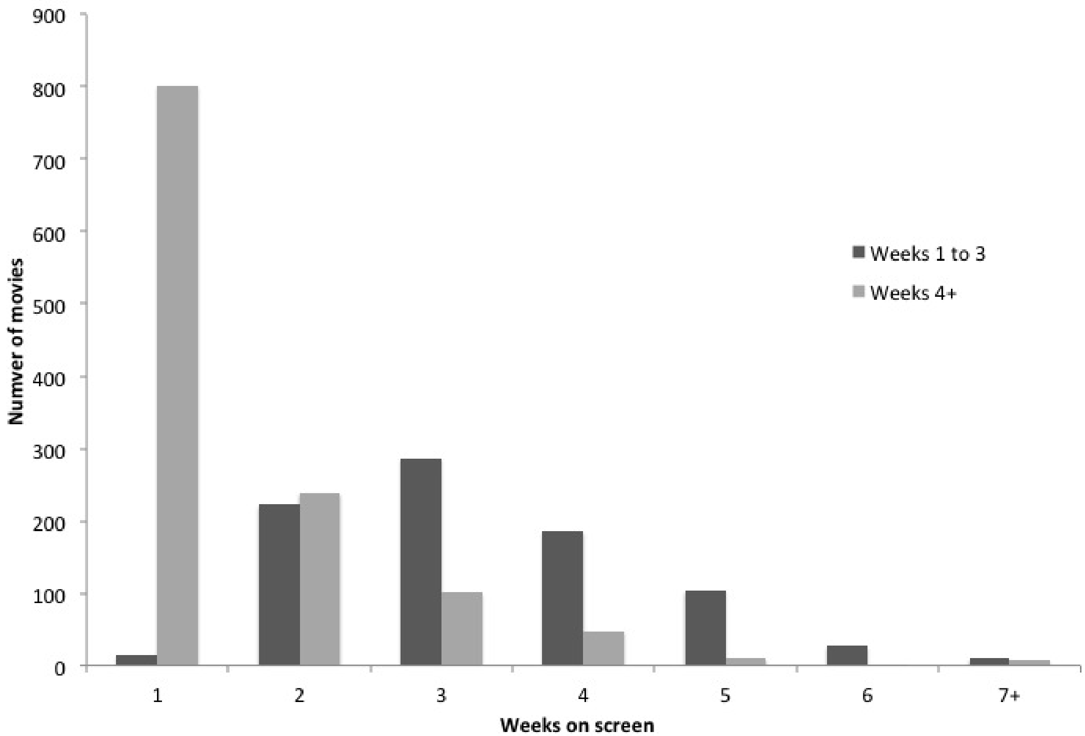

The result of contractual restraints imposed by the distributors are two-fold. First, since cinemas cannot screen more than one movie on a screen each week, the number of movies they can screen each year is limited by their screen capacity constraints. As illustrated in Chart 3, the more screens a movie theater has the more movies overall it plays, suggesting that exhibitors prefer to offer a larger variety of movies if they can. Second, since often distributors shorten the "minimum run length" period if the movies are not released on the break, movie theaters can fit more movies on their screens if they forego releasing them on the break. This is illustrated in Chart 3 by the fact that movie theaters with fewer screens release a smaller percentage of their movies on the break, which allows them to fit more movies overall. This in turn contributes to the fact that e.g. a 6-screen movie theater does not show 6 times as many movies as a 1-screen movie theater.

Chart 3: Percentage of all movies released shown by the movie theaters

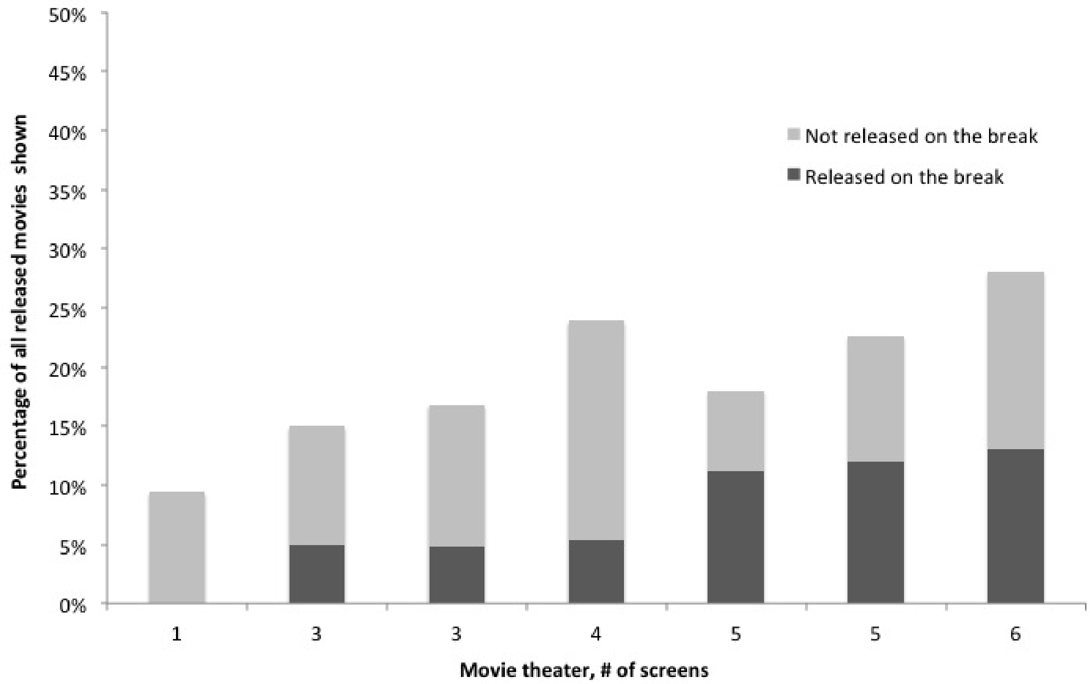

5 Estimation and Identification

My estimation strategy follows the standard (GMM) approach established by Nevo (2001), with additional micro-moments to aid estimation following Petrin (2002).

5.1 Standard BLP moments

Unlike almost all other papers using the BLP framework, I do not have an endogenous variable in my set of linear explanatory variables. As a result, setting up moments based on the unobserved product characteristic not captured by fixed effects, $\Delta\xi_{msct}$, is straightforward and yields as many moments as the number of parameters that enter the utility function linearly.

5.2 Micro Moments

In my estimation I use five sets of micro moments. The first three sets derive from information in the MPAA 2010 Theatrical Market Statistics, an annual publication put out by the Motion Pictures Association of America, while the fourth and fifth sets are derived straight from the data.

1. Moviegoing frequency, as captured by proportion of population which falls into one of four buckets: Never, Once a year, Less than once a month, More than once a month; this yields 4 moments[33]

2. Moviegoing frequency by age group (2-11, 12-17, 18-24, 25-59, 60+), relative to average moviegoing average frequency (e.g. people aged 18-24 go to the movies on average 1.7 times more often than the average person); this yields 5 moments

[33]Since this distribution reflects the nationwide moviegoing frequency of 4.1 in my estimation I only match this moment for an average of 2 out of 5 movie theaters, for which the moviegoing frequency is close to the nationwide average

3. Age composition of frequent moviegoer group (those who go to the movies more than once a month); this yields 5 moments

4. Moviegoing frequency by movie theater - this places an explicit penalty for when the model cannot capture different average moviegoing frequency (or, equivalently, overall attendance) across movie theaters in the sample; this yields 5 moments

5. Attendance at each screening by age group (2-12, 13-59, 60+) - since the additional child / adult / senior attendance split is only available for some movie theaters I utilize micro moments to capture this additional information, thus allowing information from all movie theaters to be used in estimation; this yields 3 moments.

In total I have 22 additional micro moments.

5.3 The Objetive Function

The two sets of moments that enter the GMM objective function are $\boldsymbol{G}_1(\boldsymbol{\theta})$, the standard BLP moments, and $\boldsymbol{G}_2(\boldsymbol{\theta})$, the micro moments. The population moment conditions are assumed to uniquely equal zero at the truth $\boldsymbol{\theta}_0$:

$$E[\boldsymbol{G}(\boldsymbol{\theta}_0)] = E\begin{bmatrix}\boldsymbol{G}_1(\boldsymbol{\theta}_0)\\ \boldsymbol{G}_2(\boldsymbol{\theta}_0)\end{bmatrix} = 0 \tag{14}$$

The GMM estimator then takes the form

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}} \boldsymbol{G}(\boldsymbol{\theta})'\boldsymbol{W}\boldsymbol{G}(\boldsymbol{\theta}) \tag{15}$$

Consistency GMM estimation requires that the weighting matrix, W, is a positive semi-definite matrix.

6 Results

6.1 Parameter Estimates

Table 2 presents the coefficient estimates for two different demand-side models: ordinary least squares (OLS) Logit and random coefficients with micromoments (RC + MM).

Table 2: Demand model coefficients

	OLS Logit	RC + MM
Linear coefficients		
Constant	-8.86	-10.99
	(0.247)	.
3D	0.21	0.18
	(0.025)	.
Log(week on screen)	-0.9	-0.24
	(0.011)	.
Log(Release week)	-0.24	-0.27
	(0.011)	.
Release week not important	-0.81	0.17
	(0.118)	.
Fixed Effects		
Movie	yes	yes
Movie theater	yes	yes
Year	yes	yes
Week	yes	yes
Holidays	yes	yes
Time/day of week	yes	yes
Non-linear coefficients		
Π: PG * age(2-11)		-3.10
Π: PG-13 * age(12-17)		-5.53
Π: R * age(18-24)		4.18
β_0^{MT}		-6.69
Σ		3.04
β_1^{MT}		2.61
Σ		1.86
$\Pi_{age(18-24)}$		2.74
$\Pi_{age(25-59)}$		2.01
No. of observations	56,527	56,527
R-squared	0.616	0.746
GMM objective	290.8667	20.5206
of which non-micromoment	5.5960	5.8274

6.2 Model Fit

Chart 4: Micro moment #1: Moviegoing frequency

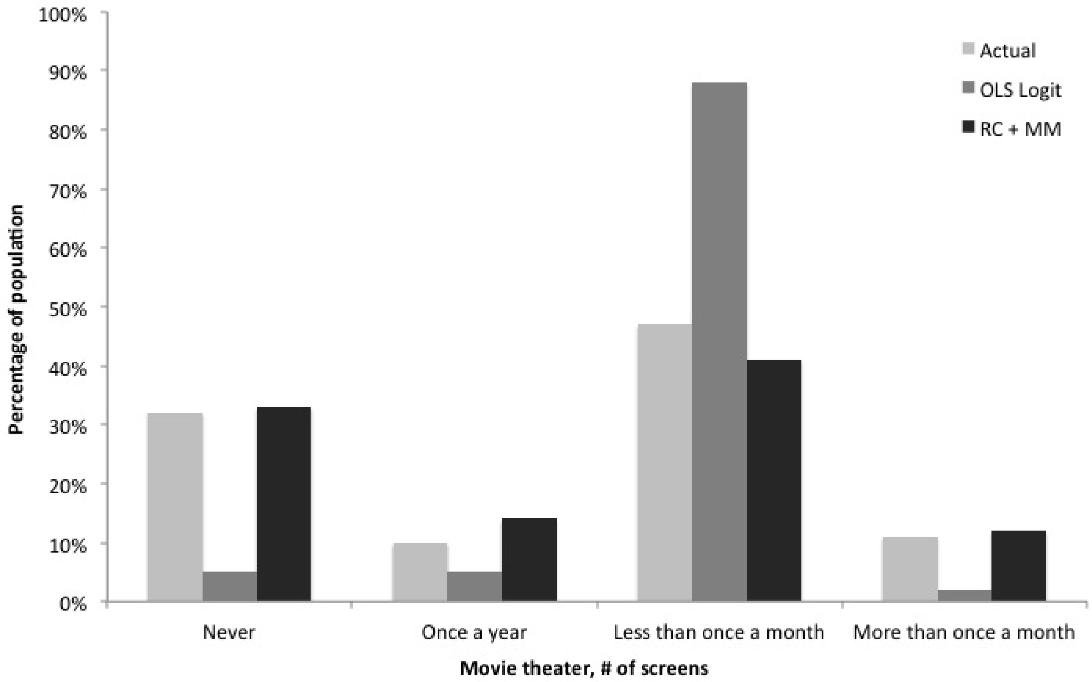

Here is the starkest showing of how much better the model with non-linear parameters does than the OLS Logit model.

Chart 5: Micro moment #2: Average moviegoing frequency, by age

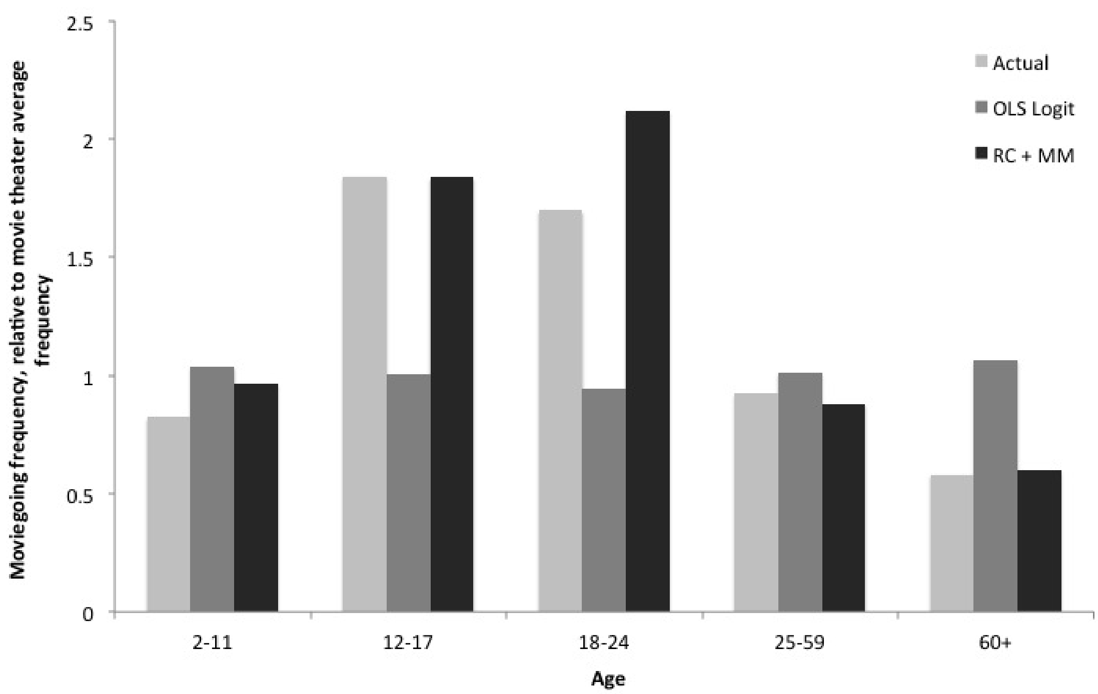

Chart 6: Micro moment #3: Composition of frequent moviegoer group

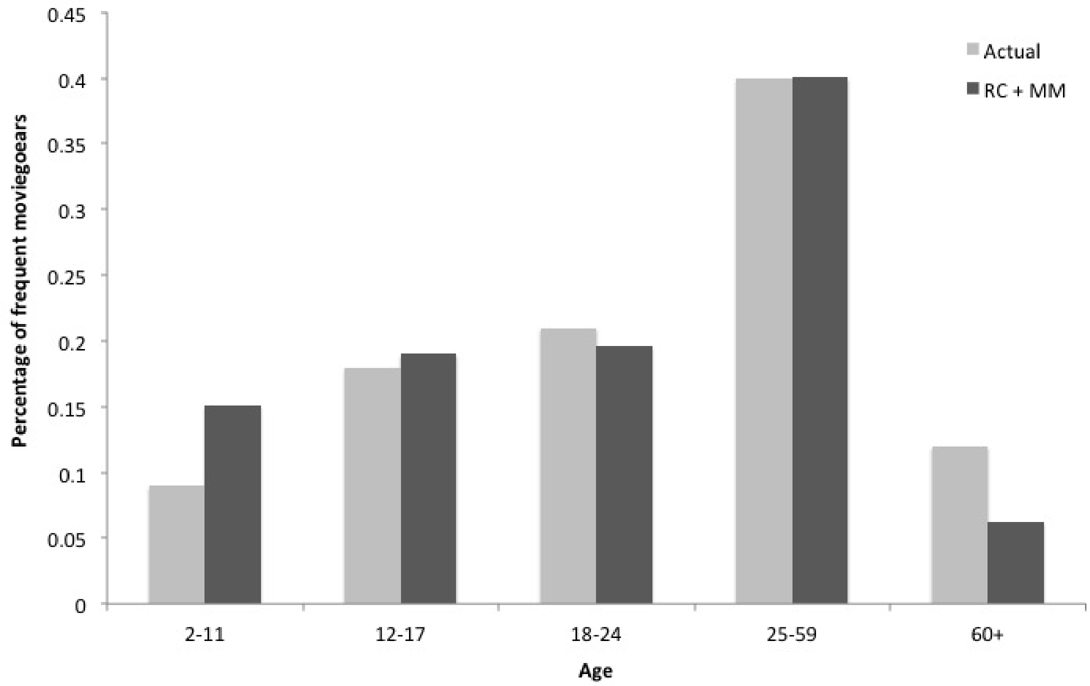

Chart 7: Micro moment #4: Average moviegoing frequency, by movie theater

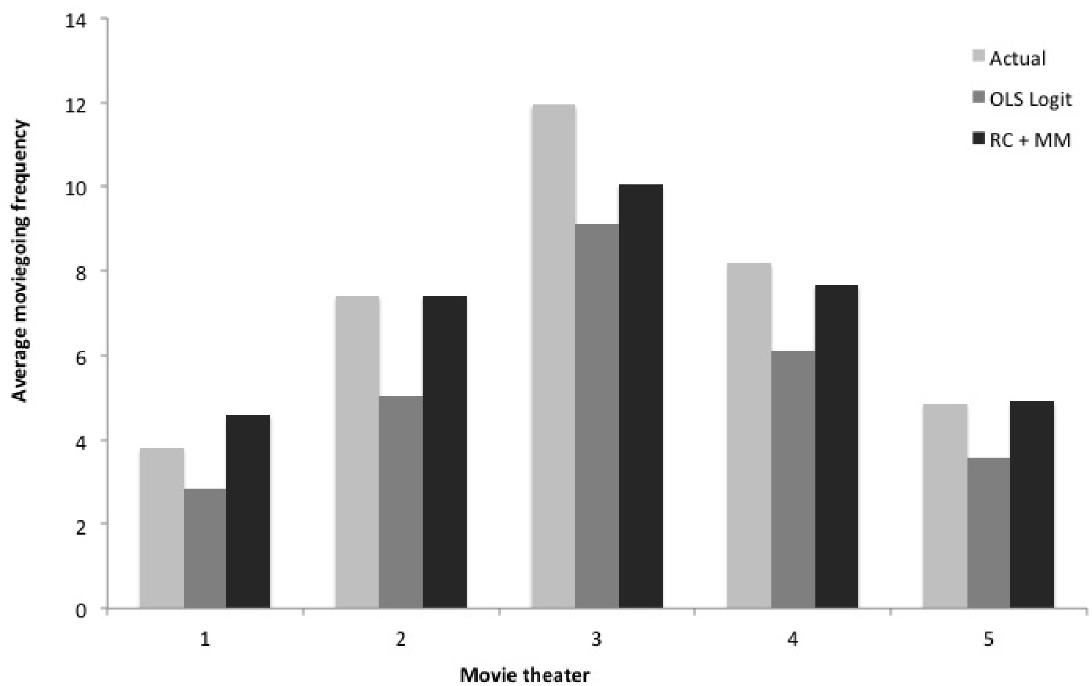

Still trying to figure out the consistently lower frequency numbers for OLS Logit model, especially that if you look at the next chart the fit of predicted attendance is very good for OLS Logit (it may have something to do with simulation error; again, more to come here)

Chart 8: Predicted/Actual attendance, by movie theater

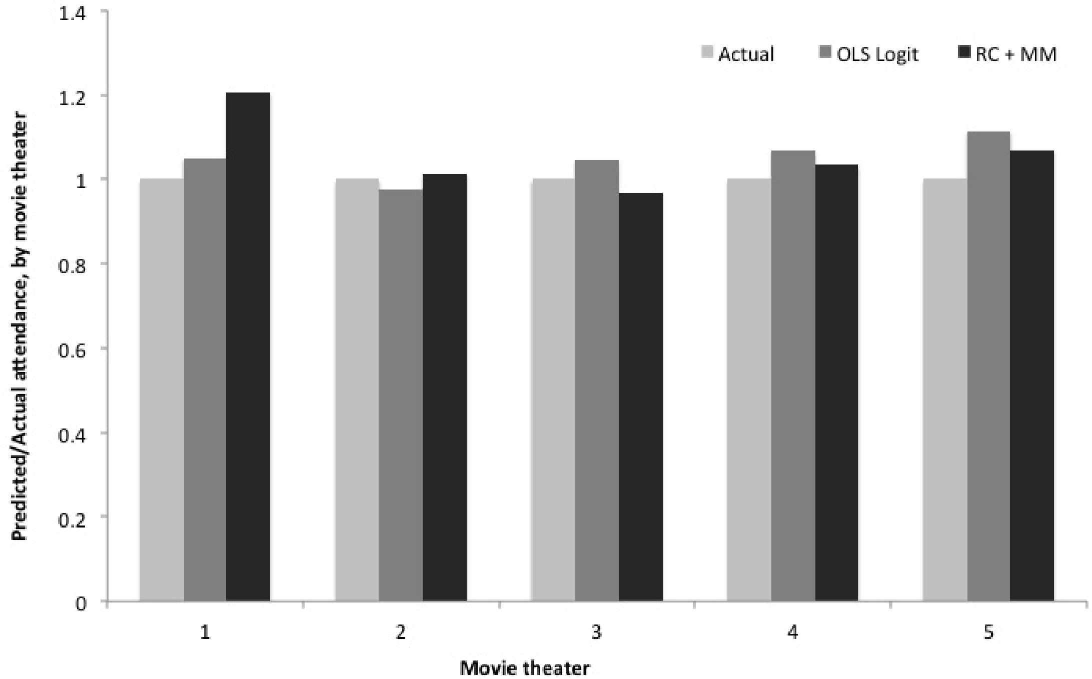

6.3 Counterfactual Results

Note: due to a lack of a final implementation of the revenue split model the results focus on attendance rather than revenues, and are the result of exhibitors maximizing the former rather than, as is intended, the latter. However, given the relatively small differences in revenue splits between movies and the relatively high importance of concession sales (which I will assume will be directly proportional to attendance) I expect the results presented below to closely foreshadow those I will obtain once I implement revenue maximization in the exhibitors problem.

Chart 9: Increase in number of movies screened, by MT, from removing contractual restraints

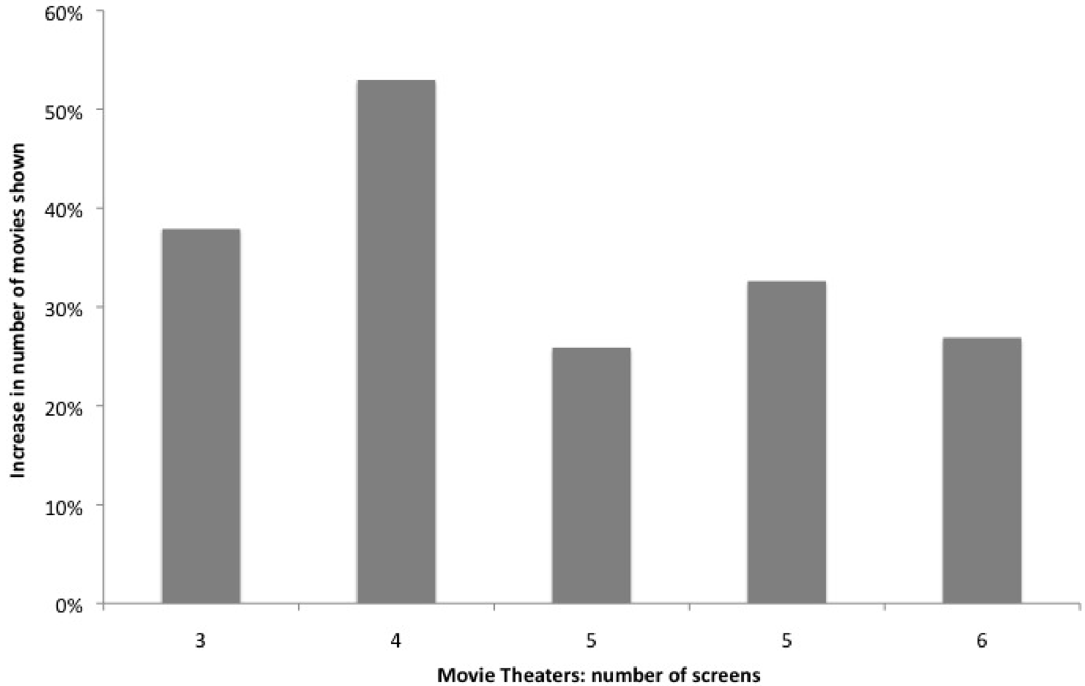

As expected, removing contractual restraints increases the number of movies screened, as illustrated in Chart 9. Smaller movie theaters experience a higher increase, which is consistent with expectations that contractual restraints combined with a limited number of screens limits their ability to profitably offer a broader range of movies.

The percentage of all movies shown released by the "majors" falls from 65% to 60%. This decrease is small, however I expect it to be larger once I implement a fully dynamic learning model. That model will give movie theaters an incentive to screen movies which they would not otherwise, on the off-chance that such movies resonate with local audiences and attract a significant number of moviegoers.

Chart 10: Increase in attendance, by MT and learning model, from removing contractual restraints

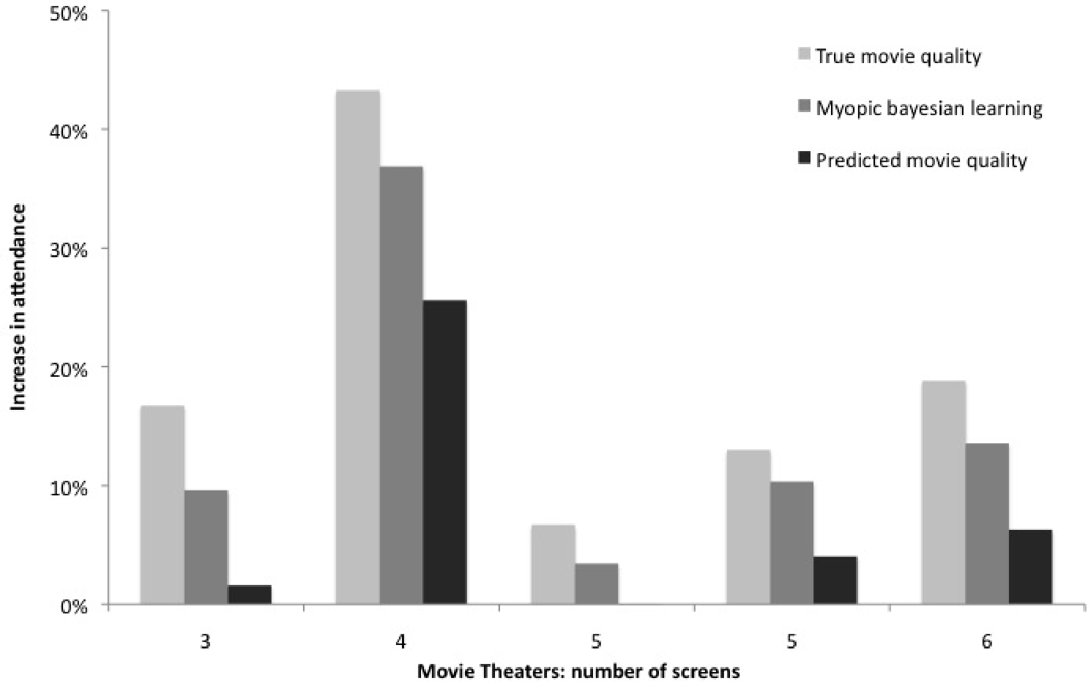

Removing contractual restraints also increases attendance, as illustrated in Chart 10. Smaller movie theaters benefit more from the removal of vertical restraints than larger ones, especially the movie theater with 4 screens. This supports the notion that offering a broader range of movies, as well as being able to schedule them freely throughout the week, allows movie theaters to profitably attract more consumers.

Also, the increase in attendance from removing the restraints depends on how well we believe movie theaters know the quality of the movies *ex ante*. The largest increase comes if movie theaters know the exact quality of each movie before they screen it, allowing them to accurately predict attendance and thus construct the schedule accordingly. On the other extreme, if the movie theaters can only rely on predicted movie quality when constructing their schedules, and do not learn a movie's quality from screening it, then the increase from removing contractual restraints, though still positive, is much smaller. Finally, if movie theaters learn movie quality in a bayesian fashion from screening them than they can use the information gathered this way to help them increase attendance almost as much as though they possessed perfect movie quality information, since the learning process happens quickly between weeks[34].

7 Conclusion

References

Ainslie, A., Drèze, X., & Zufryden, F. (2005). Modeling movie life cycles and market share. *Marketing Science*, *24*(3), 508–517.

[34]For now I only have results for a myopic bayesian learning model, but it is reasonable to expect that implementing a fully dynamic learning model will lead to a higher increase in attendance.

Asker, J. (2005). Diagnosing Foreclosure due to Exclusive Dealing.

Basuroy, S., Chatterjee, S., & Ravid, S. (2003). How critical are critical reviews? the box office effects of film critics, star power, and budgets. *Journal of Marketing*, *67*(4), 103–117.

Bork, R. (1978). The antitrust paradox: A policy at war with itself. *New York*.

Brenkers, R., & Verboven, F. (2006). Liberalizing a distribution system: the European car market. *Journal of the European Economic Association*, *4*(1), 216–251.

Comanor, W., & Frech, H. (1985). The competitive effects of vertical agreements? *The American Economic Review*, *75*(3), 539–546.

Davis, P. (2006). Spatial competition in retail markets: movie theaters. *The RAND Journal of Economics*, *37*(4), 964–982.

Einav, L. (2007). Seasonality in the us motion picture industry. *The Rand journal of economics*, *38*(1), 127–145.

Elberse, A., & Eliashberg, J. (2003). Demand and supply dynamics for sequentially released products in international markets: The case of motion pictures. *Marketing Science*, (pp. 329–354).

Eliashberg, J., Hegie, Q., Ho, J., Huisman, D., Miller, S., Swami, S., Weinberg, C., & Wierenga, B. (2009). Demand-driven scheduling of movies in a multiplex. *International Journal of Research in Marketing*, *26*(2), 75–88.

Filson, D., Switzer, D., & Besocke, P. (2005). At the movies: The economics of exhibition contracts. *Economic Inquiry*, *43*(2), 354–369.

Fumagalli, C., & Motta, M. (2006). Exclusive dealing and entry, when buyers compete. *The American Economic Review*, *96*(3), 785–795.

Gil, R., & Hartmann, W. (2007). The role and determinants of concession sales in movie theaters: Evidence from the Spanish exhibition industry. *Review of Industrial Organization*, *30*(4), 325–347.

Gil, R., & Lafontaine, F. (2009). Using Revenue–Sharing to Implement Flexible Pricing: Evidence from Movie Exhibition Contracts.

Ho, J., Ho, K., & Mortimer, J. (2010). Welfare effects of full-line forcing contracts in the video rental industry. *NBER Working Paper*.

Hotelling, H. (1929). Stability in competition. *The economic journal*, *39*(153), 41–57.

Lafontaine, F., & Slade, M. (2005). Exclusive contracts and vertical restraints: Empirical evidence and public policy. *Handbook of Antitrust Economics, Paolo Buccirossi (Ed.), Cambridge: MIT Press, Forthcoming*.

Mathewson, G., & Winter, R. (1987). The competitive effects of vertical agreements: Comment. *The American Economic Review*, *77*(5), 1057–1062.

Moul, C. (2007). Measuring word of mouth's impact on theatrical movie admissions. *Journal of Economics & Management Strategy*, *16*(4), 859–892.

MPAA (2010). Theatrical market statistics.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, *69*(2), 307–342.

Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, *110*(4).

Posner, R. A. (1976). Antitrust Law: An Economic Perspective.

Swami, S., Eliashberg, J., & Weinberg, C. (1999). SilverScreener: A modeling approach to movie screens management. *Marketing Science*, *18*(3), 352–372.

Wen, J. (2011). *Commitment and Coordination in Timing Games: An Empirical Study of the U.S. Movie Industry*. Ph.D. thesis, New York University.

Whinston, M. (1990). Tying, foreclosure, and exclusion. *The American Economic Review*, *80*(4), 837–859.