THE ECONOMIC JOURNAL

NOVEMBER 1996

The Economic Journal, 106 (*November*), 1493–1514. © Royal Economic Society 1996. Published by Blackwell Publishers, 108 Cowley Road, Oxford OX4 1JF, UK and 238 Main Street, Cambridge, MA 02142, USA.

BOSE–EINSTEIN DYNAMICS AND ADAPTIVE CONTRACTING IN THE MOTION PICTURE INDUSTRY*

Arthur De Vany and W. David Walls

Audiences discover what films they like and spread the word. Information feedback produces Bose–Einstein dynamics in the motion picture box office revenue distribution. Information cascades generate box office 'hits' and 'flops'. The revenue distribution evolves recursively over the 'run' as a mixture of the 'opening' and a stochastic competition among films. The motion picture run is decentralised, adaptive, and self-organising: semi-stationary admission prices, local information, and contingency-rich contracts match the film rental price and the supply of theatrical engagements to demand.

A film is like no other product... it only goes around once. It is like a parachute jump. If it doesn't open, you are dead.

Producer Robert Evans (Litwak, 1986, p. 84)

The hard part about understanding the motion picture industry is coming to grips with the way demand and supply operate. Film audiences make hits or flops and they do it, not by revealing preferences they already have, but by discovering what they like. When they see a movie they like, they make a discovery and they tell their friends about it; reviewers do this too. This information is transmitted to other consumers and demand develops dynamically over time as the audience sequentially discovers and reveals its demand. Supply must adapt sequentially as well, which means there must be a great deal of flexibility in supply arrangements. Pricing must be equally flexible.

The crucial factor is just this: nobody knows what makes a hit or when it will happen. When one starts to roll, everything must be geared to adapt successfully to the opportunities it presents. A hit is generated by an information cascade. If supply can ride the cascade, a superstar might be the result. A flop is an information bandwagon too; in this case the cascade kills the film. The discovery of preferences, the transmission of information, and state-contingent adaptation are the key issues around which the motion picture market is organised. This organisation is supported by adaptive contracts.

In this paper we explore theoretically and (mostly) empirically how demand

* This paper is dedicated to Ross Eckert. We would like to thank Michael Intriligator, Boyan Jovanovic, Steve Lippman, Walter Oi, Sherwin Rosen, and participants of the NBER Conference on New Products for comments on an earlier version of this paper. We alone are responsible for the paper's contents.

and supply dynamics and the path of the distribution of film revenues are related. What makes the sequential discovery of demand and adaptive supply difficult to model is the complicated distributional dynamics they can produce. If there are 50 films playing, how will information flows affect how well they do against one another? As demand unfolds, how will box office revenues change and how are these distributional dynamics linked to the opening? How do the dynamics shape the final distribution of total revenue at the end of the run? How must supply adapt to support these dynamics? What mechanisms support this adaptation? These are some of the issues considered here.[1]

We model the discovery and transmission of information and analyse the sequential dynamics of demand under various transmission mechanisms. We show that a Bayesian demand process leads to Bose–Einstein distributional dynamics for motion picture revenues. Other processes lead to distributions such as the uniform, the geometric, the Pareto and Zipf laws, and the log normal. By testing characteristics of the distribution of film revenues and its dynamic path during the course of the motion picture run, we are able to gain an understanding of the nature of the information transmission and adaptive supply dynamics in this industry. The evidence shows there is positive information feedback among film audiences that is captured by the Bose–Einstein process and that the industry's flexible supply adaptation to this feedback produces increasing returns. The Bose–Einstein process can produce information cascades (Bikhchandani *et al.* 1992), path dependent dynamics (Arthur, 1989), and superstars (Rosen, 1981), all of which are exhibited by motion pictures.

We begin in Section I with a discussion of some of the institutional features of the motion picture industry that are essential to understanding its demand and supply dynamics. The discovery and transmission of demand information is modelled in Section II and various processes and their associated revenue distributions are posed as candidates to be tested against the data. Section III tests the empirical implications of the models of information transmission with a large sample of motion picture runs in *Variety*'s Top-50. In Section IV we examine various aspects of the industry's contracts and practices in light of our empirical results and argue that the motion picture run is decentralised, adaptive, and self-organising – by using semi-stationary admission prices, local information, and contingency-rich contracts the industry flexibly matches the film rental price and the supply of theatrical engagements to demand. Section V concludes the paper.

I. THE MECHANICS OF SUPPLY

Some Examples of Motion Picture Runs and Revenues

To fix some ideas about the motion picture theatrical 'run' and how it responds adaptively to demand, consider some of the patterns relating box

[1] We are influenced in this endeavour by Brock (1993) who argues that to come to a real understanding of distributions one must model the underlying processes and dynamics that produce them.

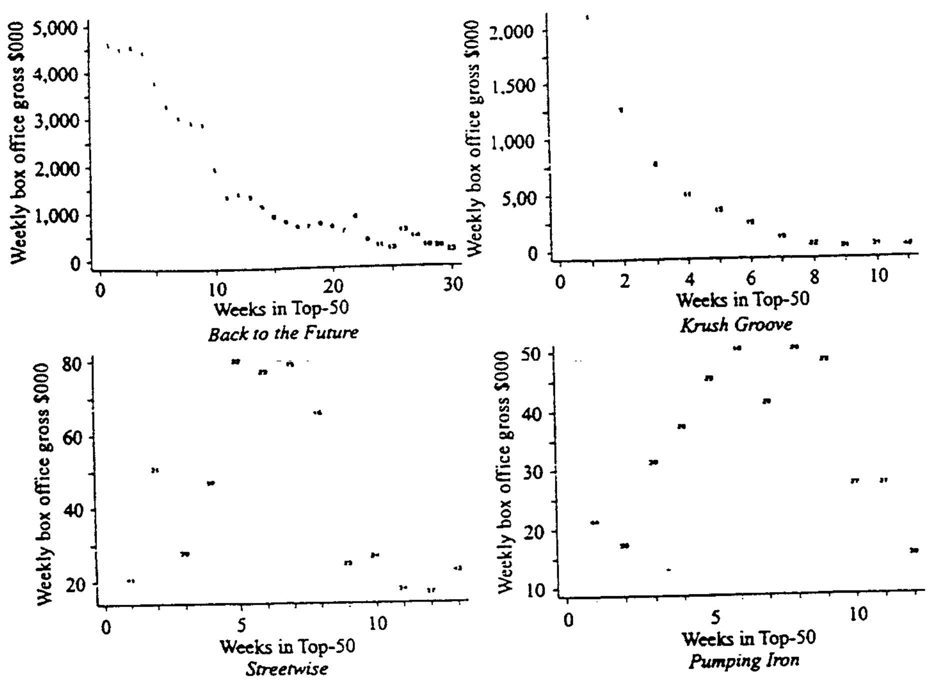

Fig. 1. Patterns of film runs. Numeric plotting symbols indicate rank in *Variety*'s Top-50.

office revenue, the rank or position of a film in the weekly *Variety*'s Top-50 and the length of the run of a film in cinemas. In Fig. 1 we show how a few films in our sample played.[2] The figure shows the rank, weekly revenue (box office gross), and week of run for four films: *Back to the Future*, *Krush Groove*, *Streetwise*, and *Pumping Iron*. The films shown in the top two panels opened strongly and played well; they had long runs and high revenues. *Back to the Future* played 30 weeks in the Top-50; it achieved rank 1 in its first week and held that rank for over 10 weeks and then declined in rank and box office revenue to the end of our sample. *Krush Groove* showed a similar pattern of a high opening revenue and near top rank followed by a uniform decline of rank and revenue. This pattern contrasts with films that build revenues over their lives, examples of which are given in the lower two panels of Fig. 1. These films opened moderately at low ranks and then began a climb in rank and revenue. *Pumping Iron*, Arnold Schwarzenegger's first film, opened at rank 44, fell to 50, and then began a steady climb to 18. It ranked in the 20s for a few weeks and then declined in rank and played off in a few more weeks. *Streetwise* followed a similar pattern – a climb in rank and revenue followed by a descent. Most movies follow neither of these patterns, but all successful ones do. The 'bombs' or

[2] The sample is described more fully in Section III. In this paper, we consider only the first run of films playing in domestic cinemas; subsequent runs, foreign runs, and video and TV runs are not considered. The domestic theatrical run accounts for about 30% of film rentals. It is an important market in its own right and still is important to launching a film.

failures open poorly (at low rank and revenue) and decline quickly, ending their run in just a few weeks.[3]

Booking and Supply

Once a film is produced, it is distributed to cinemas who 'exhibit' it for audiences.[4] The distributor chooses a release pattern – the number and location of cinemas in which the film is licensed to play or 'booked' – based on an *a priori* appraisal of demand. The size of the initial release determines the number of copies that are required for distribution to each of the cinemas. Beyond that initial supply, the number of viewers who can see the film can grow to the capacity of cinemas. Adjustments in supply are made by expanding the number of weeks the film runs in the cinemas.

Contracting

The contract to exhibit a film usually requires the cinema to show it for a minimum number of weeks; 4 weeks is a common minimum, though 6 and 8 week minimums are sometimes used. In addition, the contract contains a 'hold-over' clause that requires the cinema to continue the film another week if the previous week's box office revenue exceeds a stipulated amount. The contract also 'clears' an area near the cinema were the distributor cannot license the film to other cinemas.

Cinemas rent a motion picture when they exhibit it. The film rental is a percentage of the box office revenue which the film attracts at the cinema. The typical arrangement calls for the exhibitor to pay 90% of the cinema's box office revenue in excess of a fixed amount that is negotiated as part of the contact. In addition, the rental rate is subject to a minimum percentage on total box office revenue (without the fixed deduction). The minimum rental rate declines over the run; a typical arrangement might be a minimum of 70, 70, 60 and 40% of the total box office revenue in the first, second, third and fourth weeks, respectively. Additional weeks beyond the contract period often go at 40%, but the 90–10 split continues to hold throughout the run for reasons we will explain later.[5]

Given these arrangements, the cost to the distributor of expanding the seats available to accommodate customers is the incremental cost of extending the run at cinemas another week. This cost is the opportunity cost of booking the film in another cinema. If the prospective revenues from additional bookings at other cinemas are sufficiently high, then the print may be moved (at essentially the same marginal cost of extending the engagement at the present cinema). The exhibitor and the distributor usually prearrange in the contract

[3] De Vany and Walls (1994) show that films have hazard rates that are characteristic of highly stressed products; there is a sharp peak in the hazard at about 4 weeks.

[4] See Kenney and Klein (1983) and Blumenthal (1988) for analyses of the auction institutions under which exhibitors bid for distributors' films.

[5] See De Vany and Eckert (1991) for details and an analysis of motion picture exhibition contracts.

to extend the run whenever the film earns a high revenue in the preceding week. This is done by specifying a 'hold over' amount of revenue in the prior week that obliges both the distributor and the cinema to continue the run another week.

Adaptation

Distributors and exhibitors use box office reports to match supply to demand over time and geographically. New exhibitors try to book the film if it is drawing large audiences; they can be added in accord with contractual commitments to the initial exhibitors and the availability of prints. Distributors may produce more prints if demand is strong, though this may take some time. Supply adjusts dynamically to demand through changes in bookings at exhibitors, and by lengthening the runs at the cinemas where it is booked. As noted, the decision to extend runs is decentralised to the local cinema were the film is booked and are made weekly conditional on the revenue earned the preceding week.

A failure produces a different pattern. Poor box office revenues in early weeks fail to attract new exhibitors. Poor reviews and negative word-of-mouth information discourages additional viewers. Exhibitors try to cut their losses by dropping the film, though the exhibition contract limits this adjustment. If the film is bad enough, distributors may even permit the cinema to end the run before the contract term. In the absence of that drastic action, distributors may let exhibitors add a second feature on the marque; this is an implicit reduction in admission price, since the admission buys two films rather than one. In the worst cases, heavily promoted films like *Howard the Duck* or *Last Action Hero*, fail early leaving wounded exhibitors, distributor losses, and broken contracts in their wake.

A third form of supply adaptation is noteworthy: films not considered to be suitable for a mass audience, like *Piano*, *Smoke*, or *Hannah and her Sisters*, may be released in a few cinemas to see if they will catch on. If film-goers like a film and spread the word, then the audience will grow over time as cinema bookings and run lengths adjust to positive word-of-mouth communication.

Release Strategies

Release patterns are strategies for acquiring demand information. A wide release places a film on many screens. This strategy draws a large, simultaneous sample in many cinemas and cities. A tailored release strategy samples sequentially, starting at a few cinemas and using the information from that sample to adjust bookings if the film builds a following. The risk of releasing a film widely to many cinemas is that it may play off rapidly, say, in 4–8 weeks. If it is then dropped by cinemas, many potential viewers will miss it and the information dynamic that might make it a hit may not get under way. A film released through the small sample, sequential strategy may capture good word-of-mouth information dynamics and play 10 or 15 or even 20 weeks in a varying number of cinemas each week.

II. DEMAND DISCOVERY AND DISTRIBUTIONAL DYNAMICS

Given the above discussion of the contracts and release strategies that support an adaptive supply response, we now turn to the demand discovery process. Our aim is to examine how information discovery and transmission determine the dynamics of revenues over the course of the run and how they shape the size distribution of revenue.

Search, Communication and Demand Dynamics

Suppose for the moment that there is just one film and only one person at a time can see it. Consumers choose in random sequence whether to go to the film or not. If potential viewers share an evaluation of a film's quality, there is a common probability p that a randomly chosen person will see the film. If we let X be the number attending the film, then X is a binomial random variance such that

$$P\{X = k|p\} = \binom{x}{k} p^k (1-p)^{n-k}, \quad k = 0, 1, \ldots, n. \tag{1}$$

In the simplest case, where consumers share a common, error free evaluation of a film's quality, the film's attendance, run length, and consequently, its revenues would follow a binomial distribution.

The more interesting case is where the quality is unknown and consumers differ in their evaluation of a film's quality. Because quality is unknown and evaluations differ among viewers, p is a random variable. By conditioning on p and integrating over the binomial distribution, we get:

$$\begin{aligned} P\{X = k\} &= \int_0^1 P\{X = k|p\} f(p)\, dp \\ &= \int_0^1 \binom{n}{k} p^k (1-p)^{n-k}\, dp \\ &= 1/(n+1), \quad k = 0, 1, \ldots, n. \end{aligned} \tag{2}$$

The last equation shows that each of the $n+1$ possible outcomes is equally likely. Uncertainty in the distribution of the personal quality evaluations of the movie transforms the distribution of run length, attendance, and revenue from the binomial to the uniform distribution.

To model the information dynamics that produce the uniform distribution, we consider a sequential Bayesian decision process. A Bayesian film-buff can rely on information revealed during a film's run to refine her initial valuation to its quality. Suppose a potential film-goer has information about the opinions of other viewers; this information is not exact and it may include word-of-mouth reports from friends, movie reviews, advertising, and information about how the film is playing that can be extracted from box office reports and/or crowds at cinemas. Jovanovic (1987) has modelled this kind of information-sharing and our model is a variation on his.

The return to film-goer i is a function of the film's unknown quality which depends on viewer type and an unknown parameter θ which is common to all valuations. We assume the benefit to viewer i is

$$b_i = \theta + \epsilon_i, \quad i = 1, \ldots, N, \tag{3}$$

where the ϵ_i are the personal experiences, modelled as deviations from the common quality θ, and are assumed to be independently and uniformly distributed with zero mean. No consumer knows her valuation of the film until she sees it and experiences its quality $\theta + \epsilon_i$.

The information set available to potential viewer $n+1$ in the sequence of potential customers is $\mathbf{s}_n = (\theta + \epsilon_n, \theta + \epsilon_{n-1}, \ldots, \theta + \epsilon_0)$, where n is the number of previous trials and each element of the vector is the experienced quality of each previous trial. The expected return to i conditional on $\mathbf{s}_n$ is

$$\mathrm{E}(b_i \mid \mathbf{s}_n) = \mathrm{E}(\theta \mid \mathbf{s}_n). \tag{4}$$

If the cost of seeing the movie is common to all movie-goers and equal to c, then a randomly selected consumer on trial $n+1$ will choose to see it if

$$\mathrm{E}(b_{n+1} \mid \mathbf{s}_n) > c, \tag{5}$$

and not otherwise.

Given an unknown θ and the uniform distribution of ϵ, the distribution of quality of the movie is unknown.[6] There is a common, but unknown, conditional probability that a randomly selected person will choose to see the film. This probability is equal among all consumers who inherit the same information set because they will all have the same estimate of θ, given $\mathbf{s}_n$ and the expectations of their ϵs are all zero. Since p is equal for all trials, it can be shown (Ross, 1993) that the conditional probability that the $(r+1)$st trial is a success given k successes and $r-k$ failures in the first r trials is

$$\frac{P\{(r+1)\text{st trial is a success, } k \text{ successes in first } r \text{ trials}\}}{P\{k \text{ successes in first } r \text{ trials}\}} = \frac{k+1}{r+2}. \tag{6}$$

If the first r trials go to the movie, then the next trial will with probability $(k+1)/(r+2)$ go to the movie. The probability will evolve dynamically over the run as the discovery process unfolds. The revenue earned by the film in each week of its run (the trial period in our sample) will reveal information as to the 'true' probability p and the underlying quality parameter θ. The unconditional distribution of revenue is uniform, but the sequential dynamics leading to the final outcome will exhibit conditioning of weekly revenues on prior outcomes; weekly revenues are not independent trials because they reveal information which will cause revenues to be autocorrelated.

[6] This quality search model is equivalent to Rothschild's (1974) model of searching for the lowest price when the distribution of prices is unknown. Lippman and McCall (1976) deal elegantly with this model. In the price search model, the searcher relies on her previous trials to update; here the distribution of quality is unknown and a Bayesian searcher relies on the quality reports of the searchers ahead of her to update θ, the location parameter of her distribution.

Demand Dynamics and the Bose–Einstein Distribution

Now, consider the situation when there are m motion pictures playing on s cinema screens. We allow for the distributor's booking strategy to incorporate multiple screen releases for motion pictures with high prior expectations. An equilibrium booking pattern will be such that the number of screens per motion picture equalises *a priori* revenue expectations. Thus, by considering the distribution of movie goers over screens to be multinomial uniform, we obtain a Dirichlet prior.[7] We have only to extend the previous model a little to let the outcomes be of s types. Let there be m movies, then the types are $1, \ldots, s, s+1$, where type i is 'went to cinema i', and type $s+1$ is 'did not go to a cinema and saw no film'. Denote by X_i the revenue at cinema i. The number of type i outcomes in n independent trials, each resulting in one of the s outcomes is the vector $\mathbf{X} = (X_1, \ldots, X_s)$.

The Dirichlet is a multinominal uniform distribution whose dynamics have an urn model interpretation that is natural in the present setting. To develop the distributional dynamics associated with the conditional choice logic underlying the process with a Dirichlet prior, we follow Ross (1993) and compute the conditional probability that the $(n+1)$ st outcome is of type j if the first n trials result in x_i type i outcomes where $i = 1, \ldots, s$, and $\sum_{i=1}^{s} x_i = n$.[8] Ross (1993, pp. 118–9) shows this computation leads to

$$P\{(n+1)\text{ st is } j \mid x_i \text{ type } i \text{ in first } n, i = 1, \ldots, n\} = (x_j + 1)/(n+s). \qquad (7)$$

The probability that the $(n+1)$ st outcome is type $j = 1$ depends directly on the proportion of previous n trials that are type $j = 1$. Because p is a random variable, the successive outcomes are not independent and a Bayesian movie buff gains information about the quality of the movie by observing the previous trials.

The unconditional distribution of the vector $\mathbf{X}$ corresponding to this dynamic can be shown to be

$$P\{X_1 = x_1, \ldots, X_s = x_s\} = \frac{n!(s-1)!}{(n+s-1)!} = \binom{n+s-1}{s-1}^{-1}. \qquad (8)$$

This is the Bose–Einstein distribution. This distribution is uniform, but not in the way you would usually think of it. The usual way to think of a uniform distribution of outcomes over s cinemas would be to say that the probability

[7] If we asume that the n viewers are distributed among the s screens according to probabilities $(p_1, \ldots, p_s)$, the multinominal distribution is Dirichlet. The distribution of viewers among films $m \leq s$ showing on multiple screens will also be Dirichlet (DeGroot, 1968). The movie search problem – a search for quality with an unknown distribution – as similar to the search for price with an unknown distribution. Viewers who do not know the distribution begin with a uniform prior and adapt from there. Rothschild (1974) shows that a searcher facing an unknown price distribution has a Dirichlet prior if and only if his experience can be parameterised in exactly the way we have done it for movie quality searchers. We think this is a completely natural way to characterise information sets for Bayesian film buffs when screen bookings are set to equalise prior expectations of revenues.

[8] This model also has a well-known interpretation as an Polya's Urn Model (Ross, 1993, p. 117).

that a customer will go to any one of them is $1/s$. But, this is according to the multinominal distribution, not the Bose–Einstein distribution. In the dynamics leading to the Bose–Einstein distribution movie customers sequentially select movies and the probability that a given customer selects a particular movie is proportional to the fraction of all the previous movie goers who selected that movie.[9] This is a natural result of the fact that the probabilities are not known and sampling reveals information that causes previous selections to attract new ones.

The Bose–Einstein distribution has the property that all of the possible outcome vectors **X** are equally likely. This means the vector **X** in which the attendance at every cinema is equal to zero is as likely as one in which all n trials go to only one cinema. Every other vector is equally likely.[10] The Bose–Einstein distribution has uniform mass over a space of $s+1$-vectors of large dimension; the s-vectors are urns containing the revenues of the s cinemas and one urn collecting those who go to no film.[11] What is important about the evolution of choice probabilities under the Bose–Einstein choice logic is the way past successes are leveraged into future successes; as soon as individual differences emerge among the films, these differences can grow at exponential speed. Differences among audiences that emerge at the opening can be compounded by information feedback through the run into very large differences by the end of the run. Path dependence from the opening depends on the stochastic evolution of demand at each cinema. A broad opening at many cinemas can produce high and rapidly growing audiences; but it also can lead to swift failure if the large early crowd relays negative information. Even though a film may have high initial expectations and appear on many screens, it can disappear rapidly if initial cinema expectations are not met and negative information flows promote a declining audience.[12] The Bayesian choice model and its associated Bose–Einstein dynamics leads to genuine insights about the movie business.

Demand Dynamics and other Distributions

There are many other information and choice processes that one must consider as reasonable candidates to explain motion picture revenue dynamics. We develop a few reasonable alternatives here and indicate how they can be tested

[9] This diffusion of audiences over films is similar to the process in Brian Arthur's (1989) model, where the probability of choosing a product depends on the number who previously chose it, as well as the number who chose other products: in fact, it suggests that the Arthur model may generate the Bose–Einstein statistics.

[10] To obtain the distribution of each film's revenues, we need to consider the probability that a group of k cinemas showing film j will contain a total audience of r. Feller (1957) gives this as:

$$P\left\{\sum_{1}^{k} X_t = x_j\right\} = \binom{k+r-1}{k-1}\binom{s-k+n-r-1}{r-k} \div \binom{s+n-1}{n}.$$

[11] Consider an audience of just 1,000 film-goers and 49 cinemas plus one 'did not go' element to give a 1×50 outcome vector. There are $2{\cdot}03 \times 10^{210}$ possible outcome vectors.

[12] Notable examples are *Howard the Duck* and *Last Action Hero* two highly promoted films with high expectations and large initial cinema bookings. Both were gone from the market within a few weeks and exhibitors who booked at the opening were permitted to double-bill before the term of their exhibition contract expired.

against the Bayesian model. Continuing to denote the revenue of film i as X_i, we now indicate revenue at time t as $X_i(t)$. We can generate another distributional dynamic if we assume that the revenue grows at a rate

$$r_i(t) = \frac{1}{X_i(t)} \frac{dX_i(t)}{dt} \tag{9}$$

which may vary systematically or randomly because of challenges from other films or external shocks to the market. This process leaves the rate of growth completely free to vary over the course of a film's run. Total revenue at time t may be obtained by integrating (9) to get

$$\log X_i(t) = \log X_i(0) + \int_0^t r_i(t')\, dt'. \tag{10}$$

At week t of the run, each of the terms in (10) contributes to the total revenue: $\log X_i(0)$ is the contribution of the opening (week 0 of the run), and $\int_0^t r_i(t')\, dt'$ is the contribution of random factors. If the first term dominates, then the opening is most important; if the second term dominates, the opening is unimportant, and stochastic factors that occur during the run predominantly determine the total revenue the film will bring.

The integral term is the log transformation of the product of independent random variables $(r_i(t)\, r_j(t-1))$, and so is the sum of independent variables. If the film runs long enough, then according to the central limit theorem, the integral will be normally distributed. Hence, $X(t)$ is log normally distributed.

If the opening exerts influence, there is still a tendency for the revenue distribution to approach the log normal, but the final distribution will reflect the draw $X_i(0)$ at the opening, taken from some distribution that characterises opening revenues, and the distribution that sums the stochastic factors over the run. If we consider the distribution of all movies, it will be a mixture of the distribution of opening revenue and the distribution reflecting the stochastic evolution of demand which tends to a log normal distribution.

If we raise the random variables to a power in the distribution of (10), the resulting distribution will be normal in the log of the log of revenue. Such a distribution would be evidence of strong interactions among film-goers of the form $[r_i(t)\, r_j(t-1)]^x$.

Another distributional dynamic corresponds to a geometrical splitting of the audience among films. Suppose films interact in a complex way that expresses the idea that they competitively preempt one another in a hierarchical manner. They do this when they all seek to capture a large share of a static audience of a given size. If the most successful film takes fraction k of the audience and the next takes fraction k of the remainder, and so on, then the distribution of box office revenue will be geometric.[13] In this case, the relation between sales (X_i) of film i and its rank with rank R_i is approximately

$$X_i = k(1-k)^{R_i - 1}. \tag{11}$$

[13] Schmalensee's (1978) model of preemption in the cereal market could produce this market fractioning and a geometric share distribution, or the looser log series distribution (see the next paragraph).

Hence, geometric market shares, of which constant shares is a special case, will produce a power law distribution of revenue.[14] If the geometric market division process is not realised precisely, then the distribution will be a slightly modified geometric distribution called the log series distribution (May, 1983).

Power laws commonly are found in economics. Steindl (1965) found that the relationship between a firm's size and it's rank could be approximated by a power law. It has been shown that a power law is an implication of a model in which the growth rate of firms is independent of size (Gibrat's law), and the rate of birth/death of firms is constant.

III. EMPIRICAL BOX OFFICE REVENUE DISTRIBUTIONS

The Data

Our data are the box office revenues of *Variety*'s Top-50 motion pictures by week. The *Variety* sample is a computerised weekly report of domestic film box office performance in major and medium metropolitan market areas. These data follow the revenues of each picture for the duration of its run, starting from the picture's opening and progressing by week throughout its run, or until it drops off the Top-50 list. Each week, the pictures are ranked by order of their box office revenue. The number of screens on which each film played and other information are also given in the data. Our sample includes all the motion pictures to appear in *Variety*'s Top-50 list from 8 May 1985 to 29 January 1986; some 300 motion pictures are in the sample.

Descriptive statistics of the sample are given in Table 1. The top four grossing movies in our sample accounted for 20·5% of the total box office revenue, and the top eight grossing movies accounted for 28·1% of the total box office revenue. In descending order the eight highest grossing movies in our sample were *Back to the Future*, *Rambo*, *Rocky IV*, *Cocoon*, *The Jewel*, *The Goonies*, *The Colorado*, and *White Knights*. The four lowest grossing films accounted for only 0·0036% of the total box office revenue, and the eight lowest grossing films accounted for 0·0082% of the total box office revenue. The highest grossing film, *Back to the Future*, had a cumulative revenue of $49.2 million, while the lowest grossing film, *Matter of Importance*, had a cumulative revenue of $4,615. The median box office revenue was $311,284; the mean was $1,918,261; the S.D. at $4,528,685 was over twice the mean.

Lorenz Curves and Gini Coefficients

We initially quantify the revenue distributions through the use of Lorenz curves and Gini coefficients. These confirm that the distribution of total revenues among films and movie distributors is highly uneven. Revenues are concentrated on a few films, and these hit films are concentrated among a few distributors.

Fig. 2 plots Lorenz curves for the distribution of box office revenues across

[14] This form of the Pareto law, which uses rank instead of the number of firms above a certain size, is known as Zipf's (1949) law.

Table 1

Summary Statistics of Variety's *Box Office Sample*

Variable	Obs.	Min.	Max.	Median	Mean	S.D.
Rank at birth	350	1	50	26	25·51	16·19
Rank at weeks = 2	266	1	50	22	23·06	15·49
Rank at weeks = 3	232	1	50	24·5	24·24	15·43
Rank at weeks = 4	191	1	50	21	23·46	14·30
Rank at weeks = 5	173	1	50	24	24·28	13·39
Rank at death	300	1	50	39	36·03	11·81
Weeks at rank = 1	14	1	12	1·5	2·86	3·03
Weeks at rank = 2	27	1	4	1	1·48	0·75
Weeks at rank = 3	30	1	4	1	1·33	0·71
Weeks at rank = 4	35	1	3	1	1·14	0·43
Weeks at rank = 5	33	1	3	1	1·21	0·60
Weeks in Top-50	300	1	40	4	5·69	5·15
Showcases	2,000	1	347	13	53·12	73·25
Cities	2,000	1	191	4	7·44	8·08
Weekly revenue ($000s)	2,000	2·5	10,000	57·5	335·70	685·26

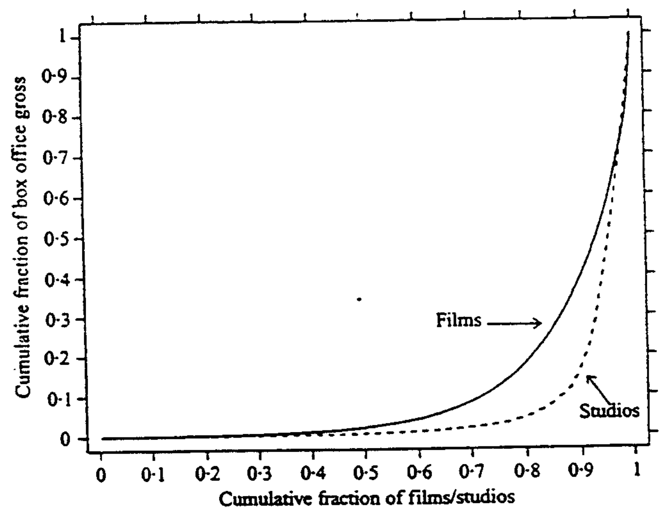

Fig. 2. Lorenz curves: distribution of revenues over films and studios.

movies and distributors. If the revenues were distributed equally, the Lorenz curves in Fig. 2 would be straight lines from the lower left corner to the upper right corner. One can readily see that both curves are far below the diagonal, which would indicate a uniform distribution, and the curve of distributor revenues is farther away from uniformity than is the Lorenz curve of film revenues.

The Gini coefficients for the distribution of revenues among films and distributors are 0·777 and 0·873, respectively; these values indicate a high

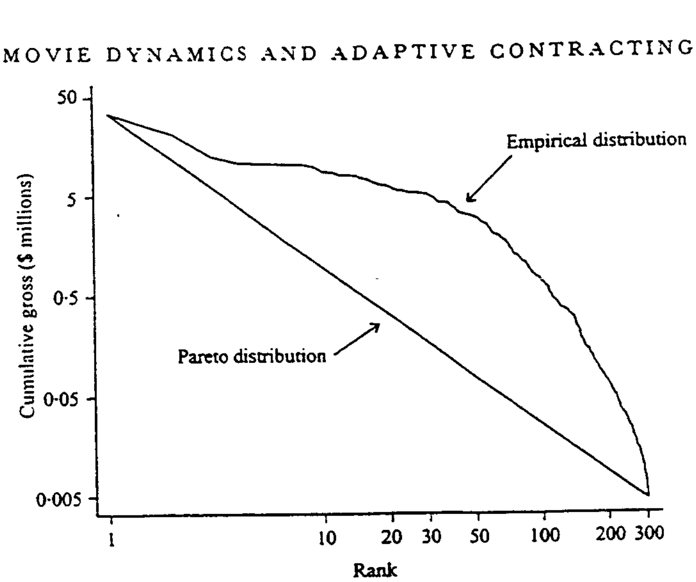

Fig. 3. Pareto distribution of cumulative revenues over films.

degree of inequality in the distributions.[15] Just 20% of the films earned 80% of box office revenues; this is such a durable relationship, it has been dubbed 'Murphy's Law' after Art Murphy, a well-known industry reporter who has reported on the business for decades. The distribution of revenues of distributors is more concentrated than film revenues: 90% of all the revenue earned by distributors was earned by just 10% of the 79 distributors in our sample.[16] Distributors of successful films tended to have several successes and those that distributed failures tended to have multiple failures or only moderately successful films. There is no evidence of distributor diversification: hits received the bulk of the box office revenue and a few distributors had most of the hits.

The Rank–Revenue Relationship

Some of the distributional dynamics we consider as candidates imply the presence of power laws between rank and revenue; others do not. We can use the rank–revenue relationship to accept or reject some of the candidate dynamics. Ijiri and Simon (1974) show that firms have autocorrelated growth when the plot of log X_i against log R_i is concave downward.[17] This means that deviations in the concave direction from a log-linear rank–revenue relationship imply autocorrelated growth.

[15] The Gini coefficient is a measure of inequality of the distribution of revenues across films and distributors. It ranges from 0 to 1, with a value of 0 indicating equality and higher values indicating less equality. It is the ratio of the area between the Lorenz curve and this line to the area above the line. It can be computed as $1-2\int_0^1 L(x)\,dx$, where $L(x)$ is the equation of the Lorenz curve. To calculate this, the Lorenz curve was first approximated by a 7th order polynomial using linear regression; the R^2 of the regression was 0·9992. The estimated Lorenz curve was integrated analytically to obtain the area beneath it.

[16] On average, there were 4·43 films released per distributor. The top four distributors made $21\frac{1}{3}$% of the films in the sample and the top eight made $37\frac{1}{3}$% of the films in the sample.

[17] We know film births and deaths are not constant (De Vany and Walls, 1994), so we already know before we test that we are unlikely to find a power law.

Table 2

Estimates of the Rank–Revenue Relationship

(Log (revenue) = $\beta_1 + \beta_2 \log(\text{rank}) + \beta_3 [\log(\text{rank})]^2 + \mu$.)

Aggregation	(1) Weekly by film	(2) Weekly by film	(3) Cumulative by film	(4) Cumulative by film	(5) Cumulative by distributor	(6) Cumulative by distributor
Regressors						
Log rank	−1·7206	0·1859	−1·9826	2·1366	−2·7222	−0·3589
	(0·0146)	(0·0435)	(0·0515)	(0·1031)	(0·0787)	(0·1772)
(Log rank)^2	—	−0·4033	—	−0·5299	—	−0·4512
	—	(0·0089)	—	(0·0130)	—	(0·0329)
Constant	16·6333	14·8405	21·7375	14·5904	21·4484	18·9070
	(0·0923)	(0·0759)	(0·2480)	(0·2004)	(0·2676)	(0·2306)
Weekly dummy variables	Yes	Yes	—	—	—	—
$\bar{R}^2$	0·8772	0·9397	0·8319	0·9744	0·9462	0·9858
Observations	2,000	2,000	300	300	69	69

Note: Estimated S.E. are in parentheses.

Fig. 3 plots the empirical rank–revenue relationship and the fitted Pareto distribution for our sample of 300 movies. It is clear from the figure that the rank–revenue relationship is downward concave, which is evidence of autocorrelated growth in film revenues. To test the rank–revenue relationship, regressions of the following form were run:

$$\log X_i = \log \alpha + \beta \log R_i + \gamma (\log R_i)^2 + \mu, \tag{12}$$

where X_i is a movie's revenue, R_i is its rank, and μ is a random disturbance with mean zero and finite variance.

If the geometric dynamic and its Pareto law describes the rank–revenue relationship, then the regression coefficient γ should be statistically indistinguishable from zero; finding that $\gamma = 0$ would be inconsistent with the hypothesis of increasing returns to motion pictures caused by information feedback. Finding that γ is less than zero would be evidence of positively autocorrelated growth in motion picture revenues. Conversely, if $\gamma > 0$ the rank–revenue distribution is convex downward and we have evidence of negatively autocorrelated growth.

Table 2 reports the estimates of the rank–revenue relationship for several levels of revenue aggregation. Columns (1) and (2) report the estimates when the rank–revenue relationship is estimated across the 50 films in each week's sample. The data from the 40 weeks of the sample were pooled and an intercept dummy was included for each week (except the base week), so there are 2,000 (40 × 50) observations. The hypothesis of a linear Pareto distribution can be rejected in favour of the alternative of downward concavity: the coefficient on (log rank)^2 is negative and statistically significant. Columns (3) and (4) report the estimates of the rank–revenue relationship where the unit of observation is the film. There were 300 films in our sample which were observed to fall from the Top-50 chart; the remaining films were censored and thus were excluded from the regression. The rank–revenue distribution clearly shows statistically significant evidence of downward concavity. The distributor is the unit of revenue aggregation in columns (5) and (6). As with the other units of aggregation, the distribution of distributor revenues by rank is also characterised by downward concavity, which is consistent with autocorrelated growth.

Revenue Distribution Dynamics

We showed that some dynamics can lead to log or log-log normal distributions. A formal test shows that the distribution of cumulative box office revenues is not log normal. Table 3 presents several test statistics for log normality, and the

Table 3
Testing Log-Normality of Box Office Revenues

Test name	Test statistic	Marginal significance level
Skewness–kurtosis	$\chi^2 = 22.560$	0.00000
Shapiro-Wilk	$Z = 3.018$	0.00127
Shapiro-Francia	$Z = 3.468$	0.00026

Note: The hypothesis of 3-parameter log normality could also be rejected at a marginal significance level of 0.02.

null hypothesis of log normality could be rejected in each case; we could also reject the hypothesis of 3-parameter log normality.

To quantify the dynamics of the revenue generation process, we have calculated the distribution of revenues across films by the week-of-run (Fig. 4). The distribution of revenues at the film's opening (week = 1) is far from log normal though the distributions appear to converge to log normality. However, we can statistically reject the hypothesis of log normality for each week-of-run revenue distribution as well as the end-of-run distribution at the 1% significance level. Thus, the opening performance is statistically a dominant factor in revenue generation, although it appears to diminish in importance near the end of a film's run.

What we find is evidence of convergence to normality of the log of the log of weekly revenue among long-lived films. Table 4 shows the marginal significance level for the null hypothesis that the log of the log of weekly cross is normally distributed. We can reject normality of log(log(revenue)) for weekly revenues across movies lasting 6 weeks or less. But, we cannot reject normality of log(log(revenue)) of films that last beyond 6 weeks. (Note that this test requires a minimum of eight observations). Normality of log(revenue) is consistent with normality of the products of random variables, which is how one would expect information to spread through interaction among member of the population of film-goers. But, normality of the log(log(revenue)) suggests a stronger form of

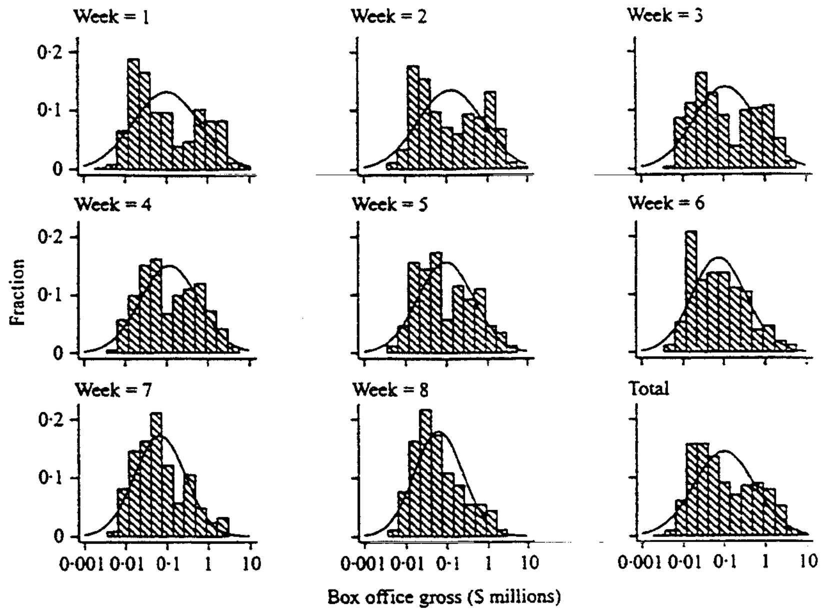

Fig. 4. Distribution of revenues by week-of-run.

Table 4
Testing Log-Log-Normality of Box Office Revenues

Week of run	Marginal significance level
1	0·0000
2	0·0000
3	0·0000
4	0·0000
5	0·0004
6	0·0082
7	0·1656
8	0·1841
9	0·0736
10	0·1983
11	0·1782
12	0·1710
13	0·3014
14	0·1330
15	0·4840
16	0·6070
17	0·4625
18	0·3060
19	0·4308

interaction for those films that survive 6 or more weeks. The interactions must be products of random variables raised to a power. The extreme upper tail of the motion picture revenue distribution is populated by films whose revenues are generated by power laws. So, beyond a run life of 6 weeks, the power laws that Brock (1993) argues will usually hold in informationally rich economic environments do in fact hold.

What Information Dynamic Accounts for the Distribution?

From our formal tests we know the distribution of box office revenue is not log normal; its departure from log normality is in a direction that indicates autocorrelated growth. There are several processes that could account for this. Films may be hierarchically dividing up a fixed audience so that revenues grow geometrically. We rejected the geometric process and the Gibrat law by formally testing the hypothesis of a geometric revenue distribution and rejecting it. Film revenues that grow at a constant or geometric rate would give a linear or nearly linear plot between the log of revenue and the log of rank and we formally rejected that too. The departure from power law behaviour is evident in the rank and revenue plots and the formal test of log-concavity of the log(revenue) versus log(rank) relationship.

We are left with two possibilities that can account for both the dynamics of revenue and the shape of the end-of-run revenue distribution. One is that movie audiences are produced by independent stochastic growth factors and an opening that has a uniform log distribution. This process will exhibit the required from log normality in the form of a mass point in the upper tail for those films where the opening is more important throughout the film's run than the stochastic factors. But, for this process, the growth in sales should be random and independent of prior growth. We know this is not true because we can formally reject independent growth. The other possibility is the one we accept for now. The best account of the data is given by the Bayesian model and the Bose–Einstein dynamics.

Information transmission and Bayesian updating give rise to the revenue distributions and their dynamics. We can account for all their features if there is uncertainty over film quality and information sharing gives us Dirichlet priors. Then Bayesian updating will produce the Bose–Einstein dynamics. If the unconditional distribution of film revenues is Bose–Einstein, then they are maximally uncertain. This is true in three senses: one, the Bose–Einstein distribution is uniform and maximally uncertain if measured by entropy, mean-preserving spreads, or stochastic dominance. Two, the space of outcomes is vast and consequently each outcome is highly unlikely. Three, the Bose–Einstein distribution is uniform but in a counter-intuitive way that masks the real uncertainty it implies. A uniform distribution of outcomes suggests that the probability that a customer will go to any one of m films is $1/m$. But, this is multinominal choice logic, not Bose–Einstein choice logic. In the information and choice dynamics that lead to the Bose–Einstein distribution customers choose movies in proportion to the previous film goers who selected that movie. When revenues are Bose–Einstein distributed, outcomes which differ in the extreme are equally likely and similar outcomes are extremely unlikely – this is the quintessential characteristic of the movie business.

IV. EXPLAINING THE INDUSTRY

How do these results explain the way the industry operates? How do its contracts and practices adjust the supply of theatrical engagements to capture the demand processes that seem to be operating here? Since the choice probabilities evolve during the run, and they are not known beforehand, dynamic optimisation in this context must essentially be adaptive. In this section, we discuss industry practices and institutions from this perspective.

Launching the Opening

Stars, and large production and advertising budgets can place a film on many exhibitor screens when it opens. This can generate high initial revenues in total, if not in each cinema. If viewers like the film and spread the word, it will earn high revenue in the following weeks. But, if viewers do not like the film, the large opening audience transmits a large flow of negative information and revenue may decline at a rapid rate. A wide release is vulnerable to negative feedback. Because a wide release lowers the gross revenue per cinema, exhibitors may drop the film sooner than they would drop a less widely released film. Thus, widely released films show more variance in the revenues and, on average, shorter run lives (De Vany and Walls, 1994). A major star-studded film may earn advance distribution fees from foreign and the post-theatrical market. The willingness of exhibitors and downstream sources of revenue like cable television, VCR distributors, pay per view and network television as well as foreign distributors to pay advance guarantees for motion pictures before their theatrical run is a major inducement for distributors to produce big budget films and promote them heavily; for such films, the cinema market (which we have studied here) can be less important than these other sources of revenues. There is a certain irony here. Some stars and films earn so much from foreign distribution rights that third world motion picture tastes may play a large role in determining the kind of films seen in US cinemas.

Adaptive Contracting

A difficult problem to solve contractually is how to keep a film on screens long enough for it to build an audience. If an exhibitor takes such a film, it is on the risk that it may build so slowly during his run that only exhibitors who show it later will benefit from information feedback. The Paramount Consent Decrees bar cinema ownership by the major distributors; they also bar franchises, extended clearances in time and location, blind-selling (booking a film before the exhibitor has seen it), and long-term contracts. Because the decrees bar long-term, exclusive showings, they make it difficult to guarantee that the first exhibitor who took the risk of introducing the film will benefit if the film later becomes a success (De Vany and Eckert, 1991). When the audience grows recursively, as the evidence suggests, then the Paramount contracting restrictions prevent risk-taking exhibitors from capturing the demand externality which they helped to make. The effect of the Paramount prohibitions is to restrict the gateway to cinema screens, allowing fewer films that would appeal to smaller and more sophisticated audiences to pass through the gate.[18]

Admission Pricing

One industry practice, widely followed though not encoded in contracts, is to use relatively inflexible admission prices to reveal excess demand and to adjust to demand by lengthening the run. Prices rarely are raised even when excess demand is revealed.[19] A relatively stationary admission price combined with a count of admissions gives a reliable signal of demand, and this signal is transmitted throughout the industry by real time reporting of box office revenues. This reporting is required in the exhibition contract and encouraged by other means as well.[20] Demand is accommodated by lengthening the run rather than raising the price and this is accounted for by the model of information transmission. If, at the first sign of excess demand, the admission price were to be increased, the number of people who would see the film in the opening weeks would fall and this would reduce the flow of information from this source to potential viewers. This lower rate of information transfer would lead to a shorter run and a lower total level of demand. The ability to extend the run makes an almost perfectly elastic supply response possible, so that price need not rise to choke off excess demand.[21]

The optimal admission price must trade off price and run length so as to maximise total revenue in the context of ignorance, not merely uncertainty, of demand. Charging a higher price will shorten the run and reduce the transfer of information; this strategy may capture only a fraction of the demand that eventually will be captured in the adaptive run so successfully used in the industry. Because motion pictures have short lives all adaptations must be made quickly and are of a short run nature.[22] Inflexible admission prices lead to a pure quantity signal and an adaptive supply response which seems ideally suited to the problem of discovering and responding to demand.

[18] Smith and Smith (1986) analyse film rental champions released before and through the 1950s, during the 1960s, and in the 1970s for evidence that the characteristics of successful films have changed. The Paramount Decree of 1949 made the vertically integrated studios divest their cinemas and altered contracting in the 1950s and after and could have changed film characteristics along the lines we describe in the text.

[19] It is not uncommon for exhibitors to restrict passes and discounts in the opening weeks, or even for the exhibition license to require it; this raises prices before demand is revealed for what usually is the period of highest demand. But, if the film is playing poorly, passes and discounts or even double billing are used to lower the effective admission price.

[20] Because film rentals are based on box office revenues, distributors will usually earn less rental if a cinema lowers its admission price. Moreover, exhibitors can reduce the film rentals even while increasing their total earnings by manipulating admission prices. A stationary price overcomes some of these disadvantages. See De Vany and Eckert (1991) and De Vany and McMillan (1993) for a discussion of how the industry deals with information and incentive issues that arise in reporting revenues.

[21] A similar mechanism is used in pricing Broadway plays. A large price increase after a successful opening chokes off the information dynamics and shortens the run, which undercompensates the artists who are paid by the week. The demand dynamics and compensation of artists makes a stationary price and adaptive run the most effective way to supply Broadway plays and compensate artists.

[22] If there were a permanent and recognised increase in the demand for motion pictures, admission prices would rise generally at cinemas throughout the industry. We are dealing here with the pricing of imperfect substitutes in the short run, a situation for which flexible prices with some form of quantity rationing, like delivery lags or queues, are well-suited. See, Gordon (1981) for a discussion of the general issues, Carlton (1978) for an analysis of delivery lags, and De Vany (1976) for an analysis of queuing.

Rentals

The contingency-rich exhibition contracts used in the industry are highly adaptive: they rely on locally generated information; they set the rental fee in a precise and non-linear way in response to demand; they share risk between exhibitors and distributors; and they create incentive for exhibitors to show films by granting a measure of exclusivity, although this is limited in scope by the Paramount Decrees. The rental price adapts precisely to the state of demand. It is usually calculated as a percentage of box office gross, as we said in Section I, so the rental varies directly with demand. But, very high demand will trigger a switch from the minimum percentage, which may be 40 or 70% depending on the week of the run, to 90% on the amount in excess of the fixed amount. The rental schedule is non-linear and has higher rates at higher grosses. Yet, in line with the declining revenues during the descent in rank and revenue that always occurs sometime in the run, the minimum percentage rental rate declines with the week of the run. This offsets the rising hazard the exhibitor faces as the run continues (De Vany and Walls, 1994).

Decentralisation

Each film's run through the market is sequential in order to exploit information dynamics. The run is self-organised because it decentralises the decision to extend the run to each cinema and uses only local information to extend or close the run at each location. The initial release is modified over time through this process and new engagements can be added subject to prior contractual obligations. These contractual features interact to adaptively capture revenue and generate strongly increasing returns from highly successful films. They help to create the Bose–Einstein distribution because they adapt sequentially to let exhibitors and the distributor ride the information cascade (up or down). When the cascade is headed in the right direction, the supply response is flexible enough for some films to ride the cascade to superstardom.

V. CONCLUSIONS

We have tied the information dynamics of motion picture audiences to the dynamics of the distribution of motion picture weekly box office revenues for the top 50 films. Our results show that weekly revenues are autocorrelated: a film which recently experienced increasing revenues is more likely to experience additional growth than a film which experienced growth in the distant past. The end-of-run or total revenue distribution for motion pictures approximates a mixture of an opening distribution and a log normal distribution, suggesting that if the film runs long enough it may acquire an independent life. Yet, the distribution never quite reaches log-normality; it has fatter tails than the log normal and mass points at the far right, where the superstars are located. Films that show extraordinary staying power exhibit power law behaviour in the extreme tail of the revenue distribution. The Bayesian information dynamics that gives the best account of the evolution of revenues is one that leads to the Bose–Einstein distribution.

The strong evidence supporting the Bose–Einstein distribution points to the stark uncertainty facing the motion picture industry; the space of outcomes in which the 50 films vying for audiences in our sample can evolve is of staggering dimension and the dynamics evolve unpredictably. Star vehicles may have higher initial expectations, but booking patterns and information dynamics produce highly unpredictable distributional dynamics and uneven revenues. The brief tenure of studio heads and the wide use of artist participation contracts is consistent with the profound uncertainty implied by our empirical findings.[23] Neither genre nor stars can guarantee success and a disaster can bankrupt a studio.[24]

Even though the motion picture industry exhibits profoundly the properties we have modelled, there are many other industries for which the model and statistical methods would seen to be appropriate: other entertainment industries such as television, music, and book publishing might be usefully modelled in the framework we develop here. In addition, the sequential dynamics of innovation and demand discovery in many industries where information sharing and Bayesian demand discovery are at work may exhibit the Bose–Einstein process. The recursive generation of demand and adaptive supply and pricing that are the primary ingredients of our approach seem to be fundamental elements of many models that focus on distributional dynamics, information diffusion, superstars, and path dependence.

The industry's organisation and contracts make sense in the uncertain and dynamic environment of the motion picture industry. The adaptive run seizes opportunities to expand supply and adjust prices sequentially using local information and decentralised decision making to do it. This adaptive response is encoded in industry contracts and institutions; it is the mechanism that lets supply follow path-dependent demand sequences to capture box office revenues and make the highly irregular distributions we reported here. The industry's market institutions and sophisticated contracts promote the information dynamics that discover extraordinary motion pictures and let supply and pricing adapt to capture the strongly increasing returns they create.

University of California, Irvine

University of Hong Kong

Date of receipt of final typescript: March 1996

References

Arthur, B. (1989). 'The economy and complexity'. In *Lectures in the Sciences of Complexity: Santa Fe Institute Studies in the Sciences of Complexity* (ed. D. Stein), pp. 713–40. Redwood City, CA: Addison-Wesley.

Bikhchandani, S., Hirshleifer, D. and Welch, I. (1992). 'A theory of fads, fashion, custom, and cultural change as informational cascades.' *Journal of Political Economy*, vol. 100, pp. 992–1026.

[23] On the tenure of studio executives, we found that over 30% of letters in a survey of nearly 400 executives were returned with no forwarding address. On participation contracts as a means of lowering studio risk, Weinstein (1995) notes that nobody would make *Forrest Gump* until Tom Hanks agreed to forego his usual fee and take a percentage of profits instead.

[24] *Heaven's Gate* ruined United Artists.

Blumenthal, M. A. (1988). 'Auctions with constrained information: blind bidding for motion pictures.' *Review of Economics and Statistics*, vol. 70, pp. 191–8.

Brock, W. A. (1993). 'Scaling laws for economists.' Working paper, University of Wisconsin at Madison.

Carlton, D. (1978). 'Market behavior with demand uncertainty and price inflexibility.' *American Economic Review*, vol. 68, pp. 571–87.

De Vany, A. (1976). 'Uncertainty, waiting time and capacity utilization – A stochastic theory of product quality.' *Journal of Political Economy*, vol. 84, pp. 523–41.

De Vany, A. and Eckert, R. (1991). 'Motion picture antitrust: the Paramount cases revisited.' *Research in Law and Economics*, vol. 14, pp. 51–112.

De Vany, A. and McMillan, H. (1993). 'Markets, hierarchies, and antitrust in the motion picture industry.' Working paper, University of California at Irvine.

De Vany, A. and Walls, W. D. (1994). 'The market for motion pictures: rank, revenue and survival.' Working paper, University of California at Irvine.

DeGroot, M. (1968). *Optimal Decision Theory*. New York: Wiley.

Feller, W. (1957). *An Introduction to Probability Theory and Its Applications*. New York: Wiley.

Gordon, R. (1981). 'Output fluctuations and gradual price adjustment.' *Journal of Economic Literature*, vol. 19, pp. 493–500.

Ijiri, Y. and Simon, H. A. (1974). 'Interpretations of departures from the Pareto curve firm-size distributions.' *Journal of Political Economy*, vol. 82, pp. 315–32.

Jovanovic, P. (1987). 'Micro shocks and aggregate risk.' *Quarterly Journal of Economics*, vol. 17, pp. 395–409.

Kenney, R. W. and Klein, B. (1983). 'The economics of block booking.' *Journal of Law and Economics*, vol. 26, pp. 497–540.

Lippman, S. A. and McCall, J. J. (1976). 'Job search in a dynamic economy.' *Journal of Economic Theory*, vol. 12, pp. 365–90.

Litwak, M. (1986). *Reel Power: The Struggle for Influence and Success in the New Hollywood*. New York: William Morrow.

May, R. M. (1983). 'Patterns of species abundance and diversity.' In *Diversity: Benchmark Papers in Ecology* (ed. R. Patrick), pp. 340–79. Stroudsburg, PA: Hutchinson Ross.

Rosen, S. (1981). 'The economics of superstars.' *American Economic Review*, vol. 71, pp. 167–83.

Ross, S. (1993). *Introduction to Probability Models*, 5th ed. San Diego: Academic Press.

Rothschild, M. (1974). 'Searching for the lowest price when the distribution of prices is unknown.' *Journal of Political Economy*, vol. 82, pp. 689–711.

Schmalensee, R. (1978). 'Entry deterence in the ready-to-eat breakfast cereal industry.' *Bell Journal of Economics*, vol. 9, pp. 305–27.

Smith, S. P. and Smith, V. K. (1986). 'Successful movies: a preliminary empirical analysis.' *Applied Economics*, vol. 18, pp. 501–7.

Steindl, J. (1965). *Random Processes and the Growth of Firms: A Study of the Pareto Law*. New York: Hafner.

Weinstein, M. (1995). 'Profit sharing contracts in Hollywood: when is the "net" "gross".' Working paper, University of Southern California Law Center.

Zipf, G. K. (1949). *Human Behavior and the Principle of Least Effort*. Reading, Massachusetts: Addison-Wesley.