GIBRAT'S LEGACY*

by

John Sutton
London School of Economics and Political Science

The Toyota Centre
Suntory and Toyota International Centres for
Economics and Related Disciplines
London School of Economics and Political Science
Houghton Street
London WC2A 2AE
Tel.: 020-7955 6674

Discussion Paper
No. EI/14
October 1996

---

* The support of the Leverhulme Trust and the Economic and Social Research Council is gratefully acknowledged. I would also like to thank Alison Hole, Reinout Koopmans, Michael Raith and three anonymous referees for their helpful comments.

Abstract

This paper traces the time series (?Growth of Firms?) tradition in the study of market structure and looks at how recent studies on entry and the size distribution of firms have modified thinking in this area.

Keywords: Robert Gibrat, market structure, *Inégalités Économiques,* the Law of Proportional Effect, firm size, industry structure.

Contact address: Professor John Sutton, STICERD, London School of Economics and Political Science, Houghton Street, London WC2A 2AE. email: j.sutton@lse.ac.uk

I. BACKGROUND

One of the most important strands in the literature on market structure begins with Robert Gibrat's Inégalités Économiques, published in Paris in 1931. Gibrat's book presented the first formal model of the dynamics of firm size and industry structure, and its lengthy subtitle confidently announced a 'new law': the Law of Proportional Effect.

Gibrat traced the origins of his thinking to the work of Jacobus Kapteyn, an astronomer who had become interested in the widespread appearance of skew distributions in various settings, especially in biology. Kapteyn's approach was to assume that underlying such distributions was a simple Gaussian process: a large number of small additive influences, operating independently of each other, would generate a normally distributed variate z (The 'Law of Laplace' in Kapteyn's

book.). An observed skew distribution of some variate x could be modelled by positing that some underlying function of x was normally distributed (Kapteyn and M.J. Van Uven (1916)). In his book, Gibrat postulated the 'simplest' such process, suggesting that the logarithm of x developed in such a fashion. This amounts to saying that the expected value of the increment to a firm's size in each period is proportional to the current size of the firm[2].

The simplest way of presenting the argument, following Joseph Steindl (1965), runs as follows. Denote the size of the firm at time t by x_t and let the a random variable ε_t denote the proportionate rate of growth between period

[2]'Size' can be measured in a number of ways, and these arguments have been variously applied to measures of annual sales, of current employment, and of total assets. Though we might in principle expect systematic differences between the several measures, such differences have not been a focus of interest in the literature.

(t-1) and period t, so that

$$x_t - x_{t-1} = \varepsilon_t x_{t-1}$$

whence

$$x_t = (1+\varepsilon_t)x_{t-1} = x_0(1+\varepsilon_1)(1+\varepsilon_2)\ldots(1+\varepsilon_t)$$

If we choose a 'short' time period, then we can regard ε_t as being 'small', justifying the approximation $\log(1+\varepsilon_t) \simeq \varepsilon_t$. Taking logs, we thus obtain

$$\log x_t \simeq \log x_0 + \varepsilon_1 + \varepsilon_2 + \ldots + \varepsilon_t$$

By assuming the increments ε_t to be independent normal variates with mean m and variance σ^2, we have from the Central Limit Theorem that as $t \to \infty$, the term $\log x_0$ will be small compared to $\log x_t$, so that $\log x_t$ is approximated by a normal distribution with mean mt and variance $\sigma^2 t$. In

other words, the limiting distribution of x_t is lognormal[3].

Gibrat first applied this to income distributions, and then to plant (establishment) sizes in manufacturing. The goodness of fit which he obtained was striking. (An example is shown in Figure 1). Gibrat presented a broad range of data of the size distribution of establishments. This data permitted comparisons across time (1896-1921), and between broad sectors of the national economy (agriculture, commerce), regional sectors (industrial establishments in Alsace-Lorraine), and some very narrowly defined industries (electrochemicals and explosives; metallurgy[4]). Gibrat's aim was to convince his

[3]More formal versions of this argument were developed in the 1950s and '60s (see below). Several of these early models combined a proportional growth hypothesis on firm size x with a constant population of firms (no entry or exit). One immediate implication of this is that the variance of the (lognormal) distribution increases over time (see footnote 8 below).

[4]Gibrat also examined the cyclical behaviour of the size distribution. He found, interestingly, that when an industry expanded, firm numbers rose slowly in

readers that this was a statistical regularity sufficiently sharp to provide a basis for serious mathematical modelling. In this he succeeded, albeit with a very long lag. Michael Kalecki's 1945 article[5] describes Gibrat's book as a 'great achievement', but as Gibrat had noted, the collection and analysis of large data sets in this area

---

comparison with output, most of the growth being captured by incumbents; but when output fell, firm numbers fell sharply, as (mostly small) firms disappeared. The cyclical behaviour of plant expansion/contraction versus entry/exit has been explored in the recent literature by Tito Boeri and Ulrich Cramer (1992) and by Steven Davis and John Haltiwanger (1992).

[5]Kalecki (1945) made a number of prescient observations. First, he noted that the Law of Proportionate Growth (applied to a fixed population of firms) implied that the variance of the distribution would increase indefinitely, a feature which seemed 'unrealistic' relative to (many) economic processes. He examined a number of modified processes, including one in which the expected rate of growth increased less than proportionately, leading to a lognormal distribution with constant variance. More interestingly, he allowed that unspecified 'economic' factors would drive the process, and examined the robustness of the law to various constraints on the process. Finally, he re-examined the empirical performance of the prediction of 'lognormality' for U.K. data on establishment size (where it worked 'fairly well') and income distributions (where a modified distribution which converged to the lognormal for large x fitted reasonably well).

involved a heavy burden of work. It was not until the 1950s and '60s that the apparent regularity of the size distribution became the focus of sustained empirical effort. By that time, a second research literature had emerged whose motivation lay in a different, but equally salient feature of market structure.

This new 'cross-sectional' literature was motivated by the observation that market structure varied in a systematic way from one industry to another, in a manner that seemed to be related to a number of industry characteristics, such as scale-economies, the role played by advertising, or the influence of R&D. The fact that the ranking of industries by some measure of concentration appeared broadly similar from one country to another lent weight to the notion that market structure was influenced by some basic 'industry characteristics' (Joe Bain (1966), Pryor (1972), Louis Phlips (1971)).

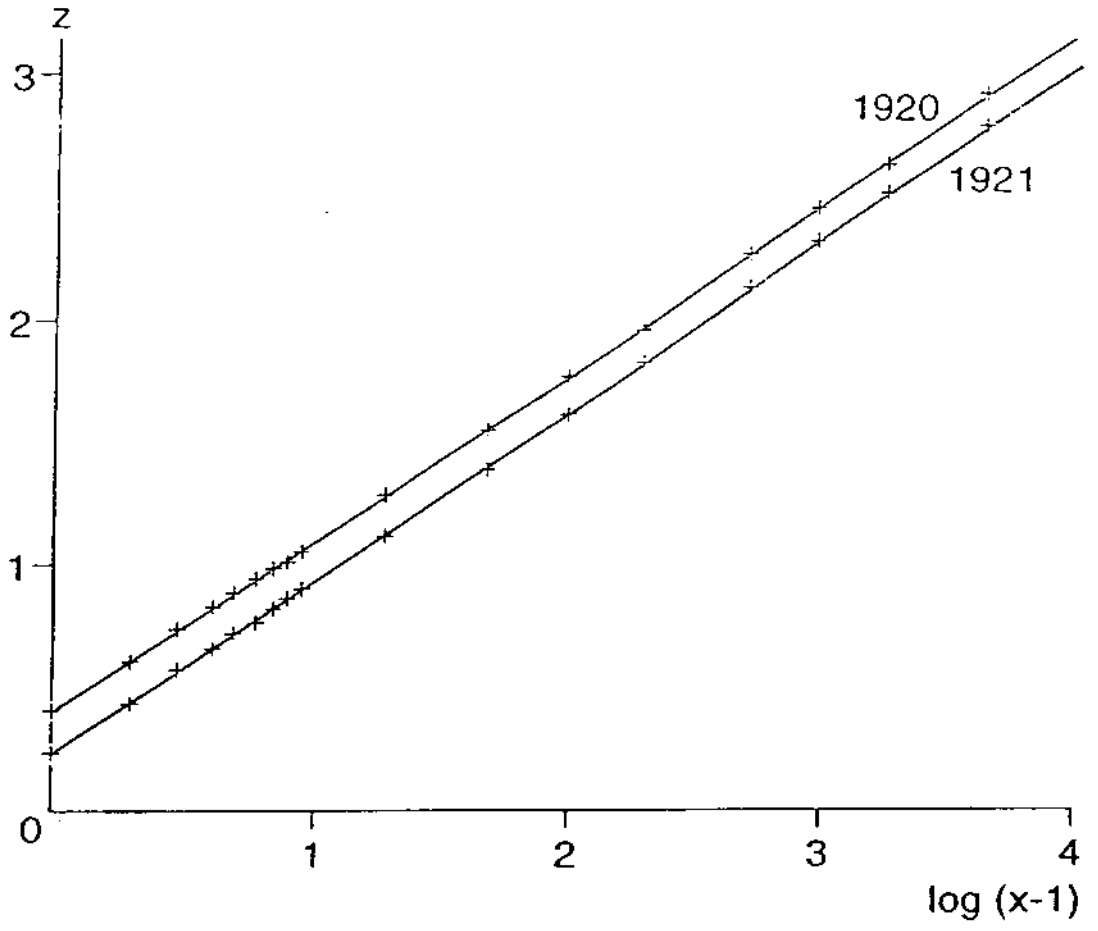

Figure 1.

Gibrat's data for French manufacturing establishments in 1920 and 1921. The horizontal axis shows log (x-1), where x is a size class, measured by employment, and the vertical axis shows z, calculated from observed values of R(z), the number of establishments of size z or greater, according to the formula for the cumulant of a normal distribution, $R(z) = \frac{1}{\sqrt{\pi}} \int_z^\infty e^{-z^2} dz$ (Gibrat (1931), p. 53 and p. 276). If the distribution of x is lognormal, the observations will lie in a straight line.

During the past fifteen years, this 'cross-sectional' literature has been re-formulated using game-theoretic models. While these models have proved to be extremely versatile in rationalizing observed outcomes, it is less easy to pin down what the theories exclude, and so to pinpoint the content of these models. This is now widely seen as a crucial issue for game-theoretic formulations of the cross-section literature. (For a companion review of the cross-sectional literature, see Sutton (1996).)

It is, however, the parallel revival of interest in the older Growth-of-Firms literature during the past decade which is the focus of interest in what follows. This revival had its roots, not in theory, but in new empirical findings which began to appear during the mid-80s. The development of the literature since then has involved a continual interplay between theoretical modelling and empirical evidence, and a shift of focus in terms of the empirical regularities which seem to be of primary interest

to researchers. At the same time, there has been a heavy emphasis on discarding the older type of purely 'stochastic' models in favour of introducing stochastic elements into standard 'maximizing' models.

The new 'maximizing' models focused on various specific settings which differed in the assumptions made on the nature of the technology, the information available to firms, and the description of the product market. Among the earliest studies was that of Lennart Hjalmarsson (1974), which examined the size distribution of plants in a homogeneous-goods industry, in the presence of increasing returns to scale at plant level[6]. The model of Boyan Jovanovic (1982) also assumed a homogeneous-goods industry, but introduced a learning mechanism which gradually revealed firm-specific efficiency differences as

[6]In a later paper, Hjalmarsson (1976) reports some tests of his model for two Swedish industries (particle board and cement) for which the assumptions made on product homogeneity and on plants' cost structure reasonable.

the industry evolved (see Section III below). Reinhard Selten (1983) considered a model in which firms incurred fixed outlays in enhancing consumers' willingness-to-pay for their respective products, which were then sold in a series of (sub-) markets of varying sizes[7].

One theme which emerges from these developments echoes the lesson learnt from the recent game-theoretic literature: reasonable 'maximizing' models can be designed in many ways, and their implications will depend on a wide range of industry characteristics, some of which will be very difficult to identify, or control for in empirical studies. In particular, there is no obvious rationale for

[7]Another important series of models, following Richard Nelson and Sidney Winter (1982) avoids strict maximizing assumptions in favour of weaker rationality requirements. This literature raises some fundamental questions as to the appropriateness of making strong rationality and informational assumptions on agents who face continuing technological change. A review of this literature, which lies beyond the scope of the present paper, will be found in Wesley Cohen and Richard Levin (1989).

positing any general relationship between a firm's size and its expected growth rate, nor is there any reason to expect the size distribution of firms to take any particular form for the general run of industries. Most authors now claim only that the distribution will be 'skew', but do not specify the extent of skewness, or the particular form which the size distribution might take. Meanwhile, empirical investigations from the 1960s onwards have thrown doubt on whether any single form of size distribution can be regarded as 'usual' or 'typical' for the general run of industries; wide differences in the form of the size distribution occur between one industry and another (Richard Schmalensee (1989), p. 994).

These developments in the Growth-of-Firms literature are very much reminiscent of what we have seen happening recently in the game-theoretic formulations of the cross-sectional literature. The lessons that might be drawn from this are explored in Section IV below.

II. THE EARLY LITERATURE

During the 1950s and '60s, a substantial class of models appeared which combined 'Gibrat's Law' with a range of ancillary assumptions. (For a review of these models, see Steindl's (1968) overview and his monograph of 1965.). Meanwhile, a growing empirical literature, associated with Peter Hart and Sigbert Prais in the U.K. and with Herbert Simon and his co-authors in the U.S., had a major influence within Industrial Organisation.[8]

---

[8]The Hart-Prais model, like that of Gibrat, implied that variance grew over time. Prais (1976) proposed this as a mechanism to explain the rate of increase in the share of the top 100 firms in U.K. industry. By measuring the variance in the year-to-year proportional growth rate, Prais could predict (ex-post) a trend rate of growth for the share of the top 100 firms for 1909 - 1970. While the predicted trend fitted the observations quite well, Prais' interpretation of this result was challenged. The stochastic growth model does not distinguish between 'internal growth' and growth achieved by way of acquisitions (or mergers). Leslie Hannah and John Kay (1977) claimed that most of the observed increase in concentration was due to mergers and acquisitions, rather than internal growth, and that once the effect of mergers and acquisitions were netted out, internal growth alone could not account for the

The generation of 'stochastic growth' models developed in this period culminated in the volume based on the papers of Simon and his co-authors (Yuji Ijiri and Simon (1977)), following which there was relatively little interest in the field for a decade. The model used by Simon and his co-authors modified Gibrat's model by incorporating an entry process according to which firm numbers rose over time as the industry grew. This avoided the implication that the variance of the size distribution increased without limit. The Simon model provides a useful point of departure in assessing the later literature. It assumes a framework in which the market consists of a sequence of independent opportunities, each of size unity, which arise over time. The best way to think of this is to imagine a number of isolated 'island' markets, each big enough to support exactly one plant. As each opportunity arises, there is

---

degree of increase in concentration. See, however, the comments on this issue in Prais (1976) and the later assessment of Hannah and Kay (1981).

some probability p that it will be taken up by a new entrant. With probability (1-p) it will be taken up by one of those firms already in the market ('active firms'). The size of any (active) firm is measured by the number of opportunities it has already taken up. There are two assumptions:

(i) Gibrat's law: the probability that the next opportunity is taken up by any particular active firm is proportional to the current size of the firm.

(ii) Entry: the probability that the next opportunity is taken up by a new entrant is constant over time.

Assumption (ii) is rather arbitrary, though it may be a reasonable empirical approximation. Simon regarded it merely as providing a useful benchmark, and presented

various robustness tests showing that 'reasonable' departures from the assumed constancy of p would have only a modest effect on the predictions of the model. The predictions are driven crucially by Assumption (i) (Gibrat's Law). What this leads to is a skew distribution of the Yule type, and Simon presented various empirical studies for large firms in the U.S. which suggested that it provided a good approximation to the size distribution of large manufacturing firms.

The goodness of fit of the size distribution provides only indirect evidence for Gibrat's Law. A second strand of the literature of the 1950s and '60s focused on the direct investigation of Gibrat's Law, by looking at the relation between firm size and growth over successive years in a panel of firms. While various studies of this kind cast doubt on the idea that proportional growth rates were independent of firm size, no clear alternative

characterisation emerged[9]. Summarizing the literature in 1980, Frederic Scherer remarked on the wide disparity between different studies, but tentatively concluded that assuming growth rates uncorrelated with initial firm size 'is not a bad first approximation,' at least for the U.S. On the other hand, most studies agreed that the standard deviation of growth rates rose less than proportionally with firm size (see Hart (1962), Stephen Hymer and Peter Pashigian (1962), Ajit Singh and Geoffrey Whittington (1968,1975) and Scherer (1980), p. 148. This observation raises the question of whether the pattern of residuals in certain fitted relationships might exhibit heteroscedasticity; see Section III.).

---

[9]While smaller firms are found to grow faster than large firms in most recent studies (see Section III below, and also Manmohan Kumar (1985), Zoltan Acs and David Audretsch (1990)), some earlier studies reported the reverse tendency (John Samuels (1965), Prais (1976)). This latter tendency in part reflects the greater role played by growth through acquisition among larger firms (Modified versions of Gibrat's Law, are described in Patrick McCloughan (1995).)

The contribution of Edwin Mansfield (1962) is of particular interest. Mansfield's study was based on samples of 'practically all' firms in three specific industries (steel, petroleum, tires) over a number of different time periods (generating ten samples in all). The author pointed out that 'Gibrat's Law' may be interpreted in different ways, depending on the way we treat firms 'disappearing' from the sample (whether by 'exit' or otherwise). Does the law relate to all firms, with the proportional growth rate of disappearing firms set at minus one? Or does it propose that the proportional rate of growth *conditional on survival* is independent of firm size?[10]

Mansfield's results rejected the first version of the hypothesis in seven of his ten samples, while the second version failed in four of his ten samples. While this led

[10]Mansfield also noted a further distinction, made by Simon, which postulated that the law held only above some minimal size of firm.

Mansfield to conclude that 'Gibrat's Law does not seem to hold up very well empirically', it nevertheless left open the possibility that Gibrat's law might still be true in yet another form. Consider the distribution of growth rates of firms that *would have resulted* if none had left the industry. Interpret Gibrat's Law as saying that this distribution is independent of firm size. Then it is possible that the measured growth-size relations could still exhibit the qualitative features observed in Mansfield's data; this depends upon the growth rates that would have been achieved by exiting firms. Now suppose that small firms with low growth rates are more likely to exit. Then the proportional rate of growth, conditional on survival, will be smaller for large firms. Whether an appeal to this 'sample censoring' effect could rescue some underlying version of Gibrat's law was one of the main questions posed in the literature of the 1980s.

III. THE NEW LITERATURE

The new literature which developed in the 1980s had two main themes. The first lay in a concern with econometric issues, and here a major focus lay in dealing with the problems of sample censoring, the specification of an appropriate functional relationship, and the problem of heteroscedasticity[11]. A central question was whether a 'failure' of (some version of) Gibrat's Law could be attributed to any of these effects. The second major theme in the literature lay in a dissatisfaction with the models of the 1950s and '60s. It seems to have been widely felt that these models might fit well, but were 'only stochastic'. The

[11]The particular problems of sample censoring is important in itself, apart from any implications which it has for the size-distribution literature. To analyse the effect of an economic policy on firms, some allowance must be made for the impact on those firms who 'disappear from the sample' when they exit the industry. Since these 'disappearing' firms may include a disproportionate number of those firms most adversely affected by the policy, it is particularly important to 'correct' for this effect.

aim was to move instead to a program of introducing stochastic elements into conventional maximizing models'[12].

The problems of measurement were addressed in three influential contributions, which led to a common view on some basic statistical regularities. Bronwyn Hall's (1987) paper was based, like much of the earlier work, on a sample of large firms (1778 publicly traded manufacturing

[12]This goal is a reasonable one, but the underlying worry that the earlier models are 'not maximising models' is misplaced. It is easy to interpret the Simon model as a maximizing model by simply regarding each market opportunity as being 'independent'. It consists, say, of the opportunity to open up a new production plant on an isolated island, whose population is sufficient to support only one plant. A maximizing model will decree that one firm enters, and will be silent on the selection mechanism that chooses the lucky firm from among potential entrants. What is striking about the 'stochastic growth models' is not their lack of 'optimizing agents', but their reliance on Gibrat's Law. If we re-cast these models as simple maximizing models within a market consisting of independent opportunities, then there seems to be no reason to posit this particular relationship between firm size and growth.

firms). Her focus of interest lay in directly tackling several econometric issues, including those noted above, and in particular the problem posed by sample selection bias. The studies of David Evans (1987a,b) and of Timothy Dunne, Mark Roberts and Larry Samuelson (1988,1989) both covered the full range of firm sizes and ages. Evans' work was based on a large dataset for U.S. manufacturing industry, constructed by the U.S. Small Business Administration using information collected by Dun and Bradstreet for its credit reports. That of Dunne, Roberts and Samuelson related to plants rather than firms, and was based on a compilation of the individual plant-level data collected in five successive U.S. Censuses of Manufactures (1963-82). Apart from the econometric issues noted above, a major focus of interest in both this study and Evans' lay in unravelling the roles played by firm age and firm size as determinants of growth.

The focus of interest in all these studies lies in estimating

(a) the probability of survival of a firm, conditional on its age, size and other characteristics, and

(b) the probability distribution describing the firm's growth rate, conditional on survival, and on its age, size and other characteristics.

Before turning to results, it is worth pausing to note an important difference in method between the studies of Hall and of Evans and that of Dunne, Roberts and Samuelson. This will require some notation.

Following Dunne, Roberts and Samuelson, we distinguish three distributions: let the random variable g' denote a proportional growth rate, and let x denote a vector of characteristics describing a firm or plant (size, age, etc.). Let $j(g' \mid x)$ denote the probability density function for g' for a firm or plant with a given set of characteristics ("the distribution of potential growth rates").

The density j (g' | x) is not observable, for some firms will exit, and these will not be a random cross-section of firms. What can be measured directly is the density of growth rates conditional on survival, labelled h (g | x), where g denotes the actual proportional growth rate, and the density of measured growth rates, labelled f (g | x), in which all exiting firms are assigned the growth rate -1[13].

In the studies of Hall, and of Evans, the technique used is to apply a standard sample selection model (for example, Takeshi Amemiya (1984)), in which a growth equation and a (profit) survival equation are estimated jointly using maximum likelihood. This allows the conditional mean of the distribution of potential growth rates j (g' | x) to be estimated. The effect of x (size and age) on the means of

[13]The relation between h and f is particularly simple in one model considered by Dunne, Roberts and Samuelson. IN this model, all firms with a growth rate below some critical value g* exit. In that case, f is obtained by truncating the tail of the distribution h below g*, and replacing it with a mass point at -1.

f (g | x) and h (g | x) can then be examined. In both studies, it was found that the tendency for proportional growth rates to decrease with firm size survives these corrections for sample selection effects.

The approach taken by Dunne, Roberts and Samuelson is different. Taking advantage of the huge size of the Census data set, the authors grouped plants into cells corresponding to successive intervals of size, and of age. Consistent estimates of the parameters of the distribution of growth rates for all plants f (g | x), and the distribution for nonfailing plants h (g | x) could be obtained, subject to the assumption that the plants within each cell are homogeneous up to a random disturbance with zero mean and constant variance (which might be 'cell-dependent'). This technique does not allow identification of the parameters of j (g' | x); on the other hand, it avoids the need for any distributional assumptions, or for assumptions on the functional form of the

growth/size/age relation. It also avoids the difficulties faced in other studies of separating out sample selection effects from heteroscedasticity and from any nonlinear effects of explanatory variables.

The studies of Evans and of Dunne, Roberts and Samuelson both permit an investigation of age as well as size effects. Both studies suggest two statistical regularities:

1. (Size and Growth):

   (a) the probability of survival increases with firm (or plant) size.

   (b) the proportional rate of growth of a firm (or plant) conditional on survival is decreasing in size.

2. (The Life Cycle):

   For any given size of firm (or plant), the proportional rate of growth is smaller according as the firm (or plant) is older, but its probability of survival is greater.

What these results indicate is that there are two effects at work in the size-growth relationship: larger firms (plants) have lower growth rates, but are more likely to survive. The combined effect of both these tendencies can be summarised by taking a cohort of firms in some size interval, and comparing the total output of these firms at the beginning of the sample period with the total output of the surviving firms at the end. This procedure is used to define the 'net' growth rate for firms of a given size class. The data presented by Dunne, Roberts and Samuelson allowed a distinction to be made between plants owned by single-plant firms and those owned by multi-plant firms. For single plant firms, the reduction in

plant failure rates with size and age was inadequate to offset the tendency for larger and older surviving plants to grow more slowly: the net growth rate *falls* with plant size and age. The opposite was true of plants owned by multiplant firms. For this group, the net growth rate of plants tended to *increase* with size and age: the fall in plant failure rate with size and age outweighed the tendency for growth rates to fall with size and age among surviving plants.

These findings prompted new interest in theoretical models of firm growth. An obvious candidate model was the recently published 'learning' model of Jovanovic (1982). In the Jovanovic model, industry demand grows over time, and a sequence of firms enters the market. Each firm has some level of 'efficiency' (its unit cost of production), but it does not know what its relative efficiency is prior to entering. Over time, the profits it achieves provide information on its relative efficiency. More efficient firms

grow and survive. Less efficient firms 'learn' of the relative inefficiency, and (some) choose to exit.

This model provides a qualitative description of a process of excess entry followed by some exit, and this was the aspect of the model which made it attractive as a vehicle for discussing the new empirical results. As to the size distribution of firms, the model said little: it would depend inter alia on unobservables, such as the initial distribution of 'efficiency levels'. This theme would be echoed throughout the next series of theoretical models.

This new interest in age-growth relations led in turn to a renewed interest in modelling the life-cycle of the industry itself, and the evolution of market structure over time[14].

[14]The life cycle of the industry had been a focus of interest since the 1960s, see for example, Raymond Vernon (1966). As the focus of interest in the literature has shifted back to industry 'life cycle' effects, less attention has been paid to firm age/survival effects. These effects, however, receive considerable attention in the literature on the

An important impetus to the discussion came with the publication of the first of a series of papers by Steven Klepper and his co-authors, based on data assembled from entries in trade journals, regarding the evolution of firm numbers over time for a wide range of narrowly defined product markets (Klepper and Elizabeth Graddy (1990), Klepper and Kenneth Simons (1993)). The feature of the data which these authors emphasise constitutes a third statistical regularity which has been influential in shaping the recent literature:

---

sociology of organisations. Carroll (1983) for example, lists thirty-two studies on groups of organisations of various kinds. Many of these studies investigate the firm-specific factors associated with individual failures. A recurring theme is that many failures reflect a bad initial judgement of market opportunities, managerial incompetence, or simply the fact that the entrant set up a business which had only modest prospects of survival, as an alternative to entering the labour market, where opportunities were poor.

3. (Shakeout):

It is frequently observed that the number of producers tends first rises to a peak, and later falls to some lower level.

The extent and timing of this 'shakeout' varies very widely across product markets. In some cases, it comes early in the life of the product, and is very sharp. An example of such a case is shown in Figure 2, which is taken from Jovanovic and Glenn MacDonald (1994).

Two types of model have been postulated in response to this observation. The first is due to Jovanovic and MacDonald, who begin by stating that Klepper's data on shakeout can not be accounted for by appealing to the 'Learning' model of Jovanovic (1982). Instead, the authors postulate a model in which early entrants employ a common technology which after some time is superseded by a new technology. The new technology offers low unit

costs, but at a higher level of output per firm (scale economies). The transition to the new technology involves a shakeout of first generation firms, and the survival of a smaller number of firms who now employ the new large-scale technology. By calibrating the model against the data for the U.S. tire industry, the authors can simulate successfully the number of firms, and the movement of stock prices over time.

Another candidate model is developed in Klepper (1993). This model combines a stochastic growth process for firms, who enter by developing some new variant ('product innovation'), with the idea that each firm may spend some fixed costs to lower its unit cost of production ('process innovation'). Assuming some inertia in sales, and some imperfection in capital markets, those firms whose current sales are larger find it profitable to devote more fixed costs to process innovation (since the fixed costs incurred are spread over a larger volume of sales). As the larger firms

cut their unit production costs, some smaller firms are no longer viable, and these exit, generating the 'shakeout'. This process is very similar to the 'escalation' process that has been studied in the recent game-theoretic literature; see Sutton (1991), Chapters 3 and 8.

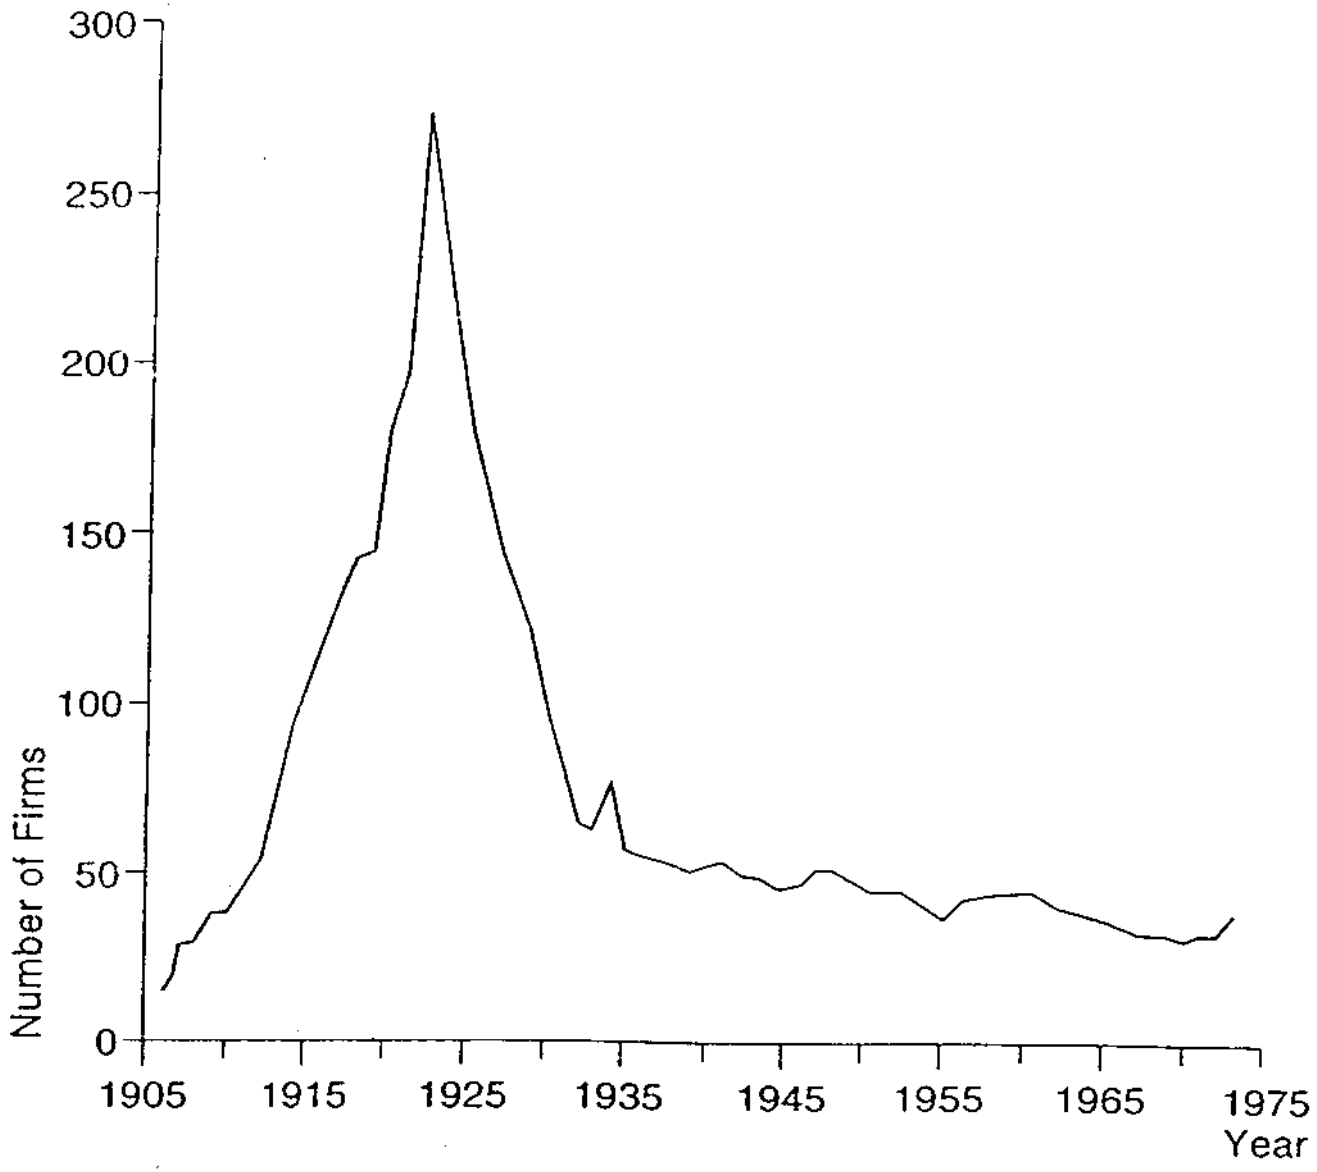

Figure 2.

Firm numbers in the U.S. tire industry. Source: The figure is a modified version of Figure 1 of Jovanovic and MacDonald (1994), which is based on Klepper (1993).

IV. THE SIZE DISTRIBUTION RECONSIDERED: A BOUNDS APPROACH

The new generation of models described in the preceding section, and in Section V below, differ from the older stochastic growth models in that the 'random growth' process has been replaced by one in which firms that differ in various attributes make different profit maximizing choices. The models remain stochastic but the source of randomness has either been pushed backward, into a description of firms' 'intrinsic efficiency differences', or forward into random outcomes emanating from R&D programs.

These recent developments, however desirable in themselves, have shifted attention away from one feature of the older stochastic models which may be crucial to a full understanding of market structure. This feature relates to the notion that the market consists of a sequence of

isolated 'opportunities', each strategically independent of the others. This appears to many readers, schooled in the importance of strategic interactions, as a rather odd and unattractive feature of these models. Yet this feature captures, albeit in a crude and extreme way, one important aspect of real markets.

Any industry will contain clusters of products or plants that compete closely. But an industry, as conventionally defined in official statistics, will usually contain more than one such cluster, in the sense that it will be possible to identify pairs, or sets, of products that do not compete directly. In other words, most conventionally defined industries exhibit both some strategic interdependence within submarkets, and some degree of independence across submarkets.

While a great deal of attention has been devoted in recent years to the analysis of strategic interactions, the role of

independence effects has received little attention. The question posed here is this: is it possible that some minimal degree of skewness in the size distribution could arise simply as a result of 'independence effects' per se? In order to address this question, it is natural to begin by thinking about the extreme case in which there are no demand side (strategic) interactions at all. The obvious way to do this abstract from strategic interactions is to think of a 'market' which consists of many independent submarkets, each of which is large enough to accommodate exactly one entrant[15]. This brings us to the setting of the Simon model.

It is worth noting that a conventional game-theoretic analysis of this situation is uninformative, in that it merely tells us that there are many pure-strategy Nash equilibria.

[15]In other words, the firm's profit function is additively separable into a number of contributions, one for each submarket.

These will include one in which the same firm enters each submarket, another in which a different firm enters each submarket, and so on.

But if we introduce the notion that each of a number of active firms has some probability of being 'selected' in each submarket, then it may turn out that certain kinds of outcomes, though possible, are 'unlikely'. Instead of imposing Gibrat's Law (which says that the probability of capturing the next opportunity increases proportionately with the firm's current size), suppose instead that the relation between firm size and the probability of capturing the next opportunity can take any form that satisfies the following constraint:

Condition 1: The probability that the next market opportunity is filled by any currently active firm is nondecreasing in the size of that firm.

Consider two businesses of different sizes. Condition 1 is violated if the smaller business is more likely to take up the next market opportunity than is the larger one. This might happen, for example, if the incremental profit realised from the new investment was smaller for the larger firm. This supposed disadvantage to the larger firm could derive either from the cost side or through 'strategic effects' on the demand side. Here, we abstract from demand side considerations, and focus on the cost side. A larger business may enjoy an advantage through economies of scope in offering several products, or in operating many plants. On the other hand, a traditional argument suggests that it will not suffer any cost disadvantage; for, if an integrated business of larger size had higher unit costs, then it should be possible to split the business into completely independent and separately managed units under single ownership, so that any such disadvantage is eliminated. This is the standard 'replication' argument for non-diminishing returns, and it

is a very appealing one. It suggests that Condition 1 might be a reasonable hypothesis for the artificial setting where firms face a sequence of 'independent opportunities'. This in itself is of little empirical interest, however, since the general run of industries will be characterised by a combination of strategic interactions within certain submarkets, and some degree of independence across submarkets. To see why Condition 1 might be of interest, we need to inquire as to how the presence of strategic interactions among firms might lead to a violation of this condition. The game-theoretic literature has afforded us a rich menu of examples in which the larger firm suffers a disadvantage in the sense that the profit per product (or plant, or unit capacity) is *decreasing* in the number of products (or plants, or units of capacity) operated by the firm. This effect has a simple intuitive interpretation: if the multi-product or multi-plant firm expands output or cuts price in order to improve the profit of one of its plants, it generates a negative externality for the other plants. In

maximizing its total profit, the firm seeks to 'internalise' this externality. This leads to higher prices and lower profits on each product or plant.

Nonetheless, empirical evidence on size-profitability relationships across businesses of different sizes within an industry suggests that the rate of return (profit) is nondecreasing in the size of the business. This suggests that firms may have some way of circumventing such strategic disadvantages where they arise. This will be the case whenever market opportunities are dispersed either geographically or in some space of 'product attributes'. If a firm that owned a number of closely clustered plants were to earn lower profit per plant, then that firm could simply expand by opening a sequence of plants in dispersed locations, thereby avoiding the strategic disadvantage. In this case, Condition 1 may still hold as an approximation, when the number of submarkets is large.

In Sutton (1995) the Simon model describedi n Section II is modified by replaicng Gibrat's Law (assumption (i)) by Condition 1, while retaining Simon's ('benchmark') assumption (ii) which posits a constant rate of entry. It is claimed that this modified model provides a good empirical description of the *least skew* distribution that we are likely to find at the 4- or 5-digit SIC level. This claim is advanced in a number of steps, as follows:

(i) In the modified Simon model, the size distribution tends to a limiting distribution that is *not less skew* than a certain reference distribution, which serves as a 'benchmark case'. (In what follows, we describe how Condition 1 leads to this result.)

(ii) This result continues to hold good in a more general model, in which there are many identical independent submarkets. No restriction is

placed on the nature of strategic interactions, or outcomes, *within* each submarket. It is assumed, however, that there are no strategic links across submarkets. (It is allowed that economies of scope may or may not exist across submarkets. If such economies are present, large firms enjoy an advantage, and this is consistent with Condition 1).

(iii) It is claimed that this more general model can reasonably be applied to the broad run of industries at the 4- and 5-digit SIC level[16]. At this level of aggregation, markets tend to contain large numbers of more or less independent submarkets.

---

[16]Strictly, to those industries in which neither advertising nor R&D are important. If advertising or R&D are important, strategic interdependence across submarkets is certainly present. However, this strategic effect runs in favour of *large* firms, as opposed to small, so it does not lead to any violation of Condition 1.

(iv) It is noted that this relationship may break down if we define markets so narrowly that all products and plants interact strategically. This allows the present interpretation to be tested by looking at a large geographical market for a physically homogeneous product, that contains many local submarkets.

The stochastic process describing firm growth is as follows: we begin at time 1 with a single firm of size 1. In each subsequent period, a new unit opportunity is taken up. With probability p, it is taken by a new entrant. With probability 1-p, it is taken up by one of the N_t firms that are already active. It can be shown that the least skew limiting distribution consistent with Condition 1 is attained in the special case where each of these N_t active firms has an *equal* probability $(1-p)/N_t$ of capturing the opportunity.

Let N_t denote the total number of firms in the market at

time t. The distribution of N_t is simply a binomial distribution with mean 1+p(t-1).

Let $n_{i,t}$ denote the number of firms of size i at time t, for i = 1, 2, 3 ... We begin by calculating the expected value of n_t, conditional on N_t. Our focus of interest lies in the behaviour of $E(n_{i,t}|N_t)$ evaluated at N_t= 1 + p (t-1), in the limit where t → ∞. It can be shown that:

$$\lim_{t\to\infty} E\left(\frac{n_{i,t}}{1+p(t-1)} \right) = p(1-p)^{i-1}$$

so that the size distribution takes the form of a geometric distribution. For empirical purposes, it is useful to approximate this by the corresponding exponential distribution. Subject to this approximation, it can be shown that the fraction of opportunities accounted for by

the k largest firms satisfies

$$C_k \geq \frac{k}{N}\left(\gamma_k - \ln\frac{k}{N}\right) \qquad (1)$$

Here C_k is the k-firm (asset) concentration ratio, N is the number of firms in the industry, and γ_k is a constant. For $k = 1$, γ_k is Euler's constant (≈ 0.577), and γ_k increases with k and converges to 1 as $k \to \infty$. (A table of values of γ_k is given in Sutton (1995b).)

We can illustrate this inequality by drawing a Lorenz curve, i.e. a relation between the fraction of (top ranking) firms k/N represented in the k-firm concentration ratio and their share of industry assets, C_k. The theory says that these points will lie above the reference curve obtained by expressing (1) as an equality.

Hence we obtain a quantitative prediction on a lower

bound to concentration[17]. This lower bound is independent of the entry rate (the proportion of opportunities captured by new entrants). The rate of entry, as noted earlier, is assumed to be constant over time, following Simon's second ('benchmark') assumption. It is shown that the result is fairly robust to empirically reasonable relaxations of this assumption. The result is also fairly robust to empirically reasonable relaxations of this assumption. It is unaffected by a number of empirically reasonable changes in the model, such as allowing opportunities to vary in size, and is invariant to different time patterns of industry growth. In particular, it is invariant to a process of shakeout which removes a random subset of plants or firms.

[17]The model does not provide any confidence interval around this bound. The model uses only an inequality constraint on the size-growth relation, and no statistical structure is placed on the distribution of the residuals $(C_k - \underline{C}_k)$.

The model does, however, rest on the assumption that the industry grows over time, possibly reaching some limiting size. It might in principle be violated in an industry which declined over time; the above assumption on size-growth relations says nothing about this.

This lower bound appears to fit reasonably well to data for various countries and time periods. An example is shown in Figure 3, which shows observed values of the 5-firm concentration ratio for all U.K. product markets in 1977[18]. Similar bounds for the U.S. and Germany are reported in Sutton (1995). The comparison between these two cases is interesting, in that the 'clouds' of points for the two countries are very different, reflecting the fact that concentration in the U.S. tends to be higher, for a given

[18]This is the lowest level of aggregation at which U.K. statistics are available, and corresponds to a level intermediate between the U.S. 4-digit and 5-digit levels. Statistics at this level have not been published for more recent years. Only sales concentration ratios, rather than asset concentration ratios, are reported.

number of firms. In spite of the wide divergence in average experience, however, the lower bound to both 'clouds' remains close to the reference curve given by the inequality (1).

This remark provides some motivation for the view that, in spite of the apparent importance of 'unobservables' in driving wide differences in 'average' outcomes, there may be a bound to outcomes which is relatively sharp. What makes such a bound interesting is the fact that much of the skewness we observe in 'typical' industries may be a reflection of the operation of independence effects per se, with greater skewness in some industries emanating in part from additional mechanisms such as the operation of scope economies across otherwise independent submarkets.

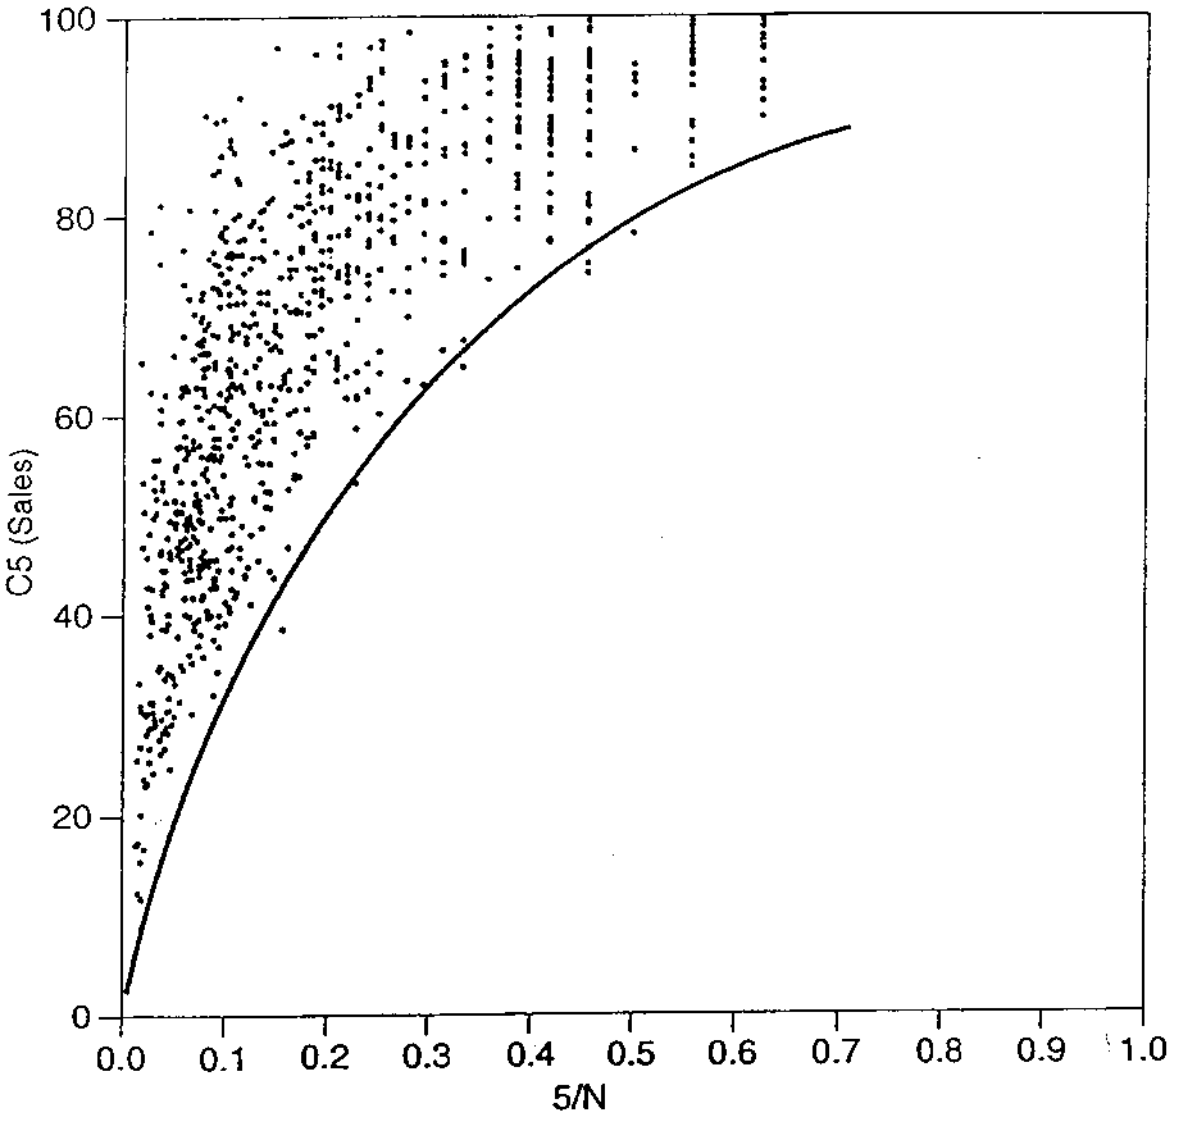

Figure 3.

Five-firm sales concentration ratios for U.K. product markets, 1977. The limiting Lorenz curve (1) for $\gamma_k = \gamma_5 = 0.90$ is shown as the solid curve lying above the diagonal. Each point corresponds to a different industry, the horizontal coordinate being 5/N where N is the number of firms, and the vertical coordinate shows the 5-firm sales concentration ratio.

If this is so, then the search for some 'average' or 'typical' shape to the size distribution may not be wise. This conclusion is very much in line with the growing pessimism among both theoretical and empirical researchers as to the reasonableness of postulating some 'typical' form for the size distribution.

V. TURBULENCE

The characterization offered by the time series models described above has been one of convergence towards some steady state of the industry, in which concentration levels and firm numbers will eventually become static. Empirical evidence on entry and exit patterns, however, emphasises that continuing entry and exit occur over the entire life of the industry. Much empirical work on entry and exit patterns has suggested a fourth statistical regularity that some authors see as particularly

noteworthy:

4. (Turbulence):

   Across different industries, there is a positive correlation between gross entry rates, and gross exit rates, i.e. the 'churning' of the population of firms is greater in some industries than others. However, most of this entry and exit has relatively little effect on the largest firms in the industry[19].

Within any one country, quite a strong correlation usually exists between entry and exit rates by industry. Paul Geroski (1991), for example, reports a correlation coefficient of 0.796 for a sample of 95 industries in the U.K.

[19]The volatility of market shares among large firms has been less widely studied. An important early study was that of Richard Caves and M. E. Porter (1977), which used the PIMS data-set. See also the recent study by Steven Davies et al. (1991).

in 1987. The most comprehensive data on this issue comes from a compilation of country studies edited by Geroski and Joachim Schwalbach (1991). The cross country comparisons afforded by this study indicate that there is at least a weak correspondence between the ranking of industries by turbulence in different countries. This is important, in that it suggests that there may be some systematic, industry-specific, determinants of turbulence levels[20]

---

[20]A major focus of interest in this literature has been on the way in which fluctuations in industry profits induce changes in the rate of entry and exit. Various authors have estimated net entry equations, following David Orr (1974), but the explanatory power of such regressions has been poor (see the review of results set out in Geroski's (1991) monograph, and the comments of Roger Clarke (1993)). A new attack on this problem has been emerging recently, following the work of Avinash Dixit and Robert Pindyck (1994) on investment under uncertainty. Here, the focus is on analysing the different thresholds associated with entry decisions, which involve sunk costs, and decisions to exit (see the review by Glenn Hubbard (1994) and the comments in Eugene Lambson (1991) below). Another strand in this literature looks at the different experience of different types of entrant, distinguishing between de novo entrants and entry into new markets by diversifying firms. The latter type of entrant tends to have a larger initial size,

These results have prompted interest in the determinants of turbulence (defined as the sum of gross entry and gross exit rates) across different industries. At least three types of influence are likely to be involved,

(a) underlying fluctuations in the pattern of demand across product varieties or plant locations;

(b) the displacement of existing technologies (modes of production) by alternatives; and

(c) the displacement of existing products by new and superior substitutes.

Of these, the first factor may be of primary importance, but while it is easy to model, it is very difficult to measure or

---

grow more rapidly, and have a lower rate of exit (see Geroski (1991), pp. 31ff. for an overview of studies on this question).

control for its influence empirically. The second and third factors pose more interesting questions in terms of modelling. Some new models have been developed recently, but these have not yet led to empirically tested claims regarding the influence of industry characteristics on the degree of turbulence.

The second factor (displacement of production technologies) has been modelled by Lambson (1991), who considers an industry facing exogenous shocks to relative factor prices, which occur at infrequent intervals. Firms incur sunk costs in building a plant using a given technology, and when factor prices change, an entrant - knowing that factor prices shift rarely - may find it profitable to enter the industry and displace incumbents. In this kind of model, the level of sunk costs incurred by firms will influence entry and exit rates, conditional on the volatility of industry demand.

The third factor listed above relates to the idea that (some) exit may be induced by entry, as new and superior product varieties displace existing products. This is the basic idea discussed in the vertical product differentiation literature. The key theoretical question is why the old varieties can't continue to retain a positive market share at some price, given that their costs of product development are sunk. Such varieties would indeed continue to survive in a 'horizontal' product differentiation model, but this is not necessarily true in a 'vertical' product differentiation model (Jean Jaskold Gabszewicz and Jacques Thisse (1980), Avner Shaked and Sutton (1983)).

This contribution to turbulence has been explored most fully by Neil Ericson and Ariel Pakes (1995). In the Ericson-Pakes model, equilibrium is characterized as the stationary state of a stochastic process, in which the fortunes of individual firms rise and fall over time. Each firm's current state is indexed by a number which can be

thought of as the relative quality of the product it offers. The vector of qualities maps into a vector of profits earned in the current period. Between one period and the next, each firm's quality will either stay the same, or rise by one unit. The firm chooses a level of R&D spending in each period, and the more it spends, the higher is the probability that its quality 'jumps'.

The model also assumes an upward drift in factor prices, or in the quality of some rival good, so that 'standing still' in terms of market share and profits requires that the firm achieve a steady rate of quality improvement over time. The authors characterise a (Markov perfect Nash) equilibrium in this setting in which each firm, taking as given the current distribution of qualities, and market parameters (factor prices), decide on an optimal level of R&D spending. The result is a steady state distribution of the vector of relative qualities. In this steady state, firms' optimal actions will reflect their relative position in the

quality spectrum. Firms with intermediate quality levels will spend on R&D in each period. Those with high levels may stop spending ("coasting"). Those with low levels may exit the industry. The distribution of relative qualities at any time is itself stochastic; and just as certain configurations may be reached at which some firms exit (their market share drops to zero), so too there will be configurations at which new entry occurs.

The strength of the Ericson-Pakes model lies in the fact that it offers an analysis of turbulence as a steady state phenomenon within a game-theoretic setting. The authors' aim of setting up as general a framework as possible carries the cost that few restrictions can be placed on the pattern of equilibrium outcomes. The authors have, however, developed a simulation package for general use, and this is already proving a useful research tool (Pakes and Paul McGuire (1994), Pakes, Gautam Gowrisankaran,

and McGuire (1993))[21].

These models developed by Pakes and his co-authors raise an interesting question as to the link between turbulence (entry and exit rates) and market share volatility among leading firms. In the stochastic growth models discussed earlier, the setting is one in which the size of the market grows, at least up to some limiting size, and the size distribution of firms converges to some limiting distribution. The entry and growth of new firms leads to market share volatility, in the sense that new firms might displace old ones as market leaders within that distribution. However, the probability that the market leader at any point is a relatively recent entrant is small.

[21]One obvious question relates to industry-specific factors influencing turbulence. The Pakes-Ericson model deals with the particular setting of an R&D-intensive industry where market share fluctuations are driven by the stochastic returns to R&D. It may be that broadening the scope of the model, by exploring different environments, may lead to testable predictions on the industry-specific determinants of turbulence.

These features seem to accord well with the empirical observations on entry and turbulence noted earlier.

In the Ericson-Pakes model, on the other hand, the focus in the numerical simulations reported to date has been on a small number of firms (fewer than ten), operating in a steady state environment where the degree of turbulence is magnified by way of stochastic returns to R&D outlays. Here, if numbers are small, the whole size distribution may swing widely over time. While there seems to be little evidence for wide swings at the 4 or 5-digit SIC level, it would be interesting to examine whether such swings are common in more narrowly defined markets where firm numbers are small.

VI. DECLINE AND EXIT

The literature discussed so far has been concerned with the growth of an industry to some steady state; in this section we turn to the final phase of industrial decline. Is the process of industrial decline associated with any systematic changes in market structure? As the industry declines, firm numbers tend to fall and reported concentration ratios show a weak tendency to rise. The latter tendency is in part, at least, a simple consequence of exit. A question that has attracted some interest lately relates to the way in which the size distribution of surviving firms evolves: is there, for example, a tendency for the largest firms in the industry to converge in size i.e. do the bigger firms shrink proportionally faster, so that the firm size distribution becomes 'less skew'?

Though this question has attracted some interest, there is little systematic evidence at the cross-industry level. One

way of checking this is to look at the ratio of, say, the 4-firm sales concentration ratio with the 8- or 20-firm concentration ratio. Figure 4 shows some scatter diagrams of the ratio of C_4 to C_8 at two widely separated dates for the set of 4-digit U.S. manufacturing industries that have experienced a large (> 40%) fall in firm numbers. The scatter shown in the figure does not suggest any tendency for the sizes of the top four firms and sizes of the next four firms to converge as the industry declines. (A similar pattern holds when we look at the ratio of C_4 to C_{20}.)

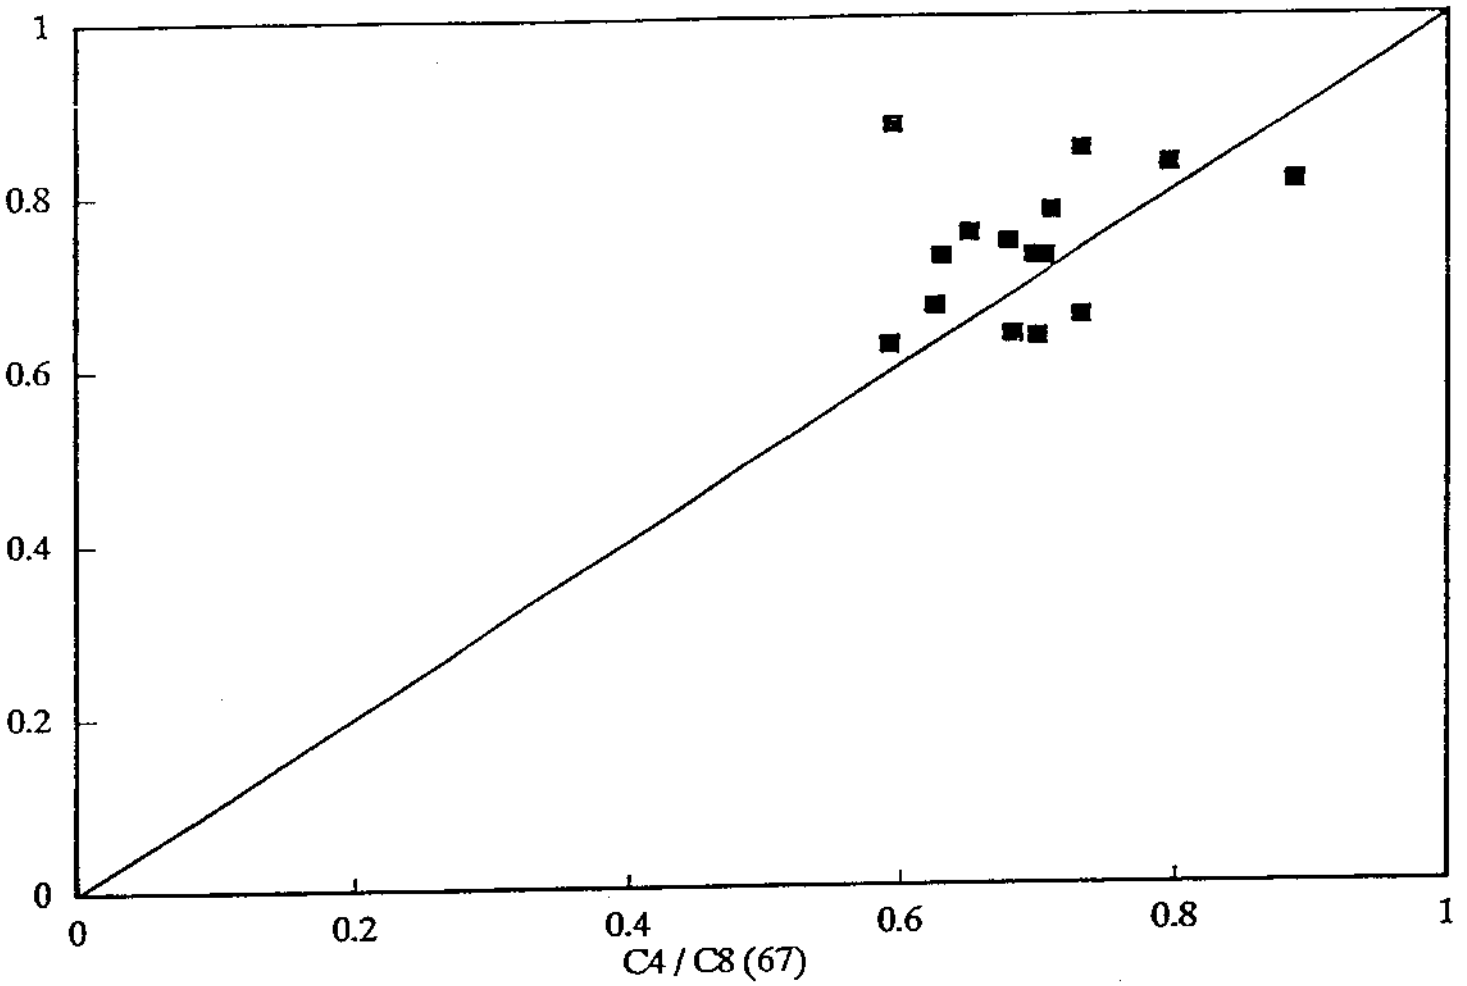

Figure 4.

The ratio of the four firm concentration to the 8-firm sales concentration ratios in 1987 (vertical axis) and 1967 (horizontal axis), for all 4-digit US Manufacturing industries that experienced a fall in firm numbers exceeding 40%. If the sizes of the top four firms and the sizes of the next four firms converge over time, then the scatter of points will lie below the diagonal.

At the theoretical level, there does not seem to be any general argument that suggests either a convergence or a divergence in the sizes of the largest firms. A rich variety of potential influences are available, and reasonable models can be constructed which lead to either outcome. Under these circumstances, it is natural to turn to a narrower domain, and to ask whether there are any types of industry for which some prediction is possible.

Pankaj Ghemawat and Barry Nalebuff (1990) consider a homogeneous goods industry with a particular kind of cost structure: unit costs are constant up to full capacity, and fixed costs are proportionate to plant capacity (as opposed to current output). Capacity can be reduced irreversibly in a continuous manner over time. The authors analyse a game in which firms reduce capacity as demand declines. They show that, along the equilibrium path of the game, the largest firm sheds capacity until it is equal in size to its nearest rival; then both these shrink together until they hit

the size of the next largest, and so on[22,23]. The authors present some striking evidence on the pattern of plant closures in the U.S. soda-ash industry, an industry for which these assumptions on product homogeneity and cost

---

[22]The intuition for this result is as follows: price in each period is determined by equating supply (total output Q) with demand; and the market demand schedule $P_t(Q)$ is falling over time. Each firm i produces to full capacity in each period, and its marginal revenue equals $Q_iP_t'(Q) + P_t(Q)$, where its output Q_i equals its current capacity. This expression is decreasing in Q_i. Therefore, it is the larger firm that has the greater incentive to withdraw a unit of capacity as market demand declines.

[23]John Londregan (1990) showed that the 'size effect' also held in a model where re-entry was possible, and Stanley Reynolds (1988) showed that it held good in the presence of cost differences, so long as these were not too large. On the other hand, if the model of Ghemawat and Nalebuff is modified by disallowing continuous shrinkage so that multiplant firms of varying size take turns in making discrete choices on the closure of (indivisible) plants, Michael Whinston (1988) shows that the tendency for large firms to close down plants first is less sharp, in that multiple equilibria now appear in the model, and it can happen along the equilibrium path that a smaller firm closes a plant ahead of a larger rival. Finally, Ingemar Dierickx, Carmen Matutes, and Damien Neven (1991) show that if decline is driven, not by a shrinkage in the consumer base, but by a diminution in consumers' willingness-to-pay, then the 'size' prediction may no longer hold.

conditions are reasonable. The observed pattern of plant closure conforms closely to the prediction that large firms close plants first. A number of later studies have looked at industries for which the assumptions of the Ghemawat-Nalebuff model are reasonable. Marvin Lieberman (1990) looks at a large sample of product markets in the chemicals sector. Controlling for plant size, Lieberman found that large multiplant firms tended to close individual plants ahead of their smaller rivals. For plants in the U.S. steel industry, however, Mary Deily (1991) found that the primary determinants of plant closure seem to have been a set of factors underlying plant profitability. Firms first disinvested in the least profitable plants, and then closed them. Controlling for plant and firm characteristics, the influence of size ran (weakly) in the opposite direction to that predicted by the theory. An earlier study by Deily (1985), however, looks in detail at the pattern of disinvestment and plant closure by the largest firms in this industry, and this suggests a less clear

pattern (Deily (1991), p. 262).

Moving beyond the special setting to which the Ghemawat-Nalebuff model applies, a number of recent studies have focused attention on the fact that the menu of relevant factors may run far beyond those normally included in the simple game-theoretic models. In looking at the U.K. steel casings industry, Charles Baden-Fuller (1989) found that many of the closing plants were not the least profitable ones. Rather, firms that were diversified and financially strong were more likely to close plants. The author's analysis of managers' views suggested that, in these firms, there were fewer internal conflicts between owners, debt holders and managers[24].

[24]It was pointed out by Jeremy Bulow and John Shoven (1978) that consideration of such decision making process 'within' the firm, (especially between equity and debt holders) might be crucial to understanding decisions on bankruptcy.

A similar theme emerges in the work of Martha Schary (1991) who appealed to a richer discussion of the determinants of exit, distinguishing between three routes by which firms 'disappeared' (bankruptcy, voluntary liquidation, or merger). Schary's paper examines the process of exit in the U.S. cotton textile industry, using a model in which firms make a series of decisions, considering each exit route in turn, in a pre-determined order. She finds that it is possible to reject a simple 'threshold profitability' model against this more complex schema. While it is probably fair to say that the specific results obtained may be sensitive to the maintained assumptions on the decision-making sequence, this study nonetheless suggests that a richer type of model may be called for in this area.

VII. SUMMARY AND CONCLUSIONS

The development of the literature over the past decade has been heavily influenced by a concern with a small number of statistical regularities (labelled 1-4 above). The aim has not, however, been to achieve some level of descriptive realism by finding a model consistent with this or any other set of statistical regularities. Rather, the role played by these regularities has been to stimulate interest in the possibility that there may be some systematic economic mechanisms at work. The emphasis in modelling has been on trying to capture features that are relevant to the working of some postulated mechanism, and to see whether the model can lead to further predictions or explanations.

One theme suggested above is that a proper understanding of the evolution of structure may require an analysis not only of such economic mechanisms, but also of the role

played by purely statistical (independence) effects, and that a complete theory will need to find an appropriate way of combining these two strands. This can be done in two ways. The current trend is towards building strategic interactions into models of the growth of firms. Along this route the same indeterminacy problem is encountered which is familiar throughout the game theory literature. An alternative way forward is to begin by tackling the indeterminacy problem by turning to a Bounds approach, as outlined in Section VI above, with a view to isolating the role played by independence effects per se. This approach may prove easier to integrate within a more general game-theoretic analysis which encompassess strategic interactions.

The evolution of market structure is a complex phenomenon and the quest for any single model that encompasses all the statistical regularities observed is probably not an appropriate goal. Yet there remain

phenomena which may well be worth encompassing in a more general theory than is currently available, and which are still poorly understood. Most notable among these are questions of the industry-specific determinants of firm turnover (turbulence) and the volatility of market shares. Another such area is that of the pattern of exit in declining industries. Nothwithstanding recent progress on these topics, many important questions still remain open.

REFERENCES

Acs, Zoltan S. and Audretsch, David B., Innovation and small firms, Cambridge, MA: MIT Press, 1990.

Amemiya, Takeshi, "Tobit Models: A Survey," Econometrics, Jan./Feb., 1984 24(1/2), pp. 3-61.

Baden-Fuller, Charles., "Exit from Declining Industries and the Case of Steel Castings," Econ J., Dec. 1989 99(398), pp. 949-61.

Bain, Joe S., International differences in industrial structure: Eight nations in the 1950s, Newhaven: Yale U. Press, 1966.

Boeri, Tito, and Cramer, Ulrich, "Employment growth, incumbents and entrants," Int. J. Ind. Org., Dec. 1992, 10(4), pp.545-565.

Bulow, Jeremy and Shoven, John, "The Bankruptcy Decision," Bell J. Econ., Autumn 1978, 9(2), pp.437-456.

Carroll, Glenn, "A Stochastic Model of Organisational Mortality: Review and Reanalysis," Soc. Sci. Res., 1983, 12,

pp. 303-329.

Caves, Richard and Porter, M.E., "From Entry Barriers to Mobility Barriers," Quart. J. Econ., May 1977, 91(2), pp. 241-67.

Clarke, Roger, "Geroski's Market Dynamics and Entry: A Review," Int. J. Econ. Bus., 1993, 1, pp. 455-459.

Cohen, Wesley and Levin, Richard, "Empirical Studies of Innovation and Market Structure," in Handbook of industrial organisation, (Richard Schmalensee and Robert Willig), 2, 1989, pp. 1059-1103.

Davis, Steven J. and Haltiwanger, John, "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," Quart. J. Econ., August 1992, 57(430), pp. 819-863.

Davies, Steven et al., The dynamics of market leadership in the U.K. manufacturing industry 1979-1986, Centre for Business Strategy Report Series, London Business School, 1991.

Deily, Mary, "Capacity Reduction in the Steel Industry,"

PhD Thesis, Harvard U, 1985.

Deily, M., " Exit Strategies and Plant-Closing Decisions: The Case of Steel," Rand J. Econ., Summer 1991, 22(2), pp. 250-263.

Dierickx, Ingemar, Matutes, Carmen, and Neven, Damien, "Cost Differences and Survival in Declining Industries," Europ. Econ. Rev., Dec. 1991, 35(8), pp. 1507-28.

Dixit, Avinash and Pindyck, Robert, Investment under uncertainty, Princeton, NJ: Princeton U. Press, 1994.

Dunne, Timothy, Roberts, Mark, and Samuelson, Larry, "Patterns of Firm Entry and Exit in U.S. Manufacturing Industries," Rand J. Econ., Winter 1988, 19(4), pp. 495-515.

Dunne, T., Roberts, M., and Samuelson, L., "The Growth and Failure of U.S. Manufacturing Plants," Quart. J. Econ., Nov 1989, 96(4), pp. 671-98.

Ericson, Neil and Pakes, Ariel, "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," Rev. Econ. Stud., Jan 1995, 62(1), pp. 53-82.

Evans, D., "The Relationship Between Firm Growth, Size and Age: Estimates for 100 Manufacturing Industries," J. Ind. Econ., June 1987, 35(4), pp. 567-81.

Evans, David, "Tests of Alternative Theories of Firm Growth," J. Polit. Econ., August 1987, 95(4), pp. 657-74.

Gabszewicz, Jean Jaskold and Thisse, Jacques, "Entry (and Exit) in a Differentiated Industry," J. Econ. Theory, April 1980, 22(2), pp. 327-38.

Geroski, Paul, Market dynamics and entry, Oxford: Basil Blackwell Ltd, 1991.

Geroski, P. and Schwalbach, Joachim, Entry and market contestability, Oxford: Blackwell Publishers, 1991.

Ghemawat, Pankaj and Nalebuff, Barry, 'The Devolution of Declining Industries, Quart. J. Econ., Feb 1990, 105(1), pp. 167-86.

Gibrat, Robert, Les inégalités économiques applications: aux inégalités des richesses, a la concentration des entreprises, aux populations des villes, aux statistiques des familles, etc., d'une loi nouvelle la loi de l'effet

proportionnel, Paris: Librairie du Recueil Sirey, 1931.

Hall, Bronwyn, "The Relationship Between Firm Size and Firm Growth in the US Manufacturing Sector," J. Ind. Econ., June 1987, 35(4), pp. 583-606.

Hannah, Leslie and Kay, John, Concentration in Modern Industry: Theory and Measurement and the U.K. Experience, London: Macmillan, 1977.

Hannah, L. and Kay, J., The Contribution of Mergers to Concentration Growth: A Reply to Professor Hart," J. Ind. Econ., March 1981, 29(3), pp. 305-13.

Hart, Peter, "The Size and Growth of Firms," Economica, Feb. 1962, 29(113), pp. 29-39.

Hjalmarsson, Lennart, "The Size Distribution of Establishments and Firms Derived from an Optimal Process of Capacity Expansion," Europ. Econ. Rev., 1974, 5, pp. 123-40.

Hjalmarsson, L., "Reply [to Vining]," Europ. Econ. Rev., April 1976, 7(3) pp. 287-92.

Hubbard, Glenn, "Investment Under Uncertainty: Keeping One's Options Open," J. Econ. Lit., Dec. 1994, 32(4), pp. 1816-31.

Hymer, Stephen and Pashigan, Peter, "Firm Size and Rate of Growth," J. Polit. Econ., Dec. 1962, 70(6), pp.556-69.

Ijiri, Yuji and Simon, Herbert, Skew distributions and the sizes of business firms, Amsterdam: North-Holland Publishing Co, 1977.

Jovanovic, Boyan, "Selection and the Evolution of Industry," Econometrica, May 1982, 50(3), pp. 649-70.

Jovanovic, B. and MacDonald, Glenn, "The Life Cycle of a Competitive Industry," J. Polit. Econ., April 1994, 102(2), pp. 322-47.

Kalecki, Michael, "On the Gibrat Distribution," Econometrica, April 1945, 13(2), pp. 161-70.

Kapteyn, Jacobus and Uven, M.J., Skew frequency curves in biology and statistics, Groningen: Hoitsema Brothers, 1916.

Klepper, Steven, and Graddy, Elizabeth, "The Evolution of New Industries and the Determinants of Market Structure," Rand J. Econ., Spring 1990, 21(1), pp. 27-44.

Klepper, S., "Entry, Exit, Growth and Innovation Over the Product Life Cycle," Working Paper, Carnegie Mellon U., Pittsburgh, 1993.

Klepper, S., and Simons, Kenneth, "Technological Change and Industry Shakeouts," Working Paper, Carnegie Mellon U., Pittsburgh, 1993.

Kumar, Manmohan, "Growth, Acquisition Activity and Firm Size: Evidence from the United Kingdom," J. Ind. Econ., March 1985, 33(3), pp.327-38.

Lambson, Eugene, "Industry Evolution with Sunk Costs and Uncertain Market Conditions," Int. J. Ind. Org., June 1991(2), 9, pp. 171-96.

Lieberman, Marvin, "Exit from Declining Industries: 'Shakeout' or 'Stakeout'?," Rand J. Econ., Winter 1990, 21(4), pp.538-554.

Londregan, John, "Entry and Exit Over the Industry Life Cycle," Rand J. Econ., Autumn 1990, 21(3), pp. 446-458.

Mansfield, Edwin, "Entry, Gibrat's Law, Innovation, and the Growth of Firms," Amer. Econ. Rev., Dec. 1962, 52(5), pp. 1023-51.

McCloughan, Patrick, "Simulation of Concentration Development from Modified Gibrat Growth-Entry-Exit Processes," J. Ind. Econ., Dec. 1995, 43(4), pp. 405-34.

Nelson, Richard and Winter, Sidney, An Evolutionary Theory of Economic Change, Cambridge, MA: Harvard U. Press, 1982.

Orr, David, "The Determinants of Entry: A Study of the Canadian Manufacturing Industries, Rev. Econ. Statist., Feb. 1974, 56(1), pp. 58-66.

Pakes, A., Gowrisankaran, Gautam and McGuire, P., "Implementing the Pakes-McGuire Algorithm for Computing Markov Perfect Equilibria in Gauss," Working Paper, Yale U., 1993.

Pakes, A., and McGuire, Paul, "Computing Markov Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model," Rand J. Econ., Winter 1994, 25(4), pp. 555-89.

Phlips, Louis, Effects of industrial concentration: a cross-section analysis for the common market, Amsterdam: North Holland, 1971.

Prais, Sigbert, The evolution of giant firms in britain, London: Cambridge U. Press,1976.

Pryor, Frederic, "An International Comparison of Concentration Ratios," Rev. Econ. Statist., May 1972, 54(2), pp.130-40.

Reynolds, Stanley, "Plant Closing and Exit Behaviour in Declining Industries," Economica, Nov. 1988 55(220), pp. 493-504.

Samuels, John, "Size and the Growth of Firms," Rev. Econ. Stud., Jan. 1965, 32(1), pp. 105-12.

Schary, Martha, "The Probability of Exit," Rand J. Econ., Autumn 1991, 22(3), pp. 339-53.

Scherer, Frederic, Industrial Market Structure and Economic Performance, 2nd ed., Boston: Houghton Mifflin, 1980.

Schmalensee, R., "Inter-Industry Studies of Structure and Performance," in in Handbook of industrial organisation, (Richard Schmalensee and Robert Willig), 2, 1989, pp. 951-1010.

Selten, Reinhard, "A Model of Oligopolistic Size Structure and Profitability," Europ. Econ. Rev., June 1983, 22(1), pp. 33-57.

Shaked, Avner, and Sutton, John, "Natural Oligopolies," Econometrica, Sept. 1983, 51(5), pp. 1469-84.

Singh, Ajit and Whittington, Geoffrey, Growth, profitability and valuation, Cambridge: Cambridge U. Press, 1968.

Singh, A. and Whittington, G., "The Size and Growth of Firms," Rev. Econ. Stud., Jan. 1975, 42(1), pp. 15-26.

Steindl, Joseph, Random processes and the growth of firms: a study of the pareto law, London: Griffin, 1965.

Steindl, J., "Size Distributions in Economics," in International Encyclopedia of Social Sciences, David L. Sills (ed.), London: Collier Macmillan, 1968.

Sutton, John, Sunk costs and market structure, Cambridge, MA: MIT Press, 1991.

Sutton, J., "The Size Distribution of Businesses, Part I," STICERD Discussion Paper No. EI/9, London School of Economics, 1995.

Sutton, John, "Game-theoretic Model of Market Structure," in D. Kreps and K. Wallis, Advances in economics and econometrics: theory and applications, (Proceedings of the Seventh World Congress of the Econometric Society), Cambridge: Cambridge U. Press, 1996.

Vernon, Rayond, "International Investment and International Trade in the Product Life Cycle," Quart. J. Econ., May 1966, 80(2), pp. 190-207.

Whinston, Michael, "Exit with Multiplant Firms," Rand J. Econ., Winter 1988, 19(4), pp. 568-88.