RAND Journal of Economics
Vol. 38, No. 1, Spring 2007
pp. 127–145

Seasonality in the U.S. motion picture industry

Liran Einav*

The observed seasonality of box-office revenues reflects both seasonality in underlying demand for movies and seasonality in the number and quality of available movies. I separately identify these aspects by estimating weekly demand for movies, using movie fixed effects, a long panel of movies' weekly revenues, and restrictions on their decay pattern. I find that the estimated seasonality in underlying demand is much smaller and slightly different from the observed seasonality of sales. The biggest movies are released at times when demand is highest, amplifying the underlying seasonality. Price rigidities in the industry may facilitate this amplification effect.

1. Introduction

■ 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 40% of the box-office revenues of the average picture. As a result, the release date of a new picture is a major focus of attention for distributors of movies. With virtually no price competition, the movie's release date is one of the main short-run vehicles by which studios compete with each other. 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. The goal of this article is to decompose the observed seasonal pattern of sales into these two components.

In economics, seasonality is typically considered as systematic "noise" that has to be adjusted for, so the analysis can focus on more fundamental patterns in the data. In some industries, however, timing decisions can be crucial for the success or failure of a business. In these cases, opening the seasonality "black box" and explicitly analyzing seasonal patterns is key to our understanding.[1] The analysis should take the underlying seasonal pattern as given, with the understanding that

* Stanford University; leinav@stanford.edu.

This article 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 thank two anonymous referees and Rob Porter, the Editor, for many useful comments. I also thank Bharat Anand, Susan Athey, Nittai Bergman, Tim Bresnahan, Estelle Cantillon, Guillermo Caruana, Richard Caves, Ken Corts, Peter Davis, Yair Eilat, Julie Mortimer, Aviv Nevo, Marc Rysman, Roni Shachar, Minjae Song, Manuel Trajtenberg, Frank Wolak, and seminar participants in many universities for helpful comments and suggestions. Insightful conversations with Jeff Blake (Sony Pictures), Ron Kastner (Goldheart Pictures), and Terry Moriarty (Hoytz Cinemas) are greatly appreciated. Earlier drafts of this article were circulated with the title of "Gross Seasonality and Underlying Seasonality: Evidence from the U.S. Motion Picture Industry." Some parts of this article draw upon initial results reported in Bergman and Einav (2001). The usual disclaimer applies.

[1] A similar approach was taken, for example, by Miron and Zeldes (1988) and Barsky and Miron (1989) in the macroeconomics literature.

this seasonal pattern will induce an equilibrium market reaction.[2] This idea has been employed for at least two decades to analyze price competition. Prices are not set exogenously; they are the sum of equilibrium marginal costs and markups. Just as we have to separate costs from markups when analyzing pricing decisions, we have to separate underlying seasonality from seasonal market reaction when analyzing timing decisions. This distinction is essential to address standard policy questions when timing decisions are an important part of the competitive environment.[3] For example, if the seasonal variation is mainly explained by changes in underlying demand, then distributors should release their movies when sales and demand are high. In contrast, if the seasonal variation is mainly explained by changes in movie quality, then movies should be released in a low sales period, when competition is soft.

The relationship between the underlying seasonality and the market reaction can be classified into two types. First, an amplification effect occurs when the market reaction to high demand increases overall demand. For example, more or better products may be introduced in high-demand seasons. Second, a dampening effect occurs when the market reaction dampens demand changes. For example, prices may increase in high-demand seasons. I discuss this distinction in more detail at the end of the article.

Movie distributors tend to release their (*ex ante*) biggest hits in the beginning of the summer and during the Christmas holiday season, and the choice set of moviegoers varies dramatically over the year. We face a classic identification problem: is the strong box-office performance of, say, Memorial Day weekend the result of higher demand, better movies, or both? One might control for variations in the set of available products in order to identify underlying demand. Such an exercise is not feasible in the motion picture industry; with weak predictors for a movie's box-office success, one cannot control for movie quality.[4]

Identification is achieved by estimating a nested logit model of weekly demand for movies. Movie quality is identified by movie fixed effects and assuming that, conditional on observables, the decay in average consumer utility from a given movie is independent of the movie's release date. The market-expansion effect is identified by variation in the number and quality of movies available in the same calendar week, in different years. This variation is assumed to be the realization of exogenous random shocks in the movie-production process. This is a reasonable assumption as long as distributors believe that underlying seasonality in demand is stable during the observation period. Given a measure of movie quality and an estimate of the market-expansion effect, the residual seasonal variation in sales is attributed to underlying demand. Section 4 discusses the basic identification assumptions, their rationale, and their limitations. In Section 5, I verify that the main results are robust under alternative specifications.

Many articles estimate demand for motion pictures.[5] Absent good predictors for the box-office success of movies, most of the literature conditions on the movie's first week revenues and investigates the effects of various variables, such as advertising, reviews, and Academy Awards, on revenue patterns over the movie's subsequent life cycle. By using movie fixed effects, I follow a similar approach, but not as the main objective of the analysis. To the best of my knowledge, this is the first article to recognize the potential difficulty in the interpretation of the seasonal effects in the industry and to address it explicitly.[6] I assemble a sample that includes all movies

---

[2] Similar logic can be applied to, say, geographically isolated markets, where the distinction would be between underlying market size (e.g., population) and observed market size (sales). In contrast to seasonal markets, in spatial markets, we often observe natural proxies for underlying market size.

[3] For articles that address these issues in price competition, see, for example, Berry, Levinsohn, and Pakes (1995), Nevo (2000, 2001), and Petrin (2002).

[4] Throughout the article I use the word *quality* to mean a measure of a movie's box office appeal or potential, independent of its release date. This may or may not coincide with the movie's cinematographic quality.

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

[6] Radas and Shugan (1998, 1999), as well as Vogel (1994), focus on the observed seasonality in box-office revenues.

released nationwide in the United States from 1985 to 1999. This panel is much longer than those used in the literature. In my article, one important source of variation is across the same calendar week in different years. In a shorter panel, the results may depend on the particular observation period. I also employ discrete-choice estimation methods to control for changes in competition and seasonal effects in a systematic way. The key identifying assumption of using the decay pattern as described above is a novel contribution.

My estimation results imply that about a third of the seasonal variation in sales can be attributed to seasonal changes in the number and quality of movies. In contrast, the industry's conventional wisdom attributes most of the seasonality in sales to demand factors. Moreover, the estimated seasonal pattern in underlying demand differs slightly from the observed pattern of sales. For example, industry revenues gradually fall from the end of July until early September, but my estimates suggest that underlying demand is stable over the summer, with a sharp decrease after Labor Day. The difference occurs because the biggest movies are typically released in the beginning of the summer. These results have important implications for the timing of release dates by movies' distributors. At the end of Section 5, I provide some crude quantitative estimates for these effects. A more rigorous analysis of the release date decisions is relegated to a companion article (Einav, 2003), in which the strategic timing game is modelled empirically.

The rest of the article is organized as follows. In Section 2 I describe the industry. Section 3 describes the data, and Section 4 presents the estimation strategy and the identification assumptions. I present the results in Section 5, which are followed by sensitivity analysis and several implications for release behavior. Section 6 concludes by discussing the potential relevance of the results to other industries.

2. Industry description

■ The motion picture industry comprises of three main players: producers, distributors, and exhibitors.[7] 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 own the theaters. The industry is dominated by the major studios that have integrated production and distribution. 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 not vertically integrated with producers or distributors.

I now examine in greater detail the process by which a movie is exhibited after its completion. The first stage is distribution. The main decisions involve setting the release date, deciding the initial scope and locations of the release, negotiating contracts with exhibitors, and designing the national advertising campaign. The two important considerations for the release date are the strong seasonal effects in demand and the competition that will be encountered throughout the movie's run. Typically, movies with higher expected revenues are released on higher (perceived) demand weekends, and there is a tradeoff between the seasonal and the competition effects. The importance of the release date is magnified because the first week accounts for almost 40% of total domestic box-office revenues on average (Figure 1). In addition, revenues in the first week are believed to create information and network effects that increase revenues in subsequent weeks.

The identity of the competing movies also matters when setting the release date. Distributors are wary of releasing a movie in close proximity to strong, popular movies. As documented in Einav (2002), distributors often change release dates in response to new information concerning release dates of similar movies. They also announce their movie's release date early, hoping that other distributors will avoid the announced date. This practice is common for movies that are thought likely to succeed.

Distributors also decide the scope of a movie's release. There are three types of release: wide release, platform release, and limited release. Wide releases are the most common with the main distributors. Screening of the movie begins in a large number of theaters, typically several

[7] Much of this section is adapted from Squire (1992), Verter and McGahan (1996), and Caves (2000).

FIGURE 1

DISTRIBUTION OF BOX-OFFICE REVENUES OVER MOVIES' LIFE CYCLE

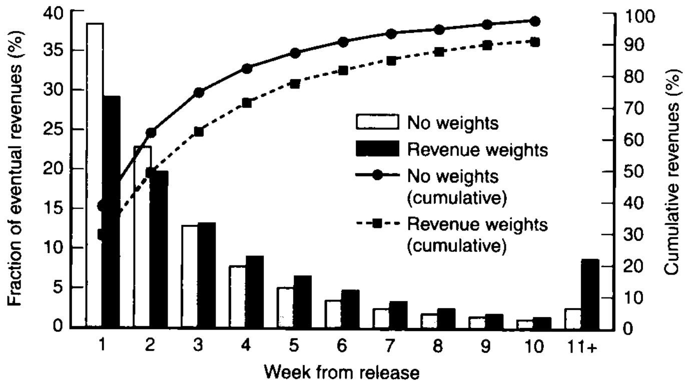

thousand, accompanied by an extensive national advertising campaign. Platform releases involve an initial release in a small number of theaters, often only in big cities, with advertising concentrated in local newspapers. The movie then expands to additional screens and to more rural areas. Distributors typically use such releases with movies that do not have obvious appeal to mainstream audiences, for example, because the movies' actors are unknown or the subject matter is difficult. The conventional wisdom is that a platform release creates word-of-mouth that is necessary for success. In most cases, there is 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.[8] Limited releases are those in which the movie is released in two or three cities without expectations of a subsequent wider release.

Contracts negotiated between distributors and exhibitors are standard. Under a typical contract, the theater pays the distributor a fixed share of the box office revenues. This share is the greater amount of either 90% of the revenues net of the theater's expenses (the "house nut") or a lower share of the gross revenues. This latter share typically starts at around 70% and then declines throughout the movie's run. The percentage decline reduces the exhibitor's incentive to terminate the movie in favor of another. On average, the distributor collects about half of the domestic box-office revenues. As the movie release date draws near, the distributor begins the national advertising campaign. The advertising campaign peaks when the movie is released (Vanderhart and Wiggins, 2001). Distributors tend to set their advertising budget as a fixed percentage of the movie's production cost.[9]

Unless specified otherwise, the exhibitor decides when to terminate the screening of a movie. There is a long-term relationship between distributors and exhibitors, however, and the distributor may punish exhibitors who prematurely terminate by denying access to future films. Typically, a theater screens a movie for six to eight weeks.

The motion picture industry has grown rapidly, with annual domestic box office revenues of more than $9 billion in 2002. On average, since 1985, more than 150 movies were released annually nationwide. Table 1 presents general trends over the sample period.

Domestic box-office revenues now account for as little as 15% of the movie's revenues (down from about 35% in the early 1980s). Additional revenues are obtained from selling screening rights to cable and television networks, from the video and DVD markets, and from the international box-office market. However, higher domestic box-office revenues are believed to increase revenues in the ancillary markets. Thus, maximizing domestic box-office revenues seems like a reasonable approximation for the objective function of distributors.

[8] An additional benefit of platform releases is risk reduction for the distributor, which can observe interim information. This is an important reason for independent producers to use such releases.

[9] Almost all marketing expenditures are funded by the distributor, who is also the one making prints (copies) of the movie. Prints involve nontrivial costs, also affecting the scope of a movie's release.

TABLE 1 Industry Trends

Year	Average Ticket Price[a]	Number of Releases[b]	Number of Wide Releases[b]	Total Box-Office Revenues[a]	Admissions (per Capita)
1985	5.54	167	103	3.51	2.69
1986	5.70	222	120	3.98	2.94
1987	5.78	220	120	4.28	3.08
1988	5.84	230	144	4.37	3.08
1989	5.40	213	115	4.92	3.73
1990	5.45	223	119	4.90	3.64
1991	5.20	224	125	4.63	3.57
1992	4.97	214	121	4.85	3.86
1993	4.83	227	144	5.09	4.14
1994	4.74	233	142	5.13	4.20
1995	4.80	246	149	5.27	4.23
1996	4.74	262	152	5.52	4.44
1997	4.81	250	150	5.90	4.63
1998	4.91	264	139	6.20	4.73
1999	5.00	244	143	6.80	5.03

[a] Ticket prices and revenues are in December 1999 U.S. dollars. Revenues include only the first ten weeks since release.

[b] Total releases is the number of all titles released every year. Wide releases is the number of titles, which reached a wide release (600 screens) within the sample. The latter is the subset of titles used in the empirical analysis throughout the article.

3. Data

■ The dataset covers all movies released in the United States between January 1, 1985 to December 31, 1999. For each movie released, the dataset includes the official release date, 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, Academy Awards, and the total run time. In addition, for the first 10 weeks following the official release date of each movie, the dataset includes weekly box-office revenues and weekly number of screens on which the movie was run.[10] The data were obtained from *ACNielsen EDI*, except for the advertising data, which was obtained from *Ad $ Summary*, published annually by *Competitive Media Reporting*.

Because of the long sample period, I deflate the box-office revenues to accommodate trends in the average ticket price and 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 interpolated from the annual ticket price schedule by assuming that prices increased linearly throughout the year. Weekly figures for the population of the United States are interpolated in a similar way, with annual figures taken from the U.S. Census. Weekly market shares for each movie are calculated by dividing weekly revenues by the weekly ticket price and by the weekly population size.[11]

The initial sample comprises 3,523 movie titles. Throughout the analysis, I restrict attention to movies that reached a wide release at some point during their run.[12] For those movies that I do not include, the average peak number of screens is 100 and the average production cost is $2.9 million. These are relatively small movies, in a different segment of the industry. I also do not use

[10] The first 10 weeks of a movie's run account, on average, for more than 90% of total box-office revenue (Figure 1).

[11] As a referee noted, this calculation underestimates the market shares of movies that cater to children and senior citizens, who purchase tickets at a significant discount. As the fraction of discounted tickets varies primarily across (and not within) movies, the use of movie fixed effects addresses this concern.

[12] Operationally, any movie that reached 600 screens on some week is included in the analysis. Corts (2001) uses a similar cutoff point. This also seems reasonable because the distribution of the peak number of screens across movies is bimodal, with 600 falling between the two modes.

in the analysis 97 titles, which went through a relatively long platform release and never reached a wide release within 10 weeks of their official release date.[13]

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.[14] I disregard weeks in which a movie was shown in a limited number of theaters. The view of the initial weeks of a gradual release as a method of advertising justifies this approach. The official and the designated release date coincide for wide releases, which is the case for more than 90% of the titles used in the analysis.

My final dataset comprises 1,956 titles and 16,103 weekly observations. These titles account for 94% of the revenues of the 3,523 titles in the original sample. On average, I observe slightly more than 8 weeks per movie rather than 10. First, about 80% of the truncated observations is attrition of movies that were terminated by exhibitors before their 10th week. Second, movies that were platform released have less than 10 weekly observations after their first wide release. Throughout the analysis, I use all observations and verify that the results are not sensitive to the truncation problem.

The industry operates on a weekly schedule. More than 75% of the movies were released on Friday (20% on Wednesday). Much of the competition is over the weekend audience, which accounts for about 70% of revenues. Theaters' termination decisions are made on a weekly basis. Consequently, release dates are tabulated at the weekly level. Holiday weekends are an important aspect of seasonality in the industry. American holidays can shift by up to one week from year to year, and it is important to adjust for the calendar differences across years. Thus, I create 56 weekly dummies; I introduce some additional weeks in some years. For example, Labor Day and Thanksgiving are always assigned the same week numbering (38 and 51, respectively) although there might be 11 or 12 weekends between them. Week 42 is then used as a "filler" in years with 12 weeks between these two holidays. Weeks 1, 6, 13, 35, 42, and 56 are those that do not appear in all years.

Table 2 provides some descriptive statistics for the final sample, along with 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 and the production cost figures show a similar pattern, while the distribution of advertising expenditure is more symmetric. Finally, Table 1 demonstrates that the industry share (annual admissions over U.S. population) has been growing since the mid 1980s.

TABLE 2 **Descriptive Statistics of Key Variables**

	Total Revenues (current $)		Population Share (%)[a]		First-Week Revenues (current $)		Production Cost (December 1999 $)		Advertising Expenditure (December 1999 $)	
Titles		1,956		1,956		1,956		1,604		1,873
Mean		34.8		2.78		10.20		29.7		8.47
Median		20.0		1.71		6.63		23.2		7.19
Standard deviation		43.5		3.16		10.80		22.6		5.77
Highest	Titanic	601	Titanic	32	The Lost World	107	Titanic	209	Toy Story	43
2nd highest	Star Wars ('99)	431	Jurassic Park	28	Star Wars ('99)	99	Waterworld	193	Titanic	35
3rd highest	Jurassic Park	357	Star Wars	27	Austin Powers 2	85	Armagedon	155	The Rookie	35
4th highest	Forrest Gump	330	Batman	23	Jurassic Park	82	Speed 2	152	Anastasia	33
5th highest	Lion King	313	Lion King	23	Independence Day	79	Tarzan	152	Forrest Gump	31

[a] Population share (box-office revenues divided by ticket price and by U.S. population) is based on the first ten weeks only.

[13] I do not use the first nine weeks of 1985, as I am missing the data on the movies released in the final weeks of 1984.

[14] For a typical movie, it is easy to see when the number of screens discontinuously jumps, marking the first week in which the movie is shown nationwide. Operationally, I define the actual release to be the first week in which the number of screens exceeds the maximum of 400 screens and 30% of the (eventual) maximal number of screens showing the movie.

FIGURE 2

SEASONALITY—DESCRIPTIVE STATISTICS

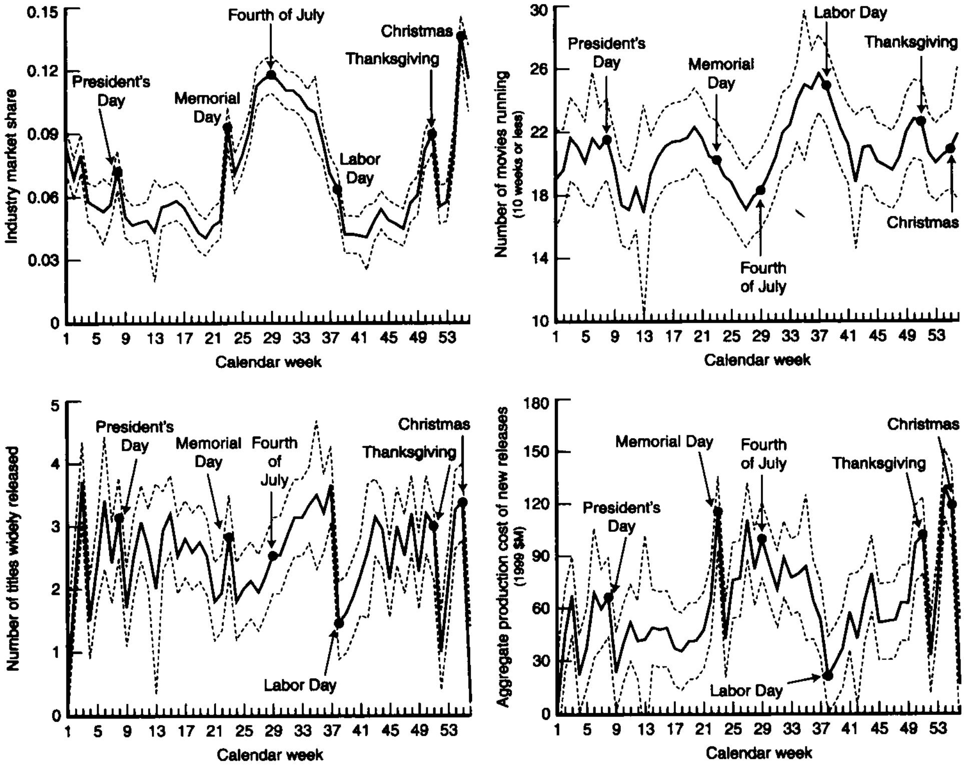

The main focus of analysis is the seasonal pattern in industry sales, which is depicted in the upper-left panel of Figure 2. This is the typical seasonal picture portrayed in the popular media. The year is thought to consist of four periods: summer (roughly, Memorial Day to Labor Day), holiday (Thanksgiving to mid January), winter/spring, and fall. The first two are generally thought of as high-demand periods, while the other two are considered weak. 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, the lower-right panel of Figure 2 shows that releases of big-budget 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. The number of movies or releases by itself has a different pattern, driven mainly by smaller budget movies releasing in different times than big-budget ones (the rest of Figure 2).

4. Specification and identification

■ **The benchmark model.** I estimate a nested logit demand model. The model is based on an individual utility function, but my aim is not to model individual decision making. Rather, this functional form allows me to separate the seasonal effects in a simple way. The nested logit model is important because a primary goal is to estimate underlying market size. The market expansion arising from an increase in the number and/or quality of movies should be estimated, not imposed. The identifying assumptions are discussed below.

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

$$u_{ijt} = \theta_j - \lambda(t - r_j) + \xi_{jt} + \zeta_{it} + (1 - \sigma)\varepsilon_{ijt}, \tag{1}$$

where θ_j is the quality of movie j, and r_j is the release week of movie j. Thus, $t - r_j$ is the number of weeks that have passed since the movie's release. ξ_{jt} is an unobserved propensity to like movie j in week t, which I interpret as a movie-week deviation from the common decay pattern. $\zeta_{it} + (1 - \sigma)\varepsilon_{ijt}$ is an individual error term. ζ_{it} does not vary across movies and is an idiosyncratic propensity of individual i to go to the movies in week t.

I allow the mean utility from the outside good (from not going to any movie in week t) to vary from week to week. Thus, the utility of consumer i from the outside good (good 0) in week t is given by

$$u_{i0t} = -\tau_t + \zeta_{it}' + (1 - \sigma)\varepsilon_{i0t}, \tag{2}$$

where τ_t is a week fixed effect. It has a negative sign in the utility for convenience. Note the difference between my specification and typical discrete-choice applications, in which the mean utility from the outside good is normalized to zero in *all* markets. The variation in the utility from the outside good across weeks captures the underlying seasonality.

I assume that ε_{ijt} (and ε_{i0t}) is distributed i.i.d. extreme value and ζ_{it} (and ζ_{it}') has a distribution that depends on σ such that $\sigma \in [0, 1]$ and the sum, $\zeta_{it} + (1 - \sigma)\varepsilon_{ijt}$ (as well as $\zeta_{it}' + (1 - \sigma)\varepsilon_{i0t}$), is a distributed extreme value (Berry, 1994; Cardell, 1997). As σ approaches zero, we obtain the simple logit model. As σ approaches one, there is no substitution between the outside good and the inside goods. As discussed later, the σ parameter is important in determining the market-expansion effect.

The nested logit predicted (weekly) market shares for each movie are given by

$$s_{jt} = \frac{\exp\left(\frac{\theta_j - \lambda(t - r_j) + \tau_t + \xi_{jt}}{1 - \sigma}\right)}{D_t^{\sigma} + D_t}, \tag{3}$$

where

$$D_t = \sum_{k \in J_t} \exp\left(\frac{\theta_k - \lambda(t - r_k) + \tau_t + \xi_{kt}}{1 - \sigma}\right) \tag{4}$$

and J_t is the set of all movies that are in theaters in week t. We can then rearrange equation (3) to obtain

$$\log(s_{jt}) - \log(s_{0t}) = \theta_j + \tau_t - \lambda(t - r_j) + \sigma \log\left(\frac{s_{jt}}{1 - s_{0t}}\right) + \xi_{jt}. \tag{5}$$

The mean utility from movie j at week t relative to that of the outside good depends on the fixed quality of the movie, the underlying seasonal effect in demand, and the number of weeks that have passed since the movie's release. I discuss these below. The fourth element, the within-industry market share, $s_{jt}/(1 - s_{0t})$, is endogenous, and an instrumental variable is required. I discuss the instruments later.

The movie fixed effect, θ_j, captures variables that do not change over the screen life of the movie, such as the cast, the plot, and the prerelease advertising campaign. A fixed-effects specification is important because observables explain little of the variation in movie quality.

The second element, τ_t, is the underlying seasonal effect in demand. The weekly dummies capture demand fluctuations that are not explained by variation in the set of available movies. Interpreting these as variation in the opportunity cost of time is natural. For example, on holiday weekends, people have more leisure time and the weekly dummies are higher. Implicitly, I assume that the mix of consumers is similar across weeks. The identification strategy compares adjacent weeks or the same calendar weeks in different years, in which case this assumption seems reasonable.

The third element, $t - r_j$, is the decay effect. The dependent variable is approximately proportional to the logarithm of revenues, so the linear term approximately represents a proportional decay in revenues, somewhat consistent with Figure 1. The decay variable captures two effects. First, most people watch a movie only once, and the potential market shrinks over time.

Second, most people prefer watching a movie earlier in its release.[15] The second effect is consistent with the individual utility specification. The saturation effect, however, requires assumptions about the correlation of idiosyncratic noise, ε_{ijt}, over time, leading to changes in the distribution of the ε_{ijt}'s due to truncation of people who have already seen the movie (Moul, forthcoming). I believe that identifying such effects requires a more detailed dataset. Thus, I view the individual utility as a convenient "reduced form."

□ **Identification.** My primary goal is to distinguish between underlying seasonality and the endogenous market reaction. Two assumptions allow me to do so. First, I assume that the decay in the mean utility is orthogonal to the movie's release date. This facilitates the identification of movie quality. Second, I stipulate that variation is exogenous across the *same* calendar week in *different* years in variables such as the number of movies released. This allows a comparison across years, facilitating the identification of the market-expansion effect. I discuss these assumptions in more detail below. For clarity, I discuss each identification problem separately. The model addresses both problems.

Movie qualities. In order to identify movie quality, one would like to observe the same movie released on different dates in a number of distinct yet similar markets. This is not feasible. First, I observe only the aggregate U.S. domestic market. Second, due to increasing returns in advertising, first-run films are predominantly released simultaneously nationwide. The identification problem is difficult because the release-date decision is endogenous. Observables account for a small fraction of the variation in quality, and controlling for observables will not alleviate the endogeneity problem.

Weekly box-office revenues provide multiple observations for each movie and are used to "extract" the average quality of the movie under certain assumptions. Identification comes from observing movies throughout their run, playing during weeks with different demand levels, and competing against different sets of movies. A higher-quality movie may be released in a week with higher demand, but it still runs in subsequent weeks, in which seasonal demand varies. A lower-quality movie released afterward is observed at the same time as the better movie. A comparison of the market shares identifies their relative qualities. Seasonal effects are identified by observing different choice sets over different weeks and over a period of 15 years.

Consider the following simple example. Suppose that a year consists of two seasons and that two movies are released each year, one every season. One movie is of quality 1, and the other is of quality $\theta_H > 1$. The market sizes in season 1 and season 2 are 1 and $M > 1$, respectively. The revenues of a movie of quality θ released in a season of size m are $m\theta$. Suppose that distributors always release their better movie in season 2. Assume first that each movie is shown for a single season. Sales are 1 in season 1 and $M\theta_H$ in season 2. In this case, M and θ_H are not identified separately.

Now suppose that each movie is shown for two seasons and that they share a common decay rate in quality, λ. For simplicity, assume also that the demands for each movie are independent. 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 1 and $M\lambda$. I have three equations in three unknowns, and all the parameters are identified. With more movies and more seasons, one can allow the decay parameter to vary across movies. The crucial identifying assumption is that the decay pattern does not vary across release dates conditional on observables.

I make several additional comments. First, the assumption concerns the decay pattern in the utility from a given movie, not the decay pattern in box-office revenues. Decay in revenues will depend on the release date, via the underlying seasonal pattern. The assumption does not allow consumers to alter relative preferences over two movies that were released at different dates, beyond the change implied by the common decay rate. Second, an important aspect of the

[15] Focus-group participants suggest that the heavy concentration of advertising just prior to release, as well as a desire of some moviegoers to be market mavens, can help explain the rapid decay. The heavy advertising just prior to release is also driven by the declining percentage structure of distribution contracts that provide incentives for distributors to front load their advertising expenditure.

decay pattern may be the exhibitors' decisions about movie termination. If exhibitors act as agents for consumers—terminating a movie only when the demand is sufficiently low—this introduces discrete drops in revenues. The national-level data smoothes out any such drops because the typical movie runs on several hundred screens and termination decisions are made separately by each exhibitor.[16]

Finally, the main concern about the identification assumption is that the decay pattern is an unobserved attribute of the movie. *Ceteris paribus*, studios may be more likely to release a fast-decay movie at the end of the high-demand summer season than in the beginning. There are several important points to note, however. First, based on multiple interviews with industry players, market participants tend to think about movies only in terms of expected box-office revenues. Variation in decay helps explain box-office revenues after the movie's run, but not as much *ex ante*. To the extent that studios are concerned with total revenues, *ex post* variation in the decay for movies should not be correlated with the release-date decision. Second, seasons vary in many ways, and the identification assumption is stronger for movies that are released within a short period, when it is more likely that decay does not vary significantly, after controlling for observables or movie quality.

Market expansion and instruments. Identifying underlying seasonality will not be possible unless we identify the market-expansion effect. We need to estimate whether more or better movies draw their customers from other movies, or by drawing people who would otherwise not go to the movies. This is captured by the σ parameter of the nested logit model. If $\sigma = 1$, there is no substitution between the inside goods (movies) and the outside good, and hence no market-expansion effect. Then observed seasonality in revenues reflects the underlying seasonality. As σ decreases, more observed seasonality is attributed to variation in the quality and number of movies, i.e., market expansion is higher.

The main feature of the data that is used for identification is the relatively long panel. I assume that the underlying seasonal pattern is stable, so that comparing the same calendar week in different years can identify market expansion. For example, suppose that more (or better) movies are released on Memorial Day 1999 compared with those released on Memorial Day 1998. Conditional on movie quality and on the assumption that underlying seasonality does not change, the difference in industry revenues identifies the market-expansion effect. Similar levels of total revenues would imply a weak market-expansion effect (σ close to 1), while higher Memorial Day revenues in 1999 will be attributed to a strong market-expansion effect (low σ). Because σ is assumed to be the same across weeks, it can be used to infer the expansion effect from changes in the number or quality of movies shown in consecutive weeks.

This identification assumption guides the choice of instruments. The benchmark regression uses the number of movies shown in a given week as the instrument. More movies are associated with more-intense competition and therefore should be negatively related to the within-industry share, the endogenous variable. These instruments are valid if they are unrelated to the error term. They may seem endogenous if more movies are released in high-demand weeks. Weekly dummy variables are included, however, and the variation that is important for identification is across the same calendar week in different years. One may still be concerned that fewer movies will be released against an extremely high-quality movie. But movie fixed effects account for this variation. The identifying assumption is that the number of movies shown in a given week is not correlated with the idiosyncratic decay pattern of the movies. In that sense, this identifying assumption is related to the previous assumption discussed above. If distributors do not account for their movie-specific decay patterns when releasing their movies, it seems reasonable that competing distributors will do the same.

Finally, if release patterns do not respond to these shocks, why should we observe any variation in release pattern? There is a natural source of variation in annual release patterns. Studios commit to new projects a year or two in advance, but these projects are highly uncertain

[16] Such drops may be more pronounced for low-revenue movies. The distribution of the weekly number of theaters over the movies' lives is quite smooth even for the smallest decile of movies.

in terms of completion time and realized quality. If studios engage in a timing game over a short release period but are reluctant to hold completed movies in inventory for long periods, there will be exogenous variation in the set of movies released within a season.

5. Results

■ **The benchmark model.** The top panel of Figure 3 presents results from the benchmark model. The seasonal pattern should be compared with the gross seasonality presented in Figure 2. To facilitate comparison, the bottom panel of Figure 3 presents results from the benchmark model without movie fixed effects. The bottom panel controls for seasonal variation in the number of movies released but not for their quality, while the top panel also accounts for movie fixed effects.

The key finding is an amplification effect. Once variation in movie quality is accounted for, the estimated seasonality is significantly dampened. The standard deviation of the estimated

FIGURE 3

ESTIMATES FROM THE BENCHMARK MODEL

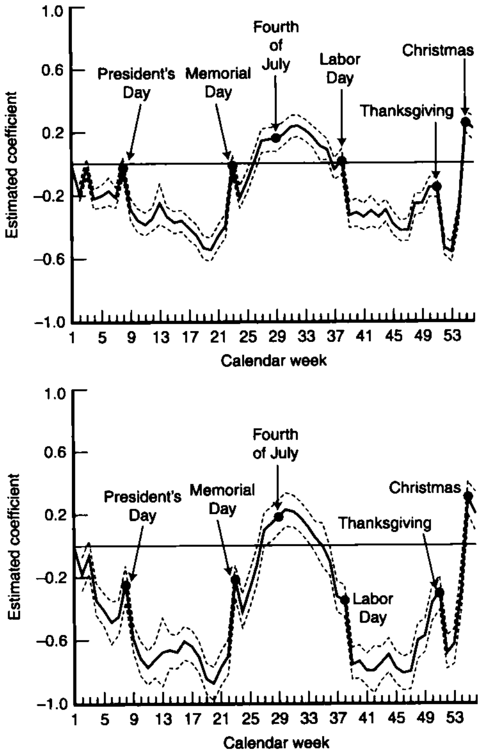

The top panel presents the estimated coefficients on the weekly dummy variables from the benchmark model. The bottom panel presents the same coefficients, when the benchmark model is estimated without movie fixed effects.

Additional results from these regressions:

	With movie FE	Without movie FE
λ	−0.220 (0.0014)	−0.163 (0.005)
σ	0.524 (0.030)	0.577 (0.011)
Number of observations	16,103	16,103
Number of titles	1,956	1,956
R^2	0.876	0.893

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).

coefficients on the week dummy variables increases from 0.238 to 0.356 when movie fixed effects are omitted. Thus, the endogeneity of movie quality amplifies seasonality by about 50%. In addition, the underlying seasonal pattern is similar to that of industry revenues (high summer and holiday seasons, low spring and fall), but there are some differences. Industry revenues are high during the big release weeks, such as Memorial Day, Fourth of July, and Thanksgiving. These weeks are estimated to have high underlying demand, but their relative demand with respect to adjacent weeks is not as high. In contrast, Labor Day, when few movies are released, is estimated to have significantly higher underlying demand than gross industry revenues would imply. More generally, underlying demand during the summer seems flat, dropping sharply after Labor Day. Industry revenues, in contrast, peak in early July and drop gradually.[17]

The estimated decay coefficient in the benchmark model is -0.22, with a very small standard error, while the market expansion parameter, σ, is 0.52. Substitution patterns differ significantly in the choice among different movies and the choice whether to go to the movies. The two coefficients imply an estimated decay of revenues of 37%, close to the observed revenue decline presented in Figure 1. The linearity assumption on the decay pattern is not important. When I allow for a flexible decay pattern by using 10 dummy variables, the resulting decay pattern is almost identical to the linear one.

The estimated movie fixed effects closely follow the total box-office revenues for each movie, with little trend over time. These coefficients capture the trend in the attractiveness of the inside good relative to the outside good. The outside good did not remain constant, with the introduction of Internet, DVD, and other alternatives to movie going. The fact that average quality did not increase does not imply that the industry is not advancing. The industry has kept up with the outside competition, albeit with increasing production costs. Regressions of the fixed effects on different movie characteristics are presented in Table 3. The results suggest that the estimated fixed effects are reasonable. All coefficients have sensible signs and magnitudes. The low R^2's also show that observable variables (some of them endogenous) explain only a small fraction of the variation in quality. Hence, the original motivation for movie fixed effects.

TABLE 3 Projection of the Estimated Movie Fixed Effects on Observables

Dependent Variable	Estimated Movie Fixed Effect[a]		
Log (production cost)	.215** (.021)	.389** (.021)	.384** (.020)
Log (advertising expenditure)	.367** (.019)	—	—
MPAA rating: PG	omitted	omitted	—
MPAA rating: PG-13	−.125** (.038)	−.135** (.043)	—
MPAA rating: R	−.154** (.039)	−.144** (.044)	—
MPAA rating: >R	−.442 (.365)	−.716* (.413)	—
Genre: Action	omitted	omitted	—
Genre: Comedy	−.052 (.035)	.029 (.039)	—
Genre: Drama	−.097** (.037)	−.021 (.040)	—
Genre: Children	−.070 (.054)	−.037 (.060)	—
Best picture nominee	.433** (.081)	—	—
Best picture award	.467** (.139)	—	—
Year fixed effects	yes	yes	no
N	1,542	1,603	1,603
Adjusted R^2	.387	.208	.193

Standard errors in parentheses.

[a] In all regressions, the dependent variable is the estimated movie-fixed effects (θ_j) from the benchmark model.

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

[17] Krider and Weinberg (1998) model the release-date timing game by assuming that seasonality follows a step function similar to the one estimated here.

FIGURE 4
MODEL FIT

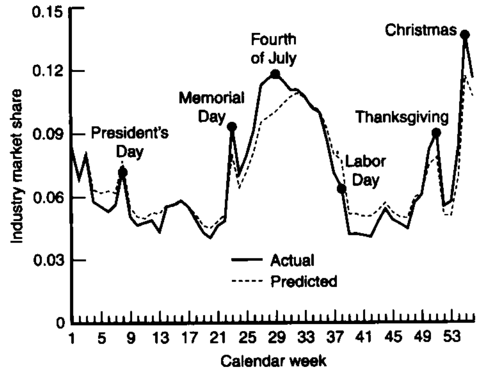

□ **Robustness.** *Model fit and weaker identifying assumptions.* Figure 4 compares the observed seasonality in industry shares with the estimated seasonality, as predicted by the model. To generate the latter, I use the point estimates to compute the implied weekly market shares for each movie, aggregate over all movies in release that week, and average over the 15 years of the data. The model replicates quite well the observed pattern. However, some patterns suggest misspecification of the benchmark model. In particular, the benchmark model underpredicts revenues in the strong-release seasons (early summer and the holiday season) and overpredicts revenues the rest of the year.

On average, movies released during the strong-release seasons decay slower than other movies. If the model does not account for differences in the decay parameter, I obtain biased estimates of underlying seasonality. In the first weeks of slowly decaying movies, underlying seasonality will be biased downward, while later underlying seasonality will be biased upward.

There are (at least) two reasons why decay rates may not be constant over the year. First, as Figure 1 indicates, bigger movies decay slower. Bigger movies are more likely to be released in the strong-release seasons. I therefore include an interaction variable between the decay parameter and the size of the movie. I tried two specifications. The first uses production costs as a proxy for the size of the movie, which does not have much additional explanatory power. As suggested in Table 3, production costs are a noisy proxy. The second specification implements a two-step procedure. I first estimate the benchmark model and then reestimate interacting the estimated movie-fixed effects from the first step with the decay parameter. Because the interaction is endogenous, I use interactions between production costs and decay as instruments. This procedure reveals that better movies decay slower, and the model fit improved. It made little difference, however, to the estimated underlying seasonality. Allowing the decay to vary with the movie genre results in similar findings. There are slightly different decay patterns for different types of movies, but little effect on the estimated underlying seasonality.

A second possible misspecification is that the identifying assumption is incorrect. Distributors may know the decay of their movies before the actual release. If so, they may release slow-decay movies early in the summer and fast-decay movies late in the summer to make the most out of the high-demand summer season. I therefore allow the decay pattern to vary over the seasons in which they were released. I allow six different seasons: winter, spring, early summer, late summer, fall, and holidays. The estimated decay coefficients reflect the fact that bigger movies are released early in the summer and during the holidays. In particular, λ is estimated to be lowest (in absolute value) for holiday releases (-0.194); higher for movies released in the winter, spring, and early in the summer (-0.206, -0.216, and -0.214, respectively); and highest for late summer and fall releases (-0.249 and -0.244, respectively). I cannot reject the hypothesis that the first four coefficients are equal, but the null that the latter two are the same as the others is rejected. Although this specification slightly improves the model fit, the estimated pattern of underlying demand does not change much. The summer coefficients are identical, while the holiday–winter periods are estimated to have somewhat higher underlying demand.

Industry trends, truncation, and Wednesday releases. In this section, I discuss other concerns and verify that the results are stable across different subsamples. First, there may be information-dissemination effects over the movie life cycle. I assume that most word-of-mouth effects occur within the first two weeks of the movie's run. The results are almost identical when the model is estimated for a subsample that includes movies only after their third week in theaters. Second, the results are virtually the same when the model is estimated for movies with a full run of 10 weeks (about 60% of the movies). This accounts for two potential concerns: (i) movies with limited or platform releases differ and (ii) movies that run less than 10 weeks in theaters bias the results. Third, the results do not change if I restrict the data to Friday releases (about 75% of the movies). A Wednesday release is likely to saturate the market for a movie faster. The (roughly) proportional decay implies that the revenue pattern of such movies shifts downward, but their decay pattern will be the same. For this reason, the estimated fixed effects for Wednesday-released movies are likely to be biased downward, but this bias should not (and does not) affect the estimates of underlying seasonality.

The main results are based on a period of 15 years. As Davis (2006) documents, the number of theaters increased during the 1990s. Capacity expansion has translated into bigger opening weekends, which may result in faster market saturation and steeper decay. I therefore estimate the benchmark model for each five-year span separately.[18] Figure 5 presents the estimates of underlying seasonality for each subperiod. Movies do decay faster later in the sample, but underlying seasonality is stable. Similarly, the results do not change much if I pool all periods together, but allow only the decay parameter to vary.

The data and the specification do not allow separate identification of the time-trend or year effects. A linear time trend is not identified.[19] A more flexible time trend will be identified only through the functional form. Similarly, year effects are identified only from movies that are released in December and continue through January. One cannot distinguish between a trend in the utility from the outside option and a trend in the average quality of movies. I can identify only the sum of the two trends, captured by the trend of the estimated movie fixed effects, which does not reveal any obvious pattern.

Instruments and market expansion. I use the number of movies in release as the instrument for the inside share. The results are similar when I use, instead or in addition, the total budget of competing movies or the total number of movies within the same genre. I also used the number of movies released within a three- and five-week window around a particular week, to account for potential endogeneity.[20] The results are fairly stable.[21]

Market expansion in the model is the result of either better or more movies. The nested logit model includes a free parameter for market expansion (σ), but it imposes an implicit functional-form assumption, relating the market-expansion effect of an increase in the number of movies and the market-expansion effect of an increase in the quality of movies. Ackerberg and Rysman (2005) discuss this restriction and suggest including the logarithm of the number of products as an additional explanatory variable. The two effects can then be separately estimated. In my application, the number of products is already an instrument. Including it as an explanatory

[18] To make the results comparable, I hold σ constant (at 0.52, as estimated in the benchmark model) for each subsample. The estimates of underlying seasonality are not sensitive to this restriction. The unconstrained σ is estimated to be 0.40, 0.74, and 0.70 for each subsample, mainly affecting the interpretation of the decay parameter.

[19] The benchmark model with a linear time trend is $\delta_{jt} = \theta_j + \tau_t - \lambda(t - r_j) + \gamma t + \sigma \log[s_{jt}/(1 - s_{0t})] + \xi_{jt}$. This can be rewritten as: $\delta_{jt} = \widehat{\theta}_j + \tau_t - (\lambda - \gamma)(t - r_j) + \sigma \log[s_{jt}/(1 - s_{0t})] + \xi_{jt}$, where $\widehat{\theta}_j = \theta_j + \gamma r_j$.

[20] Fewer movies may be shown when a particularly good movie is released. If distributors are flexible in choosing the week of release but cannot shift the date more than a few weeks, this alternative set of instruments accounts for this potential endogeneity.

[21] The instruments are weaker when interacted with a measure of the (*ex ante*) size of the movie (e.g., production costs). This leads to low estimates of σ and to somewhat different results. If this set of instruments is included in addition to the instruments used above, I retain the results reported throughout.

FIGURE 5

UNDERLYING SEASONALITY OVER TIME

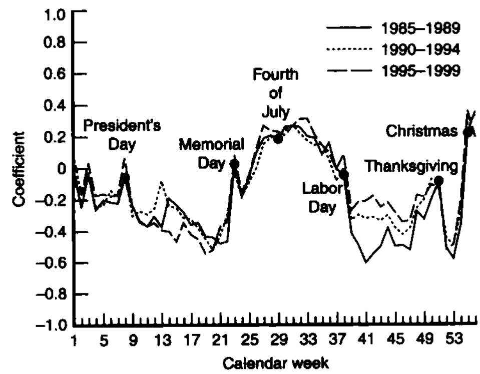

The figure presents the estimated coefficients on the weekly dummy variables from the benchmark model, when estimated for each sub-period separately. To allow comparison of the coefficients, I impose $\sigma = 0.524$ (from the benchmark model) for all sub-periods. Allowing it be separately estimated affects λ but has little effect on the estimated seasonality.

Additional results from the bottom panel regressions:

	1985–1989	1990–1994	1995–1999
λ	−0.174 (0.0016)	−0.206 (0.0014)	−0.253 (0.0015)
N	4,035	5,521	6,547
Number of titles	572	666	754
R^2	0.402	0.430	0.499

variable is possible only through functional-form restrictions or by using other characteristics as instruments (e.g., the total production costs of competing movies). Neither specification changes the results much and the estimated underlying seasonality is the same.

□ **Implication for timing decisions.** I consider the benchmark model, assuming away the disturbance term ξ_{jt}, and ignoring parameter uncertainty. Rewrite equation (5) to obtain

$$s_{jt} = \frac{\exp\left(\frac{\theta_j - \lambda(t-r_j)}{1-\sigma}\right)}{\left[\sum_{k\in J_t}\exp\left(\frac{\theta_k-\lambda(t-r_k)}{1-\sigma}\right)\right]^{\sigma}}\left[\exp(-\tau_t) + \left(\sum_{k\in J_t}\exp\left(\frac{\theta_k - \lambda(t-r_k)}{1-\sigma}\right)\right)^{1-\sigma}\right]^{-1}. \quad (6)$$

The numerator is a scaling factor, which depends on the decay-adjusted quality of the movie. The rest of the expression depends on two elements. The first element, $\exp(-\tau_t)$, is the market-size effect, increasing in the estimated week effect. The second element is the competition effect, as summarized by adding the decay-adjusted qualities of the competing movies. The competition effect is greater if there are more competing movies or if the average quality of the competing movies is higher. As σ increases, the importance of competition is greater, as movies steal more business from other movies than they gain consumers who would otherwise not go to the movies.

One can use equation (6) to characterize the problem faced by distributors when deciding movies' release dates. Because prices (and contracts) are stable and marginal costs are effectively zero, maximizing profits is equivalent to maximizing cumulative market share. The distributor of movie j chooses a release date, r_j, to maximize $\sum_{t=r_j}^{r+H} s_{jt}$, taking the release dates of competing movies as given. H is the length of the horizon taken into account by distributors.[22] I assume

[22] With the estimated fast decay, the results are not sensitive to the choice of H, as long as H is not too small. For

FIGURE 6

SEASONAL EFFECTS IN COMPETITION AND IMPLICATIONS FOR RELEASE DATES

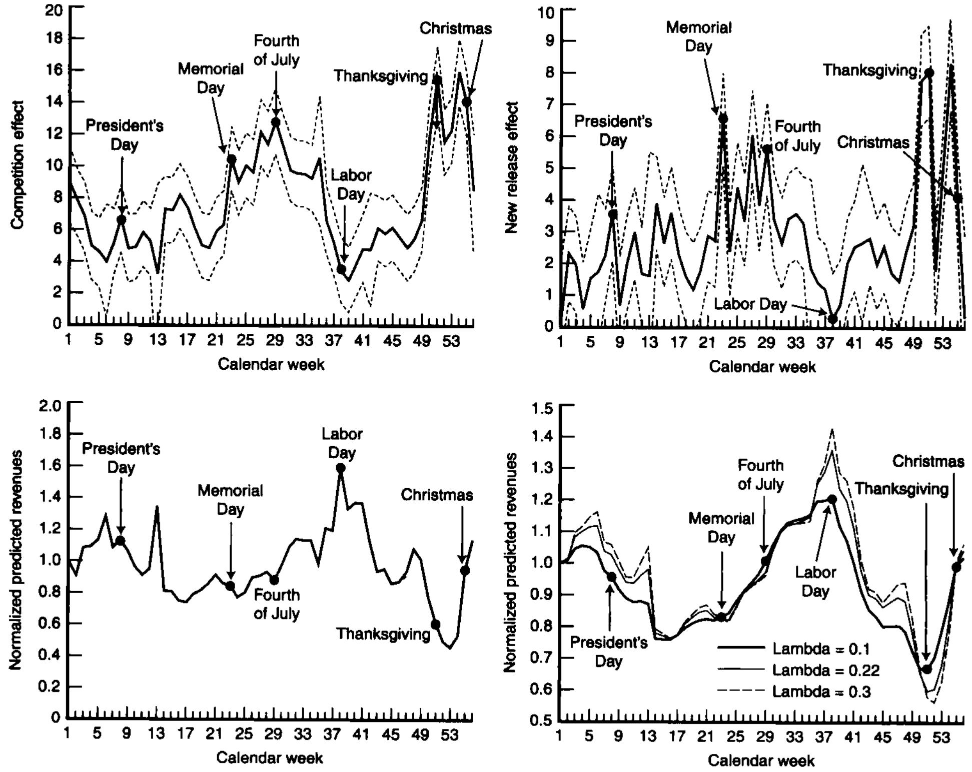

that movie quality and all other parameters are common knowledge. θ_j enters the optimization problem in two ways. The first is multiplicative and does not affect the optimal decision. The second enters through the competition effect.

Bigger movies will generally have a bigger competition effect and hence a stronger strategic incentive. The strategic game is analyzed in a companion article (Einav, 2003). For the remainder of the section, I assume that distributors are not strategic and I do not account for the competition effect. This is analogous to a price-taking assumption. Under this assumption, all distributors want to release when underlying demand is highest and competition is softest.

Figure 6 presents the implications of the point estimates from the benchmark model. The top-left panel presents the competition effect, $\sum_{k\in J_t} \exp[(\theta_k - \lambda(t - r_k))/(1 - \sigma)]$, for each week. I then project it on the weekly dummy variables. The quality of movies at the theaters peaks at the beginning of the summer and in the Christmas holiday season. The top-right panel shows the part of the competition effect due to new releases. I repeat the same exercise for movies in their first week. The seasonality in the competition effect is driven by several strong release weeks: Memorial Day, the third weekend in June, Fourth of July, Thanksgiving and the preceding week, and the week before Christmas. Labor Day, despite high estimated underlying demand, has soft competition.

The bottom part of Figure 6 combines the computed competition effect with the estimated underlying demand to produce the overall effect on demand. On the left, this exercise is done week by week. Labor Day, as well as the rest of the fall, is an attractive release period because of the soft competition during this period. However, movies remain in theaters for more than one week. The right-bottom part depicts the normalized revenues for each release date, under three different

---

the same reason, accounting for the declining distributor share of the revenues does not qualitatively affect the discussion. Discounting over several weeks is negligible.

assumptions about the decay parameter. Labor Day remains an attractive release date due to the combination of soft competition and high underlying demand. Thanksgiving is not attractive, as competition is intense and underlying demand is not very high. The estimates suggest that a small movie can double its revenues by releasing on Labor Day instead of Thanksgiving.

These exercises should be viewed with caution. First, there is a perfect-competition assumption. Bigger movies will account for their effect on competition and will not be as affected by the estimated-competition effect. A big movie cannot double its revenues by switching from Thanksgiving to Labor Day. A very big movie will account for most of the competition effect, and releasing in weeks with high underlying demand will be optimal. Second, the exercise is based on box-office revenues. There may be other important factors. For example, an early summer release allows distributors time to release the video and DVD versions before Christmas (Chiou, 2006). Similarly, a Thanksgiving release may complement related merchandise sales in the subsequent shopping season, especially for children's movies. In contrast, distributors claim that a Labor Day release may send a bad signal regarding a movie's quality. These considerations are beyond the scope of this exercise.

The next exercise allows the industry to coordinate release dates to maximize joint revenues. I assume that total movie quality is divisible and each movie comprises of many small movies of quality ε, which can be distributed over different weeks. In practice, the industry must take indivisibilities into account, so the estimates below provide an upper bound. Under this divisibility assumption, the industry chooses the number of movies released each week, n_w, subject to a constraint on the quality and number of produced movies. The industry solves the program:

$$\max_{\{n_w\}_{w=1}^{W}} R\left(\{n_w\}_{w=1}^{W}\right) = \sum_{w=1}^{W} \frac{n_w^{1-\sigma}\exp(\varepsilon)}{\exp(-\tau_w) + n_w^{1-\sigma}\exp(\varepsilon)} \text{ subject to } \sum_{w=1}^{W} n_w = N, \tag{7}$$

where N is the total number of small movies to be allocated throughout the year.

A closed-form solution does not exist, but one can examine the first-order conditions. At the optimum, each n_w must satisfy

$$\frac{\exp(-\tau_w)}{\left(\exp(-\tau_w) + n_w^{1-\sigma}\exp(\varepsilon)\right)^2}(1-\sigma)\,n_w^{-\sigma}\exp(\varepsilon) = \mu, \tag{8}$$

where μ is the Lagrange multiplier. The estimated parameters imply a positive cross-partial of the left-hand side of equation (8) with respect to τ_w and n_w. If movies were uniformly distributed over the year, joint profits would increase by shifting some movies from low-demand weeks to high-demand weeks. This is a feature of the nested logit model, consistent with the amplification effect. The number of movies released will mimic (and amplify) the estimated underlying demand.

I use the estimates from the benchmark model to construct weekly market size (τ_w) and average (over years) weekly quality of movies released, an approximation for n_w. Substituting these estimates into equation (7) yields a total industry market share of 3.94 (which is, by construction, close to the observed average industry share of 3.92). The optimal solution yields a significantly more amplified release pattern than what we observe, but improves industry market share to only 3.98. This is not surprising. First, much of the demand is explained by movie quality, and the total quality constraint bounds revenues. Second, the σ parameter is above 0.5, and any redistribution over weeks primarily redistributes revenues across movies, with little effect on total industry share.

6. Concluding discussion

■ I decompose the seasonal pattern in total box-office revenues into underlying demand and the endogenous market reaction. The latter implies a strong seasonal pattern in the number and quality of movies released. My main finding is that gross seasonality is amplified by the release decisions. Underlying demand accounts for only about two-thirds of the seasonal variation in total sales.

My results demonstrate the endogeneity of observed seasonal patterns. This endogeneity is not limited to the motion picture industry. Warner and Barsky (1995) and Chevalier, Kashyap, and Rossi (2003) document that retail prices are often lower in weekends and holidays. This will lead to an amplification effect, similar to the one I find for movies. Cooper and Haltiwanger (1993) describe a drastic change in the seasonality of automobile sales following an exogenous change in the timing of new model introductions. They demonstrate that gross seasonality in the auto industry is driven by supply-side effects. These results imply that a naive interpretation of seasonality in industry sales as seasonality in demand may be misleading. Observed seasonality is often an equilibrium outcome.

It is natural to ask why I find an amplification effect and whether it generalizes to other industries. Because ticket prices (and distribution contracts) are stable over the year,[23] the market reaction in movies is limited to changes in the number and quality of movies released. Bigger markets are likely to attract more and bigger movies. More or better movies lead to market expansion, amplifying the observed seasonal pattern. It is the rigidity in ticket pricing that drives the amplified seasonality. More-flexible pricing may lead to higher prices in high-demand seasons, limiting the amplification effect.

One should expect a similar amplification effect in industries in which choice sets vary over time but prices are stable, such as video and DVD rental and sales markets, CDs, and books. A dampening is likely in industries in which prices vary but products are fixed. In tourism and airlines, capacity constraints create higher prices in high-demand seasons, dampening overall sales. Industries such as clothing have seasonal products as well as varying prices, and the overall effect of seasonality is ambiguous. Analyzing the market response to seasonality across industries may be an interesting direction for future research.

References

ACKERBERG, D.A. AND RYSMAN, M. "Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects." *RAND Journal of Economics*, Vol. 36 (2005), pp. 771–788.

BARSKY, R.B. AND MIRON, J.A. "The Seasonal Cycle and The Business Cycle." *Journal of Political Economy*, Vol. 97 (1989), pp. 503–534.

BERGMAN, N. AND EINAV, L. "Disentangling Observed Seasonal Patterns: Evidence From the U.S. Motion-picture Industry." Mimeo, Harvard University, 2001.

BERRY, S. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics*, Vol. 25 (1994), pp. 242–262.

———, LEVINSOHN, J., AND PAKES, A. "Automobile Prices in Market Equilibrium." *Econometrica*, Vol. 63 (1995), pp. 841–890.

CARDELL, N.S. "Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity." *Econometric Theory*, Vol. 13 (1997), pp. 185–213.

CAVES, R.E. *Creative Industries*. Cambridge, Mass.: Harvard University Press, 2000.

CHEVALIER, J.A., ROSSI, P.E., AND KASHYAP, A.K. "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data." *American Economic Review*, Vol. 93 (2003), pp. 15–37.

CHIOU, L. "The Timing of Movie Releases: Evidence from the Home Video Industry." Mimeo, Occidental College, 2006.

COOPER, R. AND HALTIWANGER, J. "Automobiles and the National Industrial Recovery Act: Evidence on Industry Complementarities." *Quarterly Journal of Economics*, Vol. 108 (1993), pp. 1043–1071.

CORTS, K.S. "The Strategic Effects of Vertical Market Structure: Common Agency and Divisionalization in the U.S. Motion Picture Industry." *Journal of Economics and Management Strategy*, Vol. 10 (2001), pp. 509–528.

DAVIS, P. "Measuring the Business Stealing, Cannibalisation and Market Expansion Effects of Entry in the Motion Picture Exhibition Market." *Journal of Industrial Economics*, Vol. 54 (2006), pp. 293–321.

DE VANY, A. AND WALLS, W.D. "Bose–Einstein Dynamics and Adaptive Contracting in the Motion Picture Industry." *Economic Journal*, Vol. 106 (1996), pp. 1493–1514.

——— AND ———. "The Market for Motion Pictures: Rank, Revenue and Survival." *Economic Inquiry*, Vol. 35 (1997), pp. 783–797.

EINAV, L. *Essays in Industrial Organization*. Ph.D. Dissertation, Department of Economics, Harvard University, 2002.

———. "Not All Rivals Look Alike: Estimating an Equilibrium Model of the Release Date Timing Game." Mimeo, Department of Economics, Stanford University, 2003.

[23] See Orbach and Einav (2007) for a discussion of potential causes for the uniform prices in the industry.

ELIASHBERG, J. AND SHUGAN, S.M. "Film Critics: Influencers or Predictors?" *Journal of Marketing*, Vol. 61 (1997), pp. 68–78.

KRIDER, R.E. AND WEINBERG, C.B. "Competitive Dynamics and the Introduction of New Products: The Motion Picture Timing Game." *Journal of Marketing Research*, Vol. 35 (1998), pp. 1–15.

MIRON, J.A. AND ZELDES, S.P. "Seasonality, Cost Shocks, and the Production Smoothing Model of Inventories." *Econometrica*, Vol. 56 (1988), pp. 877–908.

MOUL, C.C. "Measuring Word-of-Mouth's Impact on Theatrical Movie Admissions." *Journal of Economics and Management Strategy*, forthcoming.

MULLIGAN, J.G. AND MOTIERE, L. "The Market for First-Run Motion Pictures." Working Paper no. 94-5, University of Delaware, 1994.

NELSON, R.A., DONIHUE, M.R., WALDMAN, D.M., AND WHEATON, C. "What's an Oscar Worth?" *Economic Inquiry*, Vol. 39 (2001), pp. 1–16.

NEVO, A. "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry." *RAND Journal of Economics*, Vol. 31 (2000), pp. 395–421.

———. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica*, Vol. 69 (2001), pp. 307–342.

ORBACH, B.Y. AND EINAV, L. "Uniform Prices for Differentiated Goods: The Case of the Movie-Theater Industry." *International Review of Law and Economics*, Vol. 27 (2007), pp. 129–153.

PETRIN, A. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy*, Vol. 110 (2002), pp. 705–729.

PRAG, J. AND CASAVANT, J. "An Empirical Study of the Determinants of Revenues and Marketing Expenditures in the Motion Picture Industry." *Journal of Cultural Economics*, Vol. 18 (1994), pp. 217–235.

RADAS, S. AND SHUGAN, S.M. "Seasonal Marketing and Timing New Product Introductions." *Journal of Marketing Research*, Vol. 35 (1998), pp. 296-315.

——— AND ———. "Services and Seasonal Demand." In T.A. Swartz and D. Iacbucci, eds., *Handbook of Services, Marketing, and Management*. London: Sage Publications, 1999.

SAWHNEY, M.S. AND ELIASHBERG, J. "A Parsimonious Model for Forecasting Gross Box-Office Revenues of Motion Pictures." *Marketing Science*, Vol. 15 (1996), pp. 113–131.

SQUIRE, J.E., eds. *The Movie Business Book*. New York: Simon and Schuster, 1992.

VANDERHART, J. AND WIGGINS, S.N. "Advertising, Quality, and Motion Picture Performance." Mimeo, Texas A&M University, 2001.

VERTER, J. AND MCGAHAN, A.M. "Coming Soon: A Theater Near You." Harvard Business School Case Study no. 9-797-011, 1996.

VOGEL, H. *Entertainment Industry Economics*. New York: Cambridge University Press, 1994.

WARNER, E.J. AND BARSKY, R.B. "The Timing and Magnitude of Retail Store Markdowns: Evidence from Weekends and Holidays." *Quarterly Journal of Economics*, Vol. 110 (1995), pp. 321–352.