
ELSEVIER
Editors:
P. Bajari
B. Caillaud
N. Gandal
Advisory Board:
J. Levinsohn
F. Malerba
R. Porter
M. Riordan
M. Slade
M. Waterson
Co-Editors:
J.P. Choi
J. Gans
A. Hortacsu
R.A. Jensen
G.Z. Jin
B. Jullien
D. Martimort
V. Nocke
A. Petrin
M. Shum
Associate Editors:
V. Aguirregabiria
R. Amir
S. Clerides
G.S. Crawford
J. Dana
U. Doraszelski
S. Fisher Ellison
J. Fox
K. Graddy
F. Khalil
K. Krishna
J. Lerner
D. McAdams
E.J. Miravete
H. Paarsch
M. Peitz
B.J. Ruffle
S. Ryan
M. Rysman
O. Shy
C. Snyder
Y. Spiegel
V. Stango
R. Stenbacka
S. Stern
G. Tan
J. Thisse
N. Vettas
R. Veugelers
Editorial Assistant:
A. Simantov-Damti
26 (4) 861–1058
(2008)
Volume 26, issue 4, July 2008
ISSN 0167-7187
INTERNATIONAL JOURNAL OF
INDUSTRIAL
ORGANIZATION
CONTENTS
Regular papers
Efficient horizontal mergers: The effects of internal capital reallocation and organizational form
S.H. Mialon 861
Pricing strategy, quality signaling, and entry deterrence
A. Utaka 878
The choice at the checkout: Quantifying demand across payment instruments
R. Borzekowski and E.K. Kiser 889
Endogenous horizontal mergers under cost uncertainty
W. Zhou 903
Assessing horizontal mergers under uncertain efficiency gains
P. Choné and L. Linnemer 913
Generic entry, price competition, and market segmentation in the prescription drug market
T.L. Regan 930
Content and advertising in the media: Pay-tv versus free-to-air
M. Peitz and T.M. Valletti 949
Retailer entry conditions and wholesaler conduct: The theatrical distribution of motion pictures
C.C. Moul 966
Biases in demand analysis due to variation in retail distribution
S. Tenn and J.M. Yun 984
(Contents continued on last page of this issue)
Available online at
ScienceDirect
www.sciencedirect.com

Available online at www.sciencedirect.com
ScienceDirect
International Journal of Industrial Organization
26 (2008) 966–983
International Journal of Industrial Organization
www.elsevier.com/locate/econbase

Retailer entry conditions and wholesaler conduct: The theatrical distribution of motion pictures ☆

Charles C. Moul *

Washington University, Campus Box 1208, St. Louis, MO 63130-4899, USA

Received 16 March 2006; received in revised form 22 August 2007; accepted 28 August 2007
Available online 6 September 2007

Abstract

I add to the empirical literature on vertical contracting and wholesaler conduct by using retailer entry conditions to infer unobserved choice variables and equilibrium responses to prices and advertising. After estimating the US demand for theatrical motion pictures from 1990–96, I apply these techniques to compare observed outcomes to predictions under various distributor-conduct hypotheses. While several caveats apply, results indicate that the hypothesis of competition among distributors fails to describe advertising levels or aggregate payments of theaters to studios. The hypothesis of some collusion among distributors, however, matches the data fairly well.

JEL classification: D3; L1; L8; Z1

Keywords: Vertical relationships; Collusion; Theatrical movie industry

Empirical research on the vertical issues between wholesaler and retailer is often held back by a lack of data on intermediate prices and other choice variables. Such issues also affect the ability to test horizontal hypotheses of wholesaler conduct and market power when only retail data are available. While wholesaler costs and prices can frequently be recovered under equilibrium and profit-maximizing assumptions for the most common case of sequential price setting, there exists another important margin of many vertical relationships, namely shelf space.[1] Slotting allowances (i.e., payments from wholesalers to retailers to secure placement for new products) have received much attention in marketing, though less in economics.[2] The infrequently observed payments along this margin clearly influence consumer behavior but do not fit readily into

☆ I would like to thank Editor Ali Hortacsu and an anonymous referee for valuable suggestions; Doug Whitford, Vice President of Film Buying for Wehrenberg Theaters, for his assistance regarding institutional details; and David Switzer for sharing his data on exhibition contracts and concessions. This paper has also greatly benefited from seminar participants at the University of Oklahoma and the University of Illinois at Urbana-Champaign. For purposes of length, supplemental material is available on a web appendix linked at http://moul.wustl.edu. All errors are my own.

* Tel.: +1 314 935 5836; fax: +1 314 935 4156.
E-mail address: moul@artsci.wustl.edu.

0167-7187/$ - see front matter
doi:10.1016/j.ijindorg.2007.08.005

[1] See Manuszak (2007) and Villas-Boas (2007) for applications of the wholesale cost inference under the sequential pricing framework.

[2] Sudhir and Rao (2006) offers a survey of the slotting allowance literature in marketing in their analysis of whether the practice is pro- or anti-competitive. Shaffer (1991) and Sullivan (1997) are prominent examples of the slotting allowance literature in economics.

the existing discussion of missing data. Furthermore, even when necessary data are available, the complexity of the retailer's decision will often prevent direct estimation of equilibrium responsiveness.

In this paper, I use entry conditions in retail to draw inferences about the equilibrium responsiveness of retailer resources (e.g., shelf space) to wholesaler choice variables and to infer unobserved wholesaler choice variables. Specifically, I use the fact that in equilibrium retailers must be satisfied with their current shelf space choices and that at the margin the retailer is indifferent between slight changes among products. Such indifference characterizes my equilibrium conditions, the total differentiation of which then yields the impacts of wholesaler choices on retailer allocation of shelf space. As an empirical illustration, I turn to the theatrical exhibition of motion pictures, a sector in which the shelf-space margin of vertical contracting is especially stark. I then apply this method to consider the question of wholesaler conduct, comparing the predictions under various hypotheses of distributor (studio) conduct to observed outcomes of vertical payments and advertising.

While the contentious vertical relationship between movie studios and exhibitors has a long and well-known history, distributor conduct and the extent of horizontal market power among studios has received scant attention. Given the concentration of theatrical movie distribution in the U.S., this omission is on its face somewhat surprising. The top eight distributors throughout the 1990s routinely received over 90% of the payments from exhibitors, with Hirschman–Herfindahl Indices between 1200 and 1550.[3] With the exceptions of Disney's entrance into distribution in 1953 and RKO's exit in 1955, the major distributors through the mid-1990s were the same as they had been in 1930s. Such concentration and stability of players suggest the possibility of at least tacit collusion, perhaps even in the face of difficulties brought on by substantial product differentiation and unstable demand conditions.

The lack of attention to the question of horizontal market power is best explained by an absence of critical price information, namely the wholesale price that theaters pay studios to exhibit a movie for a week.[4] Characteristics of the movie industry allow me to incorporate my retailer entry conditions with now common techniques in industrial organization to infer these wholesale prices under a variety of assumptions. Specifically I use structural demand estimates to back out under the different scenarios what prices would simultaneously maximize profits, clear the exhibitor market, and generate the observed outcomes. I follow a similar approach, conditional upon typical wholesale prices, to infer marginal costs of advertising and the implied expenditures on advertising. While national wholesale price information is rarely (if ever) available by movie and week, cumulative payments from theaters to studios over the theatrical run of a movie are sometimes known. Furthermore, the industry provides strong priors regarding both these prices and advertising-to-sales ratios. Comparing these observations and priors to the cumulative predictions under the various assumptions allows me to compare the validity of those assumptions and hypotheses. Exploiting data that are generally assumed to be unobserved has been used before with respect to marginal cost. Most similar to this work, Nevo (2001) estimates the demand for ready-to-eat cereal, calculates the profit-maximizing mark-ups under different conduct hypotheses, and compares the median to a crude estimate of the mark-up on a typical brand. My approach will parallel Nevo's, as identification works primarily off the level of cumulative payments rather than variation among these payments.

The basic model that I use to describe studio decisions closely follows relevant parts of the standard contract for the time period of my data. Specifically, each week studios select advertising and the percentage of exhibitor box office receipts that will be returned to the studio (i.e., the rental rate). Theater operators observe the rental rates and advertising of all available movies and decide whether or not to show the movie. The total number of exhibiting theaters for each movie is then determined by an entry condition that the marginal theater operator is indifferent between showing his movie and any other available movie at these terms, an obviously quite strong assumption that I explore later. Thus, while studios most directly care about the number of theaters showing their movies, they achieve this outcome by competing against each other in rental rates and advertising. The entry conditions in the exhibition sector imply causal relationships between the rental rate and the number of exhibiting theaters and between the level of advertising and the number of exhibiting theaters. These relationships then enter into a studio's first-order conditions to maximize profits. Differing levels of responsiveness of demand to the number of exhibiting theaters will then imply different levels of rentals rates, allowing the observed level of cumulative rental payments to distinguish between the competition and collusion hypotheses. The same approach with

[3] Figures are taken from Litman (2001), p. 177.

[4] The only economic research of which I am aware that observe these wholesale prices are Blumenthal (1988), Switzer (2004), and Filson, Besocke, and Switzer (2005).

respect to advertising yields different implied expenditures on advertising under the conduct hypotheses, and I compare these predicted advertising-to-rental ratios to the industry's conventional wisdom.

I find that rental rates implied by the simplest competitive model are lower than those suggested by industry standards and observed aggregate payments from theaters to studios. Portfolio considerations for multiproduct firms also fail to generate high enough rental rates. The collusive hypotheses of collusion among the major studios, however, do generate appropriately high rental rates and explain up to 83% of their variation. In contrast, assuming that studios are competing (regardless of the choice of other assumptions) explains no more than 58% of the cumulative rentals' variance. Similarly, competition among multiproduct firms suggests unrealistically high ratios of advertising expenditures to distributor revenues (at least 75% and upwards of 100%), while the collusion hypotheses generate much more plausible ratios (25%–30%).

The paper is laid out as follows. Section 1 discusses the history and critical aspects of the vertical contracts in the theatrical movie sector. Section 2 models first consumers and then firms, with emphasis in the latter upon how market clearing in exhibition identifies the responsiveness of the number of exhibiting theaters to wholesale prices and advertising. I discuss the data that I employ in Section 3, and Section 4 includes discussion of the identification of demand, the estimation techniques, and results. I conclude in Section 5.

1. Vertical contracting and industry details

The primary economic actors in distribution are well known, with most distributors tracing at least part of their histories back to the consolidation of the 1930s. For the period of my data (1990–1996), seven studios were considered to be the "majors": Disney, MGM/UA, Paramount, Sony (controlling Columbia and Tri-Star), 20th-Century Fox, Universal, and Warner Bros.[5] During this time, a few studios fell somewhere between these majors and the truly independent distributors. The most noteworthy were Miramax (prior to its 1993 acquisition by Disney), Orion, and Savoy.

Antitrust issues are fairly prominent in the industry's history. In 1921, the newly formed Federal Trade Commission declared the industry practice of block-booking (i.e., bundling movies for exhibitors) to be anticompetitive. The FTC renewed its litigation against the vertically integrated studios in 1928, charging the studios with monopolization through vertical foreclosure. While the studios were found guilty in 1930, the penalties were nullified as part of New Deal legislation in the early 1930s. A similar monopolization suit was filed by the Department of Justice in 1938, and, after a respite for World War II, the industry was again found guilty in 1948. The Paramount Consent Decrees resulted. Under *Paramount*, the defendant studios divested their theater holdings, legally separating theatrical distribution and exhibition. *Paramount* also affirmed the illegality of block-booking. The status of block-booking was reaffirmed by the Supreme Court in 1962, and, despite the counterarguments provided by Stigler (1968) and Kenney and Klein (1983, 2000), distribution proceeds on a movie-by-movie basis.

Distributor-exhibitor negotiations after *Paramount* eventually settled into auctions in which exhibitors provided multidimensional bids for the right to be a movie's exclusive showcase within an area. By 1990, however, these auctions had largely been replaced by direct negotiations between distributors and exhibitors.[6] Like their auction counterparts, these contracts specified payments to the distributor in three parts: a minimum guarantee, a rental rate, and a figure (called the "nut") beyond which the distributor claimed 90% of receipts. The maximum implied by these three terms was then paid to the distributor.[7] In equilibrium the rental rate is by far the most important component. For example, contracts with the largest theater chain in St. Louis, MO, from September 2001 to June 2002 showed that the payment based upon the nut was used in only 0.55% of observations and the average impact when exceeded was an additional 2 percentage points beyond the contracted rental rate. The minimum guarantee came into play in about 0.02% of observations.[8] While contracts apply for several weeks (four weeks was common for my sample), negotiated terms are allowed to vary across weeks, and declining rental rates are the norm.[9] An opening week rental rate of 70% is common for highly anticipated movies, and rental rates will typically reach 30% as a movie plays out. Given typical patterns of movie demand over time, the industry rule of thumb is that about half of box office receipts are paid to distributors as rentals. Theatrical distribution entails two primary costs: the

[5] DreamWorks began distributing movies in September 1997.

[6] Disney continued to make substantial use of auctions until 1994.

[7] See the web appendix for a figure illustrating this mapping of revenues to rental payments.

[8] See Switzer (2004), p. 41.

[9] The declining length of movies' theatrical runs as consumers are shifted to the opening week has generated a new contract, first introduced in 2001, that specifies a single rental rate that is constant over the run.

physical replication of the reels of film and their transportation to theaters (i.e., prints and shipping) and the purchase of advertising. As print expenses are essentially sunk, advertising remains as the only important expenditure under a distributor's control. It is now common for a studio to spend an amount equal to 30–50% of a movie's production cost on domestic advertising and promotion. A theatrical distributor therefore faces no important variable costs.

A common question in the theatrical movie sector is why admissions prices to consumers are constant, both across movies at a given theater and over time for a given movie. While Orbach and Einav (2007) offer some insights on the matter, no definitive answer has yet been reached. This constancy of prices has now been written into the standard distributor–exhibitor contract (exhibitors are allowed a single discount night during the week). The fixed price is unsettling from an economic perspective, but quite useful from an econometric one. I will throughout assume that admissions prices are exogenous to both distributors and exhibitors on a weekly basis.

I will work from the assumption that distributors maximize domestic theatrical weekly profits, but it is informative to consider what biases might arise if other objective functions are more accurate. The rise of video and the growing importance of foreign markets over the past decades have relegated domestic theatrical exhibition to a position of decreasing importance. An obvious conclusion is that a far-sighted distributor might sacrifice part of its share of domestic theatrical receipts if cross-market spillovers would inflate later ancillary profits. While some preliminary work has begun in this area (see Ravid 1999; Elberse and Eliashberg 2003), no definitive findings on such spillovers are yet available. If such spillovers are important, the truly profit-maximizing rental rates (advertising levels) would be less (more) than those implied by my assumption of static profit-maximization. The opposite is likely to arise if distributors follow some dynamic decision-making process, choosing rental rates to maximize their movies' cumulative theatrical profits. Higher rates (lower advertising) in this scenario are optimal because a distributor would likely recognize that, if fewer theaters show the movie, much of the demand that it is lost this period is merely deferred to later periods. There is some empirical research in the marketing literature on this front (Krider and Weinberg, 1998), but what findings there are suggest that demand is primarily lost rather than being delayed.

While market power in distribution is an obvious concern, the size of some theater chains (e.g., AMC, Regal Cinemas) suggests that market power in exhibition may also be important. My data are aggregated to the movie level, and I am unable in this paper to shed light on the extent of this retail-level market power. I therefore assume throughout the paper that exhibition is symmetric and competitive. If this competitive assumption is violated in favor of market power in the exhibition sector, my implied rental rates will be too high, as some of the economic rents are shifted from distribution to exhibition.

Finally, it is worth noting that the possibility that distributors are colluding in rental rates to exhibitors and advertising does not arise in a vacuum. Sony, Warner Bros., MGM, Paramount, and Universal were the recent focus of an Antitrust Division investigation into their joint venture in video-on-demand venture, MovieLink. A parallel civil suit by the competing service Entertainer (settled in March 2006) accused the studios of acting in concert to vertically foreclose competition. While the Antitrust Division's investigation was closed in June 2004, the area will likely remain a topic of interest.

2. The model

2.1. Consumer behavior

Like much of the literature, I model demand using the discrete-choice techniques introduced by McFadden (1978) and applied to differentiated products by Berry (1994). Specifically, I follow Bresnahan, Stern, and Trajtenberg (1997) by using a principles of differentiation generalized extreme value (PD GEV) model. The PD GEV has a parameterization similar to the nested logit without the restrictions that are generated by the (typically arbitrary) ordering of the nests of the consumer decision process. The primary advantage of this specification over estimation using random effects is its relative robustness without imposing a supply-side model (i.e., joint pricing equation) and in the absence of microlevel data such as demographics.

Discrete-choice models of demand are based upon a population of M risk-neutral consumers making choices from a set of J products each period. In my case, each consumer can choose no more than one movie in a week. Including the outside option (i.e., the option of choosing no product at all for that period) within the choice set, a consumer will purchase a product if her expected valuation from such a purchase exceeds the expected valuation from any other option available that period. Considering the entire market at time t, individual i's optimal choice j and expected valuation for choice j can be respectively expressed as

$$\begin{aligned} & j\,s.t.\ \mathrm{EU}_{itj} \geq \mathrm{EU}_{itk}, \forall k \\ & \mathrm{EU}_{itj} = \delta_{tj} + \varepsilon_{itj} \end{aligned} \tag{1}$$

where δ_{tj} is the mean expected valuation of consumers considering the movie, and ε_{itj} is the individual-specific deviation about that mean. Ties between these consumer-specific expectations occur with zero probability, and the unconditional probability of purchase is then Pr ($\mathrm{EU}_{itj} \geq \mathrm{EU}_{itk} \forall k$). As this valuation is an ordinal concept, normalizing the mean valuation of no purchase to zero ($\delta_{t0}=0\ \forall\ t$) is necessary for identification. The empirical goal is to explain variation in quantities with variables explaining mean utilities δ. Matching the literature, I will transform quantities into purchase probabilities, that is, the unconditional probability that a consumer purchases a given product during a particular week.

As is well established, the distribution of the individual-specific deviations ε determines much of the nature of aggregate demand and the resulting substitution patterns. Deviations that are independently drawn (such as those from the Type 2 extreme value distribution that generates logit probabilities) yield the familiar Independence of Irrelevant Alternatives property. Allowing such deviations to be independent across consumers but correlated for a given consumer across products can yield more intuitive substitution patterns. In his generalized extreme value model, McFadden (1978) formalizes a framework to allow for this sort of correlation. Individual-specific deviations are realized from a distribution that depends upon segmentation parameters ρ that lie on the unit interval. The econometrician chooses the potential pattern of segmentation to be estimated by selecting principles of differentiation, and each such principle is assigned its own ρ. As any dimension's ρ approaches 1, consumers give less consideration to substituting to similar products along that dimension. Conversely, as any ρ approaches 0, consumers consider only similar products along that dimension when substituting among competing products.[10] I focus upon three such dimensions for substitution within the PD GEV framework. They roughly correspond to the questions, Should I see any movie?, Should I see an action movie? and Should I see a family movie? These dimensions should capture likely effects of consumer heterogeneity in competition between movies and other entertainment options as well as between movies of different genres. While any discrete movie characteristic could be the basis for segmentation, a movie's action status and family status have the desirable traits of being relatively cleanly defined and powerful.

With this framework established to allow for more flexible substitution patterns, I turn to the components of the mean expected utility δ. I assume that consumers have preferences about travel and that, as the number of theaters showing a movie rises, the distance that the average consumer must travel to reach a theater exhibiting that movie falls. I also assume that advertising directly affects consumer utilities, which is consistent with either advertising being a complement with movie consumption or a signal of quality. Additionally, I assume that advertising derives its impact entirely from a demand complementarity with the number of exhibiting theaters. This prevents the perverse prediction that a movie receiving a great deal of advertising but showing at few theaters receives a very large number of admissions. As specifications allowing for possible impacts of lagged advertising did not substantively improve demand estimation, the models that I estimate consider only current advertising. Perhaps the most obvious characteristic of theatrical movie demand is the way in which demand decays over time. While this decay is presumably driven by a shrinking pool of inexperienced consumers, Einav (in press) shows that this decreasing demand can be sufficiently captured by using a movie's age, the number of weeks since its release. I use 1n(Age) and its square as a relatively flexible specification. The demand regressors also include other exogenous variables that vary by week and by movie.

The richness of my data allows me to utilize movie fixed effects for all movies observed at least four times. The impacts of movie-specific characteristics such as genres and the appeal of the cast and director can then be recovered by regressing the estimated fixed effects upon these variables.[11] I proxy for seasonal differences in demand by including binary variables for each month from February to December and for the eleven major holidays.[12] The week-varying variables are completed with a measure of the unemployment rate and a crude measure of average admissions price.

[10] Thus, ρ corresponds to $1-\sigma$ in the nested logit model. See the web appendix for the formal specification.

[11] Movies that I observe three or fewer times have these same movie-specific characteristics included directly in the original regression.

[12] These holidays are New Year's Day, Martin Luther King, Jr., Day, Presidents' Day, Easter, Memorial Day, Independence Day, Labor Day, Columbus Day, Veteran's Day, Thanksgiving, and Christmas.

Below is the specification for mean expected utility and the result of inverting the observed purchase probabilities.

$$\delta_{tj}(\rho) = \Lambda_j + \alpha \ln(T_{tj}) + \beta \ln(1+T_{tj})\ln(1+\mathrm{ADS}_{tj}) + \varphi_1 \ln(\mathrm{Age}_{tj}) + \varphi_2 \left(\ln(\mathrm{Age}_{tj})\right)^2 + X_{tj}\kappa + W_t\psi + \xi_{tj} \quad (2)$$

The number of theaters exhibiting movie j in week t is denoted by T_{tj} and comparable advertising by ADS_{tj}. The demand parameters to be estimated are then $\Gamma(\rho, \Lambda, \alpha, \beta, \varphi, \kappa, \psi)$.

2.2. Distributor behavior and market-clearing in exhibition

I focus upon the studio's choices of the weekly rental rate and advertising for each movie in its portfolio. This abstracts away from the multiweek nature of the contracts as well as the other details of payment (e.g., minimum guarantees, size of nut, etc.). The fact that multiweek contracts almost always include declining rental rates somewhat supports this assumption, as does the overwhelming prominence of the rental rate as the determinant of rental payments.[13] I also greatly simplify the exhibitor's profit expression by considering profits from concessions as a constant share of the admissions price.

Each week each distributor selects the rental rates (r) and advertising levels (ADS) for the movies in its portfolio. It does so recognizing that a higher rate will directly generate higher rental payments but lead fewer exhibitors to show the movie and indirectly affect the exhibitors' decisions to show other movies as well. Higher advertising generates higher box office receipts and higher rental payments, as well as increasing the number of exhibiting theaters.[14] The critical and unknown components in the corresponding profit-maximizing first-order conditions (FOCs) are then the matrices $\frac{dT}{dr}$ and $\frac{dT}{d\mathrm{ADS}}$, which include both direct and indirect impacts of distributor choices on the equilibrium number of theaters. At present, direct estimation of these relationships has been done only after imposing the restriction that the number of retailers carrying a product does not depend upon the prices, advertising, or characteristics of other products.[15] Applying the discrete choice framework that has been useful in demand estimation is not viable in this context, as the retailer profits presumably depend upon the full set of demand interactions, and thus separability does not hold. In the absence of direct estimation, characterizing these derivatives requires making assumptions on the nature of market-clearing in exhibition. To address this concern fully, I consider the admissions gain of the marginal theater to be a weighted average of the marginal impact of an additional theater and the average impact of an additional theater: $\frac{ds_k}{dT_k} = (1-\theta)\frac{\partial s_k}{\partial \delta_k}\frac{\partial \delta_k}{\partial T_k} + \theta \frac{s_k}{T_k}$. The former applies if cannibalization among theaters is non-existent ($\theta = 0$), and the latter applies if such cannibalization is total ($\theta = 1$).[16] While attempts to estimate θ precisely as a free parameter were unsuccessful, the results under the extreme hypotheses should yield some insights into how this assumption may color the final conclusions. My primary conclusions are all robust to the choice of θ through intermediate values.

The other key point in exhibition is the relevance of entry. Theaters obviously take a substantial length of time to construct and open, and an assumption of no weekly entry (or exit) might seem appropriate. In such a case, equilibrium is reached when all exhibitors are indifferent at the prospect of showing any other movie and all theaters are utilized. No solution for rental rates assuming collusion of all distributors can exist under this no-entry assumption. This no-entry assumption, however, is not the only possibility. Given the shift to multiscreen theaters that Davis (2006a) documents, superior (if unavailable) measures would be the number of exhibiting screens or seats in those auditoriums. In light of this potentially large data limitation, I also consider the opposite assumption of entry, namely that theaters enter and exit each week based upon a zero-profit condition for each movie. This is also more consistent with the substantial variation from week to week that I find in my sample's total number of theaters showing any movie. Equilibrium in both cases is

[13] This is admittedly a static approximation to a dynamic game, but my demand specification lacks an intertemporal linkage. Because my specification has movie demand decaying with age, no word-of-mouth, and advertising having only an immediate impact, my estimates yield no insight into how actions this week will carry over to demand next week. Extensions to demand that could address these matters are clearly desirable.

[14] I provide formal specifications of the objective function and first order conditions in the web appendix.

[15] See, for example, Elberse and Eliashberg (2003).

[16] Davis (2006a) provides evidence that additional theaters in a market both "steal" from existing theaters and expand the total number of consumers who see movies in theaters.

based upon a no-arbitrage condition where the marginal theater is indifferent between its choice and all other possibilities.[17]

This equilibrium condition is quite strong and deserves additional discussion. I have already discussed that block-booking (i.e., bundling of several titles) is presently illegal and does not appear in vertical contracts. The repeated interaction of theaters and studios, however, may lead to informal block-booking. In an example of this scenario, perhaps a studio deliberately limits the number of theaters showing a certain blockbuster and then charges a rental rate that is lower than it need be. That is, the studio leaves money on the table for those exhibitors lucky enough to secure exhibition rights for the movie. In return, the grateful exhibitor agrees to show some of the studio's less certain movies in the future without being compensated with lower rental rates. My equilibrium assumption in such a context would overstate the rental rate for the blockbuster and understate the rental rates for the studio's riskier releases.

A second possibility is that movie demand is substantially uncertain, even to distributors and exhibitors who have pre-screened the movie. If equilibrium rental rates are based upon expected demand, then exhibitors may be indifferent *ex ante* but not *ex post*. My equilibrium conditions, however, are inherently based upon *ex post* numbers. My model would thus overstate rental rates for movies that were upside surprises and understate rental rates for disappointments. Under either of these violations of my assumed equilibrium condition, there is no consistent bias in rental rates, and I correspondingly expect no substantial bias of rental rate magnitudes in either direction. One might alternatively argue that blockbusters and upside surprises are likely to have much longer theatrical runs then their counterparts from each scenario. Such an argument would imply that my model systematically overstates rental rates, a position that works against the collusion hypothesis.

Building upon the assumed indifference among exhibitors and recalling that θ captures the extent of cannibalization among theaters, the two entry conditions can be formally expressed as follows (week subscripts are suppressed throughout). Let there be N movies in a given week and let π_c denote average per capita concessions profit as a fraction of admissions price p. The exhibition market clears under the no-entry assumption if

$$\begin{aligned}(1-r_1+\pi_c)pM\left((1-\theta)\left(\frac{\partial s_1}{\partial \delta_1}\right)\left(\frac{\partial \delta_1}{\partial T_1}\right)+\theta\left(\frac{s_1}{T_1}\right)\right)\\ -\Phi_1=(1-r_k+\pi_c)pM\left((1-\theta)\left(\frac{\partial s_k}{\partial \delta_k}\right)\left(\frac{\partial \delta_k}{\partial T_k}\right)\right.\\ \left.+\theta\left(\frac{s_k}{T_k}\right)\right)-\Phi_k\forall\, k\neq 1\\ \sum_{i=1}^{N} T_i-\mathrm{TotT}=0 \qquad (3)\end{aligned}$$

where Φ_k denotes any fixed costs of movie k, and TotT denotes the number of available theaters. In contrast, the market-clearing condition under the free entry assumption is

$$\begin{aligned}(1-r_k+\pi_c)pM\left((1-\theta)\left(\frac{\partial s_k}{\partial \delta_k}\right)\left(\frac{\partial \delta_k}{\partial T_k}\right)+\theta\left(\frac{s_k}{T_k}\right)\right)\\ -\Phi_k=0\,\forall\, k=1,\ldots,N \qquad (4)\end{aligned}$$

If these fixed costs are specified as the sum of a weekly expectation and a disturbance ($\Phi_k=\Phi+e_k$), then it becomes clear that this fixed cost parameter is identified from the free-entry case but not the no-entry case.

Application of Cramer's Rule to these systems of equations yields the necessary expressions of $\frac{dT}{dr}$ and $\frac{dT}{d\mathrm{ADS}}$. For all cases, this application requires that quantity demanded is not too responsive to the number of theaters (e.g., solutions under the free-entry assumption require own-theater elasticities to be less than 1). The derivatives $\frac{dT}{dr}$ implied by the no-entry case depend critically and in a non-linear fashion upon all relevant rental rates, and I discuss how I implement this intermediate step in more detail under the estimation section. The derivatives implied by the free-entry case, however, are substantially simpler to recover, in that rental rates appear only as a linear coefficient after a multiplicative transformation. Substituting these expressions for $\frac{dT}{dr}$ into distributor FOCs and solving then yields the solution of equilibrium rental rates. The linearity of the free-entry case furthermore makes inference of rental rates from both the FOCs and the zero-profit conditions tractable. I therefore use feasible generalized least squares with appropriate cross-equation

[17] This no-arbitrage condition will arise in most models that consider the market to be national. It can also be approximated when the national market is broken into cities of various sizes, and theaters enter cities so long as it is profitable in expectation.

restrictions to recover those rental rates under the free-entry assumption.[18]

Interested readers can find an analytical example of a single-product monopoly with free-entry in exhibition in the web appendix. From that solution, equilibrium rental rates are increasing in per capita concessions profits and declining in a movie demand's responsiveness to the number of exhibiting theaters. Any overestimate of η^T will therefore suggest that equilibrium competitive rental rates are lower than is true and thus overstate the likelihood that competition would be rejected in favor of collusion.[19]

3. The data

3.1. Quantities, screens, and advertising

Like much of the recent empirical work in the movie-economics literature, I turn to *Variety*'s listings for revenue measures, from which I gather each week's movies. Each week *Variety* magazine publishes a ranking of the prior week's domestic (U.S. and Canada) box-office revenues for the fifty highest grossing movies.[20] The week of August 24, 1990, was the first where direct reporting by the studios fully supplanted the sampling techniques that *Variety* had previously used, and it is with this week that I begin my sample.[21] In addition to revenues, the *Variety* listings include each movie's number of exhibiting theaters. The sample ends the week of December 27, 1996, spanning 332 weeks and yielding 16,600 observations.

Revenues are broadly consistent with quantities and sometimes instructive on their own, but demand analysis requires quantities and therefore a measure of price. As explained above, the admissions price does not vary substantially across movies at a given theater, and so I turn to the National Association of Theatre Owners (NATO) annual average admissions price. While virtually no one pays this average price, it is arguably the best available when considering the domestic market as a whole. I derive quantities (Q) using a linear reconstruction from the eight annual price observations (1990–97). So long as movie selections are made by individuals paying their own admission, this price measure is reasonably appropriate, but group-purchases by a single payer (e.g., families) will clearly create distortions and make it less so. The transformation of quantities into purchase-probabilities ($s = \frac{Q}{M}$) also requires a measure of market size. I use the combined population of the U.S. and Canada for the first week of July each year and linearly reconstruct weekly counterparts.[22,23] Per capita concessions profits as a share of admissions price (π_c) are assumed to be 0.29, e.g., if the ticket price is \$5, I assume that each consumer also yields \$1.45 of concession profit to the exhibitor.[24] Without some measure of concession profits, the model will systematically understate implied rental rates and bias results towards the collusion hypothesis, but this measure is admittedly crude. It is plausible that this concessions measure should vary across genres and over the age of the movie, and future research in this particular vertical relationship may be warranted.

Other than the number of exhibiting theaters, the dominant component of demand that varies across both movies and weeks is advertising. While total advertising budgets are occasionally released after movies' runs are complete, I could find no broadly available measures of a movie's weekly advertising expenditures. The measure I use instead is newspaper advertising from *The Chicago Tribune* on that week's Sunday.[25] Sunday (along with Friday) tends to witness the highest levels of newspaper advertising. The Chicago market is sufficiently large that distributors themselves purchase movies' display ads, and consequently these ads avoid the problem of separating national distributor ads from local exhibitor ads. I measure these ads in column-inches: the product of an ad's width in columns and its height in inches.[26] This measure will capture the scale and variation of national advertising expenditures if studios spend a fixed percentage of their weekly ad budgets on newspaper advertising and a fixed percentage of that on advertising in the *Tribune*. Aside from the appropriateness of these assumptions, it is also likely

[18] While theoretically possible, using both FOCs and equalized profit conditions in the no-entry case is computationally intractable, as the estimation algorithm requires a non-linear search over all 50 rental rates for each week.

[19] Competition among distributors may also be falsely rejected if the equilibrium conditions imply responses that are much larger than those observed in practice.

[20] This Top 50 listing is fairly exhaustive. The average weekly revenue of the 50th highest grossing movie for my sample is about \$50,000.

[21] Throughout my data, weeks begin on Fridays and end on Thursdays. Dates refer to the starting Friday.

[22] U.S. Census Bureau (2001).

[23] Statistics Canada (1998).

[24] This estimate comes from Switzer (2004). It is reached from over a year of weekly data on box office, concessions sales, and concessions costs from a single theater in Arnold, MO.

[25] Presumably distorted by its heightened metropolitan nature, similar data from The New York Times Friday edition were not useful as a national ad proxy.

[26] For example, one column-inch in the Tribune corresponds to about 3 square inches.

that there exists substantial lumpiness of advertising on this measure for highly promoted movies; one-page ads are common, two-page ads are purchased occasionally, but rarely are there any ads in between the two sizes. This advertising measure thus suffers from several potential weaknesses. Results should be interpreted with all of these data caveats in mind.

3.2. Exogenous variables

The full model uses movie-specific fixed effects, but movie-specific characteristics that significantly affect demand are still critical to create effective instruments. Furthermore, second-stage regressions explaining those fixed effect estimates can yield additional confidence in the overall approach. I therefore consider several characteristics that vary across movies but are constant over time. These movie-specific characteristics come from The Motion Picture Guide and the Internet Movie Database. Movie ratings are the clearest such characteristic. I use them to classify G (general audience) and PG (parental guidance suggested) movies in a binary *Family?* variable. A similar binary *Action?* variable is based upon whether a movie is considered action or adventure in either source.

I define an actor or actress's revenue history at the time of a specific movie to be the sequence of final grosses of movies released in the prior five years in which he or she was a starring cast member. Various statistics can then be used to summarize this revenue history, and, in the case of multiple-star movies, other statistics aggregate the several revenue histories and summarize the cast appeal of the movie as a whole. Examining directors is simpler, as movies almost always have a single director.

My version of stardom refers to cast members who *star* in a movie rather than those relatively few individuals who *are* movie stars. Recognizing that distributors have an interest in informing the public about the presence of celebrities, I use movie advertising in the Friday *New York Times* as a guide. Just as theaters of the past used their marquees to advertise star presence, many *Times* ads early in a movie's run include the names of a few members of the cast.[27] I assume that these individuals are the starring cast of those movies (unless certain names are more prominent than others, in which case only the most prominent names are the movie's stars). While a few movies listed many names, most ads that list any cast display only two or three names.

I build my measure of cast appeal by summing up the revenue histories of the marquee cast and dividing by the number of movies released in the prior five years in which those cast had starred. This measure is intuitive in that it tends to be lower when a movie's stars are numerous but not very successful at the box office (e.g., Robert Altman films) and when a starring actor in a movie is prolific but not very lucrative (e.g., post-*GhostBusters* Dan Aykroyd). I measure the appeal of the director by simply summing that individual's revenue history.[28] Both these appeal variables are then scaled in billions of dollars. As an example, 1993's *Jurassic Park*'s director appeal is calculated as the total domestic box office of Steven Spielberg's three movies in the five years prior (*Hook*, 1991; *Always*, 1989; and *Indiana Jones and the Last Crusade*, 1989), or \$0.36B.

Proper incorporation of portfolio effects demands some measure of ownership, even though this ownership does not appear in the mean utility specification or the second-stage regression. I construct ten binary variables to indicate whether a movie is released by one of the notable distributors: Disney, MGM, Paramount, Sony, Fox, Universal, Warner Bros., Orion, Miramax, and Savoy King. The first seven of these distributors compose the majors, and the major-collusion hypothesis is restricted to them. Consistent with Corts (2001), all subsidiaries are collected under a single variable (e.g., Disney encompasses Buena Vista, Hollywood Pictures, and Touchstone Pictures, as well as Disney Pictures).

To this point, the only variables I have described that vary across movies and across weeks are theaters and advertising, but I also consider four exogenous variables of this type. I define *Age* as the total number of weeks that the movie has appeared in the Top 50 listings since the movie's initial entry. The binary variable *Ab?* equals one for the opening week of any movie released on a day other than Friday, and zero otherwise.[29] Academy Award announcements also fall in this category. I use the announcements to create the binary variables *OscNom?* and *OscWin?*. These variables equal one if the movie in question has been nominated or won in any of the six

[27] To maintain tractability, I limit this to printed names and not pictures of cast members. The one exception is for sequels, for which a picture of a recurring character is the equivalent of a printed name.

[28] An exception is Alan Smithee. In 1968, the Directors Guild selected that pseudonym for directors who were unhappy about the final cut of their movie for reasons beyond their control. I consider such movies to be directed by a director with no experience, and therefore they have no revenue history. Of related interest, Smithee's notoriety grew too great, and since 1999 the name has been replaced by unique pseudonyms.

[29] A movie with an abbreviated week therefore has six or fewer days in which to gross its box office for that observation.

major categories (Best Picture, Best Director, Best Actor, Best Actress, Best Supporting Actor, and Best Supporting Actress) prior to that week. The variables therefore represent the impact of the Oscar announcement and are distinct from the movie fixed effect.

3.3. Describing the demand data

The descriptive statistics of these variables can be found in Table 1. Perhaps most striking from the table are the large variances. *Jurassic Park* in its opening week of June 11, 1993, had admissions of over 19 million, while *Life and Nothing But* in the week of December 27, 1990, had 593 admissions. That same disparity is apparent with respect to the number of exhibiting theaters and advertising. *Batman Returns* opened at 3700 theaters the week of June 19, 1992, while many movies showed at a single theater. The number of weekly total theaters shows significant variation. From a time-series regression on months and holidays (not reported), total theaters were significantly higher than a non-holiday January week during April, July through September, and winter holidays, as well as increasing over the sample. As over 40% of observations have no advertising, I provide statistics conditional upon positive advertising as well as unconditional statistics. This concentration of admissions, theaters, and advertising at the top suggests that the convexity in revenues found by earlier studies applies to other aspects of the industry as well. This convexity also poses a severe challenge to estimation since the model must explain both the bulk of observations and the outliers from which most revenues come.

The binary variables capturing Oscar announcements and abbreviated weeks are straightforward, though the small number of observations for Oscar wins (304) and abbreviated weeks (81) suggests that it might be difficult to obtain precise estimates of their impacts. The typical movie in my sample has run in theaters for 7 weeks, but a few outliers running for longer dramatically inflate the mean. From the summaries of variables that capture

Table 1
Summary statistics

	Mean	S.D.	Median	Min	Max
By observation: N = 16,600					
Q (in K)	441.88	1013.53	84.90	0.59	19120.80
Theaters	580	657	297	1	3700
CT Sunday	3.86	7.27	0.00	0.00	241.00
|CT > 0 (34.29%)	11.27	8.41	9.66	0.65	241.00
OscNom?	0.12	0.62	0	0	1
OscWin?	0.02	0.18	0	0	1
Ab?	0.00	0.07	0	0	1
Age	9.00	7.82	7	1	71
CastAppeal (in $B)	0.0198	0.0245	0.0115	0	0.1419
|CastApp > 0 (60.55%)	0.0328	0.0239	0.0282	0.0001	0.1419
DirAppeal (in $B)	0.0442	0.0826	0.0067	0	0.7175
|DirApp > 0 (65.80%)	0.0672	0.094	0.0298	0.00004	0.7175
Action	0.29	0.45	0	0	1
Family	0.25	0.43	0	0	1
Major Studio?	0.75	0.43	1	0	1
By movie: N = 1602					
Ab?	0.0506	0.2192	0	0	1
CastAppeal (in $B)	0.0161	0.0223	0.0049	0	0.1419
|CastApp > 0 (55.68%)	0.0289	0.0228	0.0245	0.0001	0.1419
DirAppeal (in $B)	0.0325	0.0672	0.0019	0	0.7175
|DirApp > 0 (58.55%)	0.0555	0.0802	0.0249	0.00004	0.7175
Action	0.25	0.44	0	0	1
Family	0.20	0.40	0	0	1
Major Studio?	0.68	0.47	1	0	1
By week: N = 332					
Total theaters (in K)	28.98	3.35	28.53	21.87	39.29
Price	4.24	0.1	4.21	4.14	4.5
Unemployment	6.37	0.76	6.44	5.2	7.8
Population (in M)	287.38	5.45	287.55	277.71	296.53

movie-specific variables, one can see that about 5% of movies faced an abbreviated opening week. About 20% of the movie-week observations are categorized as family movies. Approximately 2/3 of the sample's movies are released by one of the major studios, but, given their higher percentage of observations (3/4), these movies evidently stay in the sample longer.

Disparities similar to those observed in admissions and theaters are also apparent when attention turns to the cast appeal variable. Of the observations, 44% have no cast appeal. As with the advertising variable, I provide statistics conditional upon an observation's movie having positive cast appeal. The director appeal variable is also heavily affected by outliers, with 41% of observations having no director appeal.[30] The director appeal measure seems quite plausible, suggesting that directors' past experience may be at least as important as that of the starring cast.

Lastly, I more systematically consider the potentially market-clearing variables. Simple correlations between the three (presumably) endogenous variables and age are presented in Table 2. Advertising and theaters are very positively correlated with admissions as well as with each other. All three endogenous variables also show substantial negative correlations with the age of the movie. The magnitude of these correlations suggests that multicollinearity, in addition to the use of instruments, will hamper precise estimation. Only actual estimation will reveal whether the sample's size and variation are sufficient to overcome these inherent hurdles.

3.4. Cumulative rentals

As discussed in the introduction, national weekly rental payments of exhibitors to studios are generally unavailable. The Internet Movie DataBase, however, displays cumulative rentals for 215 of the movies for which I observed the entire theatrical run. That represents about 14% of the movies for which the full run fell in my sample, but over 25% of the observations (4529 observations). Perhaps not surprisingly, these movies are not representative of the entire sample.[31] While only two-thirds of the full sample's movies were released by one of the seven major studios, over 95% of the movies for which rentals are available came out of the majors. These movies also did far better, accumulating an average of $84 M in box office receipts compared to the $21 M of the full sample. Generally, it appears that cumulative rentals are released for only commercially successful movies. These movies tend to be peak-season releases, as close to 70% were released in summer (May, June, or July) or at the year-end holidays (November or December). They are also likelier to be family movies (26%) than the overall sample (20%). The same trend is even clearer with respect to action movies (46% vs. 25%).

Table 2
Correlations

	Q	Theaters	Ads	Age
Q	1.00	0.71	0.57	−0.26
Theaters	0.71	1.00	0.58	−0.30
Ads	0.57	0.58	1.00	−0.35
Age	−0.26	−0.30	−0.35	1.00

If I were trying to match up the model's predictions for a typical movie's rentals with these observations, this unrepresentativeness would be troubling and perhaps fatal to my approach. Fortunately, the richness of the data and model allow me to consider only the relevant predictions of the model. The most succinct way to consider this cumulative rental data is as a percentage of cumulative box office receipts. The distributor of a typical movie (that released its rentals) could expect to receive 48% of the box office receipts as payment. This closely matches the 50% rule of thumb that is often used in analysis of the industry. Whether this figure is high or low can only be judged in the context of the previous model.

4. Identification, estimation and results

4.1. Identifying assumptions and estimation algorithms

The primary concern in estimating demand is that variables that are omitted by the econometrician will be substantially correlated with an included variable. In the above framework, the segmentation parameters ρ as well as the endogenous variables of theaters and advertising warrant special attention. I address identification with instruments that capture a product's competitive environment, as suggested by Berry (1994). My demand estimates based upon fixed effects then rely upon the following two identifying assumptions:

1) The individual's idiosyncratic addition to expected utility ε_{itj} is drawn from a distribution characterized by McFadden's (1978) assumptions and my own choice of segmentation dimensions.

[30] Listings of the highest ranking movies, cast members, and directors by appeal are available on the web appendix.
[31] Summary statistics for cumulative rentals can be found on the web appendix.

2) The movie-and-week-specific unobserved component of utility ξ is uncorrelated with the strength of that movie's competitive environment.

Assumption 2 supports the use of variables as instruments that will identify the segmentation parameters ρ and the coefficients of the endogenous theaters and advertising. The implicit exogeneity of movie release dates upon which this assumption relies seems at odds with the release-timing game played by studios, but this is lessened greatly when one recalls that movie fixed effects will control for the most obvious examples (e.g., highly anticipated movies released in early summer and at the end-of-year holidays).

Following the literature and especially Bresnahan, Stern, and Trajtenberg (1997), I construct seventeen instruments that capture the strength of a movie's competitive environment as well as the predictable decay in demand as the market saturates.[32] They make use of an age-of-rivals variable and age-weighted cast and director appeal measures of competing movies that are no more than four weeks old, so that a movie facing younger rivals faces a more competitive environment. These variables also exploit the action and family genre definitions in addition to studio portfolio information. Instruments that looked back more than 4 weeks provided no better fit for separate first-stage regressions with $\ln(s_M)$, $\ln(s_A)$ and $\ln(s_F)$ as dependent variables. The matrix of instruments Z then includes these 17 IVs as well as all exogenous demand regressors.

Estimation of demand under the full model follows a standard two-step process. After making a guess for ρ, I invert observed purchase probabilities to obtain the implied mean utilities δ using a contraction mapping. I then find the 2SLS point estimates and their resulting sum of squared errors. The initial guess is updated using a Nelder–Mead simplex search to find the vector of parameters that attains the global minimum of $\xi' Z(Z'Z)^{-1} Z'\xi$. I use these consistent estimates of ξ to construct $\widehat{Z'\Omega Z}$, allowing for arbitrary heteroskedasticity and serial correlation within movies using the Newey–West covariance matrix with three-week lags. I employ $\left(\widehat{Z'\Omega Z}\right)^{-1}$ as the new and efficient weighting matrix when I repeat the exercise. Standard errors for point estimates are calculated using finite perturbation of the non-linear parameters.

The remaining tasks are to transform these demand estimates into the implied rental rates and marginal costs of advertising. All derivatives of the purchase-probability function (e.g., $\frac{\partial s_1}{\partial \delta_2}$, $\frac{\partial^2 s_1}{\partial \delta_1 \partial \delta_2}$) are found using finite perturbation about the observed and inferred values. The no-entry assumption does not yield a closed-form solution for $\frac{dT}{dr}$. For these cases, I use a contraction-mapping based upon the first-order condition to yield the implied rental rates that satisfy this condition with constraints to ensure that rental rates stay on the unit interval. Estimation from various starting points all converged to the same outcome. As stated above, the free-entry assumption generates a linear system of equations in which the 50 rental rates and the expected weekly fixed cost of exhibitors are the unknown parameters. Recalling that for each week I can use both the set of FOCs and no-profit conditions, it is straightforward to create a linear system with cross-equation restrictions: $y - B*r = e$, where y is a 100×1 vector, B is a 100×51 matrix, and r consists of the 50 rental rates and weekly average fixed cost.[33] From this set-up, it is trivial to recover r as $r=(B'B)^{-1}B'y$. I then assume that the disturbances across the FOCs and the zero-profit conditions are independent and separately identically distributed, which allows me to re-estimate using generalized least squares with the efficient weighting matrix. Inferred marginal costs of advertising are derived by assuming the typical pattern of rental rates (70%–60%–50%–40%–30%–30% …) and calculating marginal benefits of observed advertising for all observations with positive advertising. Matlab code for both demand and supply estimation algorithms is available at the web appendix.

4.2. Results

4.2.1. Demand

Table 3 displays the results of the estimation of demand using the PD GEV specification. While a segmentation parameter between action and non-action movies (ρ_A) was originally estimated, results persistently indicated that such segmentation was unimportant (i.e., $\rho_A \approx 1$), and I have imposed that restriction throughout.[34] First considering the segmentation parameters ρ, the PD GEV model reveals that consumer heterogeneity is statistically important both regarding the substitutability between movies and non-movie entertainment and between family and non-family movies. The estimate of $\rho_M \approx 0.54$ is quite close to the

[32] The formal definitions for these IVs can be found on the web appendix.

[33] The y vector, for example, has as its first 50 elements the per-theater purchase probabilities s_i/T_i for that week and as the next 50 elements s_i.

[34] I continue to use the instruments based upon segmentation by Action. Results were similar but less precise when these Action-based instruments were excluded.

Table 3
PD GEV estimates of demand (n = 16600)

	b	a.s.e.
ρ_{Movie}	0.54	0.14
ρ_{Family}	0.80	0.04
log(T)	0.46	0.08
ln(1 + T)ln(1 + ADS)	0.030	0.012
ln(Age)	−0.29	0.09
ln(Age)^2	−0.068	0.018
Nom?	0.064	0.048
Win?	0.095	0.051
Abb?	−0.93	0.11
Implied % of full week	26.21%	
J-stat (Pr(χ^2 < J))	13.13	0.44
R^2 (DV = ln(s_j/s_0))	0.98	
R^2 (DV = δ(ρ))	0.95	
2nd stage movie variables		
CastAppeal	1.43	0.53
DirAppeal	0.80	0.17
Action?	−0.03	0.03
Family?	−0.13	0.03
Constant	−10.50	0.02
R^2	0.045	

Elasticities	Own-theater (η^T)	Own-advertising (η^A|A>0)
Mean	0.68	0.25
Median	0.65	0.28
Min	0.58	0.02
Max	0.91	0.35

* a.s.e. corrected for arbitrary heteroskedasticity and serial correlation of up to three weeks.

$\rho_M \approx 0.5$ of Einav (in press). The fundamental difference between the two appears to be precision, as Einav's standard error is roughly 20% the magnitude of my own. The segmentation parameter regarding a movie's Family-ness (ρ_F) is estimated much more precisely, but its economic significance is questionable $\rho_F \approx 0.8$.[35]

Advertising has a significant positive impact, a relatively rare result when using instruments and aggregate data. From the second-stage regression, the impacts of cast and director appeal are estimated to be positive and relatively precise. The fit is quite low, a problem common to most studies using pre-release exogenous variables in the industry. The average consumer dislikes family movies, but is ambivalent toward action movies. This distaste for family movies is equivalent to a movie having a director with $150 M less of box office experience. Why then do family movies continue to be made? While the typical consumer stays away, the subset of consumers with differing preferences (e.g., parents with small children) purchases. So long as a very limited number of family movies are available at any one time, family movies can be quite lucrative.

An abbreviated week is estimated to yield (on average) about 26% of the admissions that would have come from the hypothetical full week. This is quite close to the 30% result that was generated by the partial reduced form estimates of a similar model in Moul (in press) and is well beneath the 40% upper bound of a Wednesday-before-holiday release suggested by the daily averages of Switzer (2004) and Davis (2006b). Point estimates of the Oscar effects imply an average 9% increase for a nomination and a 14% increase for a win, though the underlying estimates are statistically indistinguishable from zero. Note that these estimates reflect only the direct impact of the Oscar announcements; there may be substantial indirect effects as advertising and the number of exhibiting theaters increase. I include in the web appendix the estimated seasonality of demand (which largely mimics Einav, in press) and cross-theater elasticities of demand for the five highest-grossing movies for several weeks (which illustrate the magnitudes of the segmentation parameters).

Finally, the implied own-theater and own-advertising elasticities can be found at the bottom of Table 3. The own-theater elasticities never exceed 1, and the median falls in the plausible region of $\eta^T \approx 0.65$. Own-ad elasticities are high but still plausible, with the median η^A of 0.28.[36] For comparison, the elasticity found by Nevo (2001) in the demand for ready-to-eat cereal was 0.06 (which the author considered low), and Ackerberg (2001) found an elasticity of 0.15 in the demand for Yoplait yogurt (which the author considered reasonable).

4.2.2. Supply

Before I apply these demand estimates in my supply model, it is worthwhile to consider how well the retail equilibrium conditions hold at the assumed declining rental rate schedule. Holding the rental rate and advertising level constant, one could use the estimated demand parameters and non-linear search methods to find the number of theaters that best satisfy the entry conditions, and these predictions could then be compared to actual numbers of exhibiting theaters. Such an approach requires a non-linear search of fifty

[35] Estimates using the logit restriction (available at the web appendix) suggested that both Oscar nominations and cast appeal were statistically significant negative shocks to demand. More importantly, the logit estimates also implied own-theater elasticities that were too high to implement the supply-side extension, many exceeding one.

[36] These implied elasticties ignore any indirect impact that may arise from the exhibition sector observing advertising and increasing the number of exhibiting theaters.

theater-numbers for each week of my sample, and is thus quite computationally burdensome. This exercise is the most tractable under the exhibition assumptions of no entry and perfect cannibalization.[37] Under this most straightforward case, I reached the following conclusions. First, the entry conditions do a fair job of describing the theater-numbers of the ten highest-grossing movies of a week, and the quality of the description declines as the sample expands to the full Top 50. Second, the equilibrium conditions do a systematically poor job of describing successful art-house movies. Such movies show at a small number of theaters, presumably in New York and Los Angeles, and achieve very high per-theater box office receipts even several weeks into their run when rental rates are assumed to be low. The equilibrium conditions then imply greatly expanding the number of theaters exhibiting the movie. This failure reflects two of the model's less tenable assumptions, namely, a single rental rate and homogeneous theaters. In actual contracts, the rental rate charged to theaters is dependent upon how many weeks the theater has shown the movie rather than the age of the movie.[38] The fact that theaters will differ in their innate demand for movies and probably in their matching to genres is one justification for assuming imperfect cannibalization among theaters.

Another way to approach the appropriateness of the equilibrium conditions is to consider the implied responsiveness of theaters to rental rates and advertising. The implied responsiveness of theaters to rental rates (in elasticities and semi-elasticities) can only be recovered under my method after making an assumption regarding distributor conduct. Results for the median observation are similar regardless of the cannibalization or entry assumption under the portfolio-ignoring and portfolio-recognizing competitive hypotheses. Assuming competing studios that ignore portfolio effects, a 1% increase in a movie's rental rate leads to a 1.54% decrease in the number of theaters showing it and a 10 point increase in the rental rate causes an approximately 34% decline in the number of theaters. Allowing for distributors to recognize portfolio effects mildly raises (in absolute value) the comparable elasticities and semi-elasticities to −1.6% and −36%. Similar calculations for collusion among the major studios are more dependent upon assumptions of entry and cannibalization but are substantially higher, elasticities ranging from −2.2 to −3.2 and semi-elasticities ranging from −0.46 to −0.53. Regarding advertising, I assumed the typical declining rental rate schedule and calculated the elasticities of theatrical intensity with respect to advertising. The median elasticity ranged from 1.3 to 1.6, generally similar across entry assumptions and rising as the extent of retail cannibalization decreased.

These measures of responsiveness seem large, but it is difficult to judge in the absence of prior empirical work. To offer some context, I created appropriate counterparts to the Dorfman–Steiner results for a price-setting and advertising-setting monopolist.[39] That monopolist's objective function is

$$\max_{r,\mathrm{ADS}} \Pi = rpQ(T(r,\mathrm{ADS}),\mathrm{ADS}) - \tau\mathrm{ADS} \qquad (5)$$

Such a framework predicts that a single-movie monopoly studio that faces constant marginal costs of advertising will set rental rates such that the absolute value of the own-rate elasticity of theatrical intensity will equal the inverse of the own-theater elasticity of demand $\left(/\eta_r^T/ = \frac{1}{\eta_T^Q}\right)$. Likewise, advertising should be set so that the ratio of advertising expenditures to distributor rental revenue equals the total (both direct and indirect) own-ad elasticity of demand $\left(\frac{\tau\mathrm{ADS}}{rpQ} = \eta_{\mathrm{ADS}}^Q + \eta_T^Q \eta_{\mathrm{ADS}}^T\right)$. Recalling that the mean own-theater elasticity of demand (η_T^Q) implied by my estimates was 0.68, this suggests an own-rate elasticity of theatrical intensity of about 1.5, which roughly matches the above results when portfolios are ignored. The combination of the result with respect to advertising and the earlier demand estimates are highly suggestive of portfolio effects, either legitimate or through collusion, being important. The mean own-ad elasticity of demand (η_{ADS}^Q) implied by my estimates was 0.08, so the advertising condition suggests the unlikely advertising-to-rental ratio of almost 1(=0.08+0.68 * 1.3). The only rationalization for stopping advertising when such lucrative gains remain is the concern of cannibalizing one's own portfolio.

I begin my empirical tests by examining estimated rental rates under the different hypotheses of distributor conduct and various assumptions of market-clearing in exhibition. To the extent that one "knows" the general rental rate schedule (~50% in cumulative), the question is what hypotheses can best generate these outcomes. Table 4 shows summary statistics of the implied rental rates, as well as the resulting various measures of goodness-of-fit with the observed cumulative rental

[37] The exercise under free entry requires information about weekly theater fixed costs and imperfect cannibalization requires the numerical calculation of $\frac{\partial s}{\partial \delta}$ for each iteration. Even under my preferred assumptions, the exercise required approximately one week of computing time on a 3GHz PC.

[38] On the equilibrium path of broad releases, this should not be problematic, but this is decidedly not an equilibrium-path exercise.

[39] See web appendix for details.

Table 4
Estimated rental rates and explanation of observed cumulative rentals

No Entry		mean	median	s.d.	min	max	Fit
$\theta=0$	Comp	0.40	0.44	0.10	0.10	0.61	0.53
	Port	0.43	0.46	0.11	0.12	1.00*	0.61
	MajCol	0.54	0.56	0.35	0.00*	1.00*	0.65
$\theta=1$	Comp	0.41	0.46	0.10	0.08	0.61	0.23
	Port	0.44	0.48	0.12	0.08	1.00*	0.34
	MajCol	0.52	0.52	0.35	0.00*	1.00*	0.72

Free Entry		mean	median	s.d.	min	max	E (ϕ*ρ*M)	Fit
$\theta=0$	Comp	0.40	0.44	0.10	0.09	0.53	3822	0.51
	Port	0.41	0.45	0.10	0.12	0.69	3768	0.58
	MajCol	0.58	0.60	0.11	0.15	0.92	3063	0.80
	AllCol	0.69	0.70	0.05	0.04	1.20	2575	0.64
$\theta=1$	Comp	0.41	0.45	0.10	0.08	0.54	3780	0.20
	Port	0.42	0.47	0.10	0.08	0.67	3728	0.30
	MajCol	0.59	0.62	0.12	0.13	0.79	2980	0.83
	AllCol	0.71	0.73	0.04	0.57	0.79	2470	0.70

*Percentages of implied rental rates on end point of unit interval:
θ=0, Port 0.024%=1 θ=1, Port 0.036%=1.
θ=0, MajCol 21.37%=0 & 13.69%=1 θ=1, MajCol 24.23%=0 & 12.82%=1.
Notes: Rental rate estimation in No Entry case uses only NLLS on distributor FOCs; estimation in Free Entry case uses feasible GLS on distributor FOCs and equilibrium conditions. Letting u=Rentals-PredRentals and dmr=Rentals-E(Rentals), Fit=1−(u′*u)/(dmr′*dmr).
*All estimates assume concession profitability of $\pi_c=0.29$.
*Cannibalization extent: θ=0 is none, θ=1 is smooth/perfect.

data. The top panel displays the case of no entry. Average rental rates under competitive portfolio-ignoring conduct (Comp) are about 40% under either cannibalization assumption. While such rates seem low, summary statistics could be depressed if movies have long theatrical runs (and consequently many rental rates) but cumulative rentals are largely determined by opening week rates. The explanation of variation of observed cumulative rentals is thus a better metric. Letting y denote cumulative rentals and u denote the difference between predicted and actual cumulative rentals, the measure of fit is $1-\frac{u^2}{(y-E(y))^2}$. Under no cannibalization, the predicted cumulative rentals under competitive conduct capture 53% of the observed variation, and under smooth cannibalization only 23% of the variation is matched. Assuming that distributors consider their entire portfolio of movies (Port) improves these measures somewhat, but not substantially. This lack of improvement is not surprising given the result of Corts (2001): Studios try to space out the releases of similar products to avoid direct competition among their own movies. The conduct hypothesis of collusion among the major studios ends up providing the best fit of cumulative rental payments, matching 72% of the variance under the perfect cannibalization case.

The estimates under the case of free entry in exhibition (shown in bottom panel) are quite similar to those under the no-entry case, though the fits under the collusion-of-majors conduct hypothesis are markedly higher (80% and above). This major-collusion dominates not only the competitive conduct hypotheses but also the hypothesis of collusion among all distributors, major and otherwise. Besides allowing the consideration of the all-colluding hypothesis, the free-entry case also identifies the average weekly fixed cost of operation (Φ). Averaged over the sample, the estimates under major-collusion suggest that exhibitors faced a break-even threshold of about \$3000. At the sample's average admissions price and separately assuming 60% and 30% rental rates, this suggests that each week a theater had to admit about 1000 persons early in a movie's run and 700 persons later in the run to break even. While I have no data on these particular dimensions, such figures seem high, but not implausibly so.

The statistical evidence regarding distributor conduct in the setting of rental rates then seems to favor some collusion over the competitive hypotheses. Distributors, however, set advertising as well as rental rates, and so these conduct hypotheses can also be applied to the advertising margin. Intuitively, studios raise sales with advertising increases through two routes: inducing consumers who had not been planning to see a movie to see their movie and convincing consumers to see their

Table 5
Advertising regressions, DV = ln(MBA|A>0), N=5692

	Comp		Port		MajCol		AllCol	
	b	a.s.e.	*b*	a.s.e.	*b*	a.s.e.	*b*	a.s.e.
No Entry								
θ = 1								
Int	9.31	0.08	9.13	0.08	7.85	0.06	7.32	0.07
ln(Adv)	1.05	0.04	1.08	0.04	1.20	0.03	1.33	0.03
R^2	0.15		0.15		0.12		0.15	
mean(A/R)	0.86		0.77		0.29		0.24	
θ = 0								
Int	9.67	0.08	9.51	0.08	8.02	0.06	7.47	0.06
ln(Adv)	0.98	0.04	1.01	0.04	1.21	0.03	1.35	0.03
R^2	0.15		0.16		0.13		0.16	
mean(*A*/*R*)	1.05		0.95		0.35		0.29	
Free Entry								
θ = 0								
Int	9.31	0.08	9.18	0.08	7.70	0.08	6.71	0.07
ln(Adv)	1.06	0.04	1.07	0.04	1.25	0.04	1.54	0.03
R^2	0.15		0.15		0.13		0.17	
mean(*A*/*R*)	0.88		0.79		0.29		0.23	
θ = 0								
Int	9.68	0.08	9.56	0.08	8.10	0.07	6.98	0.06
ln(Adv)	0.97	0.04	0.98	0.04	1.16	0.03	1.49	0.03
R^2	0.15		0.15		0.13		0.17	
mean(*A*/*R*)	1.03		0.94		0.34		0.26	

Weekly advertising expenditures implied by estimates are aggregated over each movie's run and then divided by sum of expected rental payments (standard rate schedule) over each run. Mean(A/R) is the average over set of movies observed at least four times.
*All std errors reflect White correction.

movie instead of another movie. Collusion among studios in advertising expenditures could therefore allow distributors to reduce such expenditures and the consequent business-stealing. Conditional upon the assumed declining rental rate schedule, the above framework and estimates readily allow the inference of the marginal benefit of advertising for each observation for which advertising is positive, and imposing the advertising first-order condition yields the marginal cost of advertising (MC_{ADS}). To consider the distributor conduct hypothesis on advertising, I next estimate the regression of log-marginal-cost as a function of log-advertising: $\ln(MC_{\mathrm{ADS}}|\mathrm{ADS}>0)=\gamma_0+\gamma_1\ln(\mathrm{ADS})+\nu$.[40] These estimates then imply weekly advertising expenditures for each observation, $TC(\mathrm{ADS})=\left(\frac{1}{\gamma_1+1}\right)e^{\gamma_0}\mathrm{ADS}^{\gamma_1+1}$. I conclude by aggregating these weekly total advertising expenditures and dividing them by the implied rentals under the assumed rental rate schedule for all movies that I observe for at least 4 weeks. While I have no strong priors as to the level of this ratio, studios must pay production and other costs from these rentals so it is unlikely that it would exceed 50%.

Table 5 displays these results under the differing assumptions on exhibition market-clearing and distributor conduct. All estimates suggest diseconomies of scale ($\gamma_1>0$), a perhaps surprising result that reinforces the earlier concern about using newspaper advertising as the measure for advertising. It generally appears that estimates under less competitive conduct hypotheses have lower γ_0 and higher γ_1 and that smooth cannibalization in exhibition generates lower advertising expenditures than the no-cannibalization case.[41] The most important robust result, however, is that the competitive conduct hypotheses imply advertising-to-rentals ratios that are implausibly high (ranging from 0.77 to 1.05). The collusive conduct hypotheses, on the other hand, imply far more plausible ratios ranging from

[40] I use OLS even though I recognize that advertising is endogenously determined. This assumption requires that there are no important unobserved components to the marginal cost of advertising that affect advertising choices. Estimates that used the earlier instruments in 2SLS were unsatisfactory.

[41] Intermediate extents of cannibalization essentially followed a linear transition between the two extremes.

0.23 to 0.35. As with rental rates, the implications are more consistent with the hypothesis of some collusion among the studios than the hypothesis that studios are recognizing the portfolio effects of their actions and competing with one another.

5. Conclusions

While I employed rather imperfect data and the retail entry conditions may be somewhat onerous here, the theatrical movie industry provided a suitable situation to test my technique for using retail equilibrium conditions to pin down retailer non-price responses to wholesaler actions. My technique allows me to draw conclusions about distributor conduct in the industry that have heretofore been prevented by insufficient data. Several caveats apply, but results are at least suggestive that the industry has found ways to engage in some form of collusion to limit payments to exhibitors and advertising that is excessive to the industry. These results are largely robust to my choice of very different assumptions that characterize retail equilibrium, a finding that makes future applications of this technique more promising.

Keeping in mind all of the caveats and the model's strong assumptions, the results' suggestion of collusion has several implications. First, it suggests that less than efficient resources are dedicated to exhibition, implying that there are too few exhibitors and leading to consumers traveling excess amounts to reach theaters. Second, the industry may have found a way to have total advertising be closer to the social optimum (which also does not value business-stealing), though this raises the possibility that too little advertising is purchased, as studios neglect to consider consumer surplus in their decisions. Last, it suggests that studios compete primarily in movie quality; if this competition is reflected in higher prices for inputs (cast, directors, etc.), then these individuals of the creative process are important beneficiaries of this collusion. Further empirical work on this margin to check the robustness of this technique and these results is clearly desirable.

References

Ackerberg, Daniel A., 2001. Empirically distinguishing informative and prestige effects of advertising. RAND Journal of Economics 32 (2), 316–333.

Berry, Steven, 1994. Estimating discrete choice models of product differentiation. RAND Journal of Economics 25, 242–262.

Blumenthal, Marsha A., 1988. Auctions with constrained information: blind bidding for motion pictures. Review of Economics and Statistics 70, 91–198.

Bresnahan, Timothy F., Stern, Scott, Trajtenberg, Manuel, 1997. Market segmentation and the sources of rents from innovation: personal computers in the late 1980s. RAND Journal of Economics 28, S17–S44.

Corts, Kenneth, 2001. The strategic effects of vertical market structure: common agency and divisionalization in the US motion picture industry. Journal of Economics and Management Strategy 10, 509–552.

Davis, Peter, 2006a. Measuring the business stealing, cannibalisation and market expansion effects of entry in the motion picture exhibition market. Journal of Industrial Economics 54 (3), 293–321.

Davis, Peter, 2006b. Spatial Competition in Retail Markets: Movie Theatres. RAND Journal of Economics 37, 964–982.

Einav, Liran, 2007. Seasonality in the U.S. Motion Picture Industry. RAND Journal of Economics 38 (1).

Elberse, Anita, Eliashberg, Jehoshua, 2003. Demand and supply dynamics for sequentially released products in international markets: the case of motion pictures. Marketing Science 22 (3), 329–354.

Filson, Darren, Besocke, Portia, Switzer, David, 2005. At the movies: the economics of exhibition contracts. Economic Inquiry 43 (2), 354–369.

Grant, E. (Eds.), 1991–97. Motion Picture Guide, The. CineBooks Publishing, New York.

Kenney, Roy W., Klein, Benjamin, 1983. The economics of block booking. Journal of Law and Economics 26 (3), 497–540.

Kenney, Roy W., Klein, Benjamin, 2000. How block booking facilitated self-enforcing film contracts. Journal of Law and Economics 43 (2), 427–435.

Krider, Robert E., Weinberg, Charles B., 1998. Competitive dynamics and the introduction of new products: the motion picture timing game. Journal of Marketing Research 35, 1–15.

Litman, Barry R., 2001. Motion picture entertainment. In: Walter, Adams, Brock, James (Eds.), The Structure of American Industry, 10th edition. Prentice Hall, Inc., Upper Saddle River, NJ.

Manuszak, Mark D., 2007. The Impact of Upstream Mergers on Retail Gasoline Markets. mimeo, Federal Reserve Board of Governors.

McFadden, Daniel, 1978. Modeling the choice of residential location. In: Karlqvist, A., et al. (Eds.), Spatial Interaction Theory and Planning Models. North Holland Pub. Co.

Moul, Charles C., in press. Measuring Word of Mouth's Impact on Theatrical Movie Admissions. Journal of Economics and Management Strategy.

National Association of Theatre Owners, 2001. Statistics: average U.S. ticket prices. The Encyclopedia of Exhibition. Available at http://www.natoonline.org/statisticstickets.htm.

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

Orbach, Barak Y., Einav, Liran, 2007. Uniform prices for differentiated goods: the case of the movie-theater industry. International Review of Law and Economics 27, 129–153.

Ravid, Abraham S., 1999. Information, blockbusters, and stars: a study of the film industry. Journal of Business 72, 463–492.

Shaffer, Greg, 1991. Slotting allowances and retail price maintenance: a comparison of facilitating practices. RAND Journal of Economics 22 (1), 120–135.

Statistics Canada, 1998. Population and dwelling counts, for Canada, provinces and territories, 1991 and 1996 censuses — 100% data. Catalogue No. 93-357-XPB. Available at http://www.statcan.ca/english/census96/table1.htm.

Stigler, George, 1968. A note on block booking. The Organization of Industry. Irwin Pub. Co.

Sudhir, K., Rao, Vithala, 2006. Do slotting allowances enhance efficiency or hinder competition? Journal of Marketing Research 18 (2).

Sullivan, Mary W., 1997. Slotting allowances and the market for new products. Journal of Law and Economics 40 (2), 461–493.

Switzer, David, 2004. The Economics of Movie Exhibition, unpublished doctoral dissertation, Washington University.

U.S. Census Bureau, 2001. Monthly estimates of the United States population: April 1, 1980 to July 1, 1999, with short-term projections to November 1, 2000. Population Estimates Program, Population Division. Available at http://www.census.gov/population/estimates/nation/intfile1-1.txt.

Villas-Boas, Sofia, 2007. Vertical relationships between manufacturers and retailers: inference with limited data. Review of Economic Studies 74 (2), 625–652.