

Vertical Restraints in the Movie Exhibition Industry*

Krzysztof Wozniak[†]

NYU Stern

November 21, 2012

Abstract

This paper analyzes vertical restraints imposed by distributors on movie theaters in the movie exhibition industry: *minimum run length* and *no screen-sharing*. These restraints, a form of *exclusive dealing*, help prevent hold-up when distributors provide exhibitors with costly analog movie copies; however this efficiency effect disappears as exhibitors switch to digital projection. The paper poses two questions: without an efficiency effect what is the welfare cost of these restraints, and is lifting the restraints an equilibrium outcome for the distributors? A structural model of industry demand and supply is estimated using a uniquely detailed panel data set of attendance and movie rental contracts collected directly from a sample of US exhibitors. Counterfactual results indicate lifting the restraints for digital movie theaters result in an increase in consumer welfare, exhibitor and distributor profits of 3.9% to 17.1%, 5.6% to 20.3% and 1.8% to 10.1%, respectively. However, despite this overall increase in profits it is not an equilibrium for distributors to lift restraints, suggesting a policy banning them would be welfare increasing.

*PRELIMINARY AND INCOMPLETE. PLEASE DO NOT RE-DISTRIBUTE WITHOUT PERMISSION

[†]E-mail: kwozniak@stern.nyu.edu

1 Introduction

Although it is a tenant of free-market thinking that giving economic agents freedom to make decisions is welfare-improving, efficiency-enhancing restrictions can be justified in markets where a *laissez-faire* approach generates inefficiencies e.g. in presence of externalities, or in cases of monopoly. Vertical restraints are a form of contractual restrictions imposed by upstream manufacturers either directly on consumers, or on downstream retailers in markets where retail is separated from manufacturing. They play a crucial role in determining the nature of competition: some, like resale price maintenance, give upstream firms influence over prices paid by consumers, while non-price restraints such as exclusive dealing, exclusive territories, or tying impact the set of products available to consumers. However, even in markets where restraints have a efficiency-enhancing effect their net impact on consumer welfare can be ambiguous, and in practice needs to be calculated using empirical analysis. This is reflected in the evolution of legal treatment of vertical restraints, with a move towards an empirical “rule-of-reason” approach to evaluating their competitive effect.¹

The aim of this paper is to conduct a structural empirical investigation of how non-price vertical restraints can be used by upstream manufacturers to impact the set of products offered to consumers by downstream retailers. The industry being studied is the movie exhibition industry, where contractual restrictions have been used to mitigate an inefficiency created when the landmark 1948 ruling in *United States v. Paramount Pictures, Inc.* which separated movie exhibition and distribution. Specifically, for a movie theater using analog projection technology to be able to screen a movie the distributor needs to provide him with a costly copy: a movie reel. This creates potential for *hold-up*: once the exhibitor² receives the movie reel he can choose to not screen the movie unless the distributor reduces the rental price. The contractual restraints employed by distributors, *minimum run-length* and *no screen-sharing*, solve the hold-up problem by guaranteeing a number of screenings for each movie. However, as the industry switches to digital projection where the costs of producing movie copies are close to zero this inefficiency is eliminated, thus removing the efficiency-enhancing justification for the use of vertical restraints.

The movie exhibition industry makes for a good setting to analyze the impact of non-price vertical restraints. First, while ticket prices vary between movie theaters and screenings they do not vary between movies, eliminating scope for vertical restraints that target price. Second, the industry is characterized by differentiated goods with uncertain *ex ante* quality, heterogeneous consumer tastes

¹Up 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 is 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)

²An exhibitor is the owner/operator of movie theaters, and can own anywhere between one (for independent exhibitors) and hundreds of movie theaters (e.g. AMC Loews, Regal Entertainment Group in the NA market)

and retailer capacity constraints - this makes the composition of the *consumer choice set* (set of products available to consumers) a crucial driver of welfare and profits. Exhibitors rent movies from upstream distributors and sell consumers tickets to movie screenings, keeping a fixed percentage of the sales proceeds. Their challenge is to offer consumers a movie schedule that best caters to their diverse tastes (a *broad choice set*) and best reflects their true valuation of movies (a *relevant choice set*). When deciding their schedules movie theaters are constrained by the vertical restraints imposed by distributors. The *minimum run length* restraint stipulates that exhibitors need to keep a given movie on screen for a minimum period of two to four weeks, depending on the movie, while the *no screen-sharing* restraint means that a given screen can only play one movie throughout the week. The scheduling impact of these restraints is twofold. First, since movie theaters have a limited number of screens at their disposal both restraints reduce the number of movies they can show each year, reducing the breadth of the consumer choice set - this is especially true for smaller movie theaters.³ Second, the *minimum run length* period prevents exhibitors from quickly replacing poorly performing movies with more attractive ones, reducing the relevance of their offering.

This paper poses two research questions. First, what is the welfare cost of the contractual restraints if they have no efficiency effects? Second, is lifting the restraints an equilibrium outcome for the distributors? Besanko & Perry (1993) suggest that adoption of exclusive dealing restraints can be a result of a prisoner's dilemma game in that while it is an equilibrium for everyone to adopt exclusive dealing each manufacturer would in fact prefer nonexclusive dealing - this hypothesis is tested using counterfactual simulations. These questions have a timely relevance to the movie exhibition industry, which is set to complete its switch to the digital projection by the end of 2013 (Geuss, 2012). In addition, findings from this paper readily translate into other vertically-separated, capacity-constrained industries such as retailing, radio/TV program scheduling, and advertising.

To answer these questions a structural model of industry demand is constructed and estimated using a unique, detailed panel data set of moviegoer attendance collected directly from a sample of US exhibitors. All movie theaters in the data set are local monopolies, which allows the analysis to abstract from competitive considerations. Consumer demand is modeled using a flexible random-coefficient logit framework that is best suited to discrete choice events such as going to the movies. The panel structure of the data set allows for the control of any ticket price variance across screenings and exhibitors, as well as for explicitly modeling consumer selection over time. The demand system is supplemented with supply-side models of exhibitor and distributor decisions: how, when faced with uncertainty about a movie's quality, exhibitors learn its true value and make optimal scheduling decisions, and how do distributors set movie rental prices? The welfare impact of lifting contractual

³The average number of movies shown each year by a movie theater in the sample is 78 out of over 500 movies released annually over the sample period (MPAA, 2010)

restraints is calculated by constructing counterfactuals where exhibitors are allowed full flexibility when deciding movie screening schedules.

Contributions and Related Literature Early empirical investigations into the effects of vertical restraints took reduced form approach (Heide et al., 1998; Slade, 1998; Chipty, 2001; Sass, 2005; Zananone, 2009). This paper contributes to the small but growing body of empirical literature which takes a fully structural approach to studying the impact of vertical restraints on market equilibria. Asker (2005) examines whether, in the Chicago beer market, exclusive dealing arrangements between upstream brewers and downstream beer distributors leads to foreclosure; he finds no evidence that it does. Brenkers & Verboven (2006) evaluate the welfare impact of enhanced competition between car dealers (downstream retailers) due to the removal of exclusive territory and exclusive dealing arrangements in the EU car market; they find removing these restraints does not lead to a significant loss in manufacturer profits. Ho et al. (2010) look at types of contracts offered in the movie rental industry by upstream distributors and contract chosen by downstream 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. Lee (2012) evaluates welfare implications of software exclusivity in the US video game industry; he finds the vertical restraints to have severely reduced industry sales and consumer welfare.

In order to fully capture substitution pattern between movies and screening times this paper takes advantage of a very detailed data set of moviegoer attendance. Previous papers have used reduced-form models to explain consumer demand using aggregate box-office revenue data (Basuroy et al., 2003; Ainslie et al., 2005; Davis, 2006; Einav, 2007; Moul, 2007); this is the first paper to use a fully structural approach to capture demand in this industry. The demand model builds on techniques introduced in Berry (1994) and Berry et al. (1995), and further developed by Nevo (2001). It also uses micro-moments to aid estimation following Petrin (2002). Due to the complexity of the demand system, in which demand for one movie screening is dependent on what other movie/screening combinations are offered, optimization algorithms such as those employed in Swami et al. (1999), Elberse & Eliashberg (2003) and Eliashberg et al. (2009) cannot be used; instead, a greedy heuristic is employed to calculate expected improvements from removing the restraints. The supply-side employs a simple take-it-or-leave-it bargaining model to determine rates charged by distributors for movie rental, as used by Ho (2009).

The findings from the paper inform the debate within theoretical literature on whether exclusive dealing restraints have a pro-efficiency or anticompetitive impact on markets.⁴ Authors such as

⁴For an overview of the literature see Ornstein (1989), Lafontaine & Slade (2005), Rey & Tirole (2007), Riordan (2008) and Whinston (2008)

Telser (1960), Marvel (1982), Klein & Murphy (1988), Besanko & Perry (1993) and Segal & Whinston (2000) show how exclusive dealing restraints can be used to promote costly investment in retailers by upstream manufacturers. On the other hand, Krattenmaker & Salop (1986), Aghion & Bolton (1987), Mathewson & Winter (1987), Rasmusen et al. (1991), Bernheim & Whinston (1998), Comanor & Rey (2000) and Fumagalli & Motta (2006) show how exclusive dealing restraints can be used to foreclose competitors. This paper also provides an empirical investigation of the *hold-up* problem, best described by Klein (1988).

In addition, the richness of the data used in this paper provides a good insight into an industry where revenue-sharing contracts are used. This contributes to a substantial literature that looks at why revenue-sharing contracts are used in the movie exhibition industry (Hanssen, 2002; Filson et al., 2005; Gil & Lafontaine, 2009), the video rental industry (Dana Jr & Spier, 2001; Cachon & Lariviere, 2005; Mortimer, 2008) and the motion picture industry in general (Chisholm, 1997; Weinstein, 1998).

Road Map The structure of the paper is as follows. Section 2 provides an overview of the movie exhibition industry and the vertical restraints imposed by distributors, while Section 3 describes the data. Section 4 provides a detailed description of the demand and supply models. Section 5 talks about estimation and identification, while Section 6 presents results of the estimation and counterfactual analysis. Section 7 concludes.

2 The Movie Exhibition Industry

The industry value chain consists of four stages: *production*, *distribution*, *exhibition* and *consumption*. The process of *production*, not analyzed in this paper, encompasses everything from the beginning until the movie is ready to be shown to paying customers. A *distributor* owns rights to the finished movie and decides on a release strategy for multiple platforms, the first of which is showing it in movie theaters.⁵ For the theatrical release the most important choice is when to release the movie, how many theaters to release it at and how much to charge theaters for screening the movie.⁶ Distributors can be split into two categories: *Majors*, the biggest studios which have their own production studios and offer a wide variety of movies, and *non-Majors*, which are smaller and offer many independently-produced movies with a narrower audience appeal.⁷ *Exhibitors* are

⁵Other media in which movies are releases are, in chronological order: movie theaters, on-demand and online services, DVD/BluRay disks and VHS cassettes, cable and network television and, increasingly, pirate channels

⁶For a more detailed discussion of distributors' decisions see Einav (2007) who investigates seasonality in movie demand, and Wen (2011), who focuses on distributors' decision when to release movies

⁷This paper uses the common definition of Majors as the "Big Six" distributors who are part of media conglomerates: Paramount Pictures (Viacom), Warner Bros Pictures (Time Warner), Columbia Pictures (Sony), Walt Disney Pictures/Touchstone Pictures (The Walt Disney Company), Universal Pictures (Comcast/General Electric) and 20th

movie theaters owners, controlling anything from a single theater to a nationwide chain of multiplexes. With few exceptions they are not vertically integrated with distributors and sign movie rental contracts on a movie-by-movie basis.⁸

2.1 Movie rental contracts

In the United States movie rental contracts between distributors and exhibitors employ a linear pricing schedule: there are no fixed fees, and each dollar of revenue from movie ticket sales is divided between the two parties on the basis of a *revenue split* that is contracted-on in advance. This means that if no one comes to see a movie the exhibitor does not need to pay the distributor anything. All one-off costs such as the cost of producing and shipping the movie reel to movie theaters and local advertising costs are covered by the distributors. Historically, contracts in the US employed a *sliding scale*, wherein the distributor’s revenue share started off high in the first week of a movie’s release and fell in subsequent weeks (Einav, 2007; Gil & Lafontaine, 2009; Gil, 2009). In recent years, however, the industry has moved toward a model with a revenue split that is constant over time, and in 2010 only 14% of movies in the sample were rented on sliding scale contracts.⁹

Exhibitors enter two types of contracts with distributors - one master contract with each distributor and separate movie rental contracts for each movie shown. The movie rental contract includes restrictions on how flexible the exhibitor can be when scheduling the movie, as well as the exact revenue split for the movie e.g. a revenue split of 60 means the distributor gets 60% of the revenue while the exhibitor gets 40%. Revenue splits differ between movies, with exhibitors “paying” more for blockbusters and less for niche and independent movies, and to a smaller extent between movie theaters. The terms also differ based on how much time has passed since the movie’s nationwide release: exhibitors face higher revenue splits if they want to release a movie *on the break* (the week of the nationwide release) than if they release it *on the second run* (usually two to three weeks after the nationwide release). The exact value of the revenue split is the result of bargaining between the distributor and the exhibitor. Contracts cover the whole period a movie is screened at the movie theater, while renegotiations are rare.

Century Fox (News Corporation); non-Majors are defined as all other distributors active in the market

⁸After 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” 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; the exception is Sony Entertainment’s ownership of Loews Theaters between 1989 and 2001

⁹The shift was driven by distributors who changed the type of contract for all their movies in the decade prior to 2010, and since this was not accompanied by any meaningful shift in the type of movies offered by distributors the type of contract offered on a movie is unrelated to the type of movie

2.2 Vertical restraints

Movie distributors employ two types of vertical restraints to influence the movies being shown by exhibitors: *no screen-sharing* and *minimum run length*.

No screen-sharing stipulates that the contracted-upon movie has the exclusive use of a screen for the duration of the contract. In practice this means that a movie theater can only screen as many movies as it has screens over the course of a week.¹⁰ This restraint applies to all movies alike.

Minimum run length is the smallest amount of time that the exhibitor is contractually obliged to screen the movie for.¹¹ The value of the restraint varies between movies, and most exhibitors need to commit to showing movies released on the break for at least two weeks.¹² While in many instances the minimum run length restraint is not binding, exhibitors say they often have to keep movies on for longer than they would without the restraint, especially if the movie turns out to have less appeal to audiences than the exhibitor expected.¹³

The best way to view these restraints is as a form of *exclusive dealing*, a vertical restraint under which the retailer can only carry products of one manufacturer.¹⁴ If one considers each screen in the movie theater as a separate retailer then the no screen-sharing clause effectively means movie distributors engage in exclusive dealing for the duration of the contract. Since the shortest period of time that distributors contract over is one week, the no screen-sharing restraint guarantees an exclusive dealing period for one week, while the minimum run length extends this period to 2-4 weeks.

No screen-sharing and minimum run length can impact exhibitors' decisions in a variety of ways. First, because exhibitors have a limited number of screens at their disposal they can only show a small fraction of all movies released. Because exhibitors have to commit to a movie for a long time, when faced with *ex ante* uncertainty about movie quality they often chose to play it safe

¹⁰The exhibitor cannot, for example, 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 to provide additional seats when releasing a movie for which it expects the first screen to sell out. Instead, when the exhibitor wants to show the movie on more than one screen the standard practice is to sign a separate contract for an additional screen, usually for the duration of the whole week

¹¹This restraint provides a floor only - the exhibitor can keep the movie for a longer period of time than that specified under minimum run length without renegotiating the contract

¹²Conversations with exhibitors suggest the usual duration of *minimum run length* for movies released on the break is 2 or 3 weeks, with few cases of 4 week requirements, while movies released at a later time have a 1-2 week minimum run length requirement. Unfortunately, exact data on the minimum run length restraint was not available since generally the restraint is a product of an unwritten understanding between the exhibitor and distributor

¹³This is in contrast to Gil & Lafontaine (2009), who reports that in the Spanish movie exhibition market the minimum run length period is rarely binding

¹⁴Examples of exclusive dealing include car dealerships which only carry one brand, gas stations, Coca-Cola's agreements with fast food restaurants and movie theaters etc.

and rent only “big” movies usually offered by major distributors because of their appeal to wider audiences and, for most part, consistently high box-office draw.¹⁵ This leaves comparatively less empty screentime available to other movies, especially for smaller movie theaters.¹⁶ Second, they may not want to take on movies with a narrower appeal, often offered by independent distributors, 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 of non-stop screenings. Finally, if the exhibitor finds *ex post* that one of the movies he’s screening is attracting 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.

2.3 Projection technology and Staggered Release Schedule

For each screen showing a movie using analog 35mm technology a movie reel has to be produced and shipped to the movie theater, the cost of which is paid in full by the distributor.¹⁷ Industry estimates and conversations with exhibitors put the cost of such a movie reel at \$1,500 (Alimurung, 2012). For a US nationwide release on over 3,000 screens this corresponds to a cost the distributor of around \$4.5M, a substantial portion of the budget for all but the biggest movies. It also compares to average weekly distributor revenues for movie theaters in the sample of between \$3,000 and \$15,000.

In order to reduce costs distributors stagger their analog release schedule. Under this scheme, distributors produce fewer movie reels than there are exhibitors wanting to screen the movie, and only give the biggest of them the chance to release the movie *on the break*. A smaller exhibitor then gets a chance to screen the movie on the *second run* after one of the first-run movie theaters is done with it. Because movie reels need to be shipped usually there is a one week period during which a movie reel is not used. Distributors offer second-run movies at lower revenues splits and, generally, with no minimum run-length restrictions.

Digital projection, first introduced in 1999, does away with the movie reel, bringing down the cost of a copy of the movie to below \$100 (for a hard drive) or even effectively zero (for digital downloads). In general, digital movie theaters gain access to all movies on the break. Although digital projectors require a substantial upfront financial outlay, this technology has been growing

¹⁵Uncertainty regarding movie quality is highlighted by Gil & Lafontaine (2009) as one of the reasons the industry uses revenue-sharing contracts

¹⁶Due to similar concerns in the retail industry an antitrust ruling case against The Coca-Cola Corporation (Gasparon & Visnar, 2005) mandated that if the company offered free coolers to small retail stores it could not require them to fill the coolers exclusively with Coca-Cola products - at least 20% of cooler space should be available to competing products

¹⁷The exact reason for this is unclear, though conversations with industry experts suggest that this was originally designed to incentivize exhibitors to take on new, untested movies

in significance, accounting for 64% of screens in 2011 (MPAA, 2011). US distributors aim to stop releasing movies in 35mm technology by the end of 2013 (Geuss, 2012).

2.4 Exhibitor decisions without vertical restraints: the Polish example

It is entirely possible that vertical restraints described above are not binding, and if they were removed exhibitors would make the same movie rental and scheduling decisions as before. Given that these restraints apply to all exhibitors in the US, to get a sense for whether this might be the case it is helpful to look at a country where exhibitors are not constrained in their movie rental and scheduling decisions. Table 1 provides summary statistics for a representative 10-screen movie theater in Lodz, Poland, where distributors do not impose vertical restraints on exhibitors, alongside summary statistics for the US movie theaters from the data set used in this paper:¹⁸

Table 1: Scheduling summary statistics, Poland vs USA

	Poland	USA
Single-movie screen-days	18.1%	85.9%
Ratio of # screenings in first vs. last day	3.94	1.03
Movies which screen for more than three weeks	44.1%	23.0%

Time period: (1: Poland) Jul 2011 - Aug 2012 (2: USA) Jan 2010 - Jan 2011

Sources: (1: Poland) one movie theater, data collected weekly over sample period from silverscreen.com.pl (2: USA) data provided by five movie theaters in the sample

It is clear from Table 1 that a big difference exists between movie schedules compiled with and without contractual restraints. While under the restraints over 85% of time a screen shows only one movie each day,¹⁹ this happens less than 20% of the time when restraints are not present, suggesting movie theaters find it beneficial to adjust their repertoire throughout the day. Looking at the ratio between number of screenings on the first vs. last day, it is also clear exhibitors in Poland adjust the schedule in time. This allows them to schedule more screenings (usually across many screens) to take advantage of the high interest early on, with sparser later screenings designed to capture audiences who were not able to see the movie earlier. This flexibility allows them to extend movie runs in time, albeit with fewer screenings per week - the number of movies which screen for more than three weeks is almost twice as high in Poland than it is in the US.

¹⁸Ideally these would be nationwide statistics, however obtaining these would have been too expensive for the purpose of this paper

¹⁹While one could expect this figure to be 100% under the contractual restraints, the divergence can be explained by additional screenings outside normal operating hours or one-off agreements with distributors to phase out heavily underperforming movies

3 Data

3.1 Data Sources and Descriptive Statistics

This paper uses two primary types of data: attendance data, which reflects the realized demand in the market, and contract data, which contains revenue split information agreed upon by distributors and exhibitors. These are supplemented with additional data sources described below.

Attendance Data Attendance data was collected directly from exhibitors, who use it for performance measurement and accounting purposes. Five movie theaters agreed to take part in this study, providing access to their attendance data at the lowest level of aggregation. Attendance figures are broken down by movie/screening time combination in the data set; in addition, data for four out of five movie theaters provides the breakdown of the attendance by type of ticket sold (*child*, *adult* and *senior*). The exact period for which the data was made available varies by movie theater - see Table 2 for data set summary statistics.

Table 2: Data set summary statistics, by movie theater

	mean	min	max
Avg. annual attendance	84,197	33,015	174,131
Avg. annual # movies	78.5	63.6	86.4
Avg. # weeks on screens (by movie)	2.9	2	3.8
Data period (mths)	35.8	14	64
Screens per theater	4.3	3	6
Market size (local population)	13,134.3	4,380	44,737
Distance to closest competitor (miles)	48.4	23	116


Sources: attendance and schedule data from movie theaters, Wikipedia, Google Maps

Each movie theater in the sample can be thought of as having a local monopoly on movie exhibition, allowing the demand and supply models to abstract from competitive considerations. The distance to the closest exhibitor is 23 miles or more for all movie theaters in the sample - although exact information of how far moviegoers are willing to travel to go to the movies is not available, conversations with exhibitors suggest none of them view the movie theaters closest to their own as direct competition.²⁰ Market size is defined as the local population of the town in which the movie theater is based.

²⁰Any bias to demand estimates resulting from not accounting for competition from other movie theaters would be largest for most popular movies, as those are the ones for which most moviegoers would be willing to travel to another movie theater for. However, these are the movies which are most likely played at all movie theaters, thus reducing this source of potential bias

Contract Data Contract data for all titles screened was made available for two out of five movie theaters in the sample, henceforth referred to as MT1 and MT2. The data contains weekly revenue split information for each movie screened, capturing the original negotiated agreement as well as any changes resulting from renegotiation. Minimum run length restraints are not observed since the industry practice is to have an unwritten agreement between the exhibitor and distributor. An analysis of observed movie run lengths, performed in Section 3.2, suggest these restraints are often binding for movie theaters in the sample. Chart 1 illustrates the values revenue splits take in the data set, while Section 8.2 in the Appendix discusses in detail how they were discretized for the purpose of the model.

Chart 1: Revenue split values



Source: data from movie theaters MT1 and MT2 in the sample, time period Jan 2009 - Jan 2011

Projection technology Of the two movie theaters for which contract data (necessary to conduct a full counterfactual analysis) was available, MT1 uses digital projection technology on its 6 screens, while MT2 uses 35mm technology on its 3 screens.

Theatrical market statistics Nationwide viewership trends, such as how often people go to the movies over the course of a year or the age composition of the frequent moviegoer group, are sourced from the annual *Theatrical Market Statistics* report release by the Motion Pictures Association of America (henceforth MPAA). These trends are used to aid estimation of the demand model.

Additional data sources In addition to the primary data sets movie characteristics were collected from The Internet Movie Database (IMDB),²¹ an online database of information related to

²¹www.IMDB.com

show business. These characteristics include the movie’s distributor, genres,²² MPAA rating, budget and Academy Awards nominations and wins. Movies’ critical ratings were taken from two sources: consumer ratings from IMDB and professional critic ratings from Metacritic,²³ an aggregator service for reviews in show business. Exhibitors provided information on ticket prices. The distribution of consumer demographics was obtained by sampling individuals from the US Census of Population.

Information not observed The minimum run-length restraint was not observed as this information is not usually written down but is rather an unwritten understanding between the exhibitor and distributor. Although this does not impact estimation of the demand model, it will require simplifying assumptions when estimating the bargaining model.

3.2 Additional Stylized Facts About the Industry

Ticket prices Movie ticket prices do not differ between competing movies, unlike for many vertically differentiated products where prices reflect quality e.g. cars, electronics, houses. Also, ticket prices do not differ over time for a given movie, not displaying the usual “skimming” behavior wherein high valuation consumers pay high prices early while low valuation consumers are targeted with price decreases later in a product’s life. Ticket prices change little over time for any given movie theater. In the sample price changes (if any) happen over New Year, with both rises (to keep up with inflation) and decrease (to stimulate demand) observed. Additionally, most exhibitors engage in modest price discrimination between age groups (child/senior discounts), time of day (matinee discounts) and 2D/3D screenings of the same movie. Ticket prices for the movie theaters in the sample are shown in Table 3.

Table 3: Ticket prices by movie theater

Movie Theater	Base	Discounts / Premia			Price Changes	
	Price	Matinee	Child/Senior	3D	Number	Magnitude
1	\$9.00	-\$2.00	-\$2.50	\$3.00	0	n/a
2	\$6.00	-\$1.00	-\$2.00	\$2.00	1	-\$0.50
3	\$7.75	-\$1.50	-\$2.25	n/a	3	\$0.75
4	\$7.75	-\$1.50	-\$2.25	n/a	0	n/a
5	\$8.00	-\$2.00	-\$2.50	n/a	1	\$0.50

Notes: prices shown at the end of the sample period, Jan 2011; price changes over the sample period

Source: information from exhibitors

Consumer heterogeneity Consumers differ significantly in how often they go to the movies, and what type of movies they watch, with consumer age being the primary driver of heterogeneity in

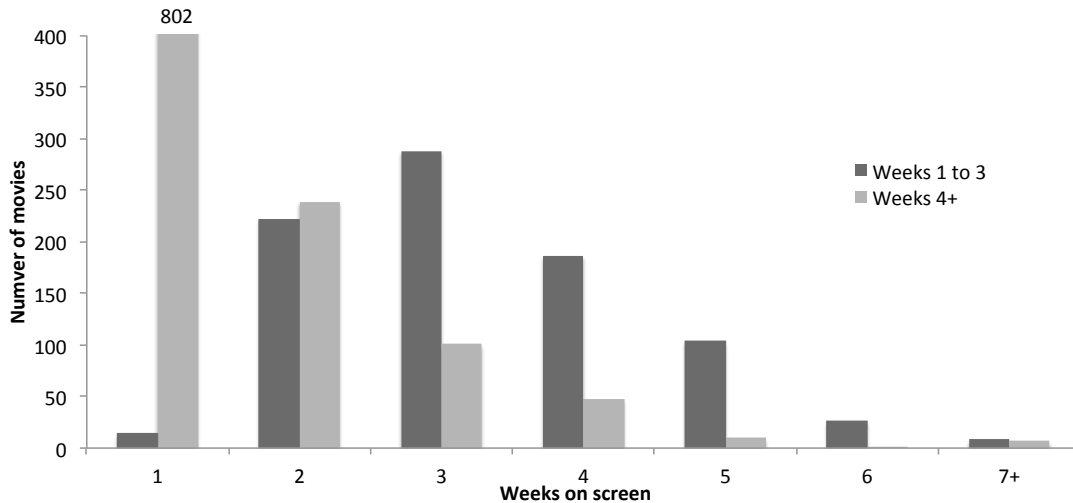
²²IMDB allows for users to add more than one genre tag to a title, leading to a multitude of genre types, thus the analysis focus on tags that describe more than 5% of the movies released in the period 2005-2010.

²³www.Metacritic.com

moviegoing patterns (MPAA, 2010). The demand model captures the impact of age on consumers' utility from going to the movie theater, different screening times and watching movies with certain genres.²⁴

Minimum run length Analyzing data in the sample by how long exhibitors keep movies on their screens illustrates the choices they face under the minimum run length restraint currently imposed on them by the distributors. If an exhibitor wants to screen a movie for one week only he has to wait until the distributor is willing to drop the minimum run length restriction later in the movie's release life - this is illustrated in Chart 2. Releasing the movie on the second run, however, means its appeal to audiences will be reduced - if an exhibitor wants to avoid this and release the movie on the break he has to comply with the minimum run length restraint.

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



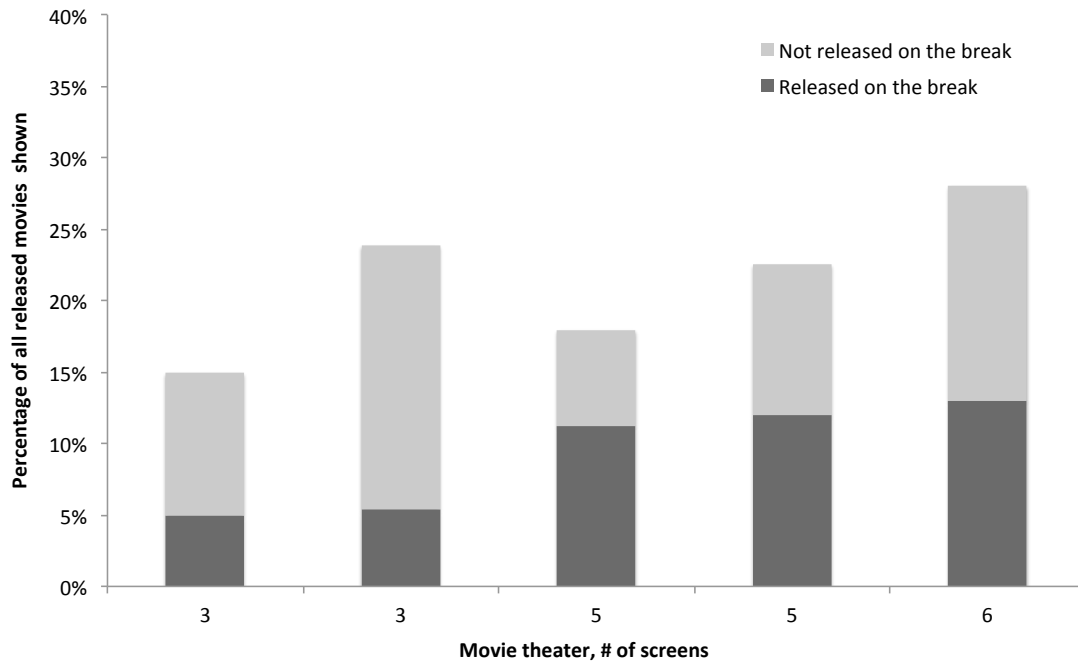
Source: schedule data from movie theaters

Number of screens constraint Under the prevailing contractual restraints cinemas cannot screen more than one movie on a screen each week, which means the number of movies they can screen each year is limited by the number of screens they have at their disposal. As illustrated in Chart 3, the more screens a movie theater has the more movies overall it plays, on average, throughout the year, suggesting rather than screen the same movie on multiple screens exhibitors chose to offer a larger variety of movies. The corollary is that to screen more movies throughout the year smaller movie theaters choose to release fewer movies on the break.²⁵

²⁴There is a tradeoff between completeness and parsimony in capturing the ways in which age impacts consumers' moviegoing decisions, and thus the model focuses on a select group of moviegoing trends most impacted by consumer age; these were selected based on correlation analysis of trends observed in the data set

²⁵Additional support to the notion that exhibitors want to screen a wide variety of movies comes from the prevalence of multiplexes. If offering a wide variety of movies was not important exhibitors could build movie theaters with fewer screens, each of which with higher capacity, to service the same number of people

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



Source: schedule data from movie theaters; National Association of Theater Owners

4 The Model

This section describes in detail the structural model which aims to explain consumers’ demand for movies, exhibitors’ scheduling choices and distributors’ movie rental pricing. The estimated model is used to determine the welfare impact of removing the contractual restraints imposed by distributors. A non-technical overview of the approach used precedes a full-detail description of the model - its aim is to provide the reader with enough background to be able to jump straight to the results and counterfactual simulations in Section 6.

Overview of Modeling Approach The structural model aims to capture three different levels of decision-making in the industry:

1. Consumers: which movies to see given what’s playing
2. Exhibitors: faced with uncertainty as to the “quality” of newly-released movies which movies to screen and when to schedule the screenings
3. Distributors: how to price movies and which movies to offer to movie theaters on the break

The primary driver of decisions at all three levels, as well as the main challenge in modeling and estimation, is movie quality. The true quality of a movie, λ_m , is not known before its release,

thus exhibitors need to make movie rental and scheduling decisions based on imperfect, exhibitor-specific expectations of movie quality, $\hat{\lambda}_{mc}$. The movie quality generating process is described by the equation:

$$\hat{\lambda}_{mc} = \underbrace{x_m^M \beta^M}_{\lambda_m} + \mu_m + \nu_{mc} \quad \mu_m \sim N(0, \sigma_\mu^2) \quad \nu_{mc} \sim N(0, \sigma_{\nu c}^2) \quad (1)$$

where $x_m^M \beta^M$ is the predicted movie quality which can be calculated *ex ante* on the basis of a movie's observable characteristics, μ_m reflects how much better/worse movie m actually is than the sum of its observable parts, and ν_{mc} captures how exhibitor c 's expectations differ from the true quality. The variables in x_m^M are: IMDB rating, dummies for MPAA ratings, the movie's budget and genre dummies. Counterfactual calculations require $\hat{\lambda}_{mc}$ and λ_m values for all movies, not only those which can be directly estimated for movies in the data set. For the rest of movies exhibitors considered these values need to be simulated based on (1) using estimated σ_μ^2 and $\sigma_{\nu c}^2$ values, although distributions of μ_m and ν_{mc} need to be truncated to reflect the fact that movies which were not chosen by exhibitors are likely to be worse than average. The challenge is that the set of $\hat{\lambda}_{mc}$ and λ_m values used in counterfactuals should be such that, if contractual restraints were *not* lifted, exhibitors would choose to screen the same movies and moviegoer attendance should closely correspond to observed attendance figures. For more details see implementation details in the Appendix.

Estimation takes place in two stages. First, a discrete-choice demand model based on Berry et al. (1995) is used to estimate parameters driving consumers' moviegoing decisions and λ_m values for movies in the data set. Second, the exhibitor and distributor models are used to back out exhibitors' *ex ante* expectations of movie quality and estimate parameters in the movie quality generating process (1). Algorithms are developed to perform counterfactual simulations, solving the exhibitor and distributor problems under relaxed contractual restraints.

Stage 1: Consumers face a discrete choice between different movie/screening time combinations offered in their local movie theater over the course of one week. The model assumes they can only watch one movie in a given week, and will see a given movie only once. A key driver of the market is that *ceteris paribus* consumer prefer to see movies closer to their nationwide release date. Tastes in movies and screening times vary within the population - in part they depend on observable characteristics such as age, while the rest of the variation is unobserved to the econometrician. Taking advantage of the panel nature of the data it is possible to estimate the true movie quality, seasonal- and long-term-trends as well as average utility moviegoers derive from different screenings *net* of ticket prices. Since ticket prices do not vary between movies of different quality they are assumed to be exogenous to consumer demand and thus, unlike in most demand estimation problems, there

is no endogeneity problem that would require the use of additional instruments. Attendance data is adjusted to best fit the model employed - see Appendix for more details.

Stage 2: The key complication faced by exhibitors is that they do not know the true movie quality λ_m when releasing movies on the break, and instead need to rely on idiosyncratic estimates, $\hat{\lambda}_{mc} = \lambda_m + \nu_{mc}$. The challenge to the econometrician in estimating $\hat{\lambda}_{mc}$ is to account for the fact that movies actually screened most likely received above-average ν_{mc} realizations, while movies not screened received below-average realizations. It is assumed that exhibitors learn the true movie quality one week after the movie's nationwide release, whether or not they screen the movie, which seems reasonable given the wealth of box office revenue information available to exhibitors today.

The second estimation stage can be split into two parts. In part 1 exhibitors' decision to screen a movie provides a floor on $\hat{\lambda}_{mc}$, while distributors' decision to charge a revenue split r_{mc} (but not higher) provides a ceiling; this is done *relative* to the expected value of the next-best alternative, $\hat{\lambda}_{cot}$, which is not identified at this stage. In part 2 a maximum likelihood estimator is constructed to calculate *absolute* values of $\hat{\lambda}_{mc}$, $\hat{\lambda}_{cot}$, as well as coefficients in the movie quality generating process: β^M , σ_μ and $\sigma_{\nu c}$.

Counterfactual simulations take the estimated consumer demand model and exhibitor movie quality expectations and ask the question: How much better/worse-off would market participants be if contractual restraints on exhibitors were relaxed? Algorithms are developed to calculate the exhibitor's schedule and distributor's revenue splits. The distributor can also choose not to release a given movie on the break if he thinks he can earn more money if he releases it in the second run. When calculating counterfactual results true movie quality is simulated for those movies that were not observed in the sample, while movie quality expectations are simulated for movies not shown by the given movie theater.

4.1 Demand: Moviegoing Decisions

Market Definition A market is defined to be a movie theater/week combination. In the geographic dimension thinking of a market as an individual movie theater comes from the focus on theaters which are local monopolies; the population of the town in which the movie theater is located defines the number of potential customers. The temporal bound of the market 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 in the temporal dimension. A time period, t , is thus one week, and a market is a cinema/week combination (c,t) .

Consumer Decision A discrete choice logit model is used to explain consumers' moviegoing decisions. Consumers choose one or none of the movie/screening time combinations available in a

given week. Consumers only ever see a particular movie once.²⁶ Their choice of a movie/screening time combination thus depends on utilities of all combinations available to them in a given week, excluding those with movies they had already seen.²⁷

Agent's Utility Function Let (m, s) index a movie m /screening time s combination offered by movie theater c in week t (e.g. Avatar screening on Friday, 9:10pm in movie theater A the week of Apr 20-26 2012). The indirect utility conditional on going to the movies for consumer i is the form:

$$u_{imsct} = u_{imt}^M + u_{imsct}^S + \Xi_{msct} + \epsilon_{imsct} \quad (2)$$

$$u_{imsct} = x_m^M \beta_i^M + \omega_{im} + I(w_{mt} = 0) \beta_1^W + w_{mt} \beta_2^W \quad (3)$$

$$+ p_{imsct}^A \beta^P + I_p^S(s) \beta_i^S + x_{msct}^{3D} \beta_c^{3D} + x_{sct}^O \beta^O \quad (4)$$

$$+ \mu_m + \mu_{sct} + \mu_w + \mu_y + \xi_{msct} \quad (5)$$

$$+ \epsilon_{imsct} \quad (6)$$

$$u_{imsct} = x_m^M \beta^M + x_m^M (\Pi D_i + \Sigma v_i) + \omega_{im} + I(w_{mt} = 0) \beta_1^W + w_{mt} \beta_2^W \quad (7)$$

$$+ p_i^A \beta^P + p_{sct}^A \beta^P + I_p^S(s) \beta^S + I_p^S(s) (\Pi D_i + \Sigma v_i) + x_{msct}^{3D} (p_c^{3D} \beta^P + \beta_c^{3D}) + x_{sct}^O \beta^O \quad (8)$$

$$+ \mu_m + \mu_{sct} + \mu_w + \mu_y + \xi_{msct} \quad (9)$$

$$+ \epsilon_{imsct} \quad (10)$$

$$\mathbf{X}_{msct} = \{x_m^M, I(w_{mt} = 0), w_{mt}, x_{msct}^{3D}, x_{sct}^O, u_{sct}, u_m, u_w, u_y\} \quad (11)$$

$$\mathbb{E}[\xi_{msct} | \mathbf{X}_{msct}, \{\iota_{it}\}_{\forall i}] = 0 \quad (12)$$

where u_{imt}^M and u_{imsct}^S capture the attractiveness of movie m and screening time s , respectively, in week t to agent i , while Ξ_{msct} and ϵ_{imsct} are unobserved portions of agent i 's utility. The unobserved term has a component which is common to all agents in the market, ξ_{msct} , and an idiosyncratic term, ϵ_{imsct} . Additive separability between u_{imt}^M and u_{imsct}^S allows consumer tastes for movies and screenings to not be correlated.

Movie m 's attractiveness to agent i is modeled as follows:

$$u_{imt}^M = x_m^M \beta_i^M + \omega_{im} + I(w_{mt} = 0) \beta_1^W + w_{mt} \beta_2^W \quad (13)$$

where x_m^M is a vector of observable movie characteristics, ω_{im} is a consumer-specific, mean-zero, time-invariant signal of whether movie m appeals to him more or less than another consumer with

²⁶While this may overestimate the quality of movies such as "Avatar" or "Titanic" which attracted lots of repeated viewers, such movies are the fringe in the data set and thus any impact this could have on estimation is small

²⁷An advantage of this approach over a nested-logit model is that it allows me 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

the same characteristics and w_{mt} the number of weeks since a movie's nationwide release (thus if a movie is released on the break $w_{mt} = 0$).²⁸ Together β_i^M and ω_{im} represent heterogeneity in consumers' movie tastes, and are crucial to modeling how consumers who choose to see a movie early on are likely to have higher appreciation of the movie than those who see it at a later time. Identification of $Var(\omega_{im})$ is made possible by the panel nature of the data set.

The attractiveness of screening time s is modeled as follows:

$$u_{imsct}^S = p_{imsct}^A \beta^P + I_p^S(s) \beta_i^S + x_{msct}^{3D} \beta_c^{3D} + x_{sct}^O \beta^O \quad (14)$$

where p_{imsct}^A is the price of admission for individual i to screening s ,²⁹ $I_p^S(s)$ are indicator variables whether screening time s is in one of the time periods p (e.g. weekday 5-8pm) and β_i^S captures the utility individual i derives from going to see a movie in period p . By making utility from screening times depend on consumers' demographic characteristics, the model is able to capture for example that children are able to see movies earlier in the day while working people can only see movies in the evenings or on weekends. Additionally, x_{msct}^{3D} captures whether a specific screening is in 3D or not, while x_{sct}^O captures factors impacting the opportunity cost of going to see a movie at time s , cinema c and in week t for example indicator variables for cultural events and holidays. In effect $x_{sct}^O \beta^O$ is the observable opportunity cost of going to the movies at this specific time and location that is common to all consumers, while the unobservable component is captured by ξ_{msct} .

Outside option If the consumer chooses not to see a movie in market (c,t) she gets

$$u_{i0ct} = \epsilon_{i0ct} \quad (15)$$

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

Unobserved product characteristic Ξ_{msct} can be broken down as

$$\Xi_{msct} = \mu_{sct} + \mu_m + \mu_w + \mu_y + \xi_{msct} \quad (16)$$

where all μ parameters can be captured by fixed effects in the estimation stage. The term μ_{sct} captures the attractiveness of screening time s that is specific to movie theater c ; it also differs over

²⁸This effect is part of what Einav (2007) captures in the "decay factor" λ - the fact that a concentrated advertising campaign and excitement of seeing the movie immediately when it is released gives people higher utility from seeing it early on. However, unlike in that paper, $\ln(w_{mt})$ 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; Ideally, the speed and shape of the decay function would be movie-specific, however given the limited number of weeks each movie spends on the screens it would be impossible to estimate

²⁹So as not to complicate notation any further, the subscript for whether the projection is in 2D or 3D, d , is suppressed

time t so as to reflect ticket price changes at c . In estimation, the screening time/movie theater/time period fixed effect which captures μ_{sct} also subsumes the price coefficient $p_{imsct}^A \beta^P$ from (14).³⁰ The term μ_m captures the unobserved quality of movie m ; μ_w represents the attractiveness of going to the movies in week $w(t)$, out of 52 weeks total in a year, and thus captures seasonality in the industry, as described by Einav (2007); μ_y captures the annual time trend in attractiveness of going to the movies (relative to, for example, seeing it on DVD/BlueRay as the release window shrinks). Finally, the leftover econometric error term ξ_{msct} captures remaining unobserved preferences e.g. a local event that draws people away from movie theater c in period st or the fact movie m does not resonate with moviegoers around cinema c in a way that is not captured by the model.

Movie quality uncertainty Substituting (13), (14), (19) and (16) into (2) yields:

$$u_{imsct} = \delta_{msct} + \omega_{im} + (x_m^M + p_{imsct}^A + I_p^S(s) + x_{sct}^O)(\Pi D_i + \Sigma v_i) + \epsilon_{imsct} \quad (17)$$

where δ_{msct} is the mean utility:

$$\delta_{msct} = \underbrace{x_m^M \beta^M + \mu_m}_{\lambda_m} + I(w_{mt} = 0) \beta_1^W + w_{mt} \beta_2^W + p_{sct}^A \beta^P + I_p^S(s) \beta^S + x_{msct}^{3D} \beta_c^{3D} + x_{sct}^O \beta^O + \mu_{sct} + \mu_w + \mu_y + \xi_{msct} \quad (18)$$

The term λ_m captures the quality of movie m that is not known *ex ante* to the exhibitor, which will feed into the learning model described in Section 4.2.

Heterogeneity in consumer tastes β_i^M and β_i^S are consumer-specific coefficients which reflect heterogeneity in movie tastes in the population. They are modeled as multivariate normal with the mean dependent on observable demographic variables and parameters to be estimated, and a variance-covariance matrix to be estimated:

$$\begin{pmatrix} \beta_i^M \\ \beta_i^S \end{pmatrix} = \begin{pmatrix} \beta^M \\ \beta^S \end{pmatrix} + \Pi D_i + \Sigma v_i, \quad v_i \sim N(0, I) \quad (19)$$

where Π captures how consumer demographics D_i impact their preferences and Σ captures idiosyncratic parameter variance between individuals.³¹

The parameter ω_{im} aims to capture heterogeneity in movie tastes which cannot be explained by observable movie characteristics. It is best thought of as an idiosyncratic variance term on the true movie quality, λ_m .

³⁰Although μ_{sct} does not differ by i , the values of i -specific discounts (presented in Table 3) are very close between movie theaters and thus will be captured by coefficients on i -specific constants in β_i

³¹The empirical implementation uses a diagonal Σ , although correlations between coefficients can easily be added

Market shares Given the choice model described above, the set of consumers who choose combination (m, s) in market (c, t) is defined as

$$A_{msct}(x_{msct}, \xi_{msct}; \theta) = \{(D_i, \epsilon_{imsct}) | u_{imsct} > u_{im's'ct} \forall m', s' s.t. (m', s') \neq (m, s) \text{ and } m \notin \iota_{it}\} \quad (20)$$

where x_{msct} and ξ_{msct} are the observable and unobservable characteristics, respectively, of combination (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_{it} the model ensures consumers see a given movie only once. 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 combination (m, s) in market (c, t) is the sum of all consumers who choose that option

$$s_{msct}(x_{msct}, \xi_{msct}; \theta) = \int_{A_{msct}} dP^*(D, \epsilon) = \int_{A_{msct}} dP^*(\epsilon) dP^*(D) \quad (21)$$

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

4.2 Supply: Exhibitor Scheduling Problem

The exhibitor problem can be split into two parts: the between-period decision, wherein the exhibitor decides on a set of movies m_{ct} to screen in each period, and the within-period decision, wherein the exhibitor decides on a schedule that will maximize profits from screening movies m_{ct} .

Between-period decision The exhibitor aims to maximize the present discounted value of expected profits $\mathbb{E}[\sum_{\tau=t}^{\infty} \beta^{\tau-t} \pi_{\tau}^* | s_{ct}, r_{ct}^*, M_{ct}^*]$, leading to the Bellman equation μ

$$V(s_{ct}, r_{ct}^*, M_{ct}^*) = \max_{m_{ct} \subset M_{ct}^*} \pi^*(s_{ct}, r_{ct}^*, m_{ct}) + \beta E[V(s_{ct+1}, r_{ct+1}^*, M_{ct+1}^*) | s_{ct}, r_{ct}^*, m_{ct}] \quad (22)$$

where s_{ct} are the state variables, M_{ct}^* are the movies offered by distributors to movie theater c in period t , $r_{ct}^* = \{r_{mct}^*\}_{\forall m}$ are revenue splits set by the distributors, π^* is the per-period expected profit function and β is the discount factor.

Within-period decision Given a set of movies m_{ct} to screen in period t the exhibitor decides when to screen them. His per-period profit given schedule z_{ct} is:

$$\pi(s_{ct}, r_{ct}^*, z_{ct}) = \sum I_{msct}(z_{ct}) \int_{A_{msct}} (p_{imsct}^A (1 - r_{mct}^*) + \pi_c^C) dP^*(\epsilon) dP^*(D) - C_c^C \quad (23)$$

where $I_{msct}(z_{ct})$ is an indicator function whether movie/screening combination (m, s) is part of schedule z_{ct} , π_c^C is the average concession profits per moviegoer in movie theater c ,³² while C_c^C is the fixed cost of keeping movie theater c open for one period.³³

Maximizing the profit function produces the optimal schedule for the set of movies m_{ct} :

$$z^*(s_{ct}, r_{ct}^*, m_{ct}) = \arg \max_{z_{ct} \in Z(m_{ct})} \pi(s_{ct}, r_{ct}^*, z_{ct}) \quad (24)$$

where $Z(m_{ct})$ is the set of all possible schedules for exhibitor c in period t given m_{ct} .

Inserting (24) into (23) produces the per-period expected profit function used in (22):

$$\pi^*(s_{ct}, r_{ct}^*, m_{ct}) = \pi(s_{ct}, r_{ct}^*, z^*(s_{ct}, r_{ct}^*, m_{ct})) \quad (25)$$

Contractual restraints The minimum run-length restraint enters the exhibitor scheduling problem through m_{ct} e.g. if exhibitor c commits in period $t-1$ to screening movie m' for 2 periods, then every combination m_{ct} in period t will contain m' .

The no screen-sharing restraint enters the exhibitor scheduling problem through the function $Z(\cdot)$ - if screen-sharing is allowed the movie theater has more flexibility with its schedule, and thus

$$Z(m_{ct})_{\text{no screen-sharing}} \subset Z(m_{ct})_{\text{screen-sharing}}$$

Movie quality expectations and learning Before a movie is released exhibitors receive a one-off signal of its quality:

$$\hat{\lambda}_{mc} = \lambda_m + \nu_{mc} \quad \nu_{mc} \sim N(0, \sigma_{\nu c}^2) \quad (26)$$

$$\lambda_m = x_m^M \beta^M + \mu_m \quad \mu_m \sim N(0, \sigma_{\mu}^2) \quad (27)$$

which depends on the true movie quality, λ_m , and an idiosyncratic term ν_{mc} . This belief is specific to each exhibitor and affects his decision whether to release a movie on the break.

The exhibitor learns a movie's true quality λ_m one week after the movie's nationwide release. This assumption reflects the fact that the exhibitor can learn a movie's true quality even if he does not screen it, by analyzing widely-available box office revenue information which captures how well the movie did in its first week of release. Moreover, by assuming learning happens independently of exhibitor's actions the estimation procedure is reduced to a static problem.

State variables State variables in the exhibitor problem include, for week t :

³²Concession sales are assumed not to vary in a meaningful way between movies and screenings and over the timeline examined, which is supported by Gil & Hartmann (2007) who find that concession sales are roughly proportional to total attendance in Spanish movie theaters

³³These costs are independent of the exhibitor's choice of movies played and the schedule, which is a reasonable assumption as long as the opening hours/days and number of screens operating are kept constant

- $\{\iota_{it}\}_{\forall i}$: the set of movies seen by consumer i up to period t ;
- $\hat{\lambda}_{mc} \forall m \in M_{ct}^*$: exhibitor c 's belief about unknown movie quality for movies he's considering playing in week t ;

4.3 Supply: Distributor Decision

Each period the distributor faces two sets of decisions: which movies to offer to an exhibitor and at what revenue splits, r_{ct}^* .

Revenue split determination Conversations with exhibitors suggest revenue splits are determined in a bargaining procedure where distributors make take-it-or-leave-it offers to exhibitors. Distributors negotiate terms for each of their movies separately,³⁴ while exhibitors negotiate separately for each movie theater they own. Negotiations happen before period t for all (c,m) combinations where $m \in M_{-c}$, the set of movies not already screened by c .

Let $\mathbf{r} = \{\mathbf{r}_{mc}\}_{\forall m,c}$ be the set of revenue splits agreed upon for all movie theaters and movies, where the revenue split time profile $\mathbf{r}_{mc} = \{r_{mct}\}_{\forall t}$ maximizes the profits movie m 's distributor derives from screening it at movie theater c :

$$\Pi_{mc}(\mathbf{r}_{mc}, \mathbf{r}_{-mc}, s_{ct}) = \mathbb{E} \left[\sum_{t=1}^{\infty} \beta^t \sum I_{msct}(z^*(s_{ct}, \mathbf{r}_{mc}, m_{ct})) \int_{A_{msct}} p_{im.sct}^A r_{mct} dP^*(\epsilon) dP^*(D) \right] - C_{mct}^M \quad (28)$$

where $z^*(\cdot)$ is the optimal schedule movie theater c arrived under its own movie quality expectations, while A_{msct} is calculated using the true movie quality for m , known to the distributor, and c 's movie quality expectations for all other movies, which are shared by the distributor.³⁵ C_{mct}^M , in turn, is the per-movie cost a distributor needs to incur to provide a copy of movie m to movie theater c in period t ; for 35mm projection technology the value depends on whether movie m is being released on the break or on the second run, while for digital movie theaters it is constant. If an agreement is not reached between movie theater c and movie m the set of revenue splits for all other movie theater/movie pairs, \mathbf{r}_{-mc} , is not renegotiated. The resulting bargaining equilibrium is:

$$\mathbf{r}_{mc}^* = \arg \max_{\mathbf{r}_{mc}} \Pi_{mc}(\mathbf{r}_{mc}, \mathbf{r}_{-mc}, s_{ct}) \quad (29)$$

Offering movies on the break (determining M_{ct}^*) For MT1, which uses digital projection technology, distributors have nothing to gain from delaying the release of movies until the second run

³⁴An extension to multi-product distributors is possible but proved intractable computationally

³⁵The assumption that the distributor knows the exhibitor's movie quality expectations is reasonable given the close relationship between movie theaters and booking agents who negotiate on the distributor's behalf

since C_{mct}^M is constant over time. It thus follows that in the second state estimation and all movies should be available to this exhibitor on their nationwide release date. Following conversations with the exhibitor C_{MT1}^M is set at \$25.

For MT2, which uses analog projection technology, distributors incur a non-trivial C_{mct}^M cost to screen a movie. The distributor’s decision on whether to offer movie m on the break to MT2 depends on what kind of release schedule the distributor is pursuing nationally for the movie, thus it is not possible to model this structurally using a non-representative sample of movie theaters. In estimation, a conservative assumption is made that apart from movies actually released on the break by MT2 all other movies were only available on the second run. Conversations with the exhibitor suggest this is a reasonable assumption, as he would screen most of the movies available to him on the break. Thus estimated model parameters are then used in the counterfactual simulations to compare how schedules, attendance and welfare would differ for MT2 if it was already a digital movie theater, *with* and *without* vertical restraints in place. In these simulations all movies are available to the exhibitor on the break, and C_{MT2}^M is set at \$25, just as for MT1.

5 Estimation, Identification and Counterfactual Calculation

Estimation proceeds in two stages: Section 5.1 describes how demand model parameters are estimated using the BLP approach (Berry et al., 1995) augmented with micro-moments à la Petrin (2002), while Section 5.2 describes how observed schedules and revenue splits are used to estimate parameters driving the movie quality generation process. Section 5.3 describes the algorithm used to calculate alternative exhibitor and distributor decisions in the counterfactuals. A non-technical overview of the main challenge of the estimation process precedes a detailed description of the model.

The challenge of estimating movie quality The main challenge of the estimation procedure is to back out true and expected movie quality values, as well as construct a way to simulate them for movies not observed in the data set. The following equation describes the relationship between predicted movie quality ($x_m^M \beta^M$) true movie quality (λ_m) and expected movie quality ($\hat{\lambda}_{mc}$):

$$\hat{\lambda}_{mc} = \underbrace{x_m^M \beta^M}_{\lambda_m} + \mu_m + \nu_{mc} \quad \mu_m \sim N(0, \sigma_\mu^2) \quad \nu_{mc} \sim N(0, \sigma_{\nu c}^2) \quad (30)$$

Counterfactual calculations require $\hat{\lambda}_{mc}$ and λ_m values for all movies available to the movie theater.

Demand model estimation, described in Section 5.1, provides λ_m only for movies in the data set. Since these movies were selected by the exhibitor on the basis of their $\hat{\lambda}_{mc}$ or λ_m values they most likely exhibit, on average, positive μ_m and ν_m values. It would thus not be correct to simply regress

λ_m values on x_m^M as the resulting estimate of β^M would be biased due to selection bias. Instead, a model is necessary which will fully capture the movie quality generation and the exhibitor movie selection processes.

The movie quality estimation procedure, described in Section 5.2, aims to do exactly that. It proceeds in two stages. First, using observed schedules and movie rental prices it backs out exhibitors' expectations of movie quality *relative* to the best alternative available each period. These calculations are carried out for a range of possible best alternative quality values, since the true value is not observed by the econometrician. Second, a maximum likelihood estimator is constructed to capture movie quality generation and the exhibitor movie selection processes. This procedure identifies the parameters β^M , σ_μ^2 and σ_{vc}^2 , the *absolute* $\hat{\lambda}_{mc}$ values for movies released on the break and calculates *absolute* upper bounds on $\hat{\lambda}_{mc}$ and λ_m values for movies not released by the exhibitor.

To calculate counterfactual scenarios $\hat{\lambda}_{mc}$ and λ_m values are simulated based on (30) and bounds estimated in the movie quality estimation procedure. The final set of estimated λ_m and $\hat{\lambda}_{mc}$ values is such that if one were to calculate counterfactuals under the current set of restrictions it result in the same schedules as those observed.³⁶

5.1 Demand Model Estimation

Primary moments The estimation strategy follows the standard GMM approach established by Berry et al. (1995). Key to this method is the inversion which, for a given set of non-linear taste parameters $\theta = (\Pi, \Sigma)$, makes it possible to back out the mean utility vector, δ_{msct} , from observed market shares by making the assumption that ϵ_{imsct} is distributed i.i.d Type I extreme value (Berry, 1994) and simulating individual-specific parameters such as ω_{im} and v_i . The inversion produces the residual:

$$\xi_{msct}(\theta) = \delta_{msct} - \lambda_m - I(w_{mt} = 0)\beta_1^W - w_{mt}\beta_2^W - p_{sct}^A\beta^P - I_p^S(s)\beta^S - x_{msct}^{3D}\beta_c^{3D} - x_{sct}^O\beta^O - \mu_{sct} - \mu_w - \mu_y$$

which is then used to construct the primary set of moments:

$$E[G_1(\theta)] = E[Z_{msct} \cdot \xi_{msct}(\theta)] \quad (31)$$

where Z_{msct} is a vector of instruments that are orthogonal to ξ_{msct} .

Instruments and price endogeneity In most demand models prices vary substantially across products, creating the problem of price endogeneity and preventing them from being included in Z_{msct} . In this application, however, the little price variation there exists can be decomposed into

³⁶Note, this calculation is not actually performed, as it requires knowledge of the minimum run length restraints for each individual movie

four categories: child/senior discounts, 3D premiums, matinee discounts and changes over time (see Table 3):

$$p_{imsct}^A = p_i^A + p_c^{3D} x_{msct}^{3D} + p_{sct}^A \quad (32)$$

where p_{ic}^A is the child/senior discount, p_c^{3D} is the premium for a 3D screening and p_{sct}^A encompasses matinee discounts and price changes over time; all three are exhibitor-specific. In estimation, all three decomposed price variables from (32) are captured entirely by other parameters in the model, thus ensuring they do not enter $\xi_{msct}(\theta)$. This removes the price endogeneity that models using the BLP framework usually have to take into account. As a result, all independent observable variables in (31) and fixed-effects are exogenous and thus are valid instruments in constructing the primary moments.³⁷

Conditionality of the primary set of moments Unlike in the standard BLP setup, the expectation expressed in (31) is not unconditional - rather, it is conditional on the selection resulting from allowing consumers to see a given movie only once. In order to make sure the expectation holds, it is important to explicitly model this selection mechanism. The model does this by explicitly keeping track of movies seen by each one of the simulated individuals, $\{\iota_{mit}\}_{\forall m,i,t}$. When deciding which movies to see in period t , individual i only considers those for which $\iota_{mit} = 0$. At the end of each period the moviegoing of each individual i is simulated, setting $\iota_{mit'} = 1 \ \forall t' \geq t$ if he decides to see movie m in period t .

Micro Moments The estimation procedure is augmented by five sets of additional micro-moments, following Petrin (2002). The first three sets are derived from information in the MPAA 2010 Theatrical Market Statistics,³⁸ an annual publication put out by the MPAA, 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 moments³⁹

³⁷The model does not explicitly take into account capacity constraints, however their impact on demand estimates is likely to be negligible, as in the sample less than 1% of screenings were 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 the sample less than 3% of screenings sell more than 75% of capacity any potential bias from this source is likely to be negligible

³⁸Since the report in its current form is only available for the years 2009 and 2010 the trends, which do not differ significantly between the two years, is assumed to apply to the whole data period 2007-2011

³⁹This moment is only matched for 2 movie theaters in the sample for which average moviegoing frequency is close to that of the nationwide moviegoing frequency of 4.1; lack of less aggregated moment information prevents the moment from being applied to the remaining movie theaters

2. Moviegoing frequency by age group (2-11, 12-17, 18-24, 25-59, 60+), relative to average moviegoing average frequency; this yields 5 moments
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 not available for all movie theaters in the data set micro moments are used to capture this additional information; this yields 3 moments.

In total there are 22 micro moments: $E[G_2(\theta)]$

The Objective Function The two sets of moments that enter the GMM objective function are $G_1(\theta)$, the standard BLP moments, and $G_2(\theta)$, the micro moments. The population moment conditions are assumed to uniquely equal zero at the true θ_0 :

$$E[G(\theta_0)] = E \begin{bmatrix} G_1(\theta_0) \\ G_2(\theta_0) \end{bmatrix} = 0 \quad (33)$$

The GMM estimator then takes the form

$$\hat{\theta} = \arg \min_{\theta} G(\theta)'W^{-1}G(\theta) \quad (34)$$

where W is a weighting matrix set to be $Z'Z$. In order to estimate standard errors the approach developed by Hansen (1982) is followed, which allows both sampling error and simulation error to be taken into account. Standard errors are clustered - see Berry et al. (2004) for further details.

Identification The estimation procedure exploits the panel nature of the dataset and employs numerous fixed effects. They include:

1. *screening time / movie theater / time interval* - fixed effects capture the observed utility from screening time $I_s^S(p)\beta^S$, the unobserved utility component μ_{sct} , as well as the price component $p_{sct}^A\beta^P$ (*time interval* is defined such that over its duration ticket prices remain constant at a given movie theater c). The primary source of identification is the time-variation in sales as movies screened change but the screening times remain constant for a given movie theater - this identifies the fixed effect for the time interval over which ticket prices at the movie theater remain constant. If ticket prices change over the sample period the panel nature of the dataset allows for the identification of an additional set of fixed effects for this movie theater

2. *movie* - fixed effects capture the true movie quality, λ_m . Identification comes from two sources: time-variation in sales at a given movie theater as movies change, and from variation in sales across movie theaters whose choice of movies screened in a given time period differs i.e. if movie m in period t was screened in one movie theater along with movies M_1 and in another movie theater along with movies M_2 such that $M_1 \neq M_2$
3. *week* - fixed effects capture the seasonal component of the unexpected utility, μ_w . Identification comes from time-variation in sales throughout the year.
4. *year* - fixed effects capture the long-term trend component of the unexpected utility, μ_y . Identification comes from time-variation in sales across the years.

The coefficient on ticket price, β^P , is not separately identified. There are three places where it enters the estimation procedure, as per (32). The first, $p_{ic}^A \beta^P$, is captured along with the component of ΠD_i associated with the constant, which combined represent the additional utility children/seniors get from going to the movies *net* of the admissions price.⁴⁰ p_c^{3D} is captured along with β_c^{3D} by the dummy variable on whether a given screening is in 3D, x_{msct}^{3D} - this can be thought of as the utility consumers get from a given movie theater's 3D screening *net* of prices charged by this movie theater. Finally, as described above, $p_{sct}^A \beta^P$ is captured by the *screening time / movie theater / time interval* - fixed effects.

Identification for β^W comes from variation in the release data for a given movie e.g. if movie m was released on the break in one movie theater but in the second run in another movie theater. If the sample included a larger selection of movie theaters it should be possible to identify a β_m^W coefficient separately for each movie, however the limited number of movie theaters in the sample prevents this. β^O is identified using time-variation in sales between weeks as the characteristics within x^O change.

In the standard BLP setup observed heterogeneity in consumer tastes is identified using variation in consumer demographics between markets. Such variation is limited, however, for the markets in the sample. Instead, the primary source of identification are the micro-moments. Micro-moments 2, 3 and 5 help identify age-specific utility from going to the movies - the components of Π in (19) that correspond to the constant. Micro-moment 5, additionally, helps identify the remaining coefficients in Π : the utility derived by different age groups from movie characteristics such as genre or MPAA rating as well as from different screening times.

⁴⁰In implementation, due to the closeness in child/senior discounts across movie theaters and the added computational complexity of estimating four additional non-linear parameters, only one parameter is estimated to capture this effect

Identification for the variance in consumer preferences comes from many sources. The first source of identification is the substitution patterns between products as these change across time periods. This helps identify the variance of ω_{im} but not the component of Σ that corresponds to screening times, as these do not vary across time periods. The second source of identification is micro-moment 1 which captures the heterogeneity in moviegoing frequency over the course of a year and thus helps identify the component of Σ corresponding to the constant. Finally, as described in Lee (2012) there is an additional source of identification that comes from the panel nature of the dataset and exploits the *self-selection* among consumers. For example, consider a world where the only potential source of heterogeneity between moviegoers is ω_{im} , and a situation where in week 2 of movie m_1 's release the exhibitor releases another movie, m_2 , with the same *mean* appeal to moviegoers accounting for attractiveness decay (i.e. $\lambda_{m_1} + \beta^W = \lambda_{m_2}$). If there is no heterogeneity in consumers' taste in movies ($Var(\omega_{im}) = 0$) than the primary driver on moviegoing decisions will be the idiosyncratic component ϵ_{imsct} , and the model will predict that, among the consumers who have not seen m_1 in week 1, an equal number will see m_1 as m_2 in week 2. If, instead, $Var(\omega_{im}) > 0$, than on average moviegoers who have not seen m_1 in week 1 will have a *lower-than-average* value of ω_{im_1} (i.e. $mean(\omega_{im_1} | \nu_{m_1 i1} = 0) < 0$), and thus in week 2 more of them will see m_2 than m_1 .

Compensating variation calculations Because the coefficient on ticket price, β^P , is not separately identified in the model, this prohibits measuring the change in consumer welfare using Compensating Variation (CV) as developed by Hicks (1939). Instead, a quasi-CV measure is developed, whose aim is to provide a conservative estimate of the monetary value a moviegoer places on the utility lost due to vertical restraints. The idea is to identify the single most valuable movie/screening combination for each moviegoer over the course of the sample period, and see how many times he would have to see it to make up for the lost utility. The quasi-CV is defined as the monetary value of those screenings.

5.2 Estimation of Movie Quality and Expectations Generation Process

This section describes the two-stage procedure used to identify $\hat{\lambda}_{mc}$ for movies released on the break, bounds on $\hat{\lambda}_{mc}$ and λ_m values for movies not screened, as well as parameters in the movie quality-generating process. The equation describing the process is reproduced below for the readers' convenience:

$$\hat{\lambda}_{mc} = \underbrace{x_m^M \beta^M}_{\lambda_m} + \mu_m + \nu_{mc} \quad \mu_m \sim N(0, \sigma_\mu^2) \quad \nu_{mc} \sim N(0, \sigma_{\nu c}^2)$$

The first stage sets up bounds on expected movie quality based on revealed schedule and pricing decisions in order to identify expected quality for movies screened *relative* to the expected quality of the best alternative. The second stage uses a maximum likelihood estimator to identify the *absolute*

value of the exhibitor's movie quality expectations, as well as parameters in (30). Due to lack of revenue split information for three of the exhibitors, the following estimation is only carried out for $c \in \{\text{MT1}, \text{MT2}\}$.

5.2.1 Stage 1: Movie quality bounds

Consider movies opening at movie theater c in period t : M_{ct}^+ . For those which are released on the break, $M_{ct}^+(w_{mt} = 0)$, the exhibitor does not know λ_m and bases his decision on his *ex ante* expectations $\hat{\lambda}_{mc}$. Knowing the exhibitor's expectations the distributor chooses a price r_{mc} . These two decisions help establish bounds on $\hat{\lambda}_{mc} \forall m \in M_{ct}^+$:

1. Lower bound $\underline{\hat{\lambda}}_{mc}$: movie m is at least as good as the best alternative from the set of movies available to the exhibitor in period t : M_{ct}^*
2. Upper bound $\overline{\hat{\lambda}}_{mc}$: movie m is not so good such that the distributor could increase r_{mc} and the exhibitor would still choose to screen the movie⁴¹

Combining the two bounds produces the equation:

$$\underline{\hat{\lambda}}_{mc}(\hat{\lambda}_{0ct}) \leq \hat{\lambda}_{mc} \leq \overline{\hat{\lambda}}_{mc}(\hat{\lambda}_{0ct}) \quad (35)$$

These bounds are expressed relative to the expected quality of the best alternative available to exhibitor c in period t : $\hat{\lambda}_{0ct}$. Which exact movie constitutes the best alternative is not known to the econometrician - it could either be a movie released this period ($w_{mt} = 0$) or a movie released in a previous period ($w_{mt} > 0$):

$$\hat{\lambda}_{0ct} = \max \left[\max_{m \in M_{ct}^-(w_{mt}=0)} (\hat{\lambda}_{mc}); \max_{m \in M_{ct}^-(w_{mt}>0)} (\lambda_m + w_{mt}\beta_2^W) \right] \quad (36)$$

where M_{ct}^- is the set of movies not released in period t .

Three cases are possible for each period t :

1. $M_{ct}^+ = \emptyset$: no movies opened in period t
2. $M_{ct}^+ = M_{ct}^+(w_{mt} = 0)$: all movies that opened in period t were also released in period t
3. $M_{ct}^+(w_{mt} > 0) \neq \emptyset$: at least one movie that opened in period t was released before period t

⁴¹ r_{mc} takes on discrete values, otherwise the upper bound would be equal to the lower bound. Also, this bound is nonexistent if r_{mc} is already at the highest value it can take (for a full discussion on discretization and bounding of r_{mc} values see Appendix

In case 1, not enough information is available regarding the expected quality of movies which were released this period. That none of them were picked up by the exhibitor may be because they are all of poor quality, but it may also be because no screens could be freed up for a new movie because of *minimum run length* restrictions on movies already being screened. In case 2 there is not enough information to identify the absolute $\hat{\lambda}_{0ct}$ value at this stage, and bounds on $\hat{\lambda}_{mt}$ are calculated relative to a range of possible $\hat{\lambda}_{0ct}$ values.. In case 3, the absolute value of $\hat{\lambda}_{0ct}$ can be set-identified relative to $\{\lambda_m\}_{m \in M_{ct}^+(w_{mt}>0)}$ ie. the fact that a movie of know quality λ_m was released provides a lower and upper bound on what the best alternative available to the exhibitor was; absolute bounds on $\hat{\lambda}_{mt}$ can also be calculated.

5.2.2 Stage 2: Maximum likelihood estimation

Maximum likelihood function Building on relative bounds for $\hat{\lambda}_{mc}$ and $\hat{\lambda}_{0ct}$ estimated in Stage 1, the second stage uses maximum likelihood estimation to identify the absolute values of these bounds, as well as estimate parameters β^M , σ_μ^2 and $\{\sigma_{\nu c}^2\}_{\forall c}$ driving the movie quality generation process (30).

The log-likelihood function is the following:

$$\begin{aligned} \ell(\sigma_\mu^2, \{\sigma_{\nu c}^2\}_{\forall c}, \beta^M, \{\hat{\lambda}_{0ct}, \lambda_{0t}\}_{\forall c, t \in T_c} | \{\lambda_m\}_{m \in M^+}, x^M, M^+, M^-) = \\ \sum_{\forall c} \sum_{t \in T_c} \ell_{tc}(\sigma_\mu^2, \sigma_{\nu c}^2, \beta^M, \hat{\lambda}_{0ct}, \lambda_{0t} | \{\lambda_m\}_{m \in M^+}, x^M, M_{ct}^+, M_{ct}^-) \end{aligned} \quad (37)$$

where T_c is the set of time periods observed for exhibitor c , and $\lambda_{0t} = \max_{m \in M^-(w_{mt}=0)}(\lambda_m)$ is the highest true movie quality value for movies not screened by any of the exhibitors.⁴²

The per-period log-likelihood function is:

$$\begin{aligned} \ell_{tc}(\cdot | \cdot) = \ln \left(P_{m \in M^-(w_{mt}=0)} \left(\hat{\lambda}_{0ct}, \lambda_{0t} | \beta^M, \sigma_{\nu c}^2, \sigma_\mu^2 \right) \right. \\ \prod_{m \in M_{ct}^+(w_{mt}=0)} \phi_\mu(\lambda_m - x_m^M \beta^M) \Phi_\nu \left(\hat{\lambda}_{mc}(\hat{\lambda}_{0ct}) \leq \hat{\lambda}_{mc} \leq \overline{\hat{\lambda}_{mc}}(\hat{\lambda}_{0ct}) \right) \\ \left. \prod_{m \in (M^+(w_{mt}=0) \setminus M_{ct}^+(w_{mt}=0))} \phi_\mu(\lambda_m - x_m^M \beta^M) \Phi_\nu(\hat{\lambda}_{0ct} - \lambda_m) \right) \end{aligned} \quad (38)$$

where $P_{m \in M^-(w_{mt}=0)} \left(\hat{\lambda}_{0ct}, \lambda_{0t} | \beta^M, \sigma_{\nu c}^2, \sigma_\mu^2 \right)$ is the joint probability that for $m \in M^-(w_{mt} = 0)$: $\lambda_m \leq \lambda_{0t} \forall m, \exists m$ s.t. $\lambda_m = \lambda_{0t}$, $\hat{\lambda}_{mc} \leq \hat{\lambda}_{0ct} \forall m, \exists m$ s.t. $\hat{\lambda}_{mc} = \hat{\lambda}_{0ct}$ if $\hat{\lambda}_{mc} \geq \max_{m \in M_{ct}^-(w_{mt}>0)} (\lambda_m + w_{mt} \beta_2^W)$.

$\ell_{tc}(\cdot | \cdot)$ accounts for all movies released nationally in period t ($w_{mt} = 0$) and

⁴²Note that since λ_m values are the same across exhibitors λ_{0t} values are also the same, unlike $\hat{\lambda}_{0ct}$ values that are exhibitor-specific

1. never screened by any exhibitor (their λ_m is unknown to the econometrician): $M^-(w_{mt} = 0)$
2. releases this period (λ_m known, $\hat{\lambda}_{mc}$ within bounds calculated in Stage 1): $M_{ct}^+(w_{mt} = 0)$
3. released in later periods (λ_m known, $\hat{\lambda}_{mc}$ below $\hat{\lambda}_{0ct}$): $M^+(w_{mt} = 0) \setminus M_{ct}^+(w_{mt} = 0)$

Computation The estimation procedure has two loops. The outside loop is a non-linear search over possible $\sigma_\mu^2, \{\sigma_{\nu c}^2\}_{\forall c}, \beta^M$ values. Given a multiple $\{\sigma_\mu^2, \{\sigma_{\nu c}^2\}_{\forall c}, \beta^M\}$ of candidate values the inner loop chooses a set of $\{\hat{\lambda}_{0ct}, \lambda_{0t}\}_{\forall c, t \in T_c}$ values which maximize the log-likelihood function as defined in (37); the set of possible $\{\hat{\lambda}_{0ct}, \lambda_{0t}\}_{\forall c, t \in T_c}$ values is restricted to align it with observed schedules.

$$\ell(\sigma_\mu^2, \{\sigma_{\nu c}^2\}_{\forall c}, \beta^M | \cdot) = \max_{\{\hat{\lambda}_{0ct}, \lambda_{0t}\}_{\forall c, t \in T_c} \in \boldsymbol{\lambda}_0} \ell(\sigma_\mu^2, \{\sigma_{\nu c}^2\}_{\forall c}, \beta^M, \{\hat{\lambda}_{0ct}, \lambda_{0t}\}_{\forall c, t \in T_c} | \cdot) \quad (39)$$

A detailed description of the estimation algorithm can be found in the Appendix.

Intuitively, identification of β^M and σ_μ^2 comes from analyzing λ_m values for movie that were screened by the exhibitors as well as λ_{0t} value for movies that were not screened. This estimation approach eliminates selection bias that would result if one tried estimating β^M using only λ_m for movies screened. Identification for $\sigma_{\nu c}^2$, in turn, comes from comparing bounds on $\hat{\lambda}_{mc}$ to λ_m for movies screened and $\hat{\lambda}_{0ct}$ to $x_m^M \beta^M$ for movies not screened (the latter also help to identify σ_μ^2). The estimation procedure is sufficiently quick that confidence intervals can be computed using a standard bootstrap, resampling at the time period level.

Implementation One of the restraints imposed on the set of possible $\{\hat{\lambda}_{0ct}, \lambda_{0t}\}_{\forall c, t \in T_c}$ values is that the best alternative cannot be better than the worst movie that was kept on by the exhibitor when he had a choice i.e. the minimum run-length restraint was no longer binding. Since the exact values of the minimum run-length restraint are not observed, an assumption needs to be made as to how long this period is. The estimation procedure is ran assuming a conservative value of three weeks for movies released within four weeks of the nationwide release date and one week thereafter, based on conversations with exhibitors. To test the sensitivity of this assumption, the estimation procedure is also ran with the former value set to two weeks, with again four weeks for the latter value.

Up to this point the exact metric exhibitors use when determining whether one movie is “better” than another has purposefully been left unspecified. Implementing a fully dynamic estimation process with forward-looking exhibitors, as described in Section 4.2, has proven computationally intractable.⁴³ To make the estimation procedure feasible a simplifying assumption is made that

⁴³Unlike most settings in which the estimation procedure is fully dynamic, exhibitors’ choice set changes over time which means the value-function is non-stationary

when making scheduling decisions exhibitors and distributors only consider profits in the current period, rather than over the whole period each movie is expected to be screened. This assumption is supported by conversations with exhibitors, who say that in the face of uncertainty about movies’ true quality they focus on first-week profits when making scheduling decisions, and only consider binding restraints once true movie quality is revealed in later periods.

Nonetheless, assuming exhibitor and distributor myopia can lead to estimation bias unless two conditions are satisfied. First, exhibitors need to *expect* that movies which are more profitable in the first week will remain more profitable than alternatives in the following weeks. If this is not true, it is possible that even though one movie is marginally more attractive in week 1 it loses greatly in the following weeks such that overall its box office revenues are lower than for the alternative. Second, if exhibitors’ ability to release blockbusters on the break is inhibited by minimum run-length restraints on other movies being screened, it is possible they strategically change their scheduling to accommodate the releases of such blockbusters. If this is true, the assumption of myopia is not justified and can lead to bias. These concerns are addressed in turn below:

The demand model, as described in Section 4.1, assumes that movies’ attractiveness decays in a linear fashion over time (component $w_{mt}\beta_2^W$ in (13)), but also that there is one-off boost to a movie’s attractiveness if it’s released the week of the nationwide release (component $I(w_{mt} = 0)\beta_1^W$). This means that assessing movies’ potential on the basis of one-week profits will unjustly favor movies released in a given week. However, the relatively small value of $\hat{\beta}_1^W$ compared to variation in movie quality (see Tables 5 and 6) suggests any potential bias from this source is likely to be negligible.

The second concern, that exhibitors change strategically don’t take on new movies with minimum run-length restraints right before the release of a “big” movie, can be assessed empirically. Table 4 reports the average proportion, across all five exhibitors, of new movies to all movies being screened in a given week. Strategic consideration for three factors is analyzed: (1) release of blockbuster movies⁴⁴ (2) high-attendance weeks⁴⁵ (3) holidays.

Overall, the results do not support the notion that exhibitors strategically do not release movies before “big” events so as not to take up the screens with movies with binding minimum run-length restraints. The only instance in which the proportion of new movies to all movies screened is lower is one week before a high-attendance week, and even then the difference is less than 1ppt. On all other measures the proportion is either unchanged, or even higher than the average.

Another way to alleviate concern that not accounting for forward-looking exhibitor behavior can bias

⁴⁴Blockbuster movies were identified subjectively, and include movies that turned out to be huge hits as well as those which promised to be huge hits but did not turn out that way, but exclude surprise hits that exhibitors could not have anticipated

⁴⁵Defined as weeks with cumulative attendance higher than 120% of the average weekly attendance

Table 4: Average proportion of new movies to all movies screened

event \ weeks before event	all weeks	restricted		
		1	2	1 - 2
Blockbuster movies	52.4%	55.0%	52.2%	53.2%
High-attendance weeks	52.4%	51.6%	54.0%	52.9%
Holidays	52.4%	58.7%	55.3%	56.3%

the estimates it to try put a sign on any potential bias. If, contrary to the simplifying assumption, exhibitors are less likely to take on new movies at times so that their binding minimum run length restraints do not get in the way of taking on new movies in the coming periods, this would bias $\hat{\lambda}_{0ct}$ and λ_{0t} estimates downward (as the exhibitor not taking on some of the outside movies would be “blamed” on these movies’ poor quality). Such a bias would be depress the simulated quality, both true and expected, in the counterfactual calculations. As a result, the increase in attendance/consumer welfare/industry profits calculated in the counterfactuals would be a conservative bound on expected improvements.

5.3 Counterfactual Calculations

This section describes the strategy for calculating counterfactuals, as well as the algorithms used to calculate $z_c^*(\cdot)$, movie theater c ’s optimal schedule, and \mathbf{r}_c^* , the set of revenue splits charged by distributors.

Counterfactual strategy The aim of the counterfactual simulations is to calculate the impact of removing contractual restraints on measures such as attendance, consumer welfare, as well as exhibitor and distributor profits. Faced with consumer demand estimates, the model solves the exhibitor and distributor problem of setting $z_c^*(\cdot)$ and \mathbf{r}_c^* so as to maximize their profits given a set of movies to choose from. While the model calculates accurate estimates of true and expected movie quality for movies screened, it only provides limited information for movies not screened in movie theater c . In order to calculate the *expected* impact of removing contractual restraints the model performs a number of true and expected movie quality simulations for movies not screened, and calculates the *average* impact across the simulations. For more details see implementation details in the Appendix.

Calculating $z^*(\cdot)$ The problem of finding an optimal schedule when attendance at one screening depends on all other screenings can be viewed as a spacial case of the maximum coverage problem which is *NP-hard* i.e. no polynomial-time algorithm is known for solving it. A typical approach to solving NP-hard problems is the heuristic approach, when a polynomial-time approximation algorithm is used that is known to work well in many cases but for which there is no proof that

it always produces the best result. The model employs one such heuristic, the greedy algorithm, which has been shown to be the best polynomial-time algorithm to solve the maximum coverage problem (Hochbaum, 1997; Feige, 1998).

The aim of the algorithm is to allocate movies to empty screening-time/screen slots, the composition of which is identical to that actually observed, so as to maximize the movie theater’s per-period profit function (23). This can either be done without restrictions or under the no screen-sharing restriction. The algorithm employed fills up the movie schedule iteratively, at each step adding the movie/screening-time combination that most increases the *combined* profits, until there are no more empty slots to fill.⁴⁶

Capacity constraints need to be accounted for explicitly when calculating $z^*(\cdot)$. The algorithm assigns screens in decreasing order of capacity, thus ensuring most attractive movie/screening time combinations chosen early on are assigned to the largest screens. For movie/screening time combinations where the predicted attendance exceeds the screen’s capacity the algorithm considers assigning another screen to the movie - doing so does not expand the consumers’ choice set, but instead raises the capacity constraint for this movie/screening time combination. After the algorithm is finished, attendance at each movie/screening time combination is capped at the combined capacity of assigned screens. The advantage of this approach is the simplicity of its implementations; however, the downside is that it may under-predict attendance. The reason for this is that by simply removing predicted attendance above a screen’s capacity, this approach ignores the fact some of the consumers who could not make it into the movie/screening combination could have considered an alternative combination. However, since this algorithm is only used in counterfactual calculations and not in estimation, this possible under-prediction does not cause bias; rather, it means calculated improvement in welfare is a conservative estimate.

Calculating r_{mc}^* This algorithm aims to find r_{mc}^* which maximizes the distributor’s profit from screening movie m in movie theater c : Π_{mc} (28). Relying on the fact that Π_{mc} is non-decreasing in r_{-mc} the algorithm starts with $r_{mc} = \min(R_{mt}) \forall m \in M_{ct}^*$, where R_{mt} is the range of values revenue splits can take.⁴⁷ It then proceeds to iteratively increase r_{mc} on each movie until no distributor finds it profitable to increase it any further.

⁴⁶A secondary algorithm can be used to go back over the schedule and considers changes to movies played, one slot at a time, searching for improvements; this algorithm was not used in the interest of computation time since it produced improvements smaller than 1%

⁴⁷The set of possible revenue split values R_{mt} depends on whether the movie is released on the break or in the second run, and in both cases the values are limited to those actually observed in the sample

6 Results

The structural model outlined above is estimated and the results are presented below. The section that follows uses the estimated model to conduct counterfactual simulations of what would be the effect of removing the contractual restraints; this is done for the two movie theaters for which revenue split information was available. The final section discusses the model fit along with various robustness considerations.

6.1 Parameter Estimates

6.1.1 Demand model

Parameter estimates from the demand system are presented in Table 5.

Table 5: Demand model coefficients

Variable	Means	St Dev	Interaction with age dummy variables			
			2-11	12-24	25-59	60+
Constant	-10.01*** (2.14)	1.00*** (0.13)		1.00*** (0.17)		
3D	0.26* (0.19)					
On the break ($w_{mt} = 0$)	0.42*** (0.09)					
w_{mt}	-0.03*** (0.01)					
MPAA: PG-13/R			-0.75 (3.16)			
Genre: Family/Animated			1.00 (2.01)			
Genre: Fantasy/SciFi/Animated						-5.75** (1.75)
Screening: Weekdays before 5pm				-3.75* (2.30)	-3.75* (2.30)	
ω_{im}		0.89*** (0.13)				
Fixed effects	movie-, week-, year-, screening time/movie theater/time interval-					
GMM objective	45.36					
Number of observations	54,785					
Number of markets	701					
Number of movie theaters	5					

Note: ***, **, and * indicate significance at the 1%, 5% and 10% levels

The results are intuitive and coefficients have the expected directionality. Consumers differ significantly in their desire to go to the movie theater, and young people aged 12-24 have an especially high valuation of going to the movies. 3D movies are more attractive to moviegoers despite their higher price, which agrees with the trend in the movie industry as a whole to produce more movies in 3D. Screening a movie on the break provides a big boost in attractiveness, while the longer it is

been since a movie's nationwide release the less attractive it is.

Age proves to be a strong predictor of consumers' tastes in movies. As expected, PG-13 and R-rated movies are less appealing to young people aged 2-11, or rather to people taking them to see the movies (technically even the youngest children can see an R-rated movie as long as they are accompanied by an adult). On the flip side, family and animated movies are more appealing to them than to older people. Seniors aged 60+ do not like Fantasy, Science Fiction or Animated movies. As expected, people aged 12-59 are less likely to go to movies before 5pm during the week because of school/work commitments. Finally, consumers differ in their preferences for individual movies, as showed in the high standard deviation of ω_{im} . This helps explain why even the "biggest" movies (e.g. Avatar) never grab the whole market to themselves.

Chart 4: Week fixed effects


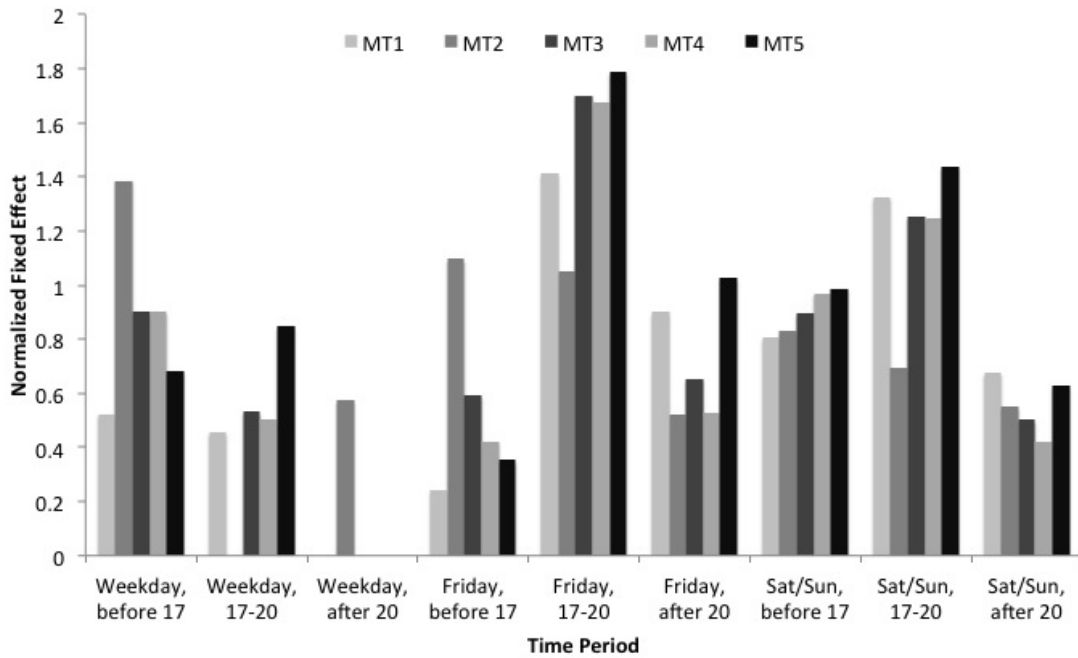


Chart 4 plots weekly fixed effects. On one hand it shows that going to the movies is more attractive during and around major holidays, which corresponds to findings in Einav (2007). On the other hand, while Einav finds that nationwide the summer months are particularly popular, the results here indicate they are in fact less attractive than the rest of the year. Most likely this reflects a difference in how small communities captured in the sample spend their summers compared to the average person in the US, possibly leaving town or simply electing to spend more time outdoor.

Chart 5: Screening time period fixed effects, by movie theater



Note: values are normalized such that, for each movie theater, the smallest value equals zero

Chart 5 plots normalized time period fixed effects, by movie theater. The first take-away is that the normalized fixed effects exhibit similar trends across different movie theaters. For all but one movie theater the *Mon-Thu, after 8pm* time period is the least attractive to moviegoers, while the *Friday, 5-8pm* and *Sat/Sun, 5-8pm* periods are the first and second most attractive overall.

Table 6: Top/Bottom 10 Movie Fixed Effects

Top 10	F.E.	BOR	Bottom 10	F.E.	BOR
		(\$MM)			(\$MM)
Pirates of the Caribbean: Dead Man's Chest	3.75	423.3	Ninja Assassin	-1.65	38.1
Cars	3.13	244.1	Saw 3D: the Final Chapter	-1.67	45.7
Wild Hogs	3.04	168.3	Death At a Funeral	-1.70	42.7
Night At the Museum	2.91	250.9	She's Out of My League	-1.71	31.6
The Chronicles of Narnia	2.82	291.7	Machete	-1.82	26.6
Over the Hedge	2.70	155.0	Skyline	-1.87	21.4
X-men: the Last Stand	2.70	234.4	The Next Three Days	-1.89	21.1
Superman Returns	2.66	200.1	The Crazyes	-1.98	39.1
Shrek the Third	2.60	320.7	Why Did I Get Married Too?	-2.13	60.1
Pirates of the Caribbean: At Worlds End	2.58	309.4	Case 39	-2.16	13.2

Notes: BOR represents the total US Box Office Revenue over the course of a movie's theatrical run

Table 6 reports the top and bottom 10 movie fixed effects, alongside each movie's national Box Office Revenue take. The first take-away is that the difference in fixed effects value between the top and bottom movie observed in the sample is considerably larger in magnitude than that for either week- or time period-fixed effects. This emphasizes the importance of screening the best movies

each week, and suggests being able to drop badly-performing movies quickly (once the minimum run length restraint is lifted) will allow exhibitors to substantially boost attendance. The movies in the top 10 were some of the the highest grossing nationwide over the sample period, and all of them topped the nationwide Box Office lists the week they opened. The bottom 10 movies all did poorly in their theatrical runs given that they all enjoyed a nationwide release.

6.1.2 Movie Quality and Expectations Generation Process

Table 7 reports results from the second stage estimation process, which identifies parameters driving the movie quality and expectations generation process (1). Standard error are calculated using bootstrapping:

Table 7: β^M estimates

Variable	Estimate	s.e.	sig.
β^M			
const	-2.00	(0.38)	***
IMDB rating	0.12	(0.05)	***
Rating: G	0.80	(0.20)	***
Budget	5.00	(0.56)	***
Genre: fantasy	0.40	(0.19)	**
Genre: sci-fi	0.40	(0.17)	***
Genre: musical	-1.00	(0.31)	***
Genre: documentary	-0.80	(0.31)	***
σ_{μ_m}	1.41	(0.39)	***
$\sigma_{\nu 1}$	0.50	(0.14)	***
$\sigma_{\nu 2}$	0.87	(0.25)	***
ℓ	1,331.02		

Notes: (a) Budget is expressed in \$100M; ***, **, * indicate significance at the 1%, 5% and 10% levels

Analyzing β^M , movies that have better “word of mouth”, as proxied by *IMDB rating*, have higher appeal to moviegoers.⁴⁸ Movies which receive a G-rating from the MPAA, meaning they are suitable for all audiences, have the highest average appeal. As expected, movies with a higher budget are, on average, more attractive to moviegoers. This captures two factors: (1) such movies are indeed better (2) a higher budget proxies for more promotional activity. Although only some of the *genre* coefficients turned out to be significant, it is clear that musicals and documentaries hold less appeal to audiences, while fantasy and sci-fi movies are more attractive *ceteris paribus*.

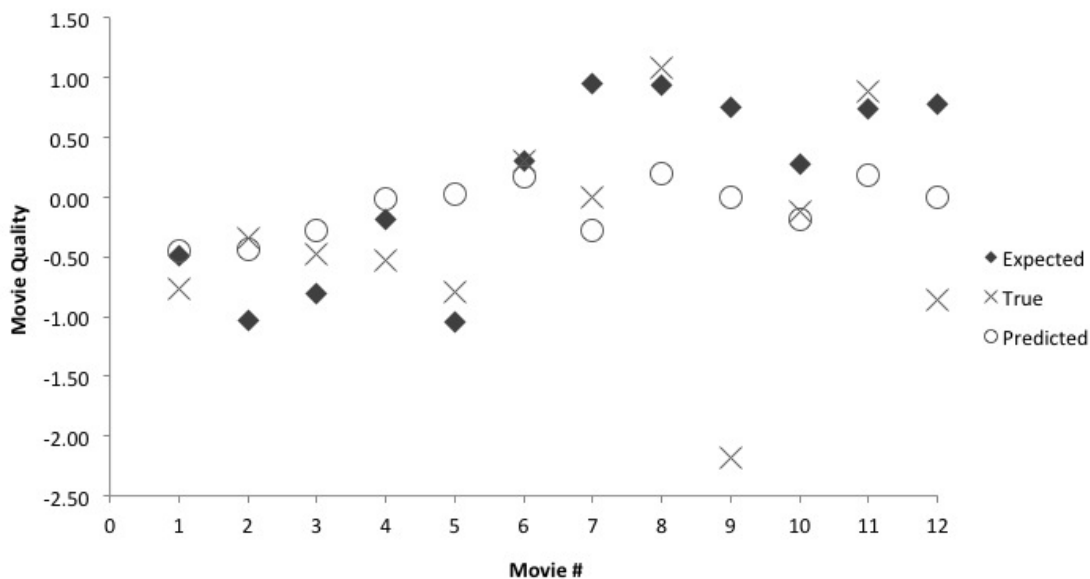
The estimated σ_{μ} coefficient is larger than either of the $\sigma_{\nu c}$ coefficients - this indicates exhibitors’ expectations of movie quality are relatively accurate compared to how much true quality varies across

⁴⁸Although in many cases moviegoer ratings on IMDB followed *after* a movie was screened, due to lack of better alternatives this variable is used as an imperfect proxy for “word of mouth”, which exhibitors are assumed to know before a movie is released

movies with the same observable characteristics x^M . In addition, the results suggest expectations made by MT1 are, on average, closer to the true movie quality than those made by MT2.

Chart 6 provides an illustration of how estimated movie quality expectations $\hat{\lambda}_{mc}$ compare to true movie quality λ_m and predicted movie quality $x^M\beta^M$. For movies which were actually screened expected and true movie qualities were estimated by the model. For movies which were not screened, these values were simulated based on the movies' predicted quality such that the expected movie quality was no higher than $\hat{\lambda}_{c0t}$, the expected quality for the best available alternative.

Chart 6: Movie quality expectations vs true and predicted values



Note: Example for MT1, period 1

MT1 has 6 screens, and in period 1 took on 6 new movies from among 12 to choose from. As expected, the expected movie quality for the chosen movies, 7-12, is higher than for movies 1-6, while the average ν_{mMT1} draw for these movies is positive, supporting the notion of positive selection among movies screened. It is interesting, however, that the true movie quality for movie 6 is in fact higher than 4 out of 6 movies which were chosen, while the true quality for movie 9 is the lowest of all movies considered. This suggests that, if the minimum run length restraint was lifted, the movie theater would greatly benefit from dropping movie 9 in period 2 and replacing it with movie 6.

6.2 Counterfactual Results

6.2.1 Welfare cost of restraints

Table 8 presents results of counterfactual simulations which measure the impact of removing the contractual restraints. The simulations were carried out for the two movie theaters for which revenue split data was available, MT1 and MT2, for four scenarios:

1. Base case, minimum run length 3 weeks
2. Base case, minimum run length 2 weeks
3. No minimum run length, screen-sharing not allowed
4. No minimum run length, screen-sharing allowed

Since the minimum run length period is not observed in the sample, the base case is split into two scenarios. In scenario 1. the minimum run length restraint imposed on all movies released on the break is the minimum of 3 weeks or the number of weeks a movie was actually screened for.⁴⁹ In scenario 2. the minimum run length period is 2 weeks; in reality the minimum run length restraints lie between those assumed in scenarios 1. and 2.. As a result, comparing results from counterfactual scenarios 3. and 4. to these in base cases 1. and 2. provides upper and lower bounds, respectively, on the improvements expected from removing the restraints.

Table 8: Results of Counterfactual Simulations
MT1

Scenario:	1. Base	2. Base	3. no screen-sharing		4. screen-sharing allowed	
Movie Theater #1 (6 screens)	abs	abs	abs	Δ	abs	Δ
Attendance (thousands)	231,872	242,777	252,019	3.8% to 8.7%	254,595	4.9% to 9.8%
# movies screened	107	127	146	15.0% to 36.4%	153	20.5% to 43.0%
% movies released on break	86.9%	88.2%	91.8%	3.6pp to 4.9pp	92.2%	4.0pp to 5.2pp
Consumer utility (utils)	1,412	1,467	1,546	5.4% to 9.5%	1,576	7.4% to 11.6%
Exhibitor profits (thousand \$)	1,189	1,277	1,349	5.6% to 13.4%	1,359	6.4% to 14.3%
Distributor profits (thousand \$)	1,008	1,023	1,041	1.8% to 3.2%	1,056	3.2% to 4.7%
Total Welfare Change				3.9% to 8.7%		5.0% to 9.9%
equiv. increase in #screens (abs)				0.4 to 0.9		0.5 to 1.0
equiv. increase in #screens (rel)				6.6% to 15.1%		8.5% to 17.1%

MT2

Scenario:	1. Base	2. Base	3. no screen-sharing		4. screen-sharing allowed	
Movie Theater #1 (6 screens)	abs	abs	abs	Δ	abs	Δ
Attendance (thousands)	57,618	60,868	64,985	6.8% to 12.8%	66,773	9.7% to 15.9%
# movies screened	94	132	188	42.4% to 100.0%	214	62.1% to 127.7%
% movies released on break	60.6%	59.1%	55.9%	-3.2pp to -4.8pp	54.7%	-4.4pp to -6.0pp
Consumer utility (utils)	4,506	4,816	5,006	3.9% to 11.1%	5,275	9.5% to 17.1%
Exhibitor profits (thousand \$)	300	317	345	9.0% to 15.1%	361	14.0% to 20.3%
Distributor profits (thousand \$)	259	273	283	3.9% to 9.5%	285	4.5% to 10.1%
Total Welfare Change				6.5% to 11.9%		9.4% to 14.9%
equiv. increase in #screens (abs)				0.4 to 0.7		0.6 to 0.9
equiv. increase in #screens (rel)				6.5% to 11.9%		9.4% to 14.9%

In all scenarios attendance figures were calculated using the estimated demand model. Every scenario was ran multiple times, each simulation based on a different set of randomly simulated $\hat{\lambda}_{mc}$

⁴⁹Although a minimum run length restraint of 4 weeks is possible conversations with exhibitors suggest it is rare and even if it is imposed, it is attached to big releases for which it is not binding

and λ_m values, and the final numbers presented are an average across all simulations. Absolute values for each scenario are reported alongside the range of change compared to the base cases.

Change in attendance Removing the minimum run length restraint only (scenario 3.) results in an increase in attendance compared to the base cases: 3.8 to 8.7% for MT1 and 6.8 to 12.8% for MT2. As expected, these values are greater for the smaller movie theater MT2 where the constraints are more restrictive. The attendance gain is driven primarily by an increase in the number of movies screened, which rises up to twofold. The equivalent increase is considerably lower for MT1, reflecting the fact that at twice the number of screens it was able to offer a broader choice set to consumers even under the prevailing restrictions. The rise in attendance is also driven by the fact that once the minimum run length restraint is lifted exhibitors are able to quickly adjust their schedules once they learn the true quality of movies.

Allowing for screen-sharing (scenario 4.) on top of removing the minimum run length restraint results in a further increase in attendance and number of movies screened, especially for MT2. However, these gains are relatively smaller than those realized in scenario 3., partly because the consumer choice set's breadth already increases greatly with just the removal of the minimum run length restraints.

Welfare impact The rise in attendance results in substantial welfare increases, though the gains are not distributed uniformly across all market participants. For both movie theaters exhibitors capture more of the incremental profits than do distributors, reflecting the fact that removing contractual restraints broadens their strategic options and increases their bargaining position. As expected, consumer welfare increases in line with attendance.

Equivalent increase in number of screens Using the model it is possible to answer the question: how many more screens would the movie theater need to generate comparable increases in total welfare under the prevailing restraints. As expected, the larger increase in total welfare would necessitate a bigger relative rise in the number of screens for MT2, though the absolute rise in number of screens is remarkably close. These are only values for the movie theaters in the sample, however a back-of-the-envelope calculation suggests keeping welfare constant while lifting the restraints would allow the closure of at least 1,600 screens.⁵⁰

6.2.2 Is Lifting Restraints an Equilibrium Outcome?

Results presented in Section 6.2.1 show that lifting vertical restraints would increase total distributor profits for digital movie theaters. However, to date distributors have continued imposing restraints on such movie theaters, even as they made up 64% of the market in 2011.

⁵⁰ Assuming the conservative bound on the equivalent number of screens, and roughly 4,000 locations in the US with 1-4 screens

The answer as to why distributors continue to impose restraints may be provided by game theory. The process of imposing/lifting restraints is formulated as a one-period game between major movie distributors, wherein each distributor decides whether to impose restraints given his expected pay-offs. Table 9 analyzes the stability of two potential equilibria: (1) no distributors impose restraints (2) all distributors impose restraints.

Table 9: Stability of potential equilibria

Potential equilibrium	1. No one imposes VRs Gain from imposing VRs		2. All impose VRs Loss from lifting VRs	
	MT1	MT2	MT1	MT2
Major #1	37.9%	22.5%	-27.9%	-20.8%
Major #2	36.3%	26.9%	-29.8%	-30.9%
Major #3	47.1%	9.5%	-46.1%	-17.6%
Major #4	36.9%	124.0%	-24.0%	-37.6%
Major #5	7.6%	104.8%	-27.4%	-36.9%
Major #6	38.1%	53.7%	-29.4%	-24.8%
Average	34.0%	56.9%	-30.8%	-28.1%

It is clear from Table 9 that “no distributors impose restraints” is not an equilibrium, as every major distributor has an incentive to unilaterally deviate. The reason for this is that by imposing restraints an individual distributor can capture a larger share of the revenues by forcing movie theaters to commit to his movies ahead of time, while the exhibitor continues to efficiently choose the best movies from amongst his competitors. On the other hand, “all distributors impose restraints” is a stable equilibrium as no Major has an incentive to deviate. The reason for this is that by lifting restraints on his movies a distributor allows them to be squeezed out by those coming from his competitors. If a movie is good it will play no longer with restraints lifted, but high expectation/low actual quality movies will be replaced sooner than if restraints were imposed.

Both these findings are true for all major distributors for both MT1 and MT2. This corresponds to the finding in Besanko & Perry (1993) that exclusive dealing restraints can be a result of a prisoner’s dilemma game, wherein exclusive dealing restraints are imposed in equilibrium even though each manufacturer would prefer nonexclusive dealing. Although it is not possible to provide a definitive answer as to why distributors impose restraints on digital movie theaters without being able to view the whole US market, these findings provide a possible explanation and suggest it would be welfare-improving for competition authorities to ban the use of these restraints for digital movie theaters.

6.3 Model Fit and Robustness

One of the central notions in the paper is that increased consumer choice will lead to a rise in the number of times consumers go to the movies throughout the year, driving leading to increased welfare and industry profits. To fully simulate the impact of such an increase in consumer choice it

is important to accurately model the way in which consumers differ between one another, as this will lead to different responses in the counterfactual scenario. Charts 7 and 8 demonstrate model fit across consumers and for specific age groups (corresponding micro-moments: #1 and #2).

Chart 7: Population composition, by moviegoing frequency

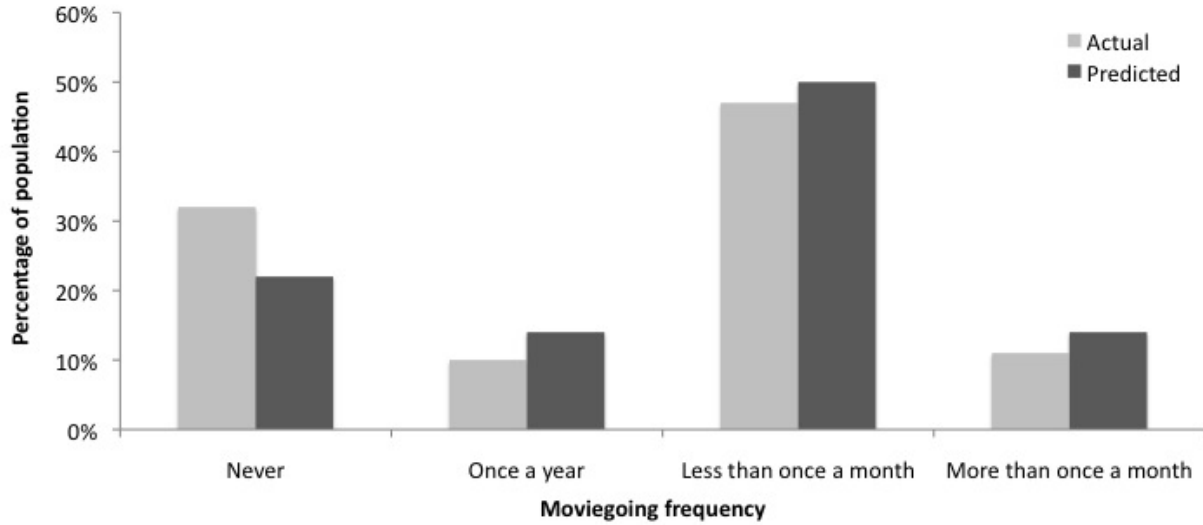

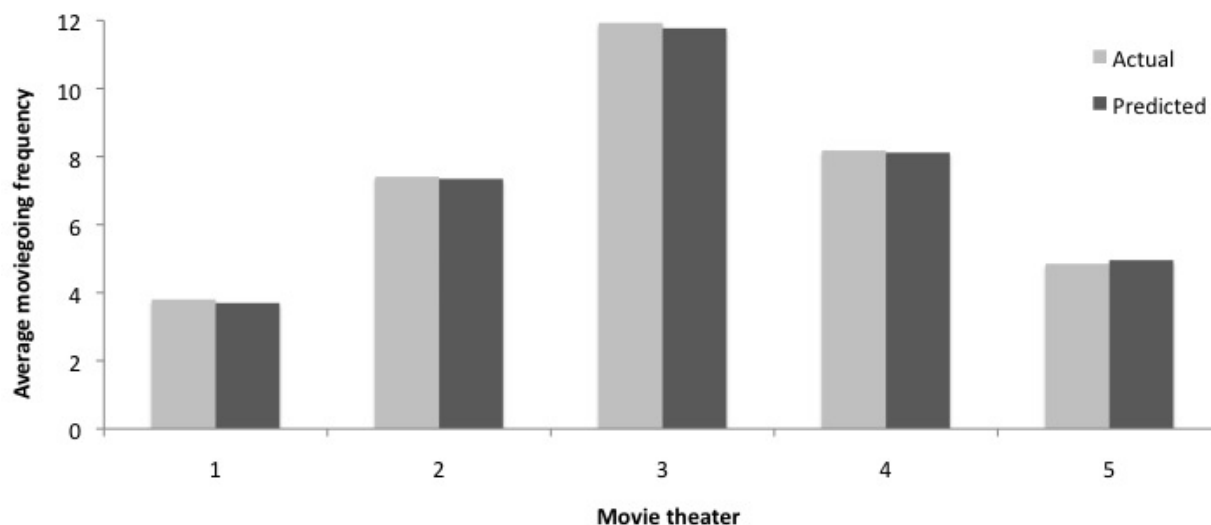


Chart 8: Moviegoing frequency by age, relative to movie theater average



Given that the attendance data set in an unbalanced panel of movie theaters with different number of screens, it is important the model captures the difference between the movie theaters, otherwise the counterfactual calculations (based on two of the five movie theaters) will be biased. Chart 9 suggests that model does a good job of capturing differences between movie theaters.

Chart 9: Moviegoing frequency, by movie theater



7 Conclusion

Counterfactual calculations show that the *minimum run length* and *no screen-sharing* restraints significantly reduce welfare and industry profits when they are imposed on digital movie theaters. Removing them would allow exhibitors to make their offering more *relevant* by quickly adjusting their schedules in response to learning the true quality of movies, as well as more *broad* by allowing them to screen more movies overall. This finding has strong implications for the US exhibition industry as it is set to fully convert to digital projection by the end of 2013.

The second major finding is that although lifting the restraints increases overall profits for distributors, it is not a stable equilibrium for all distributors to lift restraints. Counterfactual results suggest in such a situation each major distributor would find it beneficial to unilaterally deviate and impose restraints. On the other hand, the results indicate that for all distributors to impose restraints is a stable equilibrium, as unilateral deviations are profit-decreasing for all major distributors. These findings might explain why distributors continue imposing restraints on movie theaters which have adopted digital projection technology. Conversely, they also suggest it would be welfare-improving for competition authorities to ban the use of these restraints for digital movie theaters.

Given the small sample size it is important to consider how these findings translate into the whole market. The sample is non-representative in three ways: (1) movie theaters in the sample are local monopolies (2) they are all small- and medium-sized (3) the exhibitors are all independent. These limitations are addressed below.

First, conversations with exhibitors suggest movie theaters in competitive markets aim to offer a selection of movies which is no worse than their competitors'; this is borne out in the real world,

where movie theaters in direct competition to one another offer very similar schedules. This lack of scheduling complementarity suggests the combined choice set from competing movie theaters is likely to be no broader than that offered by an individual movie theater, and would thus benefit similarly from the removal of contractual restraints.

Second, since larger movie theaters can take on more movies each week it can reasonably be expected the benefits from removing contractual restraints will be smaller. This is supported in the sample, where the proportional increase in number of screens needed to generate welfare gains equivalent to removing the restraints is smaller for the larger movie theater. Nonetheless, it is impossible to conclude how much welfare gains from removing contractual restraints are reduced for large movie theaters without access to detailed attendance data for such movie theaters.

Third, conversation with exhibitors suggest movie theaters which are part of a chain (e.g. AMC Loews, Regal Entertainment) do not get preferential treatment when it comes to contractual restraints. While large chains are understood to be able to drive rental prices down for some movies, counterfactual results suggest that to exhibitors these prices are of secondary importance to movies' quality when making scheduling decisions. As such, welfare gains at movie theaters owned by large chains are not expected to be significantly different to those independently owned.

It is also important to consider any greater implications these findings have for the movie industry as a whole. The impact of removing contractual restraints on the set of movies that are *produced* is likely to be limited since movies are released over multiple channels and the theatrical box office revenue is only one factor which determines whether they are given the green light in production. Nonetheless, it is possible that of the movies that lose the most from the lifting of restraints (low quality, high expectations movies) some might never get made. Additionally, while movie theaters' locations and number of screens are taken as exogenous in the model, the removal of contractual restraints may have long-term equilibrium consequences in this respect. 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.

Beyond movie exhibition, this paper's findings confirm that in industries where products or services are sold to consumers through independent retailers non-price vertical restraints imposed by manufacturers can have a significant impact on total welfare. When only a select group of products can be offered to consumers at a given time and place, such restraints can substantially alter the composition of consumers' choice sets and thus their purchasing decisions, even if they do not impact

prices. The impact is likely to be greater when retailers regularly make decisions under uncertainty as to the appeal of new products to consumers. This has implications for industries such as radio and TV, offline and online advertising, as well as retail sales.

8 Appendix

8.1 Attendance data changes

The attendance data set is adjusted to suit the discrete choice model used to explain consumer demand.

Pooling observations In the US tickets to see movies in most movie theaters are “general admission”, meaning they do not come with seat reservations. Abstracting from screen capacity constraints, it can thus be reasonably assumed that all tickets to see a given movie at a given time are viewed by consumers as identical, even if movies are played across a couple different screens. The discrete choice model is not well-suited to handling such multiple parallel screenings of one movie, since each movie/screening combo gets its own Ξ_{msct} draw, suggesting consumers get different levels of utility from seeing one movie on different screens. The discrete choice model will thus predict higher attendance if the movie theater puts on multiple screenings of the same movie at the same time, even if one screen could have handled all the demand.

In order to avoid this problem observations are pooled for a given movie in cases where screening times are close enough to be viewed by moviegoers as identical. The cutoff for time difference between screening which are pooled together is set at 60 minutes; this value was chosen to balance two goals. On one hand, a high value is needed to pool staggered releases of one movie, e.g. 6 screenings at 10-minute intervals between 12:00am and 12:50am. On the other hand, the value has to be low enough so as to differentiate between two sequential screenings of even the shortest movies on one screen e.g. a 7:10pm and 9:00pm screenings. Overall, the number of screenings is reduced by 2.1%.

Eliminating observations The focus of this paper is on regular screenings of feature-length movies, which are the major source of revenues to exhibitors and distributors. In the data set, however, there are few observations that do not conform to this description. One way of dealing with them would be to leave them in the data set and estimate the coefficients driving their attractiveness to moviegoers. However, there are not enough observations for most of them to satisfy asymptotic requirements of the estimators. Thus the following screenings are removed from the data set:

1. Special screenings at non-standard times e.g. school trips in the morning hours
2. Free screenings e.g. summer movie series
3. Non-movie events e.g. NBA, NFL games, concerts


1.9% of observations are removed through this process .

8.2 Revenue split value changes

In order to simplify implementation and speed up model estimation and counterfactual simulation movies' revenue split values are “flattened” and discretized, while the range of possible values is limited for movies released on the break and in the second run.

Flattening Renting movies on a sliding scale is a practice the industry has been moving away from and which is likely to be discontinued in the coming years, according to exhibitors. Overall, fewer than 19% of movies in the data set were rented on a sliding scale contract, a proportion which fell to below 15% in 2010 (the last full year of observation). For these reasons all movies in the counterfactual simulated are rented on a flat rate. In estimation, so as not to eliminate observations, the revenue split values for all movies rented on a sliding scale are “flattened” i.e. converted from a sliding scale to a flat rate. The new flat rate is calculated such that the total split of box office revenues between the distributor and exhibitor over the entire observed movie run is as close as possible to that under the original sliding scale pricing schedule.

Chart 10: Revenue split values, movies on flat rate contracts released in the second run




Source: revenue split data from movie theaters

Range limitation and value discretization As illustrated in Chart 10 the majority of revenue split values for movies on flat rate contracts released in the second run falls within the 35-55 range, with only one movie rented at 60. Additionally, the vast majority of movies are rented at revenue splits which are multiples of 5. Thus, to simplify and speed up calculations, revenue split values are allowed to only take of a limited number of discrete values from the set [35, 40, 45, 50, 55]. In

estimation, revenue split values are modified to fit this set, with the new revenue split value being as close as possible to the observed value.

Chart 11: Revenue split values, movies on flat rate contracts released on the break



Source: revenue split data from movie theaters

As illustrated in Chart 11 the majority of revenue split values for movies on flat rate contracts released on the break falls within the 50-60 range. Although for this set of movies the case is not as clear-cut as for movies released in the second run, here too revenue split values are discretized such that they fall into to the set $[50, 55, 60, 65]$.

8.3 Initial moviegoing history

Consumer moviegoing history, ι_{it} , is a crucial determinant of demand as consumers in the demand model do not see the same movie twice. Since its value is not observed by the econometrician it needs to be simulated within the model, which creates the initial condition problem - what movies were seen by moviegoers prior to the first period of observation? In order to get around this problem the following approach is taken:

1. For each movie theater c the demand model is simulated for all periods $t \in T_c$ assuming $\iota_{i \min T_c} = \emptyset \forall i, c$, i.e. moviegoers had not seen any movies in periods prior to the first period of observation
2. Only time periods $t \geq \min T_c + t_{initial} \forall c$ are taken into account when forming the GMM objective function, where $t_{initial}$ is set such that $M_{c \min T_c}^+ \cap M_{c \min T_c + t_{initial}}^+ = \emptyset$ i.e. none of the

movies screened in the first period of observation were screened in period $minT_c + t_{initial}$, and any impact of the initial condition is second-order

Correspondingly, in counterfactual simulations only periods $minT_c + t_{initial}$ onwards are analyzed. Since $t_{initial}$ as defined above varies between simulations, its value is fixed at 5 such that for most simulations no movies screened in the first period of observation were screened in period $minT_c + t_{initial}$.

8.4 Movie quality estimation algorithm

The algorithm

1. Consider one possible combination of $\sigma_\mu^2, \{\sigma_{\nu c}^2\}_{\forall c}, \beta^M$ values
2. $\forall c, t$, determine λ_0 , the set of $\{\lambda_{0t}, \hat{\lambda}_{0ct}\}$ pairs which are consistent with observed schedules
3. calculate $\ell_{tc}(\cdot|\cdot)$ for all $\{\lambda_{0t}, \hat{\lambda}_{0ct}\} \in \lambda_0$
4. set $t = \min_{\forall c}(T_c)$
5. find $\{\lambda_{0t}, \{\hat{\lambda}_{0ct}\}_{\forall c} \text{ s.t. } t \in T_c\}$ multiple which maximizes $\ell_{tc}(\cdot|\cdot)$
6. check that $\hat{\lambda}_{0ct}$ is consistent with movies released in previous periods i.e. $\hat{\lambda}_{0ct} \geq \max_{t' < t} (\lambda_{0t'} + (t - t') * \beta_2^W)_{\forall c}$; if not, go to 7., else go to 8.
7. find $\left\{ \hat{\lambda}_{0ct} \in \left[\text{argmax}(\ell_{tc}(\cdot|\cdot)), \max_{t' < t} (\lambda_{0t'} + (t - t') * \beta_2^W) \right] \right\}_{\forall c}$ which maximize $\sum_{\forall c} \sum_{t' = \min_{\forall c}(T_c)}^t \ell_{t'c}(\cdot|\cdot)$
8. set $t = t + 1$
9. iterate steps 4 - 7 while $t \leq \max_{\forall c}(T_c)$, calculate $\ell(\sigma_\mu^2, \{\sigma_{\nu c}^2\}_{\forall c}, \beta^M|\cdot)$ when done
10. iterate steps 1 - 8 until $\ell(\sigma_\mu^2, \{\sigma_{\nu c}^2\}_{\forall c}, \beta^M|\cdot)$ is maximized

In point 2 inconsistent $\lambda_{0t}/\hat{\lambda}_{0ct}$ pairs are ones where:

1. $\hat{\lambda}_{0ct} > \min_{m \in M_{ct}^+(w_{mt} \geq MRL_m)}$ i.e. the best alternative cannot be better than the worst of the movies being kept on whose minimum run-length is no longer binding (that is, the exhibitor is free to replace it but chooses not to do so), where MRL is the assumed minimum-run length restraint for movie m ; this imposes restraints on $\hat{\lambda}_m$ values for movies released in period t and on λ_m values for movies released before t
2. $\hat{\lambda}_{0ct} < \min_{m \in M_{ct}^+(w_{mt} \geq 1)} (\lambda_m + w_{mt} \beta_2^W)$ i.e. in period t when at least one of the movies that opens in movie theater c was released nationally before t the best alternative cannot be better than the worst movie that was actually released; this also implies...

3. ... $\lambda_{0t'} < \min_{m \in M_{ct}^+(w_{mt} \geq 1)} (\lambda_m + w_{mt'} \beta_2^W)$ i.e. for a movie with known quality to be released it has to be better than similar alternatives in period t , and thus none of the movies released before period t can be better than this movie

These restrictions are exogenous to the estimation procedure and can be imposed by analyzing the exhibitor schedules.

The need for points 6-7 stems from the fact the, by definition,

$$\hat{\lambda}_{0ct} = \max \left[\max_{m \in M_{ct}^-(w_{mt}=0)} (\hat{\lambda}_{mc}); \max_{m \in M_c^-(w_{mt}>0)} (\lambda_m + w_{mt} \beta_2^W) \right]$$

can be driven either by the best movie released in period t that was not screened or the best movie released in previous periods that was not screened. By considering all $\{\lambda_{0t}, \hat{\lambda}_{0ct}\} \in \boldsymbol{\lambda}_0$ in point 3 the algorithm does not account for the latter component of $\hat{\lambda}_{0ct}$. If the resultant $\hat{\lambda}_{0ct} > \max_{t' < t} (\lambda_{0t'} + (t - t') * \beta_2^W)$ than it is consistent with the model. If the opposite is true, the algorithm needs to consider all candidate values between the two extremes and determine which maximizes $\ell(\cdot)$ to this point.

Implementation details In the data set over 99% of movies are screened within the first 9 weeks of their nationwide release. In order to speed up calculations and better approximate actual exhibitor behavior the algorithm only allows movies to be considered for release up to 9 weeks after their nationwide release.

References

- Aghion, P., & Bolton, P. (1987). Contracts as a Barrier to Entry. *The American economic review*, 77(3), 388–401.
- Ainslie, A., Drèze, X., & Zufryden, F. (2005). Modeling movie life cycles and market share. *Marketing Science*, 24(3), 508–517.
- Alimurung, G. (2012). Movie studios are forcing hollywood to abandon 35mm film. <http://www.lawweekly.com/2012-04-12/film-tv/35-mm-film-digital-Hollywood/>. Accessed on 09/12/2012.
- 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.
- Bernheim, B., & Whinston, M. (1998). Exclusive dealing. *Journal of Political Economy*, 106(1), 64–103.
- Berry, S. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, (pp. 242–262).
- Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, (pp. 841–890).
- Berry, S., Levinsohn, J., & Pakes, A. (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy*, 112(1), 68–105.
- Besanko, D., & Perry, M. (1993). Equilibrium incentives for exclusive dealing in a differentiated products oligopoly. *The RAND Journal of Economics*, (pp. 646–667).
- Brenkers, R., & Verboven, F. (2006). Liberalizing a distribution system: the European car market. *Journal of the European Economic Association*, 4(1), 216–251.
- Cachon, G., & Larivière, M. (2005). Supply chain coordination with revenue-sharing contracts: strengths and limitations. *Management science*, 51(1), 30–44.
- Chipty, T. (2001). Vertical integration, market foreclosure, and consumer welfare in the cable television industry. *American Economic Review*, (pp. 428–453).
- Chisholm, D. (1997). Profit-sharing versus fixed-payment contracts: Evidence from the motion pictures industry. *Journal of Law, Economics, and Organization*, 13(1), 169–201.

- Comanor, W., & Rey, P. (2000). Vertical restraints and the market power of large distributors. *Review of Industrial Organization*, 17(2), 135–153.
- Dana Jr, J., & Spier, K. (2001). Revenue sharing and vertical control in the video rental industry. *The Journal of Industrial Economics*, 49(3), 223–245.
- 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.
- Feige, U. (1998). A threshold of $\ln n$ for approximating set cover. *Journal of the ACM (JACM)*, 45(4), 634–652.
- 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.
- Gasparon, P., & Visnar, B. (2005). Coca-cola: Europe-wide remedies in fizzy drinks.
- Geuss, M. (2012). Celluloid no more: distribution of film to cease by 2013 in the us. <http://arstechnica.com/gadgets/2012/06/the-silver-screen-no-more-distribution-of-film-to-cess-by-2013-in-the-us/>. Accessed on 09/12/2012.
- Gil, R. (2009). Revenue sharing distortions and vertical integration in the movie industry. *Journal of Law, Economics, and Organization*, 25(2), 579–610.
- 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.

- Hansen, L. (1982). Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, (pp. 1029–1054).
- Hanssen, F. (2002). Revenue-sharing in movie exhibition and the arrival of sound. *Economic Inquiry*, 40(3), 380–402.
- Heide, J., Dutta, S., & Bergen, M. (1998). Exclusive dealing and business efficiency: Evidence from industry practice. *JL & Econ.*, 41, 387.
- Hicks, J. (1939). Value and capital: An inquiry into some fundamental principles of economic theory.
- Ho, J., Ho, K., & Mortimer, J. (2010). Welfare effects of full-line forcing contracts in the video rental industry. *NBER Working Paper*.
- Ho, K. (2009). Insurer-provider networks in the medical care market. *The American Economic Review*, 99(1), 393–430.
- Hochbaum, D. (1997). Approximation algorithms for np-hard problems. *ACM SIGACT News*, 28(2), 40–52.
- Klein, B. (1988). Vertical integration as organizational ownership: The fisher body-general motors relationship revisited. *JL Econ. & Org.*, 4, 199.
- Klein, B., & Murphy, K. (1988). Vertical restraints as contract enforcement mechanisms. *Journal of Law and Economics*, 31(2), 265–297.
- Krattenmaker, T., & Salop, S. (1986). Anticompetitive exclusion: raising rivals’ costs to achieve power over price. *Yale Law Journal*, 96(2), 209–293.
- 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.
- Lee, R. (2012). Vertical integration and exclusivity in platform and two-sided markets. *New York University mimeo*.
- Marvel, H. (1982). Exclusive dealing. *Journal of Law and Economics*, 25(1), 1–25.
- Mathewson, G., & Winter, R. (1987). The competitive effects of vertical agreements: Comment. *The American Economic Review*, 77(5), 1057–1062.
- Mortimer, J. (2008). Vertical Contracts in the Video Rental Industry¹. *Review of economic studies*, 75(1), 165–199.

- 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.
- MPAA (2011). Theatrical market statistics.
- Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2), 307–342.
- Ornstein, S. (1989). Exclusive dealing and antitrust. *Antitrust Bull.*, 34, 65.
- Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110(4).
- Rasmusen, E., Ramseyer, J., & Wiley Jr, J. (1991). Naked exclusion. *The American Economic Review*, 81(5), 1137–1145.
- Rey, P., & Tirole, J. (2007). A primer on foreclosure. *Handbook of industrial organization*, 3, 2145–2220.
- Riordan, M. (2008). Competitive effects of vertical integration. *Handbook of Antitrust Economics*, (pp. 145–82).
- Sass, T. (2005). The competitive effects of exclusive dealing: Evidence from the us beer industry. *International Journal of Industrial Organization*, 23(3), 203–225.
- Segal, I., & Whinston, M. (2000). Exclusive contracts and protection of investments. *RAND Journal of Economics*, (pp. 603–633).
- Slade, M. (1998). Beer and the tie: Did divestiture of brewer-owned public houses lead to higher beer prices? *The Economic Journal*, 108(448), 565–602.
- Swami, S., Eliashberg, J., & Weinberg, C. (1999). SilverScreener: A modeling approach to movie screens management. *Marketing Science*, 18(3), 352–372.
- Telser, L. (1960). Why should manufacturers want fair trade? *Journal of Law and Economics*, 3, 86–105.
- Weinstein, M. (1998). Profit-sharing contracts in hollywood: Evolution and analysis. *The Journal of Legal Studies*, 27(1), 67–112.
- 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. (2008). Lectures on antitrust economics. *MIT Press Books*, 1.

Zanarone, G. (2009). Vertical restraints and the law: evidence from automobile franchising. *Journal of Law and Economics*, 52(4), 691–700.