Seasonality and Competition in Time: An Empirical Analysis of Release Date Decisions in the U.S. Motion Picture Industry*

Liran Einav
Stanford University
leinav@stanford.edu

August 12, 2002

Abstract

Strong seasonality in demand, a short product life cycle, and the absence of any price competition make the release date of first-run movies one of the main strategic decisions taken by movies' distributors. These endogenous timing decisions, in turn, generate a strong seasonal pattern of release dates. Therefore, the observed seasonal pattern of sales is a combination of both seasonality in underlying demand and seasonal variation in the quality of movies released. In analyzing the strategic timing game over the release dates of movies one ought to take the endogeneity of release dates into account and break down the observed seasonal pattern of sales into these two component. The absence of good observable measures of product quality creates a difficulty in separately identifying the two layers of seasonal pattern. In this paper I identify these layers by estimating weekly demand for movies, using movie fixed-effects, a long panel of movies' weekly revenues, and reasonable restrictions on their decay pattern. I find that the estimated seasonal effects in underlying demand are much smaller and quite different from the observed seasonal pattern of sales. I then use the demand estimates to construct and to structurally estimate a timing game played by distributors. The results suggest that the perceived (by distributors) seasonal pattern in underlying demand that rationalizes the observed release pattern is most closely related to the observed seasonal pattern of sales and *not* to the estimated underlying demand. This implies, for example, that distributors may significantly increase their revenues by pushing some of their Memorial Day releases to August.

KEYWORDS: Motion pictures, seasonality, spatial competition, discrete choice framework
JEL classification: C13, C51, L13, L15, L82.

*This paper is a revised version of chapter 1 of my 2002 Harvard University dissertation. I am particularly grateful to Ariel Pakes for his invaluable support and comments. I also thank Bharat Anand, Nittai Bergman, Estelle Cantillon, Guillermo Caruana, Richard Caves, Ken Corts, Yair Eilat, Julie Mortimer, Aviv Nevo, Marc Rysman, Roni Shachar, Minjae Song, Manuel Trajtenberg, and seminar participants at Boston Univesity, Harvard University, MIT, NBER, NYU, Northwestern University, Princeton University, Stanford University, UC Berkeley, UCLA, University of Chicago, and Yale University for helpful comments and suggestions. I thank Jeff Blake of *Sony Pictures*, Ron Kastner of *Goldheart Pictures*, and Terry Moriarty of *Hoytz Cinemas* for insightful conversations, and *ACNielsen EDI* and *Exhibitor Relations Inc* for helping me obtain important portions of the data. Some parts of this paper draw upon an earlier draft, written jointly with Nittai Bergman. The usual disclaimer applies.

"*... a very serious game of strategy is at work–a cross between chess and chicken–which studio distribution chiefs play year round, but with increasing intensity during the summer and holiday release period.*" (The New York Times, December 6, 1999)

1 Introduction

Most empirical industry studies focus on price and quantity competition, taking other product characteristics as given. In many industries, however, prices play very little role, and competition is channeled through other product attributes. The entertainment industry is a prime example of such a case; competition among movies, television shows, or Broadway shows is on non-monetary product attributes, such as content, advertising, and time. The absence of good observable measures of product quality often complicates the analysis of competition in such industries. Moreover, this raises the need to endogenize some of the product characteristics–especially those that can be changed in the short run–just as we endogenize prices in traditional industries. This paper provides a framework in which unobserved characteristics can be handled, release decisions can be endogenized, and non-price competition can be analyzed in the context of motion pictures.

The number of Americans who go to the movies varies dramatically over the course of the year, and sometimes more than doubles within a period of two weeks. At the same time, the first week accounts for almost forty percent of the box office revenues of the average picture. The combination of these two facts makes the timing of releasing a new picture a major focus of attention for distributors of movies. With virtually no subsequent price competition, the movie's release date is one of the main short-run vehicles by which studios compete with each other.

As in many other situations of entry, product positioning, or location choice, the release date decision is a trade-off between market size and market competition. However, the motion picture industry poses an added difficulty in analyzing the release date decision. In other contexts of entry or spatial competition, proxies for market size are readily available, so variation in demand across the geographical (or other) space can be quite accurately predicted by the players and by the econometrician. This is not the case for the motion picture industry, where it is hard to find good measures of market size which are independent of the intensity of market competition, or good measures of product quality which are independent of market size. Thus, when analyzing the release date decision one has to make assumptions about the variation in demand across the different weeks, as perceived by the decision makers. This identification problem adds another focus of attention to this paper: the analysis of seasonal effects in the industry. I use this analysis to investigate the release date decisions. As it turns out from the results, not only is this identification problem faced by the econometrician, but it may be also faced by the decision makers themselves.

In many industries the set of available products is relatively fixed over time, so seasonal variation in sales is typically interpreted as variation in underlying demand.[1] Movie distributors, however, tend to release their biggest hits in the beginning of the summer and during the Christmas-winter holiday season, making the choice set of moviegoers vary quite dramatically over the year. Thus, we face a classic identification problem: is the strong box office performance of, say, Thanksgiving weekend the result of higher demand, better movies, or a combination of both? When the set of available products varies, one could potentially

[1] This is true as long as market prices do not react to the seasonal changes in demand. Once they do, a deeper analysis may be needed in order to truly identify the underlying demand.

control for the set of available products in order to identify the underlying demand. Such an exercise is not feasible in the motion picture industry; with very weak predictors for a movie's box office success, simply controlling for movie quality is not feasible.[2]

In order to properly analyze the choice of release dates, the seasonal pattern of industry sales has to be broken down into its two components: the part which is due to changes in underlying demand, fixing the choice set, and the part that is due to changes in the choice set, fixing the underlying demand. This breakdown is crucial for the release date decision. If the seasonal variation is mainly explained by changes in underlying demand then the distributor of the marginal movie should aim at releasing it in a high sales period, i.e. when demand is high. If, in contrast, the seasonal variation is mainly explained by changes in movie quality then the marginal movie would best be released in a low sales period, i.e. when competition is soft.

Accordingly, in this paper I analyze the release date decision in two separate steps. First, I break down the seasonal pattern of sales into its two components: the seasonal pattern in underlying demand and the seasonal pattern in market competition. Second, I analyze the release date decisions in light of the demand estimation results, by constructing a strategic timing game and taking it to the data. I present each step in turn.

In the first part of the paper I estimate demand for movies, focusing on the breakdown of the seasonal pattern of sales into its two components. This is done by estimating a discrete choice model of weekly demand for movies. The identification of the seasonal pattern in underlying demand can not be achieved through variables that control for movie quality, which arguably do not exist. Rather, I identify the seasonal pattern by adding movie fixed effects and assuming that the decay in the average utility to consumers from a given movie is independent of the movie's release date. Conceptually, one can think of this as estimating a measure of movie quality, and at the same time using it to control for the variation in the set of available movies. The assumptions about the decay pattern allow me to directly compare the quality of movies that overlap in their screening period, even though they were released in different weeks. This, in turn, makes it possible to compare movies with different release dates and to control for relative movie quality, leaving the residual seasonal variation explained by seasonal effects in underlying demand. Identification comes from observing movies throughout their run, playing during weeks with different demand levels, and competing against different sets of movies. While a higher quality movie may be released in a week with higher demand, it still runs for the following weeks, in which seasonal demand will vary. To do this, I exploit a rich panel of weekly box office revenues of *all* movies released nationwide in the U.S. from 1985 to 1999. The estimation results imply that about half of the seasonal variation in sales can be attributed to seasonal changes in average quality of movies. Moreover, the seasonal pattern in the estimated underlying demand is quite different from the observed pattern of sales. Roughly speaking, I find that the estimated high demand seasons are shifted forward in time, by about a month, in relation to the periods of high observed sales.

There are many other papers which have estimated demand for motion pictures.[3] In the absence of good predictors for the box office success of movies, the vast majority of the existing literature conditions on the first week movie's revenues and investigates the effects of various variables, such as advertising, reviews, and Academy Awards, on the changes in revenues pattern over the movie's subsequent life cycle. In using movie

[2] Throughout the paper we use the word *quality* to mean a measure of a movie's box office appeal or potential, independent of its release date. This, of course, may or may not coincide with the movie's cinematographical quality.

[3] See, among others, Eliashberg and Shugan (1997), Mulligan and Motiere (1994), Nelson et al (2001), De Vany and Wells (1996, 1997), Prag and Casavant (1994), Sawhney and Eliashberg (1996), Vanderhart and Wiggins (2000).

fixed effects I follow a similar approach, but only as a tool of the analysis rather than its main objective. To the best of my knowledge, this is the first paper to recognize the potential difficulty in the interpretation of the seasonal effects in the industry and to explicitly address it.[4] To do so, I assembled a sample which is much longer than those typically used in the literature (fifteen years compared to two to three years), and I employ discrete choice estimation methods that allow me to control for changes in competition and seasonal effects in a more systematic way. The key identifying assumption, using the decay pattern as described above, is a novel contribution of this paper.

In the second part of the paper I analyze the observed release date decisions. I do so in several steps. First, I describe the dynamic process which leads to the eventual release schedule using a unique data set of pre-release announcements that I have assembled. When movies enter into the production process, distributors set a broadly defined time window as a target. As production progresses and the planned release season becomes closer, the time window for release narrows, until finally a specific release date is publicly announced. This takes place, on average, about five to six months in advance. However, distributors often change the announced release dates. These changes are mainly done for strategic reasons, and in most cases do not shift the release date of the movie by more than a month in either direction. This final release date decision is the focus of this part of the paper. I take the broader release season as well as the characteristics of the finished film as exogenously given, and concentrate on the choice of the release date within the specified season.

Then, I use a reduced form analysis and a simplifying "perfect competition" assumption to find an approximation for the perceived seasonal pattern in underlying demand. By using this approximation, I show that the observed release pattern cannot be rationalized by the estimated pattern in underlying demand. However, the biggest movie in each week collects about twenty five percent of the weekly industry revenues on average, so the perfect competition assumption has to be relaxed and the strategic considerations must be accounted for. Therefore, I continue by constructing a simple timing game with private information, in which movies sequentially announce and commit to their scheduled release dates. Conditional on the private information, the game has a unique Perfect Bayesian Equilibrium. The distribution over the unobserved private information allows for a positive probability over all potential release schedules, so the model can be estimated using Maximum Likelihood. The game is simple yet general enough to allow for a wide set of potential parameterizations of either the payoff function or the commitment opportunities. Given the estimates of the demand system, I use this structure to test between different possibilities for distributors perceptions regarding the seasonal pattern in demand, when making their release decisions. The results suggest that the observed pattern in industry sales is the one that most closely resembles the perceived (by distributors) seasonal pattern in demand. This means that distributors, when making their release date decision, use the observed pattern in industry sales as a proxy for industry demand, rather than the estimated pattern in underlying demand, which is very different. These results are consistent with the spirit of the popular media that cover the industry, as well as with other industry sources (e.g. Vogel (1994)), who keep focusing on total box office revenues when describing the seasonality in the industry. In the final section of the paper I discuss various reasons that may lead to such an estimated pattern in release decisions. Among those are uncertainty, repeated game effects, bounded rationality, and agency problems.

I make two additional important remarks. First, with the vast majority of movies released on Friday and more than two thirds of the weekly box office revenues collected during the weekend, the week is the unit

[4] Other papers that deal with the seasonality of the industry look only at the aggregate seasonal pattern in box office revenues (see Radas and Shugan (1998, 2001)).

of analysis for the release date decision. Because there are only several weeks within a season, the release decision is a discrete one. This discrete nature of the decision relates this paper to the fast growing empirical literature on entry, which structurally analyzes discrete games.[5] However, in this context, the competition in time is more closely related to discrete spatial competition, and my general estimation framework most resembles the one used by Seim (2000) to analyze discrete location choice of video rental retailers.[6] This literature, however, treats all players in the industry in a symmetric fashion. An important distinction of this paper is that it allows each movie to have its own identity, which is also taken into account by its potential competitors. This additional flexibility of the model, however, does not come without cost; it makes the dimensionality problem more difficult, significantly restricting the scope of the analysis for computational reasons.

Second, the motion picture industry is not the only example in which timing decisions play a central role. Similar considerations are important in the release decisions concerning books, compact disks, and other new products, as well as in the scheduling of major events, television programming, flight schedules, and promotional sales.[7] The combination of three important characteristics, however, makes motion pictures a distinctive member of this set: (i) Ticket prices (and contracts) are roughly constant, so there exists no second-stage price competition, leaving the release date as the key strategic decision;[8] (ii) The short product life cycle increases the importance of choosing the right release date; (iii) Many new movies are released annually, making the analysis more tractable for the empiricist. Still, the methods of analysis used in this paper may prove useful for these other industries as well.

The rest of the paper continues as follows. In the next section I describe the structure of the industry and the interaction among its players. Section 3 describes the data set, and Section 4 estimates weekly demand for movies, breaking down the aggregate seasonal pattern into its two components. Section 5 analyzes the release date decision, Section 6 discusses the findings, and Section 7 concludes.

2 Industry Description

The motion picture industry has been rapidly growing in the past decades, surpassing annual domestic box office revenues of seven billion dollars in 1999. On average, since 1985, more than 150 movies were released annually nationwide. Table 1 presents the general trends in industry key variables over the sample period. I begin with a brief survey of the structure of this industry.[9]

The Motion Picture Industry is comprised of three main players: producers, distributors, and exhibitors. As can readily be understood from their names, producers are in charge of all aspects relating to the production of the movie, distributors deal with the nationwide distribution of the completed movie, and exhibitors are the owners of the theaters to which consumers go to see the movie.

[5] See, for example, Berry (1992), Mazzeo (2000), and Toivanen and Waterson (2000), as well as Reiss (1996) for a review.

[6] As will become clear later on, there are two notable differences between the second part of this paper and Seim (2000). First, I use a sequential game rather than a simultanoues move game. Second, I do not model the entry decision, which is taken as given. These simplifies the computational burden of the estimation.

[7] There are only few papers that look at competition in time. They use reduced form statistics to assess the equilibrium outcome (see Borenstein and Netz (1999), Corts (2001), and Chisholm (1999)). Goettler and Shachar (2000) construct a strategic scheduling game between television networks, but do not use it for estimation due to dimensionality limitations.

[8] The fact that ticket prices hardly varies across seasons and movies is taken as given throughout this paper. This is an interesting puzzle and is discussed in Einav and Orbach (2001).

[9] Much of the following is adapted from Caves (2000), Verter and McGahan (1996), and Squire (1992).

The structure of the industry has undergone a number of changes since Thomas Edison's invention of the Kinetoscope in 1888. As early as 1908 the industry was controlled by the Motion Picture Patents Company, otherwise known as "The Trust", which enjoyed exclusive access to the supply of raw film, and held the major patents for camera and projector equipment. Antitrust suits dissolved the Trust by 1917. Shortly afterwards the industry entered a wave of vertical integration. By the 1930's the industry was completely controlled by the "Five Majors"–Paramount, Loew's, RKO, Warner Brothers, and Twentieth Century Fox–which were fully vertically integrated, and the "Three Minors"–Columbia, Universal, and United Artists–which were not involved in exhibition. In 1946, following the antitrust case of *United States vs. Paramount Pictures*, the industry was once again disintegrated. In addition, the court prohibited a number of common practices used in the negotiations between distributors and exhibitors who were not integrated. These included "block booking"–a practice in which a group of films are licensed simultaneously–and "circuit deals"–a practice in which an exhibitor licenses a film to a group of theaters simultaneously using one contract. Later, "product splitting"–a practice in which distributors or exhibitors explicitly cooperate in splitting the market–was prohibited, and blind bidding–a practice in which exhibitors bid for films before seeing them–was partially restricted. All contracts between distributors and exhibitors are now written on a movie-by-movie and theater-by-theater basis. Thus, even theaters owned by the same exhibitor write separate contracts for the same movie.

The current industry structure is summarized in Figure 1. The industry is dominated by the major studios that have integrated production and distribution. These include such companies as Disney, MGM, Paramount, and Universal. In addition, there are a number of small independent producers ("independents") who use either the major studios' distribution division, or one of a number of independent distributors. Finally, with few exceptions, exhibitors are generally not vertically integrated with producers or distributors. Table 2 summarizes the annual market shares of the top distributors, showing that they consistently account for more than 90% of the industry revenues over the sample period. It also shows, however, that there is no clear market leader, and that the ranking of the top distributors varies from year to year. This is mainly because a few hit movies account for much of the industry box office revenues.

2.1 Movie Distribution

I now examine in greater detail the process by which a movie arrives for viewing in the theater after its completion. The first stage in this process is distribution. The main decisions included in this involve setting the release date, deciding about the initial scope and locations of the release, negotiating contracts with the exhibitors, and designing the national advertising campaign surrounding the movie.

One of the main strategic decisions made by distributors is that regarding their movie's release date. The two important considerations factored into this decision are the strong seasonal effects in the demand for movies, and the competition that will be encountered throughout the movie's run. Typically, movies with higher expected revenues are released on higher demand weekends, so there is a trade off between the seasonal and the competition effects. The importance of the release date is greatly magnified by the fact that the performance during the first week accounts for a sizeable amount of the overall performance of the movie. On average, box office revenues in the first week account for almost forty percent of the total domestic revenues (see Figure 2).[10] An additional reason for the importance of a judicious choice of a release date is the view that high revenues in the first week create information and network effects which increase revenues

[10] Furthermore, about seventy percent of the weekly revenues are collected in the weekend.

in subsequent weeks.[11]

Figure 3 presents the strong seasonality in the industry, plotting weekly average industry revenues (normalized by ticket prices and the size of the U.S. population). Major holidays such as Memorial Day, Fourth of July, Thanksgiving, Christmas, and New-Years are historically strong days for box office performance. Consistent with the figure, the conventional wisdom is that box office revenues are strong throughout the summer season and during the Christmas-winter holiday period. The period following Labor Day up to mid-November is considered to be very weak, as is the period from the beginning of March to mid-May.[12]

The identity of the competing movies is the second consideration taken into account when setting the release date. Distributors are wary of releasing a movie in close proximity to another movie with which competition will be strong due to factors such as similar genres, similar target audiences, or very disparate movie qualities. Furthermore, even once release dates are set, distributors often change them in response to new information concerning release dates of similar movies chosen by other distributors (see in more detail Section 3.2). Another strategy practiced by distributors is to announce their movie's release date early on in the hope that this preemption will result in other distributors avoiding the announced release date. This is especially common for movies with a general consensus as to their strong success.

Another decision made by the distributors is that concerning the scope of a movie's release. There are three main types of release: Wide release, Platform release, and Limited release. Wide releases, which are the most common with the main distributors, are those in which screening of the movie begins in a large number of theaters, typically several thousands, accompanied by a very extensive national advertising campaign. Platform releases involve an initial release in a smaller number of theaters, often only in big cities, with advertising concentrated more in local newspapers. In subsequent weeks the movie expands to additional screens and to more rural areas. Limited releases are those in which the movie is released in two or three cities–typically New York, Los Angeles, and occasionally Toronto–without strong expectations of the movie's potential for a wider release. Distributors typically use Platform and Limited releases with movies that are not considered to have obvious immediate appeal to mainstream audiences, for example because the movies' actors are unknown or the subject matter is difficult. The conventional wisdom is that a gradual release creates a buzz by word-of-mouth which is deemed to be necessary for the success of these movies. In most cases this is followed by a gradual expansion to a wide release within two to four weeks of the initial screening. Thus, these releases can be viewed as an alternative method of advertising prior to a movie's wide release. An additional benefit is the reduction in the risk facing the distributor due to the additional opportunity to forego a wide release upon receiving more interim information.[13]

2.2 From Contracts to Exhibition

Contracts negotiated between distributors and exhibitors are fairly standard. A typical distribution contract involves the theater paying the distributor whichever is greater, the fee calculated by the "adjusted gross receipt formula" or the fee calculated by "the gross receipts formula". According to the former, the distributor receives up to 90% of the box office revenues after a deduction for the theater's expenses (the "house nut"). According to the gross receipts formula, the distributor receives a percentage of the weekly box office revenues, with this percentage initially set to around 70% and then declining throughout the movie's

[11] To quote from Lukk (1997): "*In this business, if you are not the number one film the week you are open, you usually are never the number one*".

[12] The validity of this conventional wisdom will be later questioned.

[13] Most films produced by independent producers use limited releases.

run. This percentage decline is a mechanism to curb the exhibitor's interest in prematurely terminating the screening of a movie in favor of a new one. Finally, it should be noted that by law the contract cannot stipulate the admission price.

As the movie release date draws near, the distributor begins with the national advertising campaign. This includes a wide variety of mediums such as televised trailers, billboard advertisements, newspaper and magazine advertisements, and theater trailers.[14] The advertising campaign is meant to be structured in such a way so that its peak coincides with the release of the movie. Generally, distributors tend to use rules of thumb, setting their advertising budget as a fixed percent of the movie's production cost.

Prints (copies) of the movie are made by the distributor, each costing around $3,000, and are shipped out to the theaters shortly before the release date. For a wide release of 2,000 theaters this amounts to a cost of $6 million, a non-negligible amount compared to an average production cost of about $20 million. Thus, somewhat surprisingly, the costs of the copies are an important factor in deciding the scope of a movie's release.[15]

The movie is then screened at the theaters. Revenues are divided according to the contract described above. If not specified otherwise, the exhibitor makes the decision when to terminate the screening of a movie, thereby releasing the screen for a more successful film. However, given the long-term relationship between distributors and exhibitors in the industry, the distributor may choose to punish exhibitors that have been deemed to prematurely conclude the screening of a film by denying access to future films. Typically, a theater screens a movie between six to eight weeks. After the conclusion of first run screening, some movies are transferred to theaters specializing in second runs, which generally charge admission at around half the price of first-run theaters.

Nowadays, domestic box office revenues account for only about 15% of the movie's revenues.[16] Additional revenues are obtained from the selling of movie screening rights to cable and television networks, from the video and DVD rental markets, and from the international box office market. However, the domestic box office revenues set the tone for the ancillary markets–higher box office revenues in the domestic market lead to subsequent higher revenues. Thus, maximizing domestic box office revenues is a good approximation for the objective function of distributors.

3 Data

My analysis employs two sets of data. The first includes all the relevant data about movie revenues and movie characteristics, while the second contains information about the timing and the changes in release date announcements before the actual release. Most of the empirical analysis in this paper is based on the first set of data. The second set of data is unique and provides an important motivation for some of the main assumptions used in the analysis of the release date decision.

[14] Typically, almost all marketing expenditures (even in the local level) are funded by the distributor. Exhibitors spend very little on advertising.

[15] The industry is now in the beginning of a technological change towards digital filming and distribution. Among other consequences, this change will eliminate the cost of prints, leaving such considerations obsolete.

[16] Down from about 35% in the early 1980's.

3.1 Main Data Set

The data set covers *all* movies domestically released between 1/1/1985 and 12/31/1999 (covering 790 weeks). Compared to the coverage of two to three years prevalent in the existing literature, this wide coverage is key in allowing to pin down the seasonal effects in the industry. For each movie released, the data set includes the official release date, the total box office revenues, an estimate of the production cost (with only 85% coverage), total advertising expenditure (about 90% coverage), the distributor, the genre, the MPAA rating, data about Academy Awards, and the total run time of each movie in the sample. In addition, for the first ten weeks following the official release date of each movie, the data set includes weekly box office revenues, weekend box office revenues, and weekly number of screens on which the movie was run.[17] The majority of the data was obtained from *ACNielsen EDI*, with the exception of the advertising data which was obtained from *Ad $ Summary*, published annually by *Competitive Media Reporting (CMR)*.

Because of the long sample period, I deflate the box office revenues to accommodate trends in the average ticket price and in the total market size, the latter taken as the U.S. population. I obtain annual average ticket prices from the *Encyclopedia of Exhibition*. Weekly ticket prices are extrapolated from the annual ticket price schedule by assuming that prices increased linearly throughout the year.[18] Weekly figures for the population of the United States are extrapolated in a similar way, with annual figures taken from the U.S. Census. Weekly market shares for each movie are calculated by first dividing weekly revenues by the weekly ticket price, and then further dividing by the weekly population size.

The initial sample is comprised of 3,523 movie titles. I omit from the data set all movies that did not reach a screening of 600 screens throughout their run, which is considered to be a wide release.[19] The average distribution of an omitted movie was 100 screens, and the average estimated production cost was $2.9 million. These are therefore relatively small movies, which appear to comprise a different segment of the industry. After this deletion 2,063 movies remain. Of these, I omit additional 97 titles that went through a relatively long platform release, and never reached a wide release within 10 weeks of their official release date. I also omit from the data the observations on the first nine weeks of 1985, for which I am missing the data on the movies released in the final weeks of 1984.

For the cases of platform and limited releases, I consider the actual release date to be the first week in which the number of screens is high enough.[20] This procedure allows me to disregard weeks in which a movie was shown in a very limited number of theaters. The interpretation of viewing the initial weeks of a gradual release as another method of advertising lends itself directly to justifying this approach. Of course, for wide releases the official and the designated release date coincide, which is the case for more than 90% of the titles used in the analysis.

The final data set I use in the subsequent analysis is comprised of 1,956 titles and 16,103 weekly observations. These titles account for 94% of the revenues of the 3,523 titles that constitute the full original sample. On average I observe slightly more than eight weeks per movie, rather than ten. This is for two

[17] The first ten weeks of a movie's run account on average for more than 90% of its eventual total box office revenue.

[18] Since yearly changes in prices are quite small, this assumption is not very restrictive.

[19] Corts (2001) uses a similar cutoff point. This also seems reasonable given the distribution of the peak number of screens across movies: bimodal, with 600 falling between the two modes.

[20] For a typical movie, it is very easy to see when the number of screens discontinuously jumps, marking the first week in which the movie is shown nationwide. The exact "cleaning" procedure defines the actual release to be the first week in which the number of screens showing the movie exceeded a benchmark. The benchmark is the maximum between 400 screens and 30% of the (eventual) maximal number of screens showing the movie.

reasons. The first, accounting for about 80% of the truncated observations, is standard attrition of movies that were dropped from theaters before their tenth week. The second is a truncation of movies that were platform released, thus having less than ten weekly observations since their first wide release. Throughout the analysis I use all observations (i.e. the analog to an unbalanced panel) and verify that the results are not sensitive to the truncation problem.

Release dates of each movie in the subsample are tabulated at the weekly level. Thus, each movie's release date was associated with the week in which the movie was released. Since over 75% of the movies were released on Fridays, and an additional 20% on Wednesdays, this simplifying assumption is hardly restrictive. While there are approximately 52 weeks in a calendar year, I create 56 weekly dummies. Counting from the first week in the year, American holidays can shift by up to one week from year to year, due to the way in which they are determined.[21] Given the seasonal patterns I am looking at, which are mainly focused on holiday weekends, it is important to maintain the same week numbering for the same holiday, and adjust for the calendar differences across years. Thus, I introduce in a small number of years some additional weeks as "fillers". For example, according to this arrangement, Labor Day and Thanksgiving always fall on weeks 38 and 51, although there might be 11 or 12 weekends between them, depending on the calendar year. Weeks 6, 13, 35, and 42 are those that do not appear in all years.

Table 3 provides some descriptive statistics for the final sample, along with a few examples of the leading movies for each variable. The mean and standard deviation of the total box office revenues are $34 million and $43 million respectively. The median is less than $20 million, showing the skewness of the distribution. The corresponding population shares reveal a similar picture. It is interesting to note that the highest population share for a movie in the sample, that of Titanic, had a cumulative ten weeks market share of more than 30%. The production cost figures show a similar pattern, while the advertising distribution is more symmetric. Finally, Table 1 exhibits the progression of the industry share (annual admissions over U.S. population) and other variables over the observation period. It clearly demonstrates that the industry has been growing since the mid 1980's.

Turning to the main focus of analysis–the seasonal pattern–one should first refer to the seasonal pattern in industry sales, which is depicted in Figure 3. This is the typical picture that the popular media portrays when discussing seasonality in the industry. The year is typically thought of as consisting of four periods: the summer period (roughly, Memorial Day to Labor Day), the Holiday period (Thanksgiving to mid January), the winter/spring, and the fall period. The first two are generally thought of as high demand periods, while the other two are considered to be weak periods. The conventional wisdom is that distributors, having this in mind, scramble to release their top movies in the beginning of the high demand periods. Indeed, Figure 4 suggests that releases of top movies are concentrated around a few specific weeks of the year–Memorial Day, Fourth of July, Thanksgiving and Christmas–which fall in the beginning of the summer and in the winter holiday period. Also, it is important to note that there is no obvious similar seasonal pattern in the number of movies released, as can be seen in Figure 5.

3.2 Timing Data

The data covered in this section describe the dynamic process that leads to the eventual schedule of release dates. It is a uniquely constructed data set that contains the announcements made by movie distributors

[21] For example, Martin Luther King Day falls on the third Monday of January, usually on the third weekend of the year. However, since year 2001 started on a Monday, MLK Day fell on the second week of that year.

prior to the actual release regarding their scheduled release date. The source of the data is the monthly publication "Feature Release Schedule", which is published by *Exhibitor Relations Inc.*

In the beginning of each month, the publication lists the updated release schedule of all movies that are in the making, but have not been released as of yet. Typically, movies are first listed about twelve to eighteen months before their scheduled release. At this stage, many of the movies are in the process of casting or are in early stages of production. Thus, when first entering the monthly report, movies are generally not assigned to a specific release date. Rather, they are given a more general release season, such as "Summer 2002", "Christmas 2002", or just as "coming". As the scheduled release approaches, the release date becomes more specific, e.g. "Late Summer 2002" or "Early July 2002", converging eventually to a specific date. Occasionally, additional release information is given, such as the type of release (wide release, platform release, or limited release), the number of regional markets, and the approximate number of screens.

The data covers roughly all titles that were eventually released between 1/1/1985 and 12/31/1999,[22] a total of 3,363 titles. The analysis below uses only those that were released nationwide, applying the same definition as in the previous section. Out of the 1,956 titles covered in the main part of the data, I was able to obtain pre-release announcements data on 1,897 titles (97%). To get an idea of what the data look like, let me use *Titanic* as an example. It was first listed as "Summer 1997" in the August 1996 issue of the publication. In November 1996 the schedule became more specific – July 4, 1997. In the May 1997 issue, the movie was listed again as "Summer 1997", reflecting the concern that it would not be ready for the scheduled release date. Indeed, the next announcement pushed the movie back to the next high season, and listed it as December 19, 1997, which indeed was the eventual release date. The sequence of announcements for the 1999 release of *Star Wars: The Phantom Menace* was less eventful; it was first listed as "May 1999" in the issue of May 1998. In the September 1998 issue, the announcement became more specific, May 21, 1999, and remained the same until its actual release.[23] The *Titanic* example also illustrates how I can detect, in most cases, changes in release that are driven by production delays rather than by strategic considerations.

In what follows I describe the data from different aspects, and later in the paper I use it to motivate my estimation strategy. A more comprehensive analysis of the timing game being played is an ongoing project. Caruana and Einav (2001) provide a benchmark theoretical model that deals with the main strategic effects being considered. For computational reasons, it is hard to use the model for estimation.

As previously mentioned, I assume throughout that the general release season is exogenously given, so the strategic aspect I am interested in is related to the specific week within the season on which the movie is released. Thus, I consider as real announcements only those that carry a specific date (so announcing "Summer 2002" is like not announcing anything). Moreover, if at any point a certain title moves from a specific announcement to a more general one, I interpret this as production uncertainty (as in the Titanic example given above) and document only the sequence of "uninterrupted" specific announcements. Using this definition of announcement, I define the *time in the game* for each movie by the number of months between the report date of its first announcement and the first announced release date. An earlier entry into the game may be interpreted as trying to commit to a certain release date early on in order to scare off other movies. Figure 6 shows the distribution of the time in the game across titles. The mean for all movies is 4.84 months, with the mode being four months. Once restricted to the top movies (defined either

[22] A small number of movies were scheduled for release during the period, but were never actually released. This happened when the prospects for the movie during its production were so dire, that production was halted, or that the movie was released dicrectly to video.

[23] In fact, in the April 1999 issue the May 21 (Friday) announcement was changed to May 19 (Wednesday). However, as discussed earlier in the paper, having the week as the unit of analysis, we consider these two dates to be practically identical.

through their production costs or through their estimated quality from the benchmark model of Section 4), the mode becomes five months, and the mean is about 6.1. Thus, distributors of bigger budget movies tend to "plant their flags" earlier, taking into account the strategic impact their announcement may carry. Even within this group, the correlation between a movie's size and its time in the game is positive and statistically significant.

A major characteristic of the data is the frequent changes in the release schedule of certain movies. This is somewhat surprising, given the incurred costs of changing a release date. Such costs are incurred for several reasons, such as committed advertising slots, the implicit costs of reoptimizing the advertising campaign, reputational costs, etc. The costs become higher as the changes in release date are done closer to the scheduled release. While some of these changes are the result of unforeseen production delays, I argue that most of these changes are made for strategic reasons, and may provide some indication of unobserved characteristics of the movie, such as quality and commitment power. This is also backed by claims made by industry practitioners and the popular media, who describe the scheduling game as a kind of *war of attrition* over the best release dates.

Across all movies and announcements, more than 20% of the monthly announcements are changes in relation to the most recent announcement of the same film. Moreover, more than 60% of the movies have changed their release dates at least once. Figure 7 provides the distribution of the magnitude (in weeks) of these changes. It can be seen that the distribution is roughly symmetric (with a wider tail to the right for obvious reasons), and that the majority of the changes shift the release date by a small number of weeks; 75% of the observed changes do not shift the release date by more than a month. Both the symmetry of the distribution and its shape indicate that it is unlikely that the majority of these changes are made for exogenous non-strategic reasons, such as production delays. The likelihood of a movie changing its release date is not significantly correlated with the movie's size, measured by its production cost. However, movies with higher box office quality (as estimated in the next section) are significantly less likely to change release dates. One interpretation of this is that movies' estimated quality is originally highly correlated with their production cost, but as the shape of the finished product becomes clearer, films that turn out to be potential disappointments shift away from their previously announced release dates.

4 Demand

4.1 Specification and Identification

As discussed previously, my main objective in estimating demand is to disentangle the observed seasonal pattern in sales, and break it down into its two main components: the underlying seasonal effects in demand and the endogenous supply-side reaction. This will generate a "clean" estimate for the underlying seasonal effects in demand (and for the movie quality). Specifically, I want to distinguish between changes in demand resulting from seasonal effects and changes in demand resulting from variation in the quality of the movies released over time.

Ideally, in order to identify the two different effects I would need to observe the same movie released on different dates in a number of distinct yet similar markets. This, of course, is not feasible. First, I observe only one market: the aggregate U.S. domestic market. Second, as discussed in Section 2, even if I were to observe several regional markets separately, first-run films are predominantly released simultaneously nationwide.[24]

This problem is not trivial because the release date decision is endogenous, so the interpretation of standard procedures that correct for seasonal effects is misleading. As shown in the previous section, higher quality movies are more likely to be released on certain weekends. A simple way to separately identify the two effects would have been to control for the movie's quality. However, in the motion picture industry such controls for quality arguably do not exist, so conditional on observables the endogeneity problem remains significant.

The main feature of the data that allows me to separate these effects is the weekly box office revenues. Under certain assumptions, the weekly data allow me to obtain multiple observations for each movie, and hence to extract the average quality of the movie. Once quality is identified, one can obtain unbiased estimates for the underlying seasonal effects. Hence, I estimate *weekly* demand for movies. I do this by using the weekly data of box office revenues. Identification comes from observing movies throughout their run, playing during weeks with different demand levels, and competing against different sets of movies. While a higher quality movie may be released in a week with higher demand, it still runs for the following weeks, in which seasonal demand will vary. A lower quality movie, which is, say, released afterwards, is then observed at the same week as the better movie, enabling a direct comparison between the market shares of the two movies, thus identifying their relative qualities. Seasonal effects are identified by observing different choice sets over different weeks, and over a period of fifteen years.

To convey intuition for the identification problem and the way it is solved in this paper, let me consider the following simplifying example. Suppose that every year consists of only two seasons, and exactly two movies are released each year, one every season. One of the movies is of quality 1 and the other is of quality θ_H, with $\theta_H > 1$. The market sizes in season 1 and season 2 are 1 and M respectively, with $M > 1$. The revenues of a movie of quality θ released in season of size m are $m\theta$. Suppose that distributors always release their better movie on season 2. To understand the identification problem, assume first that each movie lasts in theaters only for a single season. Thus, the observed pattern in sales is 1 in season 1 and $M\theta_H$ in season 2. It is easy to see that in such a case only the product of M and θ_H is identified, but not each of them separately. If I had a good measure either for M or for θ_H, I could extract the other one. With no such measure, however, distinguishing between the two is impossible.

To understand the identification assumption employed in this paper, suppose now that each movie lasts in theaters for two seasons, and that they both share a common decay rate in quality, denoted by λ. For simplicity, assume also that the demand to one movie is independent of the demand to the other. In this case, I have two observations for each movie, with the sales of the movie of quality θ_H being $\lambda\theta_H$ and $M\theta_H$ in seasons 1 and 2, respectively, and the sales of the other movie being 1 and $M\lambda$. Thus, I have three equations in three unknowns, and all the parameters can be separately identified. Note also that with more movies and more seasons, one can allow the decay parameters to vary across movies. The crucial identifying assumption is that the decay pattern is a characteristic of the movie and cannot vary across release dates.

In what follows I discuss in detail the benchmark model, and then several extensions. In the next section I show that the results I am interested in are robust to these different extensions. This allows me to focus on the results from the benchmark model in the analysis of the release date decisions. While the benchmark model is simpler and more intuitive, it is important to note that incorporating the different possible extensions is

[24] This is mainly in order to maximize the effect of the mass campaign in network television. See Caves (2000) for more details. Even if it was not the case, the informational differences in the first week between markets would have made it hard to apply typical i.i.d. assumptions.

straightforward.

4.1.1 The Benchmark Model

The benchmark model follows standard discrete choice models with common coefficients for all individuals. While the model is based on an individual utility function, I should emphasize that I do not aim at actually modelling individual decision-making regarding movie-going. Rather, I use this formulation as a convenient and reasonable functional form that allows me to separate the seasonal effects in a simple way.

The utility of consumer i from going on week t to movie j is given by

$$u_{ijt} = \delta_{jt} + \epsilon_{ijt} \tag{1}$$

where δ_{jt} is the average quality of movie j at week t, and ϵ_{ijt} is an idiosyncratic individual error term. Following the literature initiated by McFadden's (1973) seminal paper, I assume that ϵ_{ijt} is distributed according to a Type 1 extreme value distribution, which allows me to integrate analytically over consumers and obtain a closed-form solution for the predicted market shares. Normalizing the utility of all consumers from the outside good (good 0) to be zero, I obtain the standard multinomial logit predicted market shares for each movie

$$s_{jt} = \frac{\exp(\delta_{jt})}{1 + \sum\limits_{k \in J_t} \exp(\delta_{kt})} \tag{2}$$

where J_t is the set of all movies that are in theaters on week t (note that given the available data, a movie that is in theaters for more than ten weeks is treated as if it is no longer in theaters). Equation (2) can then be rearranged to obtain

$$\log(s_{jt}) - \log(s_{0t}) = \delta_{jt} \tag{3}$$

The main focus then is the modelling of δ_{jt}, which is the crucial part for the identification of the seasonal effects. Still, it is important to note that the advantage of the logit model over other models that try to explain box office revenues is that here I explicitly take into account the competition in different weeks, thus explaining box office revenues adjusted for competition, which seems more plausible.

The benchmark specification for δ_{jt} is

$$\delta_{jt} = \theta_j + \tau_t - \lambda(t - r_j) + \xi_{jt} \tag{4}$$

where r_j is the release week of movie j, and the rest are estimable parameters. With this specification, the mean utility from movie j at week t is driven by three systematic elements–the fixed quality of the movie, θ_j, the underlying seasonal effect in demand (which can be thought of as the market size in the specific week), τ_t, and the decay effect, i.e. the number of weeks that have passed since the movie's release. I discuss each in turn.

The first element in the right hand side of equation (4), θ_j, is simply a movie fixed-effect which captures all the variables that do not change over the screen life of the movie, such as the cast, the plot, the pre-release advertising campaign, and any unobserved or non-quantifiable movie characteristic. As argued previously, in this specific industry using a fixed-effects specification is even more important than in other markets since observables explain very little of the variation in movie quality. A fixed-effects specification implicitly assumes that the movie's quality is indeed fixed over its screen life. This would not have been a good assumption had word-of-mouth or post-release critic reviews played an important role in determining the perceived quality of the movie. If this were the case, this would have implications on the decay pattern, which is discussed below.

The second element, τ_t, is the underlying seasonal effect in demand. Once movie fixed-effects are used in the specification, the weekly dummies capture only the fluctuations in demand that are not explained by the different set of movies that are running during each week. In the context of the utility function, the weekly dummies can be thought of as differences in the utility from the outside good across different weeks of the year. One natural interpretation is to think of differences in the opportunity cost of time, so, for example, on holiday weekends the utility from the outside good is smaller (people have more leisure time) and the weekly dummies are higher (recall that the utility from movie j is normalized by *subtracting* the utility from the outside good). Implicitly, the formulation of demand assumes that the mix of consumers is similar across weeks. Given that the main identification comes from comparing weeks that are close to each other, this assumption seems reasonable.[25]

The third element in the right hand side of equation (4), $t - r_j$, is the decay effect. The linear specification is hardly restrictive, and is used mainly to simplify notation. In the next section I show that allowing a fully flexible functional form for the decay effect results in almost a linear decay, and hence has practically no effect on the estimated seasonal pattern in demand. Note also that the dependent variable is approximately proportional to the log of the market share, so that the linear term can be thought of as a proportional decay in the box office revenues. I view the decay variable as a reduced-form way of capturing two main effects. First, most people who have watched a movie will not watch it again, so the potential market is reduced over time. Second, most people have a tendency to prefer watching a movie sooner rather than later. This is especially true in the U.S., and less so in foreign markets.[26] While the second effect may enter the utility directly, and is consistent with the individual utility specification, the first effect (no repeat buyers) deserves further discussion. To accurately model the dynamic decision one has to make assumptions about the relationship among the sets of idiosyncratic noise ϵ_{ijt} over time, and to dynamically change the distribution of the ϵ_{ijt}'s to account for truncation (of those individuals who have already seen the movie). I believe that doing so properly requires a more detailed data set than the one used in the current study. Thus, I reemphasize my view of the individual utility as a convenient functional form, rather than as an accurate description of reality.[27]

An additional effect on the decay pattern, as mentioned previously, may be the result of variation in the perceived movie quality over the movie's life cycle. This may result from word-of-mouth and critic reviews, and would imply very different decay patterns for different movies, such that some movies drop in revenues after the first week, while others increase in revenues because of good viewer reviews. While, as discussed later on, some variation in decay pattern among movies does exist, "sleepers"–a slowly building revenues pattern–is very rare. The decay pattern is quite stable across movies, and the first week revenues alone are an excellent predictor of the eventual box office revenues. Regressing the total box office revenues on the first week revenues results in an R^2 of 0.86. This is also consistent with conventional wisdom of distributors,

[25] In modelling the decision regarding the release *season*, such considerations may have been more important. The mix of consumers in the summer may be different from the one in the holiday period, and so is the general nature of movies released.

[26] Focus group participants suggest that the heavy concentration of advertising and publicity just prior to release, as well as a desire of some moviegoers to be market mavens, can help explain the rapid decay over time of mass-appeal movies. The heavy advertising just prior to release is also driven by the declining percentage structure of contracts in the industry: conditional on total admissions, distributors' incentives are to make more consumers watch the movie earlier rather than later.

[27] See also Moul (2001) for a more rigorous treatment of the saturation effect.

who claim that after they obtain box office revenues for the first Friday night in which the movie is shown, they can predict with high accuracy the eventual gross box office revenues for the film. Moreover, movies for which word of mouth seems important are generally released using a limited release pattern, so that word of mouth can spread before the movie release widens, which is the point at which the movie enters my data set. Furthermore, the reviews and trailers of most movies are published in the press before their release date. This suggests that information dissemination does not play a central role in determining box office revenues.

Finally, one may think that an important aspect of the decay pattern is the exhibitors' decision about movie replacement. As long as one is willing to assume that exhibitors act as agents for consumers–replacing a movie only when the demand for it is sufficiently low–this has no real impact. In particular, this is true for national level data, as in this case, where the typical movie runs on at least several hundred screens, so the choice of the individual exhibitor has less of an impact. With the recent increase in the number of screens in the industry, movies are dropped from theaters only when the demand for them is very low.[28]

Note also that the data and the specification do not allow me to separately identify time trend or year effects. A linear time trend is simply not identified.[29] A more flexible time trend may be estimated mathematically, but will be identified only because of the functional form, not allowing any economic interpretation. Similarly, year effects, which can be thought of in this context as a step-function time trend, would be identified only through those movies that are released in December and continue through January. This simply means that one cannot distinguish between a general trend in the utility from the outside option and a general trend in the average quality of the inside option. I can only identify the sum of the two, which is captured by looking at the trend of the estimated movie fixed-effects over time.

Let me emphasize how this specification allows me to solve the identification problem presented at the beginning of the section. I observe each movie for ten weeks only, and hence I cannot directly compare box office revenues of movies shown very far apart, as for example summer and winter movies. Rather, the identification employed is indirect. For example, movies that are released in the spring are implicitly used to compare summer and winter movies in an obvious manner: after controlling for seasonal and decay effects, I can compare summer movies with the spring ones using their overlapping period, and compare the spring movies with the winter ones. Once the movie fixed-effects are comparable, it is easy to see how the weekly dummies can be identified. It is important to note that the identification relies heavily on a sort of transitivity assumption, as discussed above. Thus, roughly speaking, the closer the weeks are to each other, the more easily they are compared, and the more confident I am in the relative magnitude of the weekly dummies. Thus, one needs to exercise more caution in comparing the estimated coefficients of Memorial Day and, say, Thanksgiving, than in comparing Memorial Day to other weeks in May or in June.

Finally, I note that the source of identification here is not as good as in a natural experiment. For example, throughout the sample period big movies have consistently been released on Thanksgiving but never on Labor Day. Hence, comparing the estimated seasonal effects in these two weeks is a more "out of sample" exercise than comparing each of them to its neighboring weeks. For the latter case, one may think that the variation in the data makes it closer to a natural experiment.

[28] Exhibitors may also shift movies from big screens to smaller ones, accomodating newly released films. Again, as long as these moves reflect the projected demand for the movie, this poses no problem for the analysis.

[29] To understand this, we can write the benchmark model and add a linear time trend to get: $\delta_{jt} = \theta_j + \tau_t - \lambda(t - r_j) + \gamma t + \widetilde{\xi}_{jt}$. This can be rewritten as: $\delta_{jrt} = \widehat{\theta}_j + \tau_t - (\lambda - \gamma)(t - r_j) + \widetilde{\xi}_{jt}$, where $\widehat{\theta}_j = \theta_j + \gamma r_j$.

4.1.2 Extensions to the Benchmark Model

There are two obvious ways by which I can extend the benchmark specification. The first is to weaken the assumption about additive separability of movie quality, market size, and decay. The second is to allow a richer specification of the error term, such as nested logit or random coefficient models (see Berry (1994) and Berry, Levinsohn and Pakes (1995)), which may be more appropriate. I discuss each in turn.

While I use dummy variables to allow a flexible functional form for each effect separately, the assumption that these three effects are additively separable is quite strong. In particular, I rule out any potential interactions between them. In considering interaction terms between the decay and the movie characteristics, it is natural to consider two major cases. The first allows different decay rates for different types of movies, while the second allows different rates of decay for movies of different quality. For example, it is well known that revenues of small movies decay faster than the revenues of big movies, as can also be seen in Figure 2. However, as long as I allow the decay pattern to depend solely on movie characteristics, which are not related to the release date of the movie, the identification mechanism discussed before goes through. Indeed, in the next section I investigate several extensions to the decay specification, and find that while the decay rate may vary across different movies, this variation does not have a significant effect on the estimated seasonal patterns in underlying demand.

Allowing interactions between the weekly dummy variables and other variables is more subtle. In considering these, one should keep in mind the original motivation, which is the identification of the seasonal effects in demand. Suppose, for example, that the decay rates are different depending on the release week. First, the economic interpretation of this is not clear. Second, this means that there is an additional seasonal effect, which enters in a multiplicative way. This would make it hard to interpret the original weekly dummy variables, so I rule out such interactions. Similarly, it is hard to interpret and identify interactions between the quality of the movies and the weekly dummies. These would suggest that certain weeks are more favorable to, say, bigger movies than to smaller ones. This is hard to justify. I do investigate, however, interactions between horizontal differences in movie characteristics, namely genres, and the weekly dummies. These may be generated by a different mix of consumers at different weeks. In the next section I show that, with the exception of children's movies, these are not very important.

A second line of extensions is to consider richer substitution patterns. As is well known, the multinomial logit allows for the substitution pattern among movies to depend only on their size (market share). In the context of this paper, this implies that, say, a children's movie released on a certain week is going to shift consumers away from all other movies showing on the same week proportionally to the movies' weekly market share. One would think that a more realistic model should predict that more consumers will be shifted away from other children's movies than from horror movies. In general, it seems reasonable to assume that movies of a certain genre are closer substitutes of movies of the same genre. To deal with this, one may consider either nested logit specifications, where the nests depend on the genre of the movie, or random coefficients models, with random coefficients on genre dummy variables. I investigate several specifications in the next section. It is important to understand, however, that the weekly dummy variables capture first order effects in demand, while these suggested specifications correct for the cross substitution, i.e. second order effects. Thus, I expect that the results I focus on, namely the seasonal effects in underlying demand, will be more sensitive to alternative specifications of the decay effects than to alternative specifications of the error term. Still, these alternative specifications may prove important when analyzing the release date decision, for which the substitution pattern is crucial.

4.2 Results

4.2.1 The Benchmark Model

I move towards the estimation of the benchmark model in two steps. As will become clear later, I do so in order to illustrate that the fixed effects, and not other modelling assumptions, are those that drive most of the results. First, I estimate the model without the movie fixed effects, i.e. I estimate the following equation: $\delta_{jt} = \tau_t - \lambda(t - r_j) + \xi_{jt}$. Unlike the aggregate industry share, with its seasonal variation presented in Figure 3, the unit of analysis here is the individual movie. Thus, this specification can be thought of as controlling for the number of movies released each week, without controlling for their quality. Another difference is the dependent variable, which is now the mean utility of the movie, namely $\delta_{jt} = \log(s_{jt}) - \log(s_{0t})$. Thus, compared to the aggregate industry share, I also control for the competition by subtracting $\log(s_{0t})$. These two differences should work in opposite directions, and hence I expect to observe a similar picture to the one we see in the aggregate. The results are presented in Figure 8. Indeed, the estimation results are quite similar to the seasonal pattern of the industry share, with the summer and the holiday periods being quite strong compared to other times of the year. Still, some differences should be noted. These seem to be generated by the high number of movies released on Memorial Day and the small number of movies released in Labor Day (see Figure 5). Thus, the dummy variable for Memorial Day, on which many movies are released, loses a lot of its magnitude (compared to the neighboring weeks), and the Labor Day dummy variable, for which very few distributors release their movies, becomes relatively higher.

Figure 9 presents the results of estimating the full benchmark model, in which movie fixed-effects are incorporated. Now we can see a strikingly different picture from the one we saw before. First, the magnitude of the estimated seasonal pattern in demand is much smaller compared to the one estimated without movie fixed effects; the standard deviation of the estimated coefficients on the week dummies goes down by almost 50% (from 0.37 to 0.20), suggesting that seasonal differences in the quality of the movies released account for about half of the observed seasonal pattern in revenues. Second, the shape of the seasonal pattern becomes quite different, and is shifted forward by one to two months. In particular, the early summer period obtains relatively low coefficients, the coefficient on Labor Day becomes surprisingly high, and the Thanksgiving (and the proceeding weeks') coefficients are very low. These estimates are very different from the conventional wisdom of the industry. The conventional wisdom is reflected in the release decisions, with studios fighting over Memorial Day releases and trying to avoid Labor Day releases. The estimated underlying demand suggests a very different seasonal pattern.

The estimated decay coefficient in the benchmark model is about -0.46, with a very small standard error. This translates into about a 37% decline in the movie's appeal every additional week it is in theaters. This mimics quite closely the observed decline in the movies' box office revenues over weeks, as presented in Figure 2. Figure 10 presents the estimated decay when I allow for a flexible pattern. It shows that the flexible form gives rise to a function which is almost linear, so we can see why the linearity assumption is hardly restrictive. All other results of the paper would not be altered if I used a flexible pattern rather than the linear one.

The estimated movie fixed effects, not surprisingly, closely follow the total box office revenues for each movie (R^2 of 0.89). Figure 11 presents the trend in the estimated fixed effects over the sample period. This trend captures the trend in the attractiveness of the "inside good" relative to that of the outside good. Clearly, the outside good did not remain constant over the long sample period, with the introduction of Internet, DVD, and other alternative opportunities to moviegoing. Thus, the fact that the average quality did not increase over time does not necessarily imply that the industry is not advancing. It shows that the industry is able to keep up with the outside competition, albeit with steadily increasing production costs. Finally, simple regression of the fixed effects on different characteristics of movies is presented in Table 4. Besides being of some interest on their own, the results suggest that the estimated fixed effects are quite reasonable. All coefficients have sensible signs and magnitudes. The low R^2's in the table also show how little of the variation in quality can be explained by the observable variables (some of which are endogenous anyway), which relates to the original motivation for using movie fixed effects.

4.2.2 Robustness tests

The results from the benchmark model are quite different from conventional wisdom, making it even more important than usual to verify that they are robust to different assumptions and to various changes in the specification. First, I show that the estimated seasonal pattern in demand is not driven by certain subsections of the sample. Figure 12 and Figure 13 present the estimation results of the benchmark model on subsamples of the data. The first figure shows estimates of the benchmark model for different genres, and the second shows the estimates for different sub periods of the sample period. While the estimated parameters are not the same as in the basic specification, they all share roughly the same pattern. Across periods, the shape of the seasonal pattern remains similar, with its magnitude becoming larger over time. Across genres, all three genres of action, drama, and comedy show similar patterns to the aggregate one, with children's movies being a notable exception. For children's movies, the seasonal pattern in demand peaks higher and earlier in the summer period. This actually makes sense and is consistent with the conventional wisdom that the underlying seasonal effects in demand should be stronger for children's movies, and more sensitive to school holidays.

Next, I address concerns regarding the fixed effects and whether they are indeed fixed over the ten week observation period of each movie. It may be argued that over time people learn about the movie, thus making it look better (or worse) after its initial release. To accommodate this concern, I assume that after two weeks in theaters informational diffusion is minor, so the fixed effects are fixed only from then on. I then estimate the model on a subsample that includes all movies only after their second week. The estimated seasonal pattern is very similar to the one obtained by the benchmark model.

Truncation may also be an issue. Recall that while I am "supposed" to observe each movie for ten weeks, on average I have about 8.5 observations per movie. I address this concern in several ways. First, I estimate the model on a "balanced panel", by dropping all movies for which I have less than ten observations. Second, I use the full unbalanced panel, but let the decay coefficient vary, depending on the number of weeks the movie has survived in theaters. In both cases I find very little effect on the estimated seasonal pattern in underlying demand.

Allowing the decay pattern to vary across types of films does not have much effect as well. I investigate this by interacting the decay with the genre, as well as with the size of the movie, knowing that bigger movies decay slower. Although, not surprisingly, some of the interaction coefficients obtain significant values, neither case has a significant effect on the pattern of the estimated coefficients on the weekly dummy variables. This is also the case with interacting decay and season dummy variables. Interaction of decay and week dummy variables does change the pattern, but, as discussed in the previous section, this exercise makes the interpretation of the coefficients on the weekly dummies unclear, because much of their effect is captured by the week-specific decay pattern.

Finally, I investigate some natural extensions by using richer specifications of the error term. I allow a nested logit specification (see Berry (1994) for details) either by nesting the movies by their genre or by nesting all movies together, with the outside good as a separate nest. These specifications result in relatively low estimates of the parameter that captures the importance of the nests (the σ parameter in Berry (1994)). Therefore, the nested logit results almost reduce back to the simple logit, i.e. the benchmark model. Similar results were obtained by experimenting with simple versions of random coefficient models that allow for random coefficients on a set of genre dummy variables. Again, just as in the other robustness tests discussed above, these results do not imply that measures of proximity among films have no effect on substitution patterns. The results only suggest that these richer substitution patterns do not have a major effect on the estimates of the underlying seasonal pattern in demand.

5 The Release Pattern

The estimation of the demand system poses a somewhat surprising result–the release pattern does not seem to correspond to the estimated seasonality in underlying demand. For example, we observe big movies released early in the summer while the demand estimates suggest that demand is higher only later in the summer, not before July. Thus, in this section I analyze the release date decision using the estimates obtained in the previous section, trying to get a handle on the seasonal pattern in demand, as perceived by distributors when making their release date decisions.

I begin by using a simplifying "perfect competition" assumption to construct an approximation that allows me to analyze the release date decisions in a reduced form way. This type of analysis suggests that the observed release pattern cannot be rationalized by the estimated pattern of underlying demand.

However, given the skewness of movie qualities, market power in the industry is very important, and the release date decision is taken more strategically, as is also suggested by the pre-release data described in Section 3.2. To address this, I continue by relaxing the perfect competition assumption and constructing an estimable timing game. The game is then used to estimate the perceived seasonal pattern in demand that rationalizes the observed release pattern.

The estimates lead us to a similar conclusion as those obtained in the perfect competition case. I find that the release date decision cannot be rationalized by the estimated demand. However, the estimates from the strategic game are driven from within-season variation (rather than from between-season), and also allow me to investigate other effects, such as the importance of strategic considerations in the industry.

Throughout this section I make many simplifying assumptions for computational reasons. Two of them are important enough to discuss early on. First, I assume away uncertainty, although it is an important feature in the industry. I treat movie quality as known in advance to all players in the market. I plan to incorporate uncertainty into the analysis in future work. Second, I ignore the fact that studios distribute a portfolio of movies and therefore, when making release decisions, internalize their effect on other movies in their portfolio. Relaxing this assumption and incorporating it into the analysis does not pose any conceptual problems, but is computationally very intensive, making the decision space much bigger. I discuss in Section 6 how relaxing these two assumptions may affect the results.

Finally, I assume throughout that the distributors' objective is to maximize the cumulative undiscounted market share of the movie. This seems to be a decent approximation of reality. Reasonable alternatives to this assumption are unlikely to change any of the results significantly.[30]

[30] For example, one may discount the revenues in accordance with the declining revenue share of distributors over the movie's run.

5.1 Release Date Under Perfect Competition

As discussed above, the release pattern does not seem to correspond to the estimated seasonality in underlying demand. To make this argument in a more formal way, let me consider the benchmark model, assume away the disturbance term ξ_{jt}, and ignore the uncertainty around the estimated parameters. In such a case, I can divide both parts of equation (2) by $\exp(\tau_t)$ to obtain the following:

$$s_{jt} = \frac{\exp(\theta_j - \lambda(t - r_j))}{\exp(-\tau_t) + \sum_{k \in J_t} \exp(\theta_k - \lambda(t - r_k))} \tag{5}$$

Equation (5) has a simple and intuitive interpretation. The numerator in the right hand side of the equation is just a scaling factor, which depends on the decay-adjusted quality of the movie. The denominator has two additive elements. The first element, $\exp(-\tau_t)$, can be thought of as the *market size effect*, and is decreasing in the estimated week effect. The second element can be thought of as the *competition effect*, which in this model is simply summarized by adding up the exponents of the qualities of the competing movies, adjusted for their decay.[31] The competition effect may be larger either because there are more competing movies or because the average quality of the competing movies is higher, or both. As we saw in Figure 5, the number of movies released each week does not vary a lot over the year, so the competition effect varies mainly because of the seasonal difference in the quality of movies released.

Using this simple model, the optimal release date for each movie is given by

$$r_j^* = \arg\max_{r_j} \sum_{t=r_j}^{r+H} s_{jt} = \arg\max_{r_j} \sum_{t=r_j}^{r+H} \frac{\exp(\theta_j - \lambda(t - r_j))}{\exp(-\tau_t) + \sum_{k \in J_t} \exp(\theta_k - \lambda(t - r_k))} \tag{6}$$

where H is the length of the horizon taken into account by distributors.[32] Equation (6) shows in a simple way the trade-off faced by distributors between the market size effect and the competition effect. Note that I assume that θ_j is known and taken as given by distributors. Therefore, the release decision is only driven by the magnitude of the denominator, with smaller being better. On the one hand, distributors prefer to release their movie in a more favorable week (higher τ_t), but on the other, they prefer to compete against less and lower quality movies (smaller competition effect). Using the demand estimates obtained in the previous section, Figure 14 presents the estimated seasonal pattern in the competition effect. Given the fast decay, the pattern in the competition effect is mainly driven by the traditional pattern in release decisions, which is presented in Figure 15.[33]

The surprising result is that if I use the point estimates of the coefficients, I obtain a significantly *positive* correlation coefficient (0.4 with t-statistic greater than 3) between the market size effect, $\exp(-\tau_t)$, and the competition effect, $\sum_{k \in J_t} \exp(\theta_k - \lambda(t - r_k))$. This implies that, given the historical release pattern, there is no real trade-off–better weeks are in general associated with weeks of softer competition. Thus, taking this model as given, it seems evident from this simple exercise that the estimated seasonal pattern in underlying demand is not the pattern according to which distributors make their release decisions.

[31] For example, in a nested logit model the competitive effect becomes more complicated, giving more weight to movies of the same nest.

[32] With the estimated fast decay none of the results are very sensitive to the choice of H, as long as H is not very small.

[33] It is interesting to note how closely Figure 15 mimics Figure 4, despite the relatively small correlation between production costs and quality.

Hence, I try to find the perceived seasonal pattern in demand that rationalizes the release pattern. An approximation is constructed in the following way. I take the estimated movie fixed effects and decay parameter as given, and I search for the seasonal pattern in demand that would exactly offset the seasonal pattern of the competition effect.[34] Such an approximation would make the denominator in equation (5) constant across weeks, thus making movies indifferent between releasing in one week or another. I argue that this is an approximation to the perceived seasonal effect in demand because such seasonal patterns make the observed release pattern a Nash Equilibrium of a release game, under a few simplifying assumptions. For this to be a Nash Equilibrium, movie quality has to be continuous and divisible. A movie of a certain quality is actually thought of as infinitely many movies of low quality that add up to that quality of the original movie, with each one of these infinitely many movies distributed separately. This implies that distributors act non-strategically and take the denominator as given. This corresponds to a price taking assumption in standard markets.[35] This sort of assumption is very similar to those that give rise to "perfect competition" in supply-demand models.

To see why these assumptions are sufficient, one can check equation (6) and verify that the perfect competition assumption implies that all movies solve the same optimization problem. The movie's own quality enters multiplicatively into the objective function, and the perfect competition assumptions imply that the denominator is taken as given. It is easy to see that in such a case a sufficient statistic for the release decision is the sequence of denominators, as given by equation (5), across the different weeks. In this perfect competition world, all movies look at the same sequence. In particular, all movies solve the following problem:

$$r^* = \arg\max_r \sum_{t=r}^{r+H} \frac{\exp(-\lambda(t-r))}{x_t} \tag{7}$$

where

$$x_t \equiv \exp(-\tau_t) + \sum_{k \in J_t} \exp(\theta_k - \lambda(t - r_k)) \tag{8}$$

Thus, in order to observe movies released throughout the year on all possible weeks, and not concentrated in specific weeks, it can be shown that a stationarity assumption implies that the sequence of x_t's must be constant (see proof in the Appendix). This simply means that distributors are indifferent between the different release dates. It implies that the market size effect is offset completely by the competition effect, i.e. a correlation coefficient of -1, rather than $+0.4$ as I found before.

5.2 Implication for Revenues

At this point, a natural question that comes to mind regards the amount of moey left on the table by the industry, conditional on the demand estimates. In this section I will provide two approximations for an answer to such a question, which mainly provide upper bounds to this amount. Both approximations rely on the divisibility (of movie quality) assumption discussed in the previous section. The first approximation calculates this amount from the perspective of a single movie, taking the release pattern of the rest of the

[34] This exercise may seem puzzling, given that the estimaed underlying demand and the other parameters of the model were estimated together. However, what makes the exercise reasonable is the fact that the estimated pattern in competition effect is robust to changes in the estimated underlying demand.

[35] Moreover, the divisibility implies that equating the denominator across weeks is feasible mathematically, something that would not have been possible with the "lumpy" movies.

industry as given, while the second approximation looks at this question from the industry-wide perspective, as if all distributors were to collude and "reshuffle" the movie releases, conditional on the set of movies available to them.

The first approximation relies on the perfect coimpetition model described above. Figure 16 depicts the implied values of the objective function in equation (7). It uses the point estimates of the week dummies and the average (over years) of the competition effect, for different values of the decay parameter. This can be thought of as the value for the objective function for a marginal movie, taking the historical release pattern as given. The figure suggests that, in contradiction to the conventional wisdom in the industry, releasing on Labor Day or in the end of the summer is quite attractive, because the competition effect is low and underlying demand is high. This is especially true for movies which are expected to decay fast, and hence not to be affected much by the weak period of September and October. The estimates suggest that the box office revenues can be dramatically affected by the choice of release date. For example, a release in the "shopping period" (between Thanksgiving and Christmas) is estimated to account for about half of the revenues compared to a release of the same movie two months later. Again, as I cautioned before, one should be careful in the interpretation of these comparisons, given that the identification is better for weeks that are closer to each other, and the predictions suggested above are far "out of sample" exercises. Moreover, it is easy to see that movies of higher quality, which are going to affect the denominatoir in the objective function, will be less affected by shifting around their release dates, thus the amplitude of Figure 16 will be scaled down. This is why we should think of the figure as an upper bound on the potential percentage increase in revenues by unilateral changes of release dates.

Another measure to assess the efficiency loss of the industry is to estimate how much additional revenues could the industry make using the estimated underlying demand, by using a release pattern that maximizes the total revenues. Again, using the divisibilty assumption discussed earlier, one can think of a total quality of movies that has to be distributed over the weeks of the year. In practice, the industry must take the indivisibilities of movie qulities into account. This would imply another constraint on the maximization problem, making it clear why the calculation described below provides only an upper bound for the industry revenues. Having these assumptions in place, the industry has to solve the following program:

$$\max_{\theta_1,\theta_2,...} R(\theta_1,\theta_2,...) = \max_{\theta_1,\theta_2,...} \sum_{w=1}^{W} \frac{\exp(\theta_w)}{\exp(-\tau_w)+\exp(\theta_w)} \text{ subject to } \sum_{w=1}^{W} \exp(\theta_w) = \overline{\Theta} \tag{9}$$

where θ_w is the total movie quality allocated to week w, and $\overline{\Theta}$ is the total movie quality that has to be allocated throughout the year. The solution to this program is

$$\exp(\theta_w^*) = \left[\frac{\overline{\Theta}+\sum_w \exp(-\tau_w)}{\sum_w \sqrt{\exp(-\tau_w)}}\right]\sqrt{\exp(-\tau_w)} - \exp(-\tau_w) \tag{10}$$

and

$$R(\theta_1^*,\theta_2^*,...) = W - \frac{\left(\sum_w \sqrt{\exp(-\tau_w)}\right)^2}{\overline{\Theta}+\sum_w \exp(-\tau_w)} \tag{11}$$

I then use the average weekly industry market share as a representative year, simulate from the estimated distribution of the τ parameters, and calculate the optimal revenues by backing out the implied value of $\overline{\Theta}$ for each simulation. This gives me a distribution of what I think of as the upper bound of the amount of money the industry as a whole is leaving on the table, if it were to collude. Compared to actual annual market share[36] of 3.68, the distribution of the optimal industry market share has a mean of 4.22 with a standard error of 0.053. In other words, by reshuffling release dates, the industry can increase its revenues by about 14.4% (with a standard error of 1.44%). Part of this increase is due to collusion, and the other part is due to the potential misperceptions of the true underlying demand for movies (the τ's).

5.3 An Estimable Strategic Timing Game

The perfect competition assumption employed in the previous section is by no means an accurate description of reality. The top grossing movie every week accounts, on average, for about twenty five percent of the industry's weekly box office revenues. Therefore, in order to relax this assumption, one has to take into account the strategic considerations by movies' distributors. This can only be done through formulating a timing game and making an equilibrium assumption. This is what I do in this section.

Many industry practitioners describe the timing game as a war of attrition–distributors "plant their flags" at the big holiday weekends, and those that lose the game eventually shift away. This would predict that there are more changes *away* from big weekends than switches *into* big weekends. Figure 17 plots the average release effect as suggested by the initial announcements, compared to the one implied by the final announcements. A war of attrition story would imply that initial patterns should reflect more concentration in the big weekends. The pattern for Memorial Day and the Fourth of July is somewhat consistent with this story, but the opposite picture is revealed in the holiday period. Thus, this simple observation does not provide strong evidence for the war of attrition story. Rather, the patterns in the data seem to be more consistent with a story according to which there is initially an equilibrium based on estimated qualities, and later on, as the estimates of the qualities become more accurate, the equilibrium configuration changes.[37] Consequently, in the rest of this section I model the game played after the qualities of all movies can be (relatively) accurately predicted. I assume that the qualities of all movies are common knowledge, and search for the release schedule that results in equilibrium.

The discreteness of the strategy space coupled with the heterogeneity among movies creates difficulty in estimation.[38] Typically, for such games, a Nash Equilibrium assumption leads to a multiplicity of equilibria, which requires further assumptions in order to be taken to the data. There are different ways to solve this problem. First, one can find a property that holds in all equilibria and rely on it for estimation (e.g. Berry (1992)), just as we rely on first order conditions in continuous strategy spaces. Alternatively, one can add more structure to the game that will generate uniqueness (e.g. Mazzeo (2000) and Seim (2000)). I do the latter by moving from a simultaneous-move game to a sequential one. In what follows, I let each player move once in a pre-specified order, and use a Perfect Bayesian Equilibrium concept. This generates a probability measure over all possible outcomes of the game, which can then be estimated using Maximum Likelihood. The pre-specified order is assumed to be known to all players, but may be unobserved by the econometrician.

The sequential move assumption seems very realistic in the present context, where one can observe that the release date announcements of different movies are done sequentially. However, one may question the commitment value of such announcements, and hence the sensitivity of the solution to the order of moves.

[36] To convert to dollar amount one has to multiply the annual market share by the size of the US population and by the average ticket price.

[37] It is interesting to note that Figure 17 is consistent with the sentence cited at the beginning of the paper – the release changes seem to be more important in the summer and in the holiday seasons, where one can see some changes in the competition effect between the initial announcements and the final ones. For the rest of the year these two graphs almost coincide.

[38] See Reiss (1996) for a review.

Indeed, in a related paper, Caruana and Einav (2001) develop a perfect information dynamic model that investigates the commitment opportunities in a similar setting, where the announced release date creates a deadline for the game. Their model is much less sensitive to the order of moves, allowing the players to change their minds later on. Applying this model in estimation is an ongoing project, with the main difficulty being the computational burden of solving the model for a large number of players and big strategy spaces.

The general setup for the estimation is as follows. I choose several time windows ("seasons") within the year. The length of each season, K, is mainly restricted by computational limitations, as will become clear later on. I take the set of movies that were released within the season as exogenously given. The motivation for this assumption comes from the pre-release timing data described in Section 3.2–distributors decide far in advance that a certain movie is scheduled for around, say, Memorial Day, but only later on, decide about the specific week on which it is released. The release week must lie within the season to which Memorial Day belongs. This assumption is consistent with the evidence provided previously, which shows that most changes of previously announced release dates do not shift the date by much. Thus, for any given season I have a set of N players, those that eventually released their movie in that season. For each player, the action taken is to release the movie on either one of the K weeks that lie within the season.

Thus, with K weeks and N players, each player can choose one of K different release dates, so I have K^N potential release configurations (equilibrium outcomes). Therefore, it should be clear that computational limitations restrict the choice of N and K. Clearly, the main strategic players are the big movies. The lower quality movies have a minor effect on the demand for the other movies, and hence on other movies' release decisions. Thus, the lower quality movies can be thought of as non-strategic players. Therefore, I choose the identities of the N players in each season to be the top quality movies within the season, where the quality measure is given by the point estimates from the demand system. These movies play against each other, conditioning on the observed release dates of all other movies.[39] Given that all movies remain in the market for longer than one week, not only do the "active" players (top quality movies) condition on the release pattern of the lower quality movies within the season, but they also condition on the release pattern of all movies in the adjacent seasons. While conditioning on the release dates of movies from the preceding season is sensible, it is questionable whether it is valid to assume that movies can condition on the release dates of movies in the following season. I justify this assumption by the high decay coefficient, which implies that the effect of movies that are released more than one week apart is relatively small, and hence has little effect on strategic considerations. Still, I use these other movies and their observed release dates to calculate the counterfactual box office revenues.

Given a season s, denote the different players by $j \in J_s = \{1, 2, ..., N\}$, and their observed release dates by r_j^{obs}. These movies' strategy space is given by $M_s = \{m, m+1, ..., m+K-1\}$, where m denotes the first market (week) of the season. Conditional on the estimates from the previous section, one can calculate the counterfactual box office revenues for each movie, given the release configuration (r_j, r_{-j}) of the N players. The total payoffs for movie j is $\pi_j(r_j, r_{-j}; \gamma, \eta) = \eta \widehat{\pi}_j(r_j, r_{-j}; \gamma) + \epsilon_{jr_j}$, where γ is a vector of payoff parameters, η is a normalization parameter, and ϵ_{jr} is an i.i.d (across release dates and movies) draw from a type 1 extreme value distribution. The vector of ϵ_{jr}'s is assumed to be *private information* of movie j. The assumption that it is private information is somewhat restrictive, but greatly simplifies the analysis. The

[39] A reasonable approach would be to assume that the order of moves is dictated by the ordering of the movie qualities, the biggest movie playing first. This implies that the smallest movies condition their decisions on the release dates of the bigger ones, but not vice versa (which is what we do in this paper). Given the computational restrictions, such an approach would rely on the decisions of the small, less strategic, players. For these movies it is not clear that we need the "highpowered" structural game for estimation. Rather, given that they have no real strategic effect, we can estimate each movie's decision separately.

private information can be thought of as non-strategic considerations that may make a certain movie more likely to be released on a certain week, regardless of release dates of other movies. The magnitude of the estimated η provides an indication for the explanatory power of the model. The higher η is, the more we can explain the decisions by the estimated payoffs rather than by the random noise. An insignificant estimate of η implies that the model for the payoffs has no explanatory power. Note that this specification leads to simple logit probabilities. Conditional on the other players' decisions, r_{-j}, movie j is released on week r_j^* with the following probability:

$$\Pr(r_j^*|r_{-j}) = \frac{\exp(\eta\widehat{\pi}_{jt}(r_j^*, r_{-j};\gamma))}{\sum\limits_{r_j \in M_s} \exp(\eta\widehat{\pi}_{jt}(r_j, r_{-j};\gamma))} \tag{12}$$

The game is played sequentially, with each movie moving once according to a pre-specified order, which may be unknown to the econometrician. The solution concept is a Perfect Bayesian Equilibrium (PBE). Note, however, that the payoffs of each movie depend only on the release date of the other movies, but not on their realizations of ϵ_{jr_j}'s. Thus, each movie's strategy depends only on the past actions of players who moved previously, but not on their exact "types". This significantly simplifies the solution, since there are no incentives for movies to alter the prior used by their opponents about the realized ϵ_{jr_j}'s they face. Consequently, given the pre-specified order of play, the game can be solved backwards in a simple way.

Given N players, let the order of play be given by a permutation $o \in \mathcal{P}_N$, such that $o(m) = j$ implies that player j is the m^{th} player who moves. Let $prev(j) = \left\{o^{-1}(k)|o^{-1}(k) < o^{-1}(j)\right\}$ denote the set of players who play before player j. I solve the game backwards as follows. The last player to move, $o(N)$, conditions on the other players' decisions, and hence I can use equation (4) to see that $r^*_{o(N)}$ is chosen with probability

$$\Pr(r^*_{o(N)}|r_{-o(N)}) = \frac{\exp(\eta\widehat{\pi}_j(r^*_{o(N)}, r_{-o(N)};\gamma))}{\sum\limits_{r_j \in M_s} \exp(\eta\widehat{\pi}_j(r_{o(N)}, r_{-o(N)};\gamma))} \tag{13}$$

Now, going backwards, I can update the continuation values for all players by integrating out player $o(N)$'s decision, i.e.

$$\widehat{\pi}_j^{(N-1)}(r_{-o(N)}; r^o_{-J}, \gamma, \eta) = \sum_{r_{o(N)} \in M_s} \Pr(r_{o(N)}|r_{-o(N)})\widehat{\pi}_j(r_{o(N)}, r_{-o(N)}; r^o_{-J}, \gamma) \tag{14}$$

and

$$\pi_j^{(N-1)}(r_j, r_{-j};\gamma,\eta) = \eta\widehat{\pi}_j^{(N-1)}(r_j, r_{-j};\gamma) + \epsilon_{jr_j} \tag{15}$$

The key is that the ϵ_{jr_j}'s are drawn independently of the decision, so they can be taken out of the sum, allowing me to use equation (4) again, with the modified payoffs. This procedure can be done iteratively up to the player who moves first, i.e.

$$\Pr(r_j^*|r_{prev(j)}) = \frac{\exp(\eta\widehat{\pi}_j^{o^{-1}(j)}(r_j^*|r_{prev(j)}))}{\sum\limits_{r_j \in M_s} \exp(\eta\widehat{\pi}_j^{o^{-1}(j)}(r_j|r_{prev(j)}))} \tag{16}$$

and

$$\widehat{\pi}_k^{(o^{-1}(j)-1)}(r_{prev(j)};\gamma,\eta) = \sum_{r_j \in M_s} \Pr(r_j|r_{prev(j)})\widehat{\pi}_k^{(o^{-1}(j))}(r_j, r_{prev(j)};\gamma,\eta)\ \forall k \in prev(j) \tag{17}$$

Thus, the above implies a probability measure over all potential outcomes of the game. In particular, given an order o, the likelihood of observing a release schedule r^* is given by $\Pr(r^*|o) = \prod_{j=1}^{N} \Pr(r_j^*|r_{prev(j)}^*, o)$, and given a probability measure over all possible permutations, the unconditional likelihood of observing an outcome r^* is given by

$$f(r^*) = \sum_{o \in \mathcal{P}_N} \Pr(o) \prod_{i=1}^{N} \Pr(r_j^*|r_{prev(j)}^*, o) \tag{18}$$

A simple yet general way to specify the probability measure over the permutations is to assign a commitment measure for each player, given by $\mu_j = X_j\beta + \zeta_j$, where ζ_j is distributed i.i.d extreme value, and X_j is a vector of observed characteristics of movie j, which affect its distributor's commitment power. The order of moves is then dictated by the commitment measure μ_j. This implies that the probability of an order o is given by

$$\Pr(o) = \prod_{m=1}^{N-1} \frac{\exp(X_m\beta)}{\sum\limits_{k \notin prev(o(m))} \exp(X_k\beta)} \tag{19}$$

This concludes the estimation model. Given S distinct and independent release seasons, the model can be estimated using Maximum Likelihood. In the next section I estimate the model, with a specification that focuses on the seasonal pattern in demand as perceived by distributors.

5.4 Specification and Results

The estimation model specified in the previous section provides the framework for analysis. For the model to be taken to the data, I still need to specify the particular functions and the model parameters. In doing so, I am guided by two main considerations. First, the computational burden of the estimation is significant, so I need to reduce it as much as possible. This is done by reducing the dimensions of the game, by setting $N = 3$ and $K = 5$, and by restricting the number of parameters, which crucially affects the number of iterations in the search for the maximum. Second, I would like to continue focusing on the seasonal pattern. In particular, the reduced form analysis in Section 5.1 suggests that the estimated seasonal pattern in demand is not consistent with an equilibrium choice of the observed release dates. Thus, I use the timing game to test which seasonal pattern is most consistent with the perceived pattern in demand.

For estimation, I choose four annual release seasons, which are all centered around a dominant release date. These are Presidents Day, Memorial Day, the Fourth of July, and Thanksgiving. Each season includes the dominant week, and two weeks before and after, adding up to five weeks in each season ($K = 5$). Thus, I have a total of 60 seasons, on which the estimates are based. Within each season, I take the top three movies ($N = 3$) according to their estimated quality, out of all the movies released in the season. The number of movies in each season is between 6 to 17 with a mean of 11.2 and standard deviation of 2.34. The quality of the top three movies that are used in estimation accounts for 44% to 91% (with a mean of 66%), as a fraction of the total quality of all movies in the season. Thus, the top three movies account for the majority of the industry box office revenues in the season in which they are released.[40]

[40] One may argue that the top three movies are likely to take into account their effect on the smaller movies in the season, when they make their release date decisions. Computational burden prevents me from allowing this more explicitly. However, if I allow only the lower three movies in a season to play strategically, conditioning on the decisions of the better movies, as suggested by the argument above, the parameter estimates are quite similar to those I report below.

I continue with the logit specification of profits. I use the point estimates from the demand analysis, but allow for varying seasonal patterns in demand, focusing on the perceived seasonal pattern in underlying demand. Thus the specification for payoffs is

$$\widehat{\pi}_j(r_j, r_{-j}; \gamma) = \sum_{t=r_j}^{r_j+H} s_{jt}(r_j, r_{-j}; \gamma) = \sum_{t=r_j}^{r_j+H} \left[\frac{\exp(\theta_j - \lambda(t - r_j))}{x(\gamma) + \sum_{k \in \widetilde{J}_t} \exp(\theta_k - \lambda(t - r_k))} \right] \tag{20}$$

where $x(\gamma)$ is the market size effect, which depends on the model parameters and will be specified later on, and all other parameters are given by the demand estimation of the benchmark model. H is the length of the period that is taken into account by distributors, when making their release decision. The choice of H is guided by computational limitations, so I choose $H = 2$, thereby restricting distributors to base their decisions on the first three weeks after release. $\widetilde{J}_t$ is the set of movies that play on week t, which depends on the observed release dates of the non-competing movies, as well as on the strategies, r_j and r_{-j}, of the competing movies.

The general specification of the market size effect, $x(\gamma)$, namely the perceived seasonal pattern in underlying demand, is

$$x(\gamma) = \sum_{i=1}^{W} w_i \exp(-\tau_t^{(i)}) \tag{21}$$

where the w_i's are a set of weights, such that $w_i \geq 0$ and $\sum_{i=1}^{W} w_i = 1$, and $\{\tau_1^{(i)}, \tau_2^{(i)}, ..., \tau_{56}^{(i)}\}$ is a set of estimates for the seasonal pattern in underlying demand. I use three potential candidates for the seasonal pattern in demand: (i) the estimated underlying demand from the benchmark model; (ii) the observed seasonal pattern of industry sales; and (iii) the perfect competition estimates, which are produced by having the denominator of equation (7) averages to a constant.

Finally, there are a few technical details that I list. First, I normalize all candidates of market size sequence to have the same magnitude as the estimated underlying demand, so they are somewhat consistent with the logit market shares and the estimated fixed-effects. Second, to constrain the set of weights, w_i's, to add up to one, I set

$$w_l = \frac{\exp(\gamma_l)}{\sum_{i=1}^{W} \exp(\gamma_i)} \text{ for } l = 1, 2, ..., W \text{ with } \gamma_L = 0 \tag{22}$$

and use $\gamma_1, \gamma_2, ..., \gamma_{L-1}$ as the estimated parameters.

I do not model the choice of the order of moves, mainly to save on parameters. With the parameter estimates of the demand system, it turns out that the equilibrium outcome is not very sensitive to the order anyway, so the β parameters are not well identified by the data.

The results of several specifications are presented in Table 5. First, when I estimate only η, it can be seen that the estimated underlying demand has very little explanatory power, which is not the case when the seasonal pattern in industry sales is used. The estimated η is basically zero when I use only underlying demand, with hardly any change in the likelihood function compared to a purely randomized decision making. When I use the seasonal pattern in industry sales η becomes higher and significant, helping to explain the release pattern. When both of the seasonal patterns are used, one can see again that the underlying demand has no impact, obtaining a weight of zero, without changing the likelihood function compared to the case in which only the pattern in industry sales is used. Finally, when I also use the perfect competition approximation, it adds a little explanatory power to the model, with a weight coefficient that is roughly equal to the one on the pattern of industry sales. This is mainly because the two patterns are strongly correlated (with a correlation coefficient of 0.78), making it hard to distinguish between them.

Therefore, it turns out that the conclusions from the reduced form analysis still hold; the estimated seasonal pattern in demand cannot rationalize the observed release pattern. In general, the strategic considerations and the exact modelling of the timing game suggest that the seasonal pattern in total sales rationalizes the observed release pattern a little better than the one suggested by the perfect competition approximation. I conclude that the results suggest that the observed release pattern is more consistent with the seasonal pattern in total industry sales, with some deviations that are in line with the perfect competition pattern. These deviations can be explained by different aspects that are assumed away in my model, such as uncertainty, portfolio considerations, signalling, and repeated game effects. I discuss these in the next section.

6 Discussion

The evidence provided in this paper suggests that seasonality in underlying demand for motion pictures is much flatter and quite different from the seasonality of observed industry revenues. Moreover, it is shown that release date decisions, following the conventional wisdom in the industry, are taken under the implicit assumption that the observed seasonality in revenues is a good proxy for underlying demand. The results suggest that this assumption is incorrect.

While the results of this paper are, in part, based on out-of-sample prediction, and hence are more suggestive than conclusive, they clearly demonstrate the endogeneity problem that creates difficulty with the interpretation of the seasonality in industry sales. Similar difficulties may be present in other industries, in which, for example, sales and price cuts are prevalent in high demand seasons. A full understanding of the market mechanism would require estimating how much of the increased sales is due to higher demand, and how much is explained by more intense price competition.[41]

Taken at face value, the results of this paper call for a reassessment of the current release pattern of first-run movies. There are several different ways of interpreting these results. The first line of interpretations is consistent with the assumptions of optimizing behavior by studios. There are several features, which are not captured by the simple timing game I estimate, that may make movies concentrate on high demand weeks, more so than predicted by the model. First, uncertainty may play an important role. While a complete information equilibrium may have the movies spreading out over the different weeks, uncertainty may result in more movies being released on the better weeks. Under uncertainty these holiday weekends may look ex-ante more attractive for movies that would have otherwise stayed away, knowing they are not as good. This over-concentration may prove inefficient ex-post, but may be optimal ex-ante. Second, repeated game effects (see Chisholm (1999)) may also change the static "value" of releasing on a certain week. If releasing on, say, the Fourth of July makes the same distributor more likely to capture the same week in future years, this may make a Fourth of July release more attractive than it is estimated to be. Third, the fact that distributors release more than one movie, and hence take into account the effect of one release on the sales of other movies may make them spread their own movies more than would be otherwise predicted. Finally, while the demand estimates are quite robust to alternative specifications of the substitution patterns, these may be more important for the release decision.

[41] Cooper and Haltiwanger (1993) provide an example for the drastic change in the seasonality of automobile sales following an exogenous change in the timing of new model introductions.

Most of the explanations suggested above may help to explain why we observe a concentration of big budget movies released on big holiday weekends. In a simplistic world of two movies and two potential release dates, this would suggest that the simple model predicts each movie to release on a different week, while the above extensions may lead to both movies releasing on the week with the higher demand. None of the explanations, however, seems to reconcile the differences in the seasonal pattern between actual release dates and the estimated underlying demand. In particular, given our estimated underlying demand, the fact that hardly any movies are released on Labor Day is a striking empirical puzzle; within the two movies-two dates intuitive framework, this would imply *both* movies releasing on the *lower* demand week.

Thus, I would also suggest a different line of interpretation, a behavioral one. First, movie distributors may err in their assessment of the underlying demand in the industry. After all, the "learning from experience" argument, according to which economic agents cannot err for a long period, may not work in the motion picture industry. To learn from experience, distributors have to first obtain enough of it and then be able to use it properly. In particular, the information about the seasonal pattern comes only once a year, and the high uncertainty about movie quality makes inference difficult regarding the separation between the underlying demand and the movie quality effects. With each movie having its own identity, a control experiment of releasing the same (or very similar) movie in different dates is not feasible. For decades the industry has followed the same release pattern, according to which big hits are released on big weekends, such as Memorial Day, the Fourth of July, Thanksgiving, and Christmas. Thus, there are no natural experiments which make it easy to distinguish between higher movie quality and higher underlying demand. Any deviation from the "predicted" seasonal pattern, e.g. successful movies in October, is typically interpreted by industry observers as an extremely good movie in the wrong season rather than as a decent movie in a mediocre season.

Even if distributors are fully rational, conservatism may lead them to stick to the traditional release pattern. The results presented in this paper are very "out of sample", and one may think that industry practitioners cannot rely on such analysis in their decision making, given the high stakes involved. This conservatism may be magnified if we think of the institutional context and the potential agency costs in the industry. Top directors and actors do not want to see their films fail because of a poor marketing decision. Thus, considering the traditional release pattern in the industry, they frequently lobby for a traditionally good choice of release date. Distributors are likely to be conservative and satisfy these requests, rather than risk their jobs, reputation, and potentially future business, because of a movie failure that would be attributed to a poor release decision. They would rather stick to the traditional release pattern, according to which they can be adequately evaluated by the market.[42] This is not the only example where the motion picture industry seems to be conservative and to follow tradition. Other examples include the current uniform ticket pricing policy in the industry (see Einav and Orbach (2001) for a discussion), or the massive capacity expansion that took place in the 1990's and has recently led many of the largest chains of movie theaters to file for bankruptcy.

7 Conclusion and Further Research

This paper investigates the seasonal pattern in the U.S. motion picture industry, and its effects on release date decisions. The analysis is done in two steps. First, I break down the seasonal pattern in industry sales into two components: one attributed to underlying demand and the other attributed to the endogenous

[42] This is in the spirit of "you cannot be fired for buying IBM". For a formal treatment, see Zweibel (1995).

market reaction, which, in turn, creates a strong seasonal pattern in the available choice set of movies. I find that underlying demand accounts for about half of the seasonal variation in total sales. Moreover, it is estimated that the seasonal pattern of the estimated underlying demand is quite different from the observed seasonal pattern of industry sales.

In the second part of the paper I investigate the release date decision, which is one of the main strategic vehicles in the competition between movies' distributors. I construct a strategic timing game and take it to the data. I find that release date decisions are best rationalized by a perceived (by distributors) pattern in demand which is closely related to the seasonal pattern in industry sales, and not to the estimated underlying demand. This model suggests that distributors of movies could have done better by changing the release pattern of their movies. For example, they may consider releasing some of their big budget movies later in the summer or in January, rather than around Memorial Day, when underlying demand has yet to peak.

The results of this paper demonstrate the importance of endogenizing non-price product attributes, and exemplify the effects it has on decision making. They also highlight the difficulty of interpreting seasonality in sales, and the need to break it down into its two components: the underlying seasonality in demand and the endogenous market reaction. I believe that this is an important point that should be taken seriously in any analysis of seasonality.

The paper provides a general framework for analyzing strategic interaction of discrete nature. Taken at face value, the results of this analysis suggest that bounded rationality may affect not only individuals but also well established firms operating in a complex environment.[43] One may also find that the results of this paper provide some evidence for theories regarding managerial conservatism and herd behavior. For the motion picture industry, this implies that distributors prefer to release on traditional release dates, rather than to experiment with "new technology" and, for example, release good movies on Labor Day.

The current paper only provides a starting point in the analysis of timing effects in the industry, and may be extended in different directions. First, one may relax some of the major assumptions made in the model, and in particular incorporate uncertainty and movie-portfolio considerations into the analysis. Other extensions include bringing the exhibitors into the analysis, and more explicitly modelling their decisions regarding movie replacement. In an ongoing project, the structure of the timing game is planned to allow for a richer set of dynamic interactions, a la Caruana and Einav (2001). It would be interesting to see how much difference a richer game structure makes, compared to the simplified sequential game employed in this paper.

Finally, the unique and detailed data about the pre-release announcements, which is described in this paper, provide a unique opportunity to observe a complete play of the equilibrium path of a dynamic game, rather than only an equilibrium outcome as we usually use with typical data sets. A deeper analysis of the commitment opportunities and the strategic changes of announced release dates is the topic of another ongoing research project.

[43] Goettler and Shachar (2000) reach a similar conclusion for network television programming decisions.

References

[1] Berry, S. (1992), "Estimation of a Model of Entry in the Airline Industry," *Econometrica*, 60, 889-917.

[2] Berry, S. (1994), "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 25(2), 242-262.

[3] Berry, S., Levinsohn, J., and Pakes, A. (1995), "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841-890.

[4] Borenstein, S., and Netz, J. (1999) "Why Do All the Flights Leave at 8 a.m.? Competition and Departure-Time Differentiation in Airline Markets," *International Journal of Industrial Organization*, 17, 611-640.

[5] Caruana, G., and Einav, L. (2001), "A Theory of Endogenous Commitment," manuscript, Harvard University.

[6] Caves, R. E. (2000), *Creative Industries*, Harvard University Press

[7] Chisholm, D. C. (1999), "The War of Attrition and Optimal Timing of Motion-Picture Releases," manuscript, Lehigh University.

[8] Cooper, R., and Haltiwanger (1993), "Automobiles and the National Industrial Recovery Act: Evidence on Industry Complementarities," *Quaretrly Journal of Economics*, 108(4), 1043-1071.

[9] Corts, K. S. (2001), "The Strategic Effects of Vertical Market Structure: Common Agency and Divisionalization in the U.S. Motion Picture Industry," *Journal of Economics and Management Strategy*, 10(4), 509-528.

[10] De Vany, A., and Walls, W. D. (1996), "Bose-Einstein Dynamics and Adaptive Contracting in the Motion Picture Industry," *Economic Journal*, 106, 1493-1514.

[11] De Vany, A., and Walls, W. D. (1997), "The Market for Motion Pictures: Rank, Revenue and Survival," *Economic Inquiry*, 35, 783-797.

[12] Einav, L., and Orbach, B. Y. (2001), "Uniform Prices for Differentiated Goods: The Case of the Movie-Theater Industry," Harvard Olin Discussion Paper No. 337.

[13] Eliashberg, J., and Shugan, S. M. (1997), "Film Critics: Influencers or Predictors?" *Journal of Marketing*, 61(2), 68-78.

[14] Goettler, R., and Shachar, R. (2000), "An Empirical Analysis of Spatial Competition," *RAND Journal of Economics,* forthcoming.

[15] Lukk, T. (1997), *Movie Marketing: Opening the Picture and Giving it Legs*, Los Angeles, Silman-James Press.

[16] Mazzeo, M. J. (2000), "Product Choice and Oligopolistic Market Structure," manuscript, Northwestern Univeristy.

[17] Mcfadden, D. (1973), "Conditional Logit analysis of Qualitative Choice Behavior," in P. Zarembka, eds., *Frontiers of Econometrics*, New York, Academic Press.

[18] Moul, C. C. (2001), "Saturation and The Demand for Motion Pictures," manuscript, Washington University.

[19] Mulligan, J. G., and Motiere, L. (1994), "The Market for First-Run Motion Pictures," Working Paper No. 94-5, University of Delaware.

[20] Nelson, R. A., Donihue, M. R., Waldman, D. M., and Wheaton, C. (2001), "What's an Oscar Worth?" *Economic Inquiry*, 39(1), 1-16.

[21] Prag, J., and Casavant, J. (1994), "An Empirical Study of the Determinants of Revenues and Marketing Expenditures in the Motion Picture Industry," *Journal of Cultural Economics*, 18, 217-235.

[22] Radas, S., and Shugan, S. M. (1998), "Seasonal Marketing and Timing New Product Introductions," *Journal of Marketing Research*, 35, 296-315.

[23] Radas, S., and Shugan, S. M. (2001), "Services and Seasonal Demand," manuscript, Univeristy of Florida.

[24] Reiss, P. C. (1996), "Empirical Models of Discrete Strategic Choices," *American Economic Review*, 86(2), 421-426.

[25] Sawhney, M. S., and Eliashberg, J. (1996), "A Parsimonious Model for Forecasting Gross Box-Office Revenues of Motion Pictures," *Marketing Science*, 15(2), 113-131.

[26] Seim, K. (2000), "Geographic Differentiation and Market Structure: The Video Retail Industry," manuscript, Yale University.

[27] Squire, J. E. (1992), eds., *The Movie Business book*, New York, Simon and Schuster.

[28] *The New York Times*, various articles.

[29] Toivanen, O., and Waterson M. (2000), "Empirical Research on Discrete Choice Game Theory Models of Entry: An Illustration," *European Economic Review*, 44, 985-992.

[30] Vanderhart, J., and Wiggins, S. N. (2001), "Advertising, Word-of-Mouth, and Motion Picture Performance," manuscript, Texas A&M Univeristy.

[31] Verter, J., and McGahan A. M. (1996), "Coming Soon: A Theater Near You," Case Study No. 9-797-011, Harvard Business School.

[32] Vogel, H. (1994), *Entertainment Industry Economics*, Cambridge, Cambridge University Press.

[33] Zweibel, J. (1995), "Corporate Conservatism and Relative Compensation," *Journal of Political Economy*, 103(1), 1-25.

Appendix

Claim 1 *Given the assumption made in Section 5.1, if distributors are indifferent between all weeks and the seasonal effects and release pattern are stationary then the sequence of denominators is a constant.*

Proof. All we need to show is that cycles in the denominator are not possible. To see this, let me reformulate the argument using the following notation. Let the number of weeks in the year be W, and the sequence of denominators be given by $\{x_0, x_2, ..., x_{W-1}\}$. The fact that distributors are indifferent between all weeks implies that

$$\sum_{t=k}^{k+H} \delta^{t-k} y_{t \bmod W} = \sum_{t=l}^{k+H} \delta^{t-k} y_{t \bmod W} = c \;\forall k, l$$

where $\delta \equiv \exp(-\lambda) < 1$, $y_i = \frac{1}{x_i}$, and c is a constant.

Now, observe that:

$$c = \sum_{t=k}^{k+H} \delta^{t-k} y_{t \bmod W} = y_k + \delta \sum_{t=k+1}^{k+H+1} \delta^{t-(k+1)} y_{t \bmod W} - \delta^{H+1} y_{(k+H+1) \bmod W} = y_k + \delta c - \delta^{H+1} y_{(k+H+1) \bmod W}$$

Rearranging, we obtain that $c(1-\delta) = y_k - \delta^{H+1} y_{(k+H+1) \bmod W}$. We can repeat the same procedure replacing k with $k+H+1$ to get $c(1-\delta) = y_{k+H+1} - \delta^{H+1} y_{(k+2(H+1)) \bmod W}$ and so on, until $k = (k+n(H+1)) \bmod W$. Such n is guaranteed to exist, with $n \leq W$. Simple algebra allows me to obtain $y_k - y_{(k+H+1) \bmod W} = \delta^{n(H+1)} \left(y_k - y_{(k+H+1) \bmod W}\right)$, which can be possible only if $y_k = y_{(k+H+1) \bmod W}$. This guarantees that any pair of y's that are exactly H weeks apart from each other has the same denominator.

This almost completes the proof. The final step is to observe that:

$$c = \sum_{t=k}^{k+H} \delta^{t-k} y_{t \bmod W} = y_k + \delta y_{k+1} + \delta^2 \sum_{t=k+2}^{k+H+2} \delta^{t-(k+2)} y_{t \bmod W} - \delta^{H+1} y_{(k+H+1) \bmod W} - \delta^{H+2} y_{(k+H+2) \bmod W}$$

but given that any pair of y's that are exactly H weeks apart from each other has the same denominator we can substitute y_k and y_{k+1} for $y_{(k+H+1) \bmod W}$ and $y_{(k+H+2) \bmod W}$ respectively, to obtain $c = (1-\delta^{H+1})(y_k + \delta y_{k+1}) + \delta^2 c$. Rearranging, we have that $y_k + \delta y_{k+1} = \frac{(1-\delta^2)c}{1-\delta^{H+1}}$. Now set $c' \equiv \frac{(1+\delta)c}{1-\delta^{H+1}}$, and we are back to the original problem with $c = c'$ and $H = 1$. Thus, we can repeat the same procedure to obtain that every adjacent y's are equal. *Q.E.D.* ■

Table 1: Industry Trends

This table shows the trends in key industry variables over the sample period. It can be seen that the industry has been steadily growing, with the number of titles increasing by 40-50% and the revenues and admissions almost doubling from 1985 to 1999.

year	Avg Ticket Price (12/1999 $)	Population (millions as of July 1)	Titles (all)	Titles (wide)	Revenues[a] (12/1999 $B)	Annual admissions (per person)
1985	5.54	237.92	167	103	3.51	2.69
1986	5.70	240.13	222	120	3.98	2.94
1987	5.78	242.29	220	120	4.28	3.08
1988	5.84	244.50	230	144	4.37	3.08
1989	5.40	246.82	213	115	4.92	3.73
1990	5.45	249.44	223	119	4.90	3.64
1991	5.20	252.13	224	125	4.63	3.57
1992	4.97	255.00	214	121	4.85	3.86
1993	4.83	257.75	227	144	5.09	4.14
1994	4.74	260.29	233	142	5.13	4.20
1995	4.80	262.77	246	149	5.27	4.23
1996	4.74	265.19	262	152	5.52	4.44
1997	4.81	267.74	250	150	5.90	4.63
1998	4.91	270.30	264	139	6.20	4.73
1999	5.00	272.88	244	143	6.80	5.03

Table 1: Industry Trends

[a] Revenues include only first ten weeks since release

Table 2: Market Shares and Ranking of Top Distributors

This table presents the annual market shares (top panel) and rank (bottom panel) of the leading distributors. It can be seen that the identity of the top distributors, as well as their overall market share (around 90%), were quite stable across the sample period. However, their individual market shares and ranking were not stable. This should be mainly attributed to the fact that the market shares are strongly driven by a few hit movies every year. Therefore, the (somewhat random) allocation of these top movies dictates the annual market shares.

	'85	'86	'87	'88	'89	'90	'91	'92	'93	'94	'95	'96	'97	'98	'99
Buena Vista (Disney)	4.1	10.7	14.5	18.2	14.0	15.1	14.1	18.2	16.2	19.7	19.5	21.1	14.4	17.4	17.2
Warner Brothers	18.3	11.1	13.1	11.1	17.3	12.0	14.2	20.5	19.2	16.8	17.2	16.2	11.4	11.4	15.0
Paramount	7.7	20.1	17.9	15.5	13.7	14.1	12.8	10.1	10.0	14.4	9.6	13.2	12.5	14.1	12.1
Columbia	9.5	9.5	4.8	2.0	7.5	5.4	9.1	12.9	11.9						
Tristar	10.5	7.8	7.1	6.3	8.6	8.2	11.3	6.6	6.5						
Sony										10.0	13.6	10.9	20.7	11.7	9.5
Universal	15.7	8.6	7.7	10.8	16.2	13.6	11.9	11.1	14.1	11.8	13.0	8.6	10.2	6.0	13.5
20th Century Fox	10.1	8.9	9.4	11.5	6.7	14.1	10.2	14.5	11.7	9.0	8.1	12.9	11.0	11.1	11.0
New Line	0.4	0.5	1.9	2.0	1.5	4.5	4.1	1.6	3.6	6.9	6.2	4.9	5.7	8.2	4.2
MGM/UA	8.6	5.0	4.5	9.7	5.1	3.0	2.5	1.3	1.9	3.1	6.7	4.6	2.8	3.1	4.4
Miramax			0.1	0.1	0.9	0.7	1.2	0.9	1.6	3.6	2.5	4.4	5.2	5.1	3.3
Orion	5.1	6.8	8.3	7.2	4.4	6.3	6.3	0.3	0.5	0.4	0.1	0.7	0.1		
Dreamworks													1.6	7.3	4.3
Cannon	3.3	3.1	1.5	1.1	0.8	0.03	0.13		0.02	0.01					
Others	6.8	7.9	9.3	4.5	3.2	3.0	2.1	1.9	2.7	4.4	3.5	2.5	4.2	4.5	5.5
C4	54.5	51.4	54.9	56.3	61.3	56.9	53.1	66.2	61.4	62.6	63.3	63.4	59.2	54.6	57.8
C8	85.4	83.5	82.8	90.3	89.2	88.9	90.0	95.6	93.2	92.1	93.9	92.5	91.2	87.3	87.1

	'85	'86	'87	'88	'89	'90	'91	'92	'93	'94	'95	'96	'97	'98	'99
Buena Vista (Disney)	9	3	2	1	3	1	2	2	2	1	1	1	2	1	1
Warner Brothers	1	2	3	4	1	5	1	1	1	2	2	2	4	4	2
Paramount	7	1	1	2	4	2	3	6	6	3	5	3	3	2	4
Columbia	5	4	8	10	6	8	7	4	4						
Tristar	3	7	7	8	5	6	5	7	7						
Sony										5	3	5	1	3	6
Universal	2	6	6	5	2	4	4	5	3	4	4	6	6	8	3
20th Century Fox	4	5	4	3	7	3	6	3	5	6	6	4	5	5	5
New Line	11	11	10	9	10	9	9	8	8	7	8	7	7	6	9
MGM/UA	6	9	9	6	8	10	10	9	9	9	7	8	9	10	7
Miramax			12	12	11	11	11	10	10	8	9	9	8	9	10
Orion	8	8	5	7	9	7	8	11	11	10	10	10	11		
Dreamworks													10	7	8
Cannon	10	10	11	11	12	12	12		12	12					

Table 2: Market Shares and Ranking of Top Distributors

Sony acquired Columbia Pictures and Tristar Pictures in 1994, creating Sony Pictures.
All figures are based on the first ten weeks of the films.

Table 3: Descriptive Statistics of Key Variables

This table presents descriptive statistics of key variables. There are a few interesting facts to note. First, by comparing the mean to the median it can be seen how the distribution of movie revenues and production cost is skewed. This is not so great for advertising expenditure. Second, the cumulative population share for the top grossing movie, Titanic, is about one third. Third, two of the top five advertisers (and five of the top ten) are children's movies. This may suggest that children's movies advertise not only the film, but also its accompanying industry (toys etc.).

	Total Revenues (current U.S. million dollars)		Population Share* (revenues normalized by price and population)		Week 1 Revenues (current U.S. million dollars)	
Number of Obs.	1,956		1,956		1,956	
Mean	34.8		2.78%		10.2	
Median	20.0		1.71%		6.63	
Standard Dev.	43.5		3.16%		10.8	
Top 5	Titanic	600.8	Titanic	31.75%	The Lost World	107.4
	Star Wars (1999)	430.9	Jurassic Park	28.20%	Star Wars (1999)	99.3
	Jurassic Park	357.1	Star Wars	26.65%	Austin Powers 2	84.7
	Forrest Gump	329.7	Batman	23.17%	Jurassic Park	81.7
	Lion King	312.9	Lion King	22.51%	Independence Day	79.2

	Production Cost (Dec 1999 U.S. million dollars)		Advertising Exp (Dec 1999 U.S. million dollars)	
Number of Obs.	1,604		1,873	
Mean	29.7		8.47	
Median	23.2		7.19	
Standard Dev.	22.6		5.77	
Top 5	Titanic	208.7	Toy Story	42.5
	Waterworld	193.1	Titanic	35.1
	Armageddon	154.7	The Rookie	34.5
	Speed 2	152.2	Anastasia	32.5
	Tarzan	151.9	Forrest Gump	30.9

Table 3: Descriptive Statistics of Key Variables

* Industry share is based on first ten weeks only

Table 4: Projection of the Estimated Movie Fixed-Effects over Observables

This table presents a set of regressions in which the dependent variable is the estimated movie fixed-effects from the benchmark model, and the regressors are sets of movie characteristics. It should be noted that we should not interpret these regressions as a causality. This is true in particular with respect to the advertising expenditure (some of which is spent after the movie has been released), and the Academy Awards (which are more likely to go to movies that did well at the box office). Still, the estimated parameters are quite sensible, and suggest that the estimated fixed-effects do capture movie quality. Note also that the low R^2's in all regressions motivates the original argument that fixed effects are needed in the absence of good predictors of movie quality.

Log(Production Cost)		0.454^{**} (0.043)	0.416^{**} (0.046)	0.813^{**} (0.043)	0.724^{**} (0.042)	0.813^{**} (0.042)	0.727^{**} (0.041)
Log(Advertising Expenditure)		0.751^{**} (0.040)	0.522^{**} (0.039)	-	-	-	-
MPAA Rating	PG-13	-0.247^{**} (0.080)	-0.312^{**} (0.085)	-0.267^{**} (0.088)	-0.315^{**} (0.089)	-	-
	R	-0.302^{**} (0.082)	-0.408^{**} (0.086)	-0.285^{**} (0.090)	-0.349^{**} (0.091)		
	>R	-0.819 (0.757)	-0.981 (0.805)	-1.377 (0.854)	-1.369 (0.865)		
Genre	Comedy	-0.101 (0.074)	-0.166^{*} (0.078)	0.065 (0.081)	0.006 (0.081)	-	-
	Drama	-0.201^{**} (0.077)	-0.289^{**} (0.081)	-0.046 (0.083)	-0.103 (0.084)		
	Children	-0.141 (0.112)	-0.224 (0.119)	-0.076 (0.123)	-0.124 (0.124)		
Best Picture	Nominee	0.865^{**} (0.168)	0.995^{**} (0.179)	-	-	-	-
	Award	0.957^{**} (0.288)	1.210^{**} (0.306)				
Year Dummies		yes	no	yes	no	yes	no
No. of Observations		1,542	1,542	1,603	1,603	1,603	1,603
Adjusted R^2		0.374	0.286	0.197	0.171	0.190	0.161

Table 4: Projection of The Estimated Movie Fixed Effects over Observables

In all regressions, the dependent variable is the estimated movie fixed effects from the benchmark model.
Recall that the dependent variable in the benchmark model is roughly proportional to the logarithm of market share, so the coefficients on production cost and advertising expenditure can be interpreted as elasticities.
Standard errors in parentheses.
* Significant at the 5% level, ** Significant at the 1% level.
[a] Base category for MPAA rating is unrestricted and PG movies, [b] Base category for genre is Action movies.

Table 5: Timing Game Estimation Results

This table presents the results from a set of specifications of the timing game. In all specifications I use a fully random order of moves (equal probability for each permutation). One can see that the estimated underlying demand has very little explanatory power, while the industry sales pattern significantly helps in explaining the release decisions. When the perfect competition series is incorporated (column (4)), it takes some of the explanatory power from the industry sales pattern. This is not surprising, however, given the strong colinearity between the two series (correlation coefficient of 0.78). The addition of the perfect competition series adds some, but not much, explanatory power to the model, as can be seen by the value of the likelihood function.

	(1)	(2)	(3)	(4)
η (coefficient on profits)	0.37 (1.72)	5.30*** (1.76)	5.30*** (1.83)	5.26** (2.34)
γ (underlying demand)	0 (imposed)	-	-32.74 (34.43)	-28.97 (30.23)
[implied weight]	[1]		[0]	[0]
γ (industry sales)	-	0 (imposed)	0 (imposed)	0.24 (3.32)
[implied weight]		[1]	[1]	[0.56]
γ (perfect competition)	-			0 (imposed)
[implied weight]				[0.44]
order	random	random	random	random
$\log(L)$	-289.65	-282.12	-282.12	-281.34

Table 5: Timing Game Estimation Results - Parameters

Standard errors in parentheses.

Given the specification, which uses exponentials to restrict the weights to be between zero and one, whenever the implied weights are close to the boundary (to zero or one) the likelihood is very flat, and the calculated standard errors are almost meaningless. In such cases the implied weight is almost certainly at the boundary, whatever the exact value of γ is.

* Significant at the 10% level, ** Significant at the 5% level, *** Significant at the 1% level.

Figure 1: The Structure of the Motion Picture Industry

This figure sketches the structure of the industry. As can be readily understood from the names, motion pictures are produced by the producers, distributed to theaters by the distributors, and exhibited to the public by the exhibitors (theaters). Generally, production and distribution are vertically integrated within the structure of the major studios (e.g. MGM, Paramount, or Disney), while exhibitors operate separately. Some of the lower budget films are produced by independent producers, and distributed by either the distribution arm of the major studios or by independent distributors.

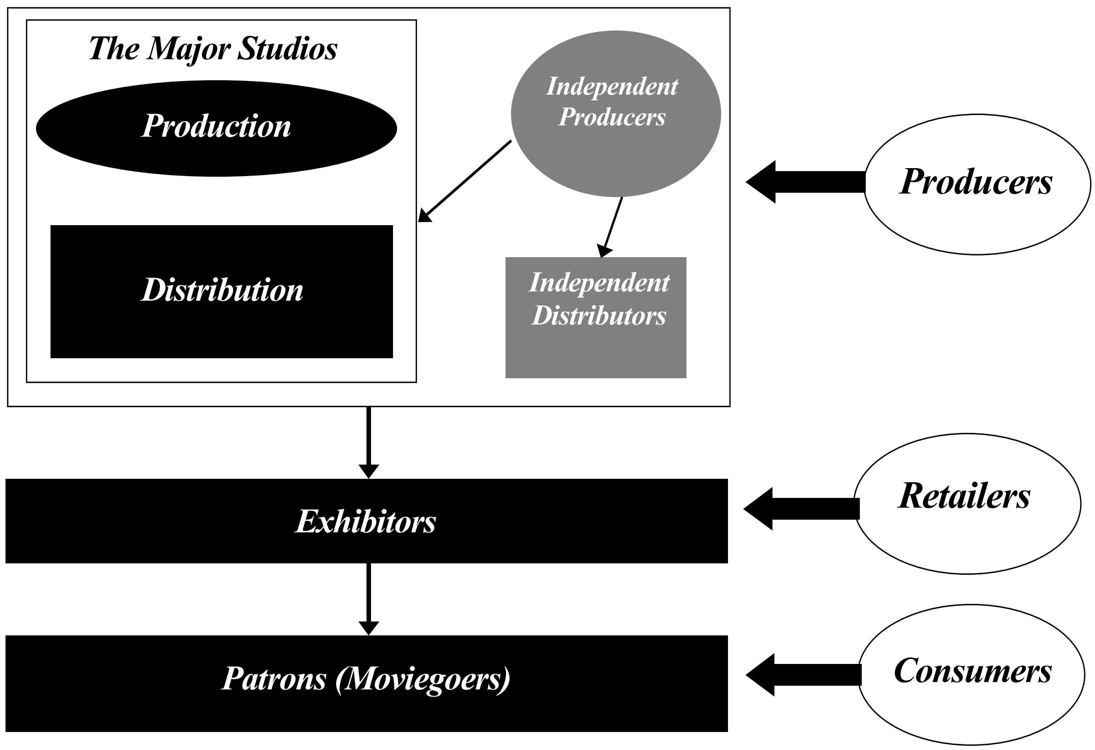

Figure 1: The Structure of the Motion Picture Industry

Figure 2: Distribution of Movie's Box Office Revenues over Its Life Cycle

This figure shows the distribution of total box office revenues over the movie's life cycle. The bars stand for the week-by-week share, while the lines stand for the cumulative share as of the end of the corresponding week. It can be seen that most of the revenues are concentrated in the first few weeks, with the first week accounting, on average, for almost 40% of the eventual box office revenues, and the first four weeks accounting for about 80% of them. Once I weight the averages by the gross box office revenues of the different films (white bars and dashed line), the distribution is less skewed and has a wider tail, suggesting that revenues of bigger movies decay slower. It is still the case, however, that much of the revenues are collected early on – almost 30% of the revenues are collected in the first week, and 50% in the first two weeks.

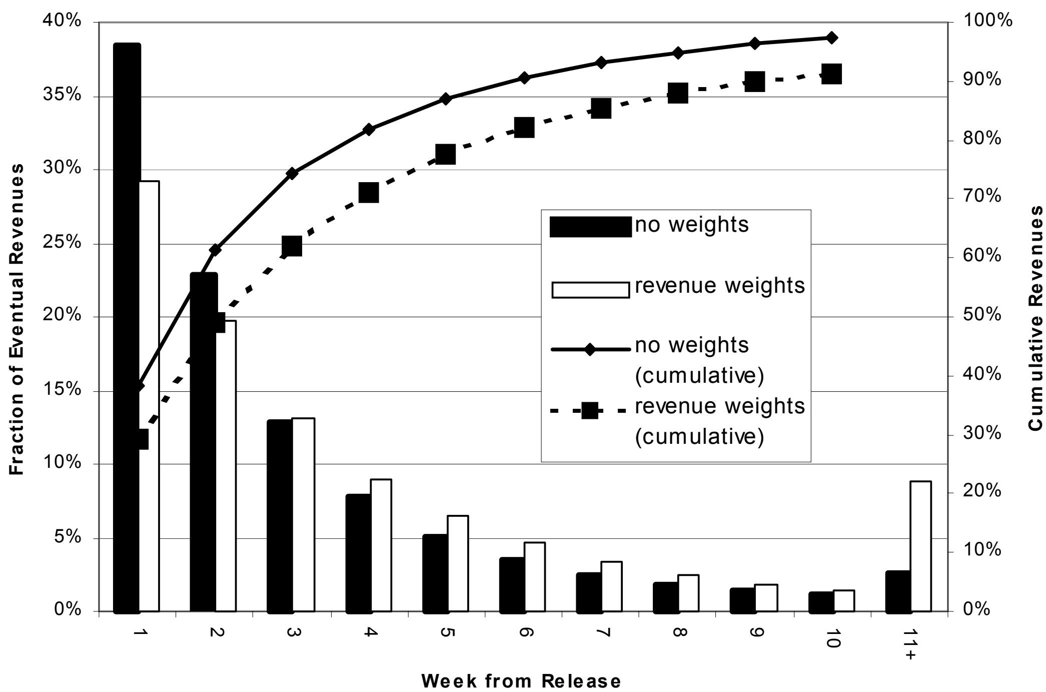

Figure 2: Distribution of Movie's Box Office Revenues over Its Life Cycle

The "no weights" series calculates the simple mean of these percentages for all movies in the data by calculating weekly percentages for each film separately, and then taking averages of these percentages over all movies. The "revenue weights" series calculates a weighted average, where the weights are proportional to the total box office revenues of each movie, so that higher weights are given to bigger movies. This weighted average is equivalent to adding up total revenues by week, and normalizing them by the total box office revenues in the sample.

Figure 3: Seasonal Effects in Total Admissions

This figure shows the seasonality in total industry sales. The vertical line ("industry share") is the industry's weekly revenues, normalized by average ticket price and by the U.S. population. Thus, it can be thought of as the per capita number of movies seen each week (an industry share of, say, 0.1 implies that one of every ten people in the U.S. goes to the movies in the corresponding week). The figure shows the industry shares, averaged over the fifteen years of the sample period.

The figure clearly demonstrates the perceived seasonal effects in the industry. The year has two strong periods, the long summer period (Memorial Day to Labor Day) and the Christmas-Winter holiday period. The Spring and the Fall are typically considered very weak periods for the industry, and the drop after Thanksgiving is generally seen as a "shopping period".

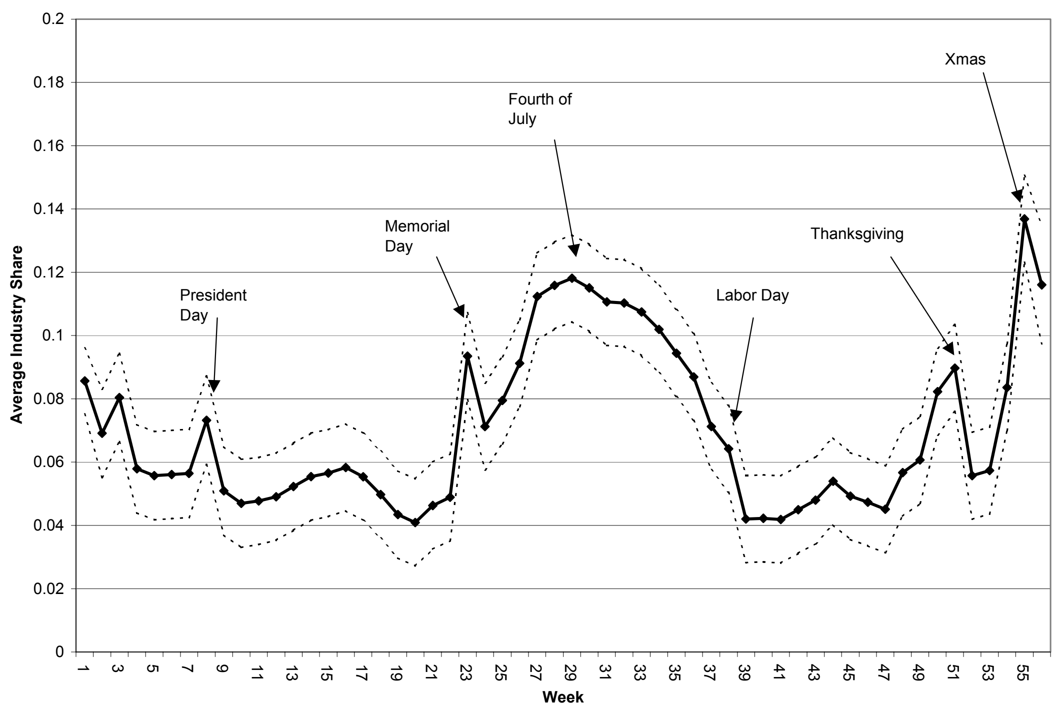

Figure 3: Seasonal Effects in Total Admissions

The dashed lines stand for deviations of two standard errors.
The number of weekly dummies is 56 to account for the timing variation in U.S. holidays across the years (see Section 3.1).

Figure 4: Seasonal Effects in Total Production Cost Released

This figure shows the seasonality in release decisions. In the absence of readily available measures of movie quality, I use the production cost (in December 1999 $M) as a proxy for quality (later on, in Figure 15, I present a similar figure based on the estimated movie quality). The figure is constructed by adding up the production cost of all movies released in a certain week, then taking averages over the fifteen years of the sample period.

It is clearly demonstrated that the release pattern follows some seasonal patterns. Big budget movies are generally released early in the summer and before Christmas, with high concentration on a few big holiday weekends – Memorial Day, the Fourth of July, Thanksgiving, and Christmas. In contrast, Labor Day, viewed as a symbol of the end of the attractive summer season, has only a few releases of low budget movies.

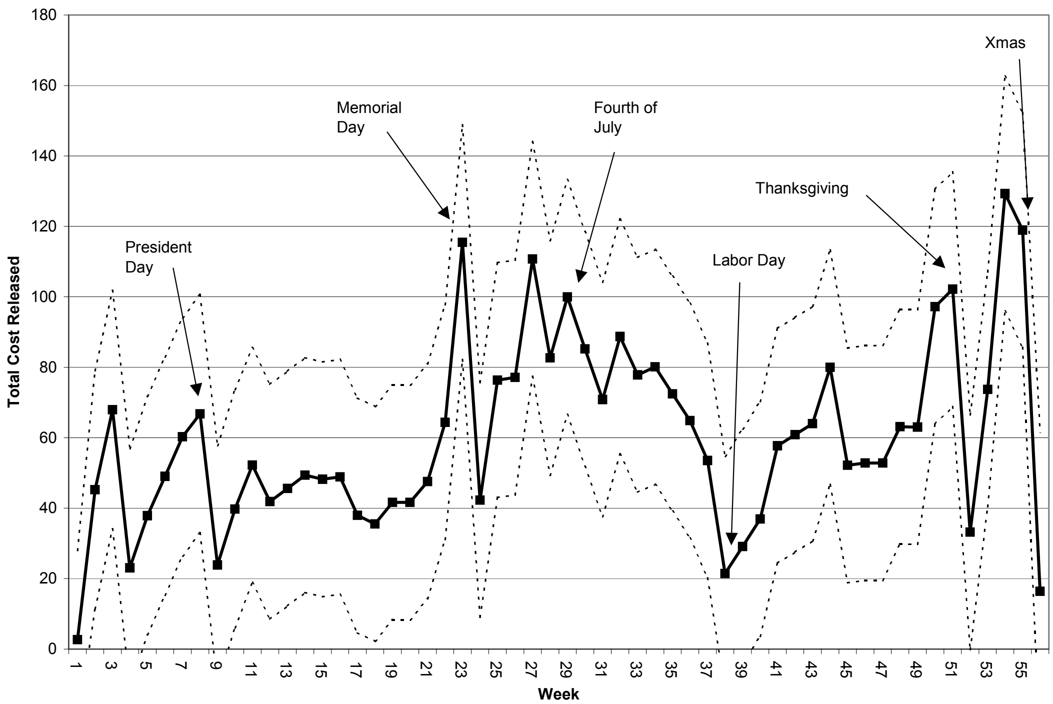

Figure 4: Seasonal Effects in Total Production Costs Released

The dashed lines stand for two standard errors from the average.
The number of weekly dummies is 56 to account for the timing variation in U.S. holidays across the years (see Section 3.1)

Figure 5: Seasonal Effects in Number of Titles Released

This figure shows the seasonality in the number of movies released. The figure is constructed by counting the number of movies released in a certain week, then taking averages over the fifteen years of the sample period. With the exception of a few notable examples, there does not seem to be a strong seasonal pattern in the number of releases. This helps us to understand the previous figure, Figure 4, as driven mainly by seasonality in movie quality and not by seasonality in the number of movies released.

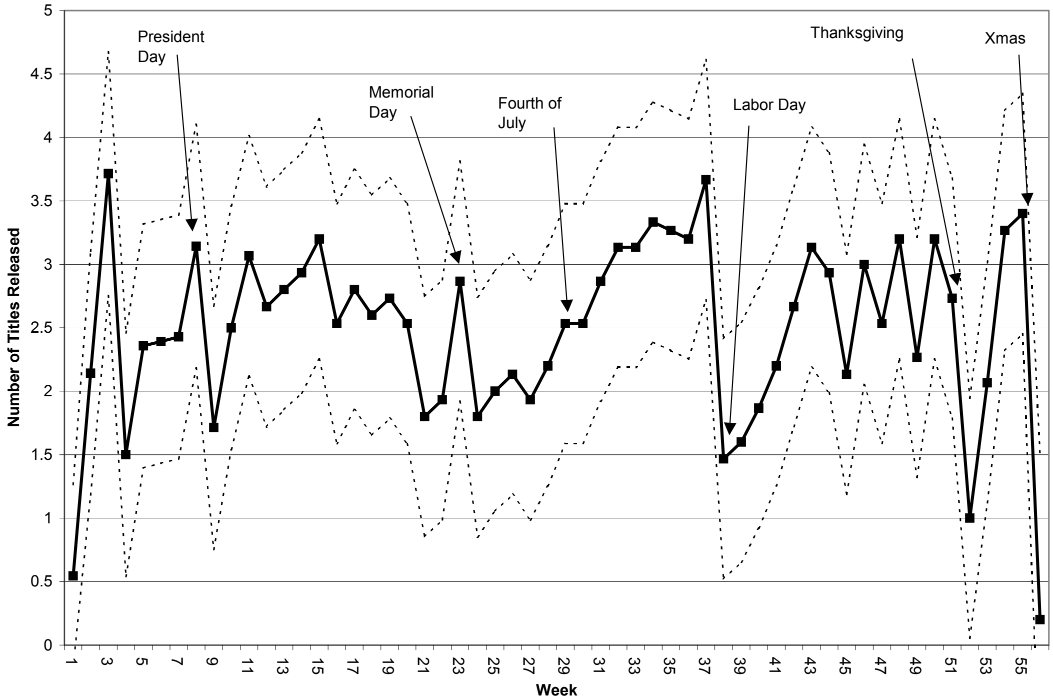

Figure 5: Seasonal Effects in Number of Titles Released

The dashed lines stand for two standard errors from the average.
The number of weekly dummies is 56 to account for the timing variation in U.S. holidays across the years (see Section 3.1)

Figure 6: Distribution of Timing of the Initial Release Date Announcement

This figure shows the distribution of the timing of the initial release announcement. The initial release announcement is considered to be the earliest month from which the movie *continuously* announces a specific release date. I then calculate the difference (in months) between the timing of the initial announcement and the initial release date specified. Note that this initial release date may be changed later on. However, given that most of the changes generally do not shift the release date by more than few weeks, the distinction between the actual release date and the release date initially announced is negligible for this figure.

The white and grey bars present the same distribution, considering only big movies. For big movies I use the top 20 movies, measured by either the production cost or by the estimated quality from the demand system, every year. This gives us 300 movies over all, compared to 1,897 movies used for the general distribution. The figure demonstrates that bigger movies tend to announce their release date earlier than the average movie.

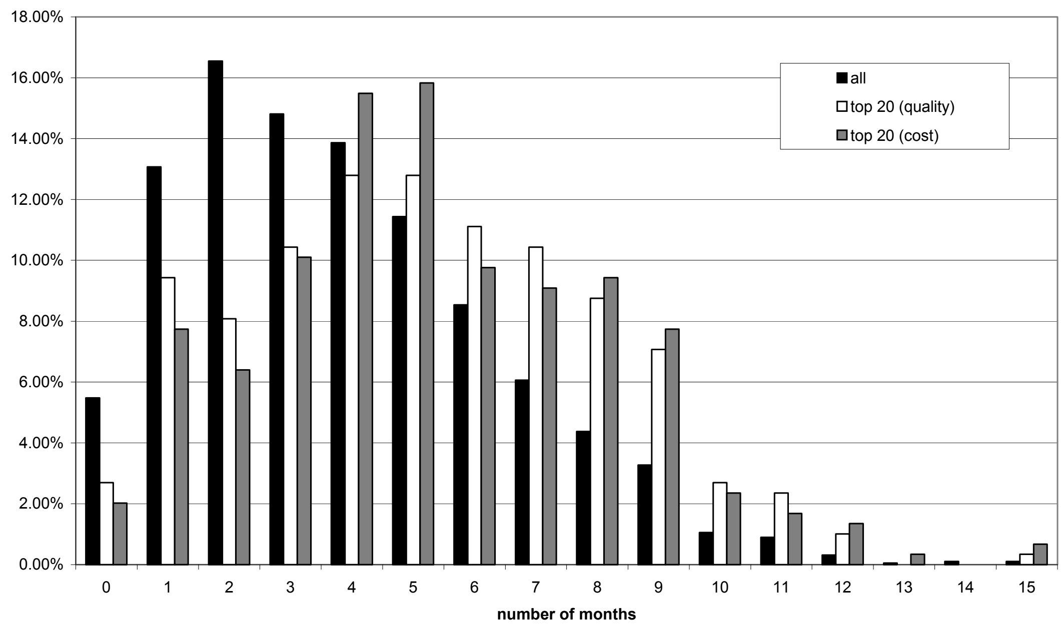

Figure 6: Distribution of Timing of the Initial Release Date Announcement

Figure 7: Distribution of the Magnitude of Switches in Announced Release Dates

This figure shows the distribution of the magnitude of switches of announced release dates. A switch is a change in the announced release date compared to the most recent announcement of the same movie (which is generally made a month before). The distribution is taken over all movies and announcements, and tabulates the difference (in weeks) between the new announcement and the previous one. A difference of zero implies no change (which is the case for about 80% of the announcements). A positive difference implies a shift forward of the release date, and a negative difference implies making the release date earlier than announced before.

This figure provides two main insights. First, the distribution is roughly symmetric (with a somewhat fatter tale in the positive part, for obvious reasons). Second, the majority of the changes shift the release date by a small number of weeks. These two observations suggest that these changes are done mainly for strategic reasons, and not because of exogenous factors, such as production delays.

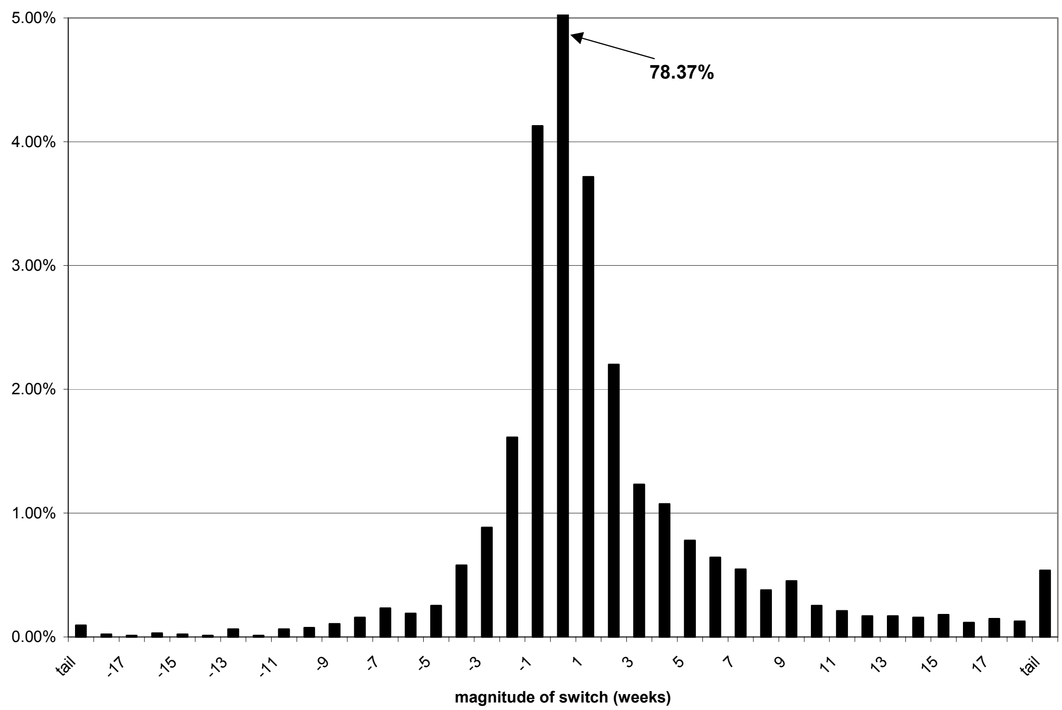

Figure 7: Distribution of the Magnitude of Switches in Announced Release Dates

Note that the bar at zero is out of scale, and accounts for 78% of the announcements. This implies that one of every five release date announcements is a change compared to previously announced release date.

Figure 8: Estimated Seasonal Effects in Demand without Controlling for Movie Fixed Effects

This figure plots the estimated coefficients on the weekly dummy variables, when movie fixed-effects are *not* incorporated into the specification, i.e. $\delta_{jt} = \tau_t - \lambda(t - r_j) + \xi_{jt}$. The estimated seasonal pattern roughly follows the seasonal pattern in industry sales, presented in Figure 3. It clearly shows the two high seasons, the summer and the winter holiday period. One can observe, however, some small differences, such as a decrease in the relative effects of Memorial Day and Thanksgiving, and an increase in the relative effect of Labor Day. These differences are mainly driven by the fact that the unit of analysis is now the individual movie, rather than the industry. In addition, note that this figure's similarity to the general pattern in industry sales suggests that adding a constant decay does not by itself drive the results of the benchmark model, presented next.

Figure 8: Estimated Seasonal Effects in Demand without Controlling for Movie Fixed Effects

The dashed lines stand for deviations of two standard errors.
The estimated decay coefficient is -0.391 with a standard error of 0.0039.
The number of observations: 16,103 (1,956 titles).
The number of weekly dummies is 56 to account for the timing variation in U.S. holidays across the years (see Section 3.1)

Figure 9: Seasonal Effects in Underlying Demand

This figure plots the estimated coefficients on the weekly dummy variables from estimating the benchmark model, i.e. $\delta_{jt} = \theta_j + \tau_t - \lambda(t - r_j) + \xi_{jt}$. There are two major differences, compared to the estimates from a model that does not control for movie quality (Figure 8). First, the variation in the estimated coefficients is much smaller (the standard deviation of the 56 point estimates is 0.20 compared to 0.37), suggesting that about half of the seasonal variation is explained by variation in quality. Second, the seasonal pattern is quite different, shifted forward by about a month.

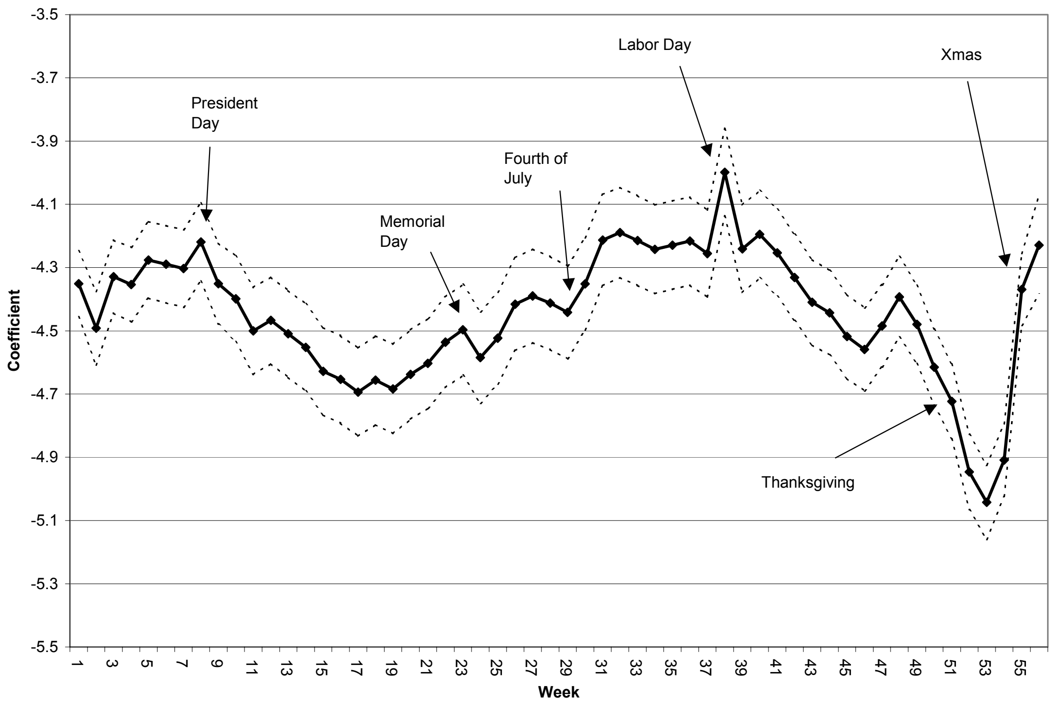

Figure 9: Seasonal Effects in Underlying Demand

The dashed lines stand for deviations of two standard errors.

The estimated decay coefficient is -0.460 with a standard error of 0.0019.

The number of observations: 16,103 (1,956 titles).

The number of weekly dummies is 56 to account for the timing variation in U.S. holidays across the years (see Section 3.1)

Figure 10: Estimated Decay Pattern

This figure shows the decay pattern in two cases, when I allow for a flexible functional form, using ten different dummy variables, and when I impose linear decay. It is easy to see that the linearity assumption is hardly restrictive. Using the flexible function form changes almost nothing in the estimates of the other parameters of the model.

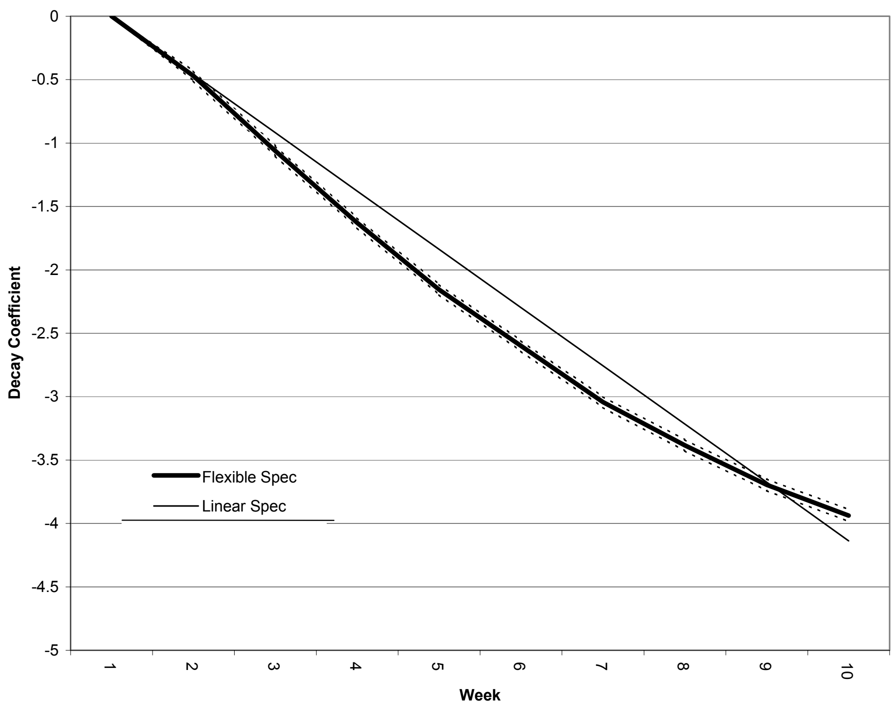

Figure 10: Estimated Decay Pattern

The dashed lines stand for deviations of two standard errors.

Figure 11: Trends in Estimated Movie Quality

Given that a general trend in movie quality cannot be identified in the model, this figure shows the trend in estimated movie fixed-effects, which capture the general trend in quality over the sample period. Note, however, that the quality is measured with respect to the outside good, which is arguably increasing over the sample period. Thus, the relative flatness of the graph does not imply that movies' quality does not improve, but only that it does not improve faster than the rate of improvement in the outside good. Note also that the number of titles released is increasing over time (see Table 1), which may account for the slight downward trend in the mean and median quality. This trend only disappears once the top movies are accounted for.

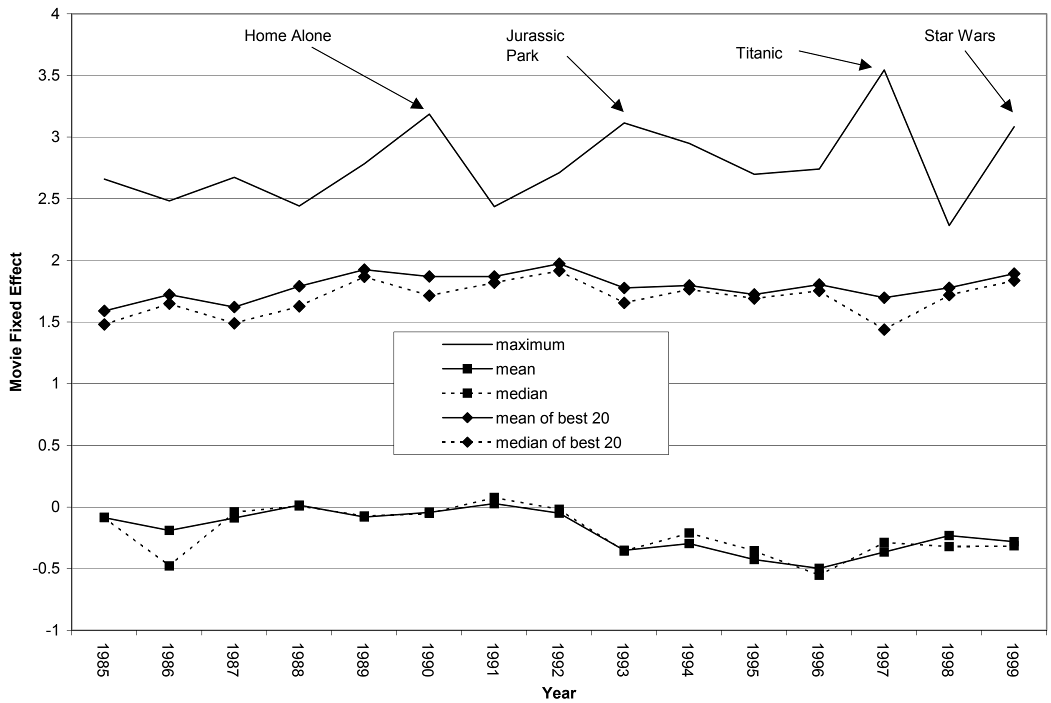

Figure 11: Trends in Estimated Movie Quality

Figure 12: Underlying Demand by Genre

This figure presents the estimated underlying demand when estimated for each genre separately, as if there is no substitution between different genres. It can be seen that the pattern is roughly similar for all genres, suggesting that the overall results are not driven by one genre in particular. The children's movies do show some difference in the estimates, with the summer being higher compared to the rest of the year. This is, in fact, encouraging and follows the conventional wisdom, according to which we expect children to have higher demand in the summer season, a pattern less clear for adults.

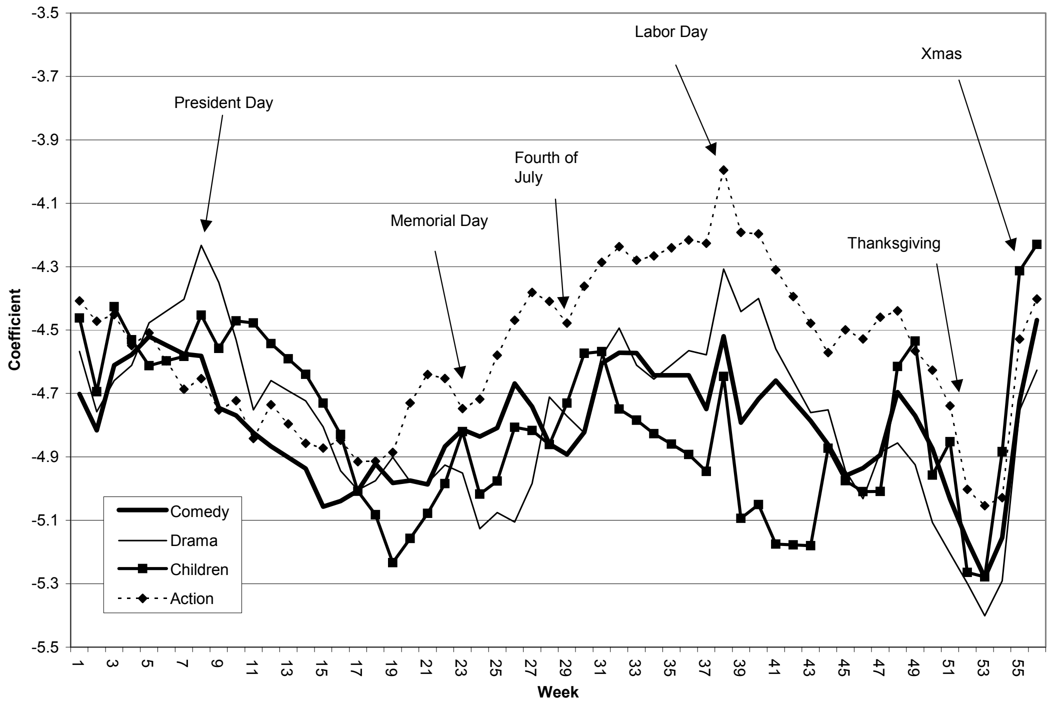

Figure 12: Underlying Demand by Genre

Additional important results from these four regressions:

	Action	Comedy	Drama	Children
Decay Coefficient	-0.484	-0.443	-0.449	-0.415
(standard error)	(0.0035)	(0.0033)	(0.0044)	(0.0057)
n (titles)	5,131 (627)	5,777 (697)	3,386 (416)	1,809 (216)

The number of weekly dummies is 56 to account for the timing variation in U.S. holidays across the years (see Section 3.1)

Figure 13: Underlying Demand by Period

This figure presents the estimated underlying demand when estimated for each five-year span separately. It can be seen that the pattern is roughly similar for all periods, although there is an obvious trend over time. First, the seasonal variation in underlying demand increases over time. Second, the decay rate becomes significantly faster over time.

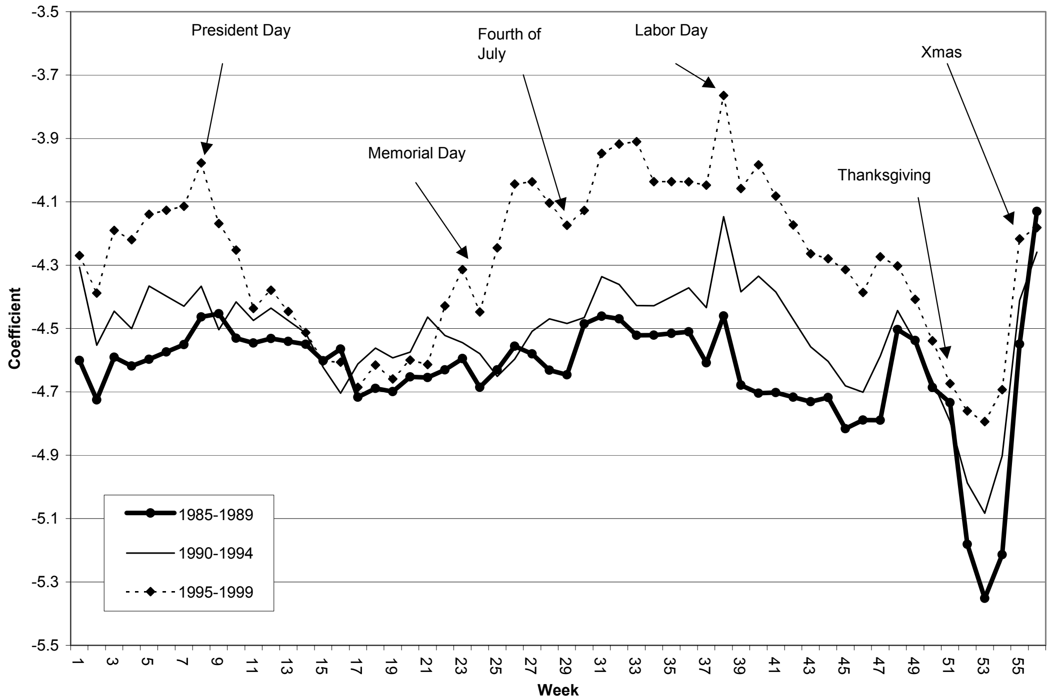

Figure 13: Underlying Demand by Period

Additional important results from these three regressions:

	1985-1989	1990-1994	1995-1999
Decay Coefficient	-0.362	-0.430	-0.532
(standard error)	(0.0032)	(0.0030)	(0.0031)
n (titles)	4,035 (572)	5,521 (666)	6,547 (754)

The number of weekly dummies is 56 to account for the timing variation in U.S. holidays across the years (see Section 3.1)

Figure 14: Seasonal Effects in Competition

This figure plots the estimated competition effect ($\sum_{k \in J_t} \exp(\theta_k - \lambda(t - r_k))$) averaged over the fifteen years of the data. The competition effect is defined in equation 5 and discussed in Section 5.1). The figure shows that the seasonal pattern in competition effect closely follows the seasonal pattern of industry sales (as shown in Figure 3). Clearly, the pattern of competition effect is driven by the pattern in release dates, which is shown in the next figure (Figure 15).

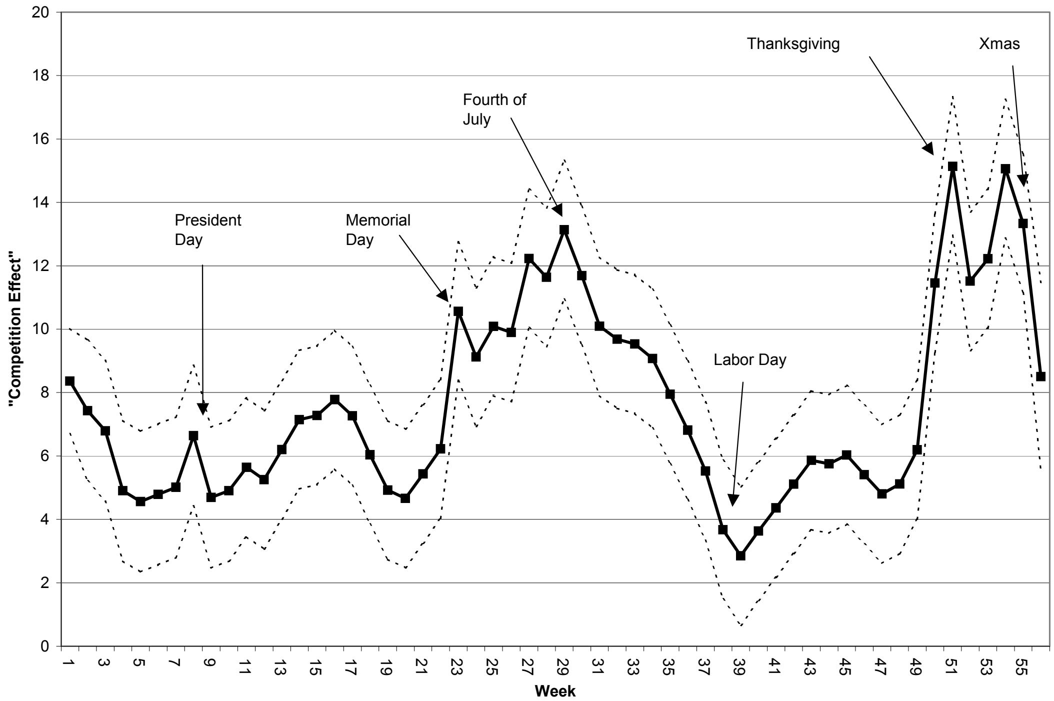

Figure 14: Seasonal Effects in Competition

The dashed lines stand for two standard errors from the average (ignoring the implicit standard error that comes from coefficient uncertainty).
The number of weekly dummies is 56 to account for the timing variation in U.S. holidays across the years (see Section 3.1)

Figure 15: Seasonal Effects in New Releases

This figure plots the estimated release effect, which is defined as the competition effect for only those movies that were released on the same week, i.e. $\sum_{k \in J_t, r_n = t} \exp(\theta_k)$. It is then averaged over the fifteen years of the data. The release effect is quite similar to Figure 4, which plots the seasonal pattern in total cost released. This similarity is quite remarkable, given the fact that the production cost is a very weak predictor of movie quality (see Table 4).

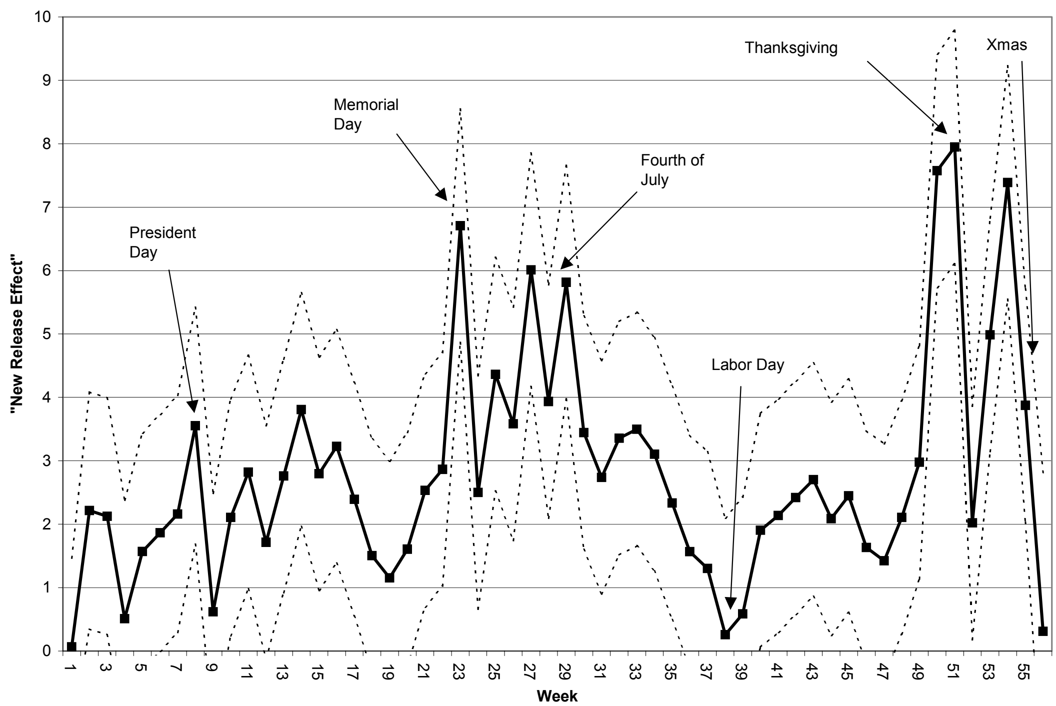

Figure 15: Seasonal Effects in New Releases

The dashed lines stand for two standard errors from the average (ignoring the implicit standard error that comes from coefficient uncertainty).

The number of weekly dummies is 56 to account for the timing variation in U.S. holidays across the years (see Section 3.1)

Figure 16: Predicted Normalized Revenues for the Marginal Movie, by Release Date

This figure uses Equation (6) in order to calculate the predicted box office revenues from releasing a certain movie at different weeks. The profits are then normalized by the predicted profits of releasing at week 1. The predicted profits are calculated using the point estimates of the seasonal effects in demand (reported in Figure 9) and the average estimate of the competition effect (reported in Figure 14). I let the decay parameter, λ, obtain different values around its point estimate. The figure implies that the marginal movie should best release in the end of the summer or in January, when underlying demand is relatively high, and competitoin is quite soft.

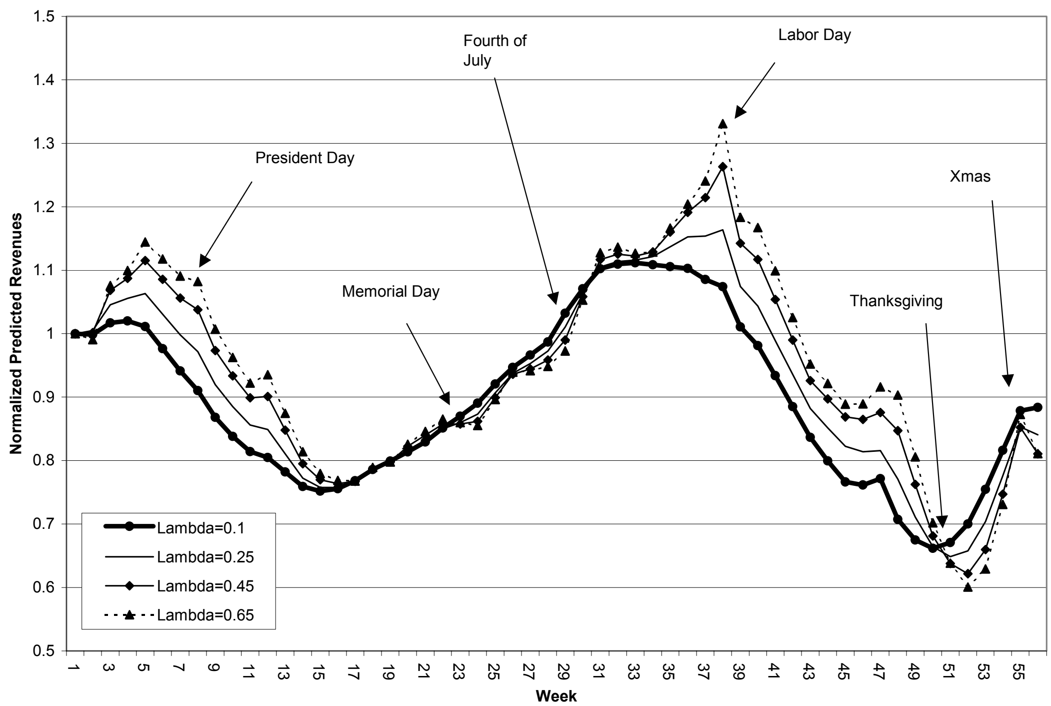

Figure 16: Predicted Normalized Revenues for the Marginal Movie, by Release Date

The number of weekly dummies is 56 to account for the timing variation in U.S. holidays across the years (see Section 3.1)

Figure 17: Evolution of Release Effect from Initial to Final Announcement

This figure plots the estimated release effect ($\sum_{k \in J_t,, r_k = t} \exp(\theta_k)$), averaged over the fifteen years of the data, twice. First, it plots the actual release effect implied by the final release date decision. Second, it plots the same effect, as implied by the initial release announcement by each movie. The two lines are quite similar, despite the frequent changes in announced release dates. This suggests that the pattern of changes cannot be described by, for example, moving away from big weekends to small weekends (which would have been consistent with stories of a war of attrition). The figure is more consistent with the idea of the changes standing for different movies swapping (not necessarily simultaneously) release slots.

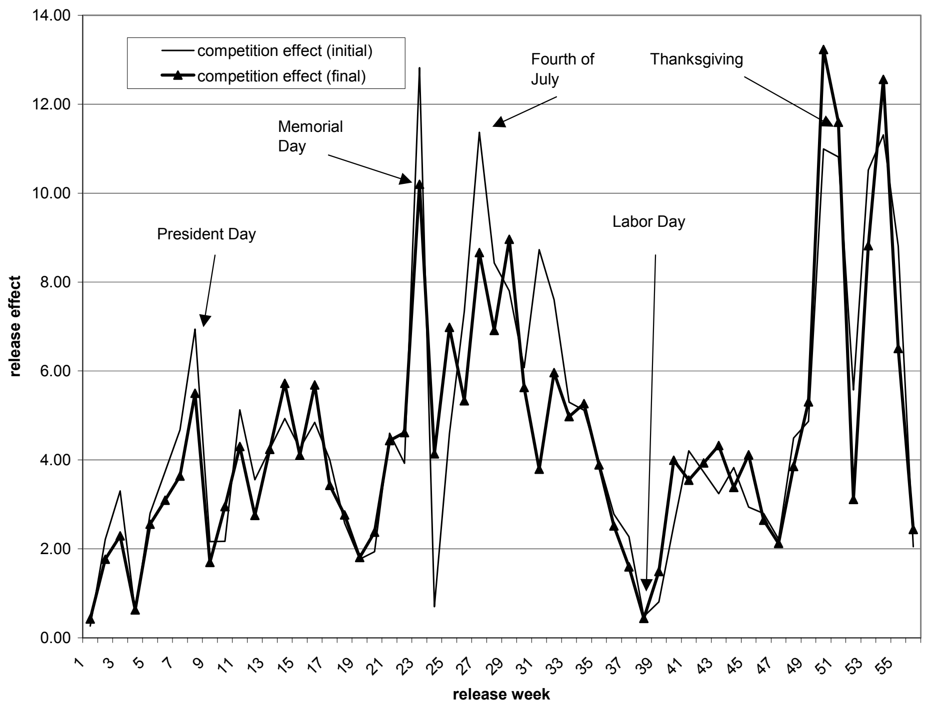

Figure 17: Evolution of Release Effect from Initial to Final Announcement

The number of weekly dummies is 56 to account for the timing variation in U.S. holidays across the years (see Section 3.1)