

Market Share Dynamics and the ‘Persistence of Leadership’ Debate

John Sutton*
London School of Economics**

Table of Contents

1. Introduction.....	1
2. The Main Idea	5
3. A Model	8
4. The Data	16
5. The Scaling Relationship	17
6. Passage Times I	21
7. Passage Times II	28
8. Digging Deeper	30
9. Conclusions	37
References	40
Appendix 1	41
Appendix 2	47

EI/37
May 2004

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

* The financial support of the Economic and Social Research Council is acknowledged. I would like to thank Ciara Whelan for her research assistance, and Yoshiro Tamai, Daisuke Tsuruta and Kuniyoshi Saito for helping with data collection.

** STICERD, London School of Economics, Houghton Street, London WC2A 2AE, UK, Tel: 020 7955 7716, Fax: 020 7955 6951, Email: j.sutton.@lse.ac.uk

© The authors. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Abstract

This paper introduces a novel analysis of the classic “persistence of leadership” question, and applies it to a newly constructed dataset for Japanese manufacturing. The analysis rests on an appeal to an empirical “scaling relationship” between current market share and the variance of changes in market share. This relationship provides a powerful “model selection criterion” for candidate models of market share dynamics. It also makes it feasible, even in small datasets, to test directly for the properties of the “first passage times” corresponding to loss of leadership.

Keywords: market share, industry dynamics, scaling, Japanese economy

JEL: L10, L60

1. Introduction

For how long does a typical ‘market leader’ in an industry maintain its position? This question has attracted continuing attention in the I.O. literature over the past generation. Two rival views have emerged. The first, associated *inter alia* with Alfred Chandler (1990), asserts that leadership tends to persist for a ‘long’ time. The rival view, sometimes labelled ‘Schumpeterian’, emphasises the transience of leadership positions; an explicit version of this view is spelt out in Franklin Fisher’s (1983), model of ‘leapfrogging competition’.

The central problem with this debate is that no benchmark is proposed relative to which the duration of leadership might be judged ‘long’ or ‘short’. Thus, if it is observed that the typical market leader stays in place for 20 years this can be interpreted as ‘long’ by writers in the first group, and as ‘short’ by those in the second. This point has not gone unnoticed by contributors to the literature; an unusually full and frank acknowledgement of the difficulty is set out by Mueller (1986),¹ who notes that his conclusion as to the degree of persistence rests on a subjective judgement.

This paper introduces a formal model of market share dynamics, and uses it to provide a benchmark case, corresponding to a ‘neutral’ situation in which neither positive (‘Chandleran’) effects or negative (‘Schumpeterian’) effects are present. This model provides a natural benchmark against which empirically observed patterns of persistence can be gauged.

¹ Mueller’s study relates to profit rates, while the present paper relates to market shares; but the present point applies equally to both measures.

What degree of persistence should we expect on the basis of theory? Game-theoretic models offer little guidance on this question. The issue turns on the following consideration: suppose the market share gap between the leader and its (nearest) rival narrows, then will this induce an increase or a decrease in effort by the leader relative to the rival? The factors that may influence outcomes here are numerous. One ('Chandlerian') view emphasises the role played by the 'dynamic capabilities' of firms. On this view, market leadership is a correlate ('signal') of superior capability, which is a slowly changing attribute. This suggests a story in which a short-run narrowing of the market share gap between leader and rival will tend to be followed by a reverse movement as the gap reverts to the level corresponding to the firm's relative capabilities. Another important factor relates to the details of the underlying technology, as represented by a stochastic mapping from R&D to product quality. If, for example, this mapping takes the (special) form used by Ericsson and Pakes (1995) for example, then the leading firm may find it optimal to cease investing in R&D ('coasting') even though this leads to a greater probability of being leapfrogged by its rival.

Given the rival perspectives on the issues, how can we define a useful benchmark case? One way forward is to begin with the question: if the gap between the leader and its (nearest) rival narrows, does this induce a tendency for a further narrowing, or a tendency for a widening? The benchmark case proposed here is that in which neither of these tendencies is present; instead, market share dynamics follow a simple random walk (or first-order Markovian process). Relative to this benchmark, we can consider two kinds of bias, one of which ('Chandlerian') leads to longer persistence of leadership, while the other ('Schumpeterian') leads to shorter persistence.²

² The discussions of these two schools of thought in the literature do not admit of any sharper definition of their respective positions.

An idea that forms an important motivation for this exercise lies in the classic observation of Feller (1950), to the effect that passage times in Markovian processes tend to be extremely long relative to what we might expect intuitively. Feller identifies this as the most surprising feature to emerge from the study of stochastic processes. In the light of this, it seems natural to inquire into the degree of persistence that we would get in a simple Markovian model; for much of the discussion the literature pre-supposes that a ‘long’ duration of leadership must imply that some ‘economically interesting’ mechanism is at work that accounts for this persistence. What Feller’s insight suggests, is that looking for such explanations may be inappropriate. Even if leader and laggard are equally lucky or equally capable, then we will still see leadership persist for what appears intuitively to be a ‘long’ time; and for reasons which are more a matter of arithmetic than economics.

The idea that some kind of Markovian model might offer a useful first approximation in modelling market dynamics is not new; indeed, within the different but related ‘growth of firms’ literature it has a substantial history, beginning from the seminal contribution of Little (1962) and Little and Rayner (1966).³ Yet such models are often thought of as being unsatisfactory, on the grounds that they do not treat changes in firms’ shares as an outcome of strategic interactions (maximizing behaviour) in marketing, R&D, etc. but rather as the outcome of ‘stochastic shocks’. Here, I defend the usefulness of such models on the following grounds: while traditional discussions between and among ‘Chandlerians’ and Schumpeterians’ tacitly assume that there is some single mechanism driving (high or low) levels of persistence, the central message of the game-theoretic

³ The ‘tests’ used in that literature have been based on examining correlations between growth rates over successive periods. What is novel in the present paper, relative to that kind of representation, is the direct examination of the statistics of ‘first passage times’ (see below).

literature in this area is that we should not expect any single mechanism to play a dominant and systematic role in driving market share dynamics. Many patterns of interaction may emerge between a leader and its rivals, and these patterns will reflect inter alia the beliefs of agents as to rivals' likely responses to their actions. The 'beliefs of agents' are among the several industry characteristics that may influence outcomes, but which are notoriously difficult to measure, proxy or control for in empirical studies (Harris (1994)). This point is developed in Section 8 below, where we examine the pattern of market share dynamics in selected industries. What emerges from these examples is:

- a) Very different patterns may arise across industries with apparently similar characteristics,
- b) Major shifts in the pattern of dynamics may occur within an industry over successive time periods.

What this suggests is that, while it might be possible to build a satisfactory 'structural' model of market share dynamics for a single industry, or even a group of cognate industries, it is helpful in looking across the general run of industries to begin by examining the data against the background of a more modest, low level representation of the kind proposed here.

2. The Main Idea

The main idea underlying the method of analysis proposed here lies in exploiting two key features of the empirical data, which permit a very simple representation of the stochastic process driving the pattern of market shares.

The analysis of market share dynamics poses, in general, two serious challenges. First, since market shares add to unity, shocks to different firms' shares are interdependent. Second, the (distribution of the) size of shocks to each firm's share might be expected to depend inter alia on that firm's current share. This implies that an appropriate model might be one in which the distribution of shocks to each firm's share would need to be conditioned on the full vector of market shares in the current period. The role of the two empirical features of the data on which the present method of analysis rests is to permit a much simpler representation of the underlying stochastic process.

The first feature of the data on which we rely is that, for all but four (highly concentrated) industries among the 45 industries in the dataset, the shocks to the market shares of the industry's leading firms display an extremely low degree of correlation, so that we may impose, as a reasonable approximation, a model of 'independent shocks'.

The second feature of the data on which we rely is that it exhibits a simple 'scaling relationship' between a firm's market share and the variance (or standard deviation) of its change in market share.

The nature of this scaling relationship is as follows: the variance σ^2 of the change Δm in a firm's market share m , measured in percentage points, increases in direct proportion to m ; equivalently, the standard deviation of the fractional change in m , i.e. $\Delta m/m$, falls proportionally with $1/\sqrt{m}$.

The method of analysis used in what follows takes advantage of this feature of the data. Essentially, it allows us to characterize the size distribution of annual shocks to market shares within each industry by reference to a pooled sample of all observations (avoiding the need to condition directly on current market share, a procedure which would not be practicable using the 'small' dataset involved here).

Taking these two features together, the most basic 'persistence of leadership' question, i.e. that of analysing the time elapsed until the market leader is overtaken by any specific rival, can be handled by reference to well-known properties of a (simple) random walk. By appealing to the standard properties of first passage times for such processes, we can achieve a considerable simplification in the analysis.⁴

Before turning to empirical matters, it may be helpful to begin by setting out an illustrative theoretical model. It is important to note, however, the empirical analysis which follows rests solely on a direct appeal to the two features of the data just

⁴ While the earlier literature has tested the null-hypothesis of 'neutral' or 'first order Markovian' property on which we focus below, it has done so by looking at (low-power) tests involving comparisons of Δx_t and Δx_{t+1} . A test of this standard kind for the present dataset indicates no significant correlation(s) of this kind, over any timescale. By focussing directly on the statistic of interest (the first passage time), we can arrive here at a more powerful and direct test of the hypothesis.

A more fundamental problem with the standard approach of examining changes in each firm's sales as an (independent) stochastic process, as is done in the 'growth of firms' literature, is that this approach is unsuited to examining the 'persistence of leadership' question: the counterhypothesis against which the null is tested, is that the sales of each firm form independent higher order Markov processes. However, the economically interesting counterhypothesis in the 'persistence of leadership' setting are ones in which changes in the sales, or shares, of the firms depend inter alia on the current difference in shares between the leader and its (nearest) rival(s), and this cannot be captured using the standard methods.

mentioned, and does not depend upon the particular model presented below. The reason for introducing the model is to provide an intuitive explanation for three points which might otherwise seem puzzling. These are:

- i. the idea that market share shocks may be ‘approximately independent’ in industries where concentration is low;
- ii. the scaling relationship It is natural to ask whether this relationship has any theoretical basis. An examination of the various standard product differentiation models indicates that the only type of model that appears to exhibit this feature is a multi-product firm model that combines a vertical product attribute of the standard kind with a horizontal attribute of the locational (Hotelling) type. In particular, this form of scaling relationship does not arise either in ‘single attribute’ quality models, whether of the ‘vertical product differentiation’ type (Sutton (1991, 1998)) or of the ‘stochastic quality jump’ type used by Ericson and Pakes (1995) in their model of market share dynamics. Intuitively what drives the present scaling property is the idea that a large firm receives shocks of the same absolute size, but that the expected number of such shocks occurring in a given time interval increases in direct proportion to the firm’s size. In order to provide a framework for the analysis that follows, we begin by introducing a (deliberately simple) model of this kind.
- iii. The model introduced here is a non-strategic one in which changes in market shares are driven by exogenous shocks to product quality. The motivation for introducing a non-strategic model in this context lies in the argument that the appropriate strategic model(s) would be highly industry specific, a point on which we elaborate in the final section below. This raises the question: what

of strategic influences that do not depend on highly specific industry characteristics, but operate robustly across the general run of industries? It is well known that certain systematic strategic effects operate to place a lower bound on the level of concentration that is sustainable as an industry equilibrium (for example, Sutton (1991, 1998)). How does this square with the notion that market shares may fluctuate over time, at least once some ‘lower bound’ to industry concentration is respected? The point is addressed below (footnote 9).

3. A Model

The model is a standard circular road model, in which firms offer products that are differentiated by ‘location’. The products are located evenly around the circumference of a circle of unit diameter. Each (active) firm owns a subset of these products. For simplicity, we confine attention to the case where no firm owns two adjacent products; this allows us to obtain a simple characterization of a Nash equilibrium in prices (it coincides with the price equilibrium for single product firms).

We associate with each product a quality index u . Consumers are located uniformly along the circle, the total size of the population of consumers being normalized to unity. Each consumer buys exactly one unit of one of the goods on offer, the supplier being chosen to maximize the consumer’s utility,

$$U(p, u) = u - p - td$$

where p is the price of the chosen good and t is the (constant) unit cost of transport along the circle.




Figure 1

In what follows, the range of u will be restricted so as to ensure that the ‘marginal consumers’ defining the left and right hand boundaries of product i ’s clientele will lie between product i and its immediate neighbours. We can then write down the conditions defining the distance from firm i to the marginal consumer on its right, which we label d_i , as follows:

$$p_i + u_i + td_i = p_{i+1} + u_{i+1} + t(1/N - d_i)$$

$$\text{whence } d_i = 1/2N + [(p_{i+1} - p_i) - (u_{i+1} - u_i)]/2t \quad (1)$$

The distance from firm i to the marginal consumer on its left, denoted by $1/N - d_{i-1}$ (see Figure 1), is calculated in the same way, viz.

$$1/N - d_{i-1} = 1/2N + [(p_{i-1} - p_i) - (u_{i-1} - u_i)]/2t \quad (2)$$

Adding (1) and (2) we obtain the quantity sold by firm i, viz

$$q_i = 1/N + [(p_{i+1} + p_{i-1} - 2p_i) - (u_{i+1} + u_{i-1} - 2u_i)]/2t \quad (3)$$

Setting cost to zero and writing the profit of firm i as $p_i q_i$, we differentiate with respect to p_i to obtain the optimal reply (reaction function) of firm i, viz.

$$p_i = t/2N + (p_{i+1} + p_{i-1})/4 - (u_{i+1} + u_{i-1} - 2u_i)/4 \quad (4)$$

Given our assumptions that no firm owns two adjacent goods, and that the range of the quality index is restricted so as to ensure that the marginal consumer always lies between the product and its closest neighbour, it follows that the optimal reply (reaction function) for each firm is to set the price of each of its products in accordance with equation (4) i.e. the firm's profit function is additively separable into a number of functions, corresponding to the profit earned from each product.

In the special case where all the u 's are zero, the set of equations defined by (4) collapse to those of the standard circular road model: there is a symmetric Nash equilibrium in prices, in which all firms set the same price $p = t/N$, as can be confirmed by inspection of (4).

Our focus of interest lies in examining the manner in which exogenous shocks to (relative) quality levels of individual products impinge on the sales of the firm.

It is shown in Appendix 1 that a unit shock to the quality of product i , given equilibrium price responses by all firms, leads to a rise in the quantity (sales volume) of product i of $1/\sqrt{3}$ units, and a fall in the sales volume of each other product. These impacts on the sales of other products decline geometrically as we move away from product i ; for the k -th product to the right or left of product i the change in sales volume is $-(2 - \sqrt{3})^k / \sqrt{3}$. Given our normalization of the total size of the population of consumers to unity, the (change in) quantity sold by a firm equals its (change in) volume market share.⁵

We do not restrict the pattern of shocks to qualities in what follows. In each (short) period, there is a small probability p that a single shock to the quality of some one (randomly chosen) product occurs, the size of the shock being drawn from some distribution $f(\Delta u)$. This leads to a geometrically declining series of shocks to the product in question, and its neighbours. We confine attention throughout the case where the number of products is large; and given the geometrically declining size of impact, we approximate by neglecting all shocks beyond a certain radius viz. the ℓ th product on the left to the ℓ th product on the right.

So far we have ignored the possibility that a shock to the quality of product k might bring it outside the range $[0, \bar{u}]$. We treat this by setting $u_{t+1} = \bar{u}$ if $u_t + \Delta u \geq \bar{u}$, and $u_{t+1} = 0$ if $u_t + \Delta u \leq 0$.

⁵ We work for convenience in terms of volume market shares. The results for market shares by value are similar, subject to an approximation.

When a product quality falls to 0, we treat this as an ‘exit’ event.⁶ We assume that such an event is followed by the entry of a new product by some firm, at initial quality 0. The probability that the new product is entered by firm j is set equal to the proportion of products currently owned by firm j .⁷

We will not be directly concerned in what follows with the long run steady state properties of the model;⁸ here, it suffices to remark that a firm’s expected market share, conditional on its having n_j out of N products, equals n_j / N .

We consider the impact on the pattern of market shares of a quality shock that affects a single randomly chosen product. In what follows, we focus on the largest and second largest firm in the industry; their respective numbers of products are denoted as $n_1 / N = p_1$ and $n_2 / N = p_2$, and we denote by $p_3 = 1 - p_1 - p_2$ the combined share of products owned by all other firms.

⁶ This representation of exit events is chosen purely for convenience; a more sophisticated model would involve a consideration of the sunk cost incurred in entering a product, and would involve the determination of an optimal threshold u^* at which a product would be deleted.

⁷ Again, this feature of the model is chosen in order to bring the model into line with the empirically observed ‘scaling’ property.

⁸ In a model of this type, there will, once entry and exit are modelled as optimizing decisions, be a lower bound to the level of concentration (specifically, to the market share of the largest firm; see Sutton (1991, 1998)). This comes about as follows: suppose we allow firms to choose the quality of their products optimally, subject to some fixed cost schedule. Then, if the number of firms becomes sufficiently large, so that the maximum market share falls below some critical level, it be optimal for one firm to deviate, either by raising the quality of (at least one) of its products, so as to capture a greater market share.

The idea behind the present model is that the number of firms that are active in the market has been arrived at by some earlier (unmodelled) process of entry, and it is not so large as to violate the ‘lower bound to concentration’. The focus of interest here lies in asking, how do market shares fluctuate within the region permitted by these bounds?

It is intuitively clear that, in examining the behaviour of the market share gap $m_1 - m_2$ (or the gap $p_1 - p_2$ which coincides with the expected value of $m_1 - m_2$), that there are two polar cases of interest, viz. where p_3 is large, so that firms 1 and 2 are ‘small’ and where p_3 is close to zero, so that $p_1 \cong 1 - p_2$. In the latter case, there is close (negative) correlation between changes in the market shares of firm 1 and firm 2. In the former case, this correlation is close to zero and we can approximate shocks to $m_1 - m_2$ by treating m_1 and m_2 as independent. An empirical examination of the present dataset indicates that the correlation between Δm_1 and Δm_2 is very close to zero (see Section 5 below). With this in mind, we focus on the case where p_3 is large, where we may analyse the impact of a single unit shock to the quality of some randomly chosen good by representing the probability that firm 1 (or firm 2) receives the associated quantity shock of order k as p_1 (or p_2 respectively), and ignore all multiple events. In this case, the expected change in m_i can be approximated as:

$$\sum_k s_k p_i = p_i \sum_k s_k$$

where s_k is the change in quantity (volume market share) for a product deriving from a unit shock to the quality of a product at the k -th location to the right, or left associated with a shock of order k , and p_i is the share of products owned by firm i , and where the sum is taken over $k = -\ell, \dots, -1, 0, 1, \dots, \ell$.

Now consider any (discrete) distribution of quality shocks: let f_j denote the probability that a shock of size Δ_j occurs. Then, recalling that the derived quantity changes are

directly proportional to the size of the quality shocks, the variance of changes to m_i can be represented as

$$\begin{aligned}\text{var}(\Delta m_i) &= \sum_j \sum_k p_i f_j(\Delta_j s_k)^2 \\ &= p_i \sum_j \sum_k f_j(\Delta_j s_k)^2\end{aligned}$$

Noting that the double sum in this last expression is a constant, the variance of Δm_i is proportional to p_i , which we can proxy empirically by m_i .

It follows that the standard deviation of changes to market shares satisfies

$$\sigma(\Delta m_i) \cong \text{constant} \cdot \sqrt{m_i}$$

It follows that, if we replace the market share m_i by $\sqrt{m_i}$, then for small changes we may write

$$\Delta \sqrt{m_i} \cong \frac{1}{2\sqrt{m_i}} \Delta m_i$$

whence

$$\sigma(\Delta \sqrt{m_i}) \cong \text{constant}$$

so that we have a measure of volatility that is constant over m_i . We can now construct an industry-specific measure of the degree of volatility by pooling all observations of $\Delta\sqrt{m_i}$ for all firms over some period; whence we define the volatility measure

$$\text{vol} = \sigma(\Delta\sqrt{m_i})$$

In the case under consideration, where Δm_1 and Δm_2 are treated as independent, we may now proceed as follows. Denote by $g(\Delta\sqrt{m_i})$ the (symmetric) p.d.f. of (small) changes to $\sqrt{m_i}$. Define the ‘gap’ between firm 1 and firm 2 as

$$g \equiv \Delta\sqrt{m_1} - \Delta\sqrt{m_2}$$

Given the independence of changes to m_1 and m_2 , we may model the evolution of g as a simple random walk whose increments are drawn from the distribution $g \circ g$. If the distribution of shocks to $\sqrt{m_1}$ is, for example, normal with standard deviation σ , then changes to $\sqrt{m_1} - \sqrt{m_2}$ are normal with standard deviation $\sqrt{\sigma^2 + \sigma^2} = \sqrt{2}\sigma$. We can therefore normalize by defining the gap

$$g = \frac{\sqrt{m_1} - \sqrt{m_2}}{\sqrt{2} \cdot \sigma(\sqrt{m_1})}$$

whose evolution can be modelled as a simple random walk, whose movements are drawn from the standard normal distribution $N(0,1)$. In the next section, we follow this procedure with one modification; the distribution of shocks is better represented as a t -distribution, and we modify the procedure slightly to reflect this.

4. The Data

The dataset consists of annual observations of market shares for leading firms in 45 narrowly defined industries in Japanese manufacturing over the 25 year period 1974-1999. (Appendix 2). These data were compiled using the annual volumes published by Yano Company. This source covers a large number of industries, but occasional changes in coverage and presentation occur, and it was possible to construct fairly long and consistent series only for these 45 industries. Specifically, it was possible to compile a history of 23 years or longer for each of the 45 industries. The starting data for this history is 1974 for the large majority of industries, but it is between 1975 and 1977 in a small number of cases. The main tests described in the next section are carried out by reference to the 22-year history of these 45 industries. A series of interviews with selected companies was used to check issues of interpretation and reliability of data.

Data of this kind would be very difficult to compile for a broad cross-section of industries in other countries; the availability of the Yano data was a primary reason for focussing on Japan. The second, equally important, reason for this focus lies in the rarity of mergers and acquisitions. For U.S. or U.K. data, for example, it would be difficult to study the distribution of first passage times over an extended time period without having to confront the confounding influence of M&A events. In the present data-set, only one merger involving 'leading' firms occurs over the 25 year period in these 45 industries.

The level of aggregation in this dataset corresponds roughly to the 5-digit SIC classification for the U.S. The industries include, for example, margarine, photographic film, beer and cash registers.

The number of firms included varies across industries, the typical case being half a dozen or so. Excluded firms generally have very small shares. Their exclusion does not affect the computation of first-passage times, since if one of these firms grows to become a leading supplier, it is incorporated in the data-set. There are no instances in which such a ‘newly entered’ firm overtakes the market leader during the period covered by the data.

5. The Scaling Relationship

We begin with a descriptive account of some basic features of the data.

- a. We begin by taking the top two firms in some reference year (year 5), and we examine, which we label hereafter ‘firm 1’ and ‘firm 2’ respectively. We examine the annual change in market share for firm 1, versus the change for firm 2, in each year. The resulting scatter for the pooled sample of all industries is shown in panel (a) of figure 2. The correlation coefficient is 0.01, indicating that the data is well represented by the first limiting case described in the preceding section. To explore this further, the exercise was repeated by excluding successive groups of industries, using as a criterion the combined market share of the top two firms in the reference year. Only when the critical value of this combined market share was set to exclude all but 4 industries did a clear negative correlation appear. With this in mind, we proceed to explore the full dataset using the model based on the first limiting case, as developed in the preceding section. (Excluding these 4

industries from the analysis which follows has no material effect on our conclusions).⁹

- b. To investigate the relationship between current market share, and the distribution of changes in market share, a pooled sample of all annual observations was formed, and partitioned into groups (bands) by market share, i.e. all pairs $(m_t, \Delta m_t)$ for which $m_i \leq m_t \leq m_{i+1}$ fall in group i , and so on. For each band, the standard deviation of Δm_t was estimated. Finally a regression of $\ln \sigma(\Delta m_t / m_t)$ against $\ln m_t$ illustrated in Figure 3, yields a slope of -0.53 , which is not significantly different to $-1/2$, and which suggests that the data is well represented by the first limiting case (as opposed to the second limiting case) of the model set out earlier.
- c. To investigate the distribution of the size of shocks to market share (which is not restricted within the above model), we may take advantage of the scaling relationship to examine the distribution of $\Delta \sqrt{m_t} / \sqrt{m_t}$, which should be independent of m_t . This indicates that the distribution is represented by a t -distribution with a coefficient of about 1.3, i.e. of the form¹⁰ $f(x) = \frac{a}{(1+x^2)^{1.3}}$. This is illustrated in Figure 4. This description does not however extend to the tails of the distribution; there are no observations outside a range of about 3 standard deviations from the origin, a point to which we return below.

⁹ It might seem surprising *prima facie* that the lack of correlation holds even in moderately concentrated industries. This may reflect the fact that, in some industries, the two leading firms do not compete ‘head-to-head’, so that their gains (or losses) of market share impinge more on lower ranked firms, than on each other.

¹⁰ By repeating this exercise for subsets of the data corresponding to different bands of m_t , it is confirmed that the form of this distribution does not vary noticeably with m_t , as expected.

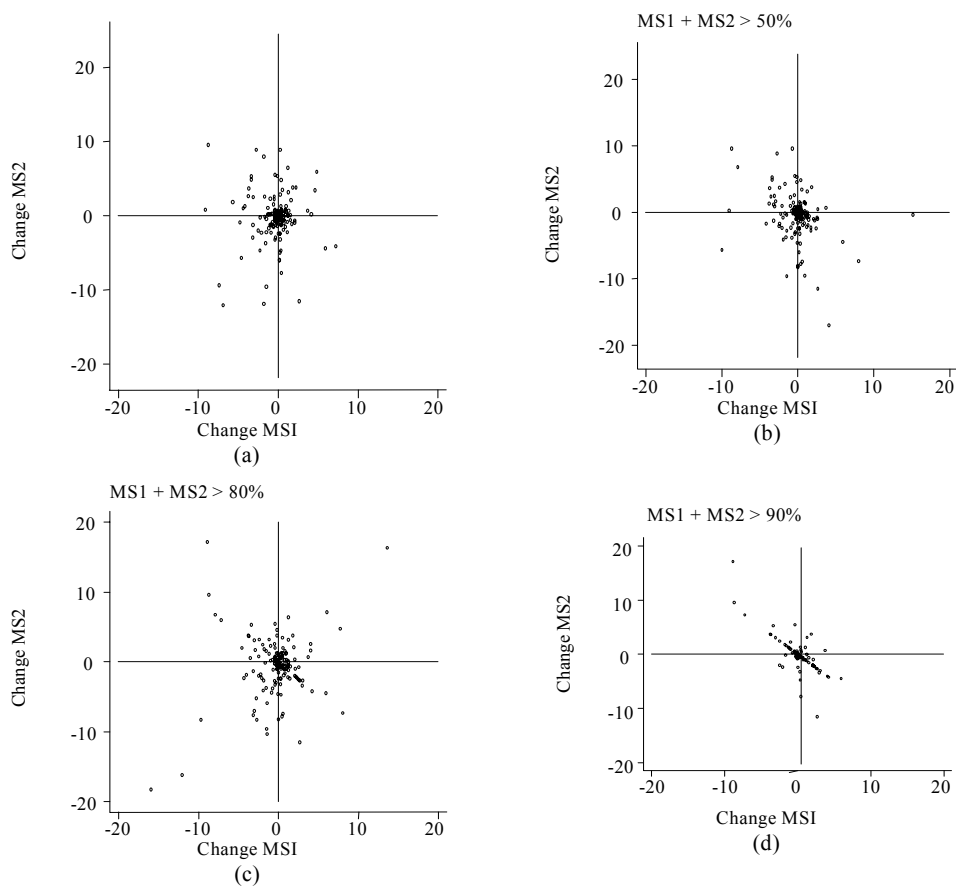



Figure 2. The annual change in market share for the top ranking firm (horizontal axis) versus the change for the second ranking firm (vertical axis). (The two firms are those ranked 1 and 2 in year 5 of the dataset).

Panel (a) shows the data for all 45 industries, while panels (b), (c) and (d) show data for those industries in which the combined market share of the top 2 firms in year 5 exceeded 50%, 80% and 90% respectively.

d. Given the scaling property, it is natural to begin the investigation of passage times by focussing attention on the top two firms, viz: we take some reference date $t_0 = 0$ and label firms in descending order of market share at that date. We now examine the first date at which the market share of firm 2 exceeds that of firm 1 (labelled t_{12} in what follows).

The advantage of beginning with an analysis of t_{12} is two-fold. First, we may take advantage of the theoretical results developed above to reduce the problem to the study of a simple random walk, thus allowing us to draw some standard results for passage times. Second, this allows us to place an upper bound on the passage time for loss of leadership.¹¹

Figure 3: The Scaling Relationship



¹¹ It might seem natural to begin by checking whether changes in the market share gap between the two firms exhibit any (positive or negative) serial correlation over successive years. A series of checks, using different time periods (lags), indicated no significant correlation of this kind. This, however, does not exclude more subtle forms of departure from the null hypothesis explored here, as was noted in footnote 4 above.




Figure 4. The form of the distribution of $\Delta\sqrt{m_1} / \sqrt{m_1}$

6. Passage Times I

This section looks at two methods of testing in relation to the passage time t_{12} . The theoretical model of Section 3 leads to an approximate representation in which $\sqrt{m_1} - \sqrt{m_2}$ can be modelled as a simple random walk in which the jump in value from t to $t+1$ is represented by a draw from some distribution $g \circ g$. To control for the different levels of market share volatility across industries, we represent the gap as $(\sqrt{m_1} - \sqrt{m_2}) / \sigma$, where σ is an industry specific volatility parameter.

Our way of predicting the distribution of the passage time t_{12} , therefore, would be to begin by estimating the volatility parameter $\sigma(\sqrt{m_i})$ for each industry, and using this

estimate to ‘normalize’ the size of the initial gap $(\sqrt{m_1} - \sqrt{m_2})$. To do this, we use the data for some initial period $[0, t_0]$ to estimate $\sigma(\sqrt{m_i})$; take the gap at the end of this period as the ‘initial gap’, and predict the probability for each industry of a crossing t_{12} during the remaining time period $[t_0, T]$. Finally, these probabilities are summed across industries to obtain a predicted number of crossings in these 45 industries over the period $[t_0, T]$.

The disadvantages of this method are:

- i. It ‘uses up’ several years of data in estimating $\sigma(\Delta\sqrt{m_i})$;
- ii. Even if σ is constant over the entire period, the precision of the estimate obtained from some short initial period may be low.
- iii. It depends upon the empirically estimated form of $f(\Delta\sqrt{m_i})$. While a sharp characterization of this distribution is possible, a difficulty lies in specifying the range of observations, i.e. the tails of the distribution. Results may be sensitive to the specification used here, and this involves some arbitrariness.

This method was implemented using the first 5 years of data to estimate the industry specific volatility parameter σ . However in view of problem (iii), the results are not reported here (they are broadly consistent with the results reported below). Instead, two alternative methods are used, as follows:

The first method of testing takes advantage of the properties of the simple random walk, in order to avoid the need to estimate the volatility parameter. This property is as follows: let t_0 be defined as the first date at which m_1 crosses m_2 , so that the gap at this date

equals (approximately) zero. Now take the interval from t_0 onwards. Divide this into two equal sub-intervals, the second sub-interval being $[t_0 + (T-1-t_0)/2, T]$ or $[t_0 + (T-t_0)/2, T-1]$ according as t_0 is odd or even respectively.

If the gap $\sqrt{\Delta m_1} - \sqrt{\Delta m_2}$ follows a random walk, then the probability that a crossing t_{12} occurs during the second sub-interval equals $\frac{1}{2}$. This result holds independently of the distribution of shocks, and so of the volatility parameter. The intuition is as follows: if we make the process more volatile, then the probability that the gap drifts upwards (or downwards) by a large amount during the first subinterval rises; but its probability of returning thereafter from a distant value increases to the same degree.

We can interpret the null hypothesis being investigated here in terms of the circular road model of section 3: we infer from the equality of market shares at rate t_0 that the two firms have an equal number of products at that date (i.e. each has the same expected number of products conditional on its observed share, viz, it owns a fraction m_i of the products). The dynamics of market share then follows a random walk, with no (positive or negative) drift.

Under this null hypothesis, it follows that, when we observe a crossing at time t_0 , the probability that we will observe a crossing in the ‘second sub-interval’ defined above is $\frac{1}{2}$. What if the hypothesis fails? Say, for example, that firm 1 had some underlying ‘capability’ superior to that of firm 2, and that this will make it (more) likely that firm 1 will pull ahead, and stay ahead, of firm 2 in the future. The presence of such a bias in favour of firm 1 would cause the stochastic process describing the market share gap to

exhibit positive drift. This is consistent with the presence of (a reduced number of) crossings of the firms' market shares, but the observation of a crossing does not imply an equality between the firms' future fortunes; the expected number of crossings in the second sub-interval is now less than $1/2$.¹²

The main disadvantage of this test is that it requires us to discard all industries in which no crossing occurs prior to the last two years of the data; this leaves us with only 18 industries out of 45. The number of second period crossings in the set of 18 industries equals 5, as compared with an expected level of 9 under the null hypothesis.¹³ The null hypothesis is rejected at the 5% level (one-tail test, see footnote 12).

This suggests that, once a firm moves ahead of its rival(s), the degree of persistence of leadership may be greater than predicted under the null.

Turning to the second method, we again appeal to the scaling property to justify the pooling of all observations of $\Delta\sqrt{m_1} - \Delta\sqrt{m_2}$ for each industry for the full period $[1, T]$, and the modelling of the evolution of $\sqrt{m_1} - \sqrt{m_2}$ as a simple random walk. Now, however, we use the set of pooled observations of $\Delta\sqrt{m_2}$ for all firms and for all periods, within each industry, to predict the distribution of the first passage time for that

¹² To see what is involved here, it is useful to ask what analogous argument would be for the gap in scores in a basketball game. (I am grateful to Barry Nalebuff for suggesting this analogy.) At time 0, the teams' scores are equal, but the abilities of the teams will, in general, differ. Only if abilities (scoring possibilities) are equal, does the present model apply. If abilities differ, the gap in scores follows a random walk with (positive or negative) drift, and while scores may at some time(s) coincide, the probabilities of a second period crossing is less than $1/2$.

¹³ A deviation of 4 or more from the expected level of 9 will occur with probability 4.8% under the null hypothesis. Thus we just fail to reject the null at the 5% level (2-tail test).

industry. We then sum the probability of observing a crossing in each industry by time T in order to arrive at the expected number of crossings over our full set of industries.¹⁴

The advantage of this procedure is that it avoids the need to fit some distribution to the observations of market share changes, and it allows us to focus directly on the issue of interest, i.e. whether the evolution of the market share gap exhibits some subtle form of behaviour that distinguishes it from our Markovian benchmark.

The results of the procedure are shown in Table 1 and Figure 5. The predicted values, and the associated 95% confidence interval, are obtained by simulating the random walk (Monte Carlo estimates).

The results are shown for a series of alternative starting dates (year 0, 5, 10 and 15, respectively).¹⁵ The observed values show a tendency to lie below predicted values. They lie within a 95% confidence interval in two of the four cases, while falling outside in two.

¹⁴ This procedure is equivalent, under our independence assumption, to modelling $\Delta\sqrt{m_1} - \Delta\sqrt{m_2}$ as a simple random walk.

¹⁵ The results predicted values for these alternative starting dates were obtained by drawing market share shocks from the pooled sample for all periods. The procedure was repeated using the sample of shocks for the corresponding time period (period 6 to final period, etc.). The results were closely similar to those shown in Figure 5.

Years after initial year (year 5)	Actual number of crossings	Predicted number of crossings	
		Mean	95% Confidence Interval
1	4	2.6	0.2 - 6.0
2	4	4.5	1.0 - 8.3
3	6	5.9	2.0 - 9.2
4	8	7.1	2.9 - 10.7
5	9	8.2	3.7 - 11.9
6	9	9.1	4.5 - 13.0
7	9	10.0	5.2 - 14.0
8	10	10.8	6.0 - 14.9
9	10	11.5	6.5 - 15.7
10	10	12.1	7.1 - 16.4
11	11	12.7	7.6 - 17.0
12	11	13.3	8.2 - 17.6
13	11	13.8	8.6 - 18.0
14	11	14.3	9.1 - 18.6
15	11	14.7	9.4 - 18.9

Table 1 Actual and Predicted Crossing Times for the Two Leading Firms, over 45 Industries.


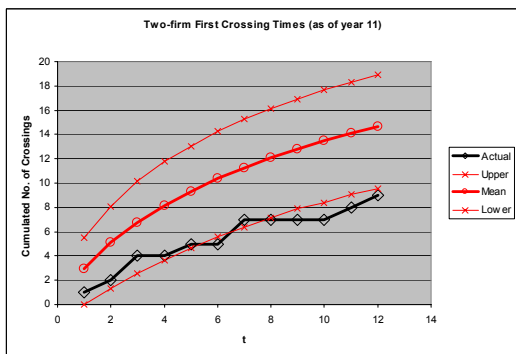

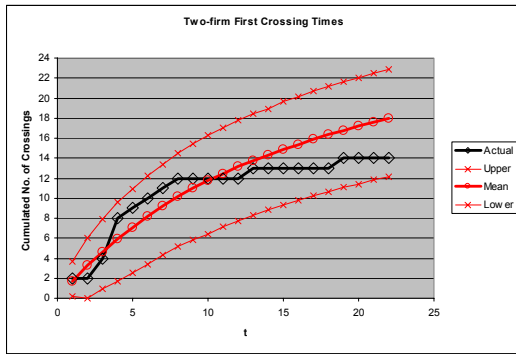


Figure 5 The top panel shows the number of crossing of the leading firm by its closest rival in the initial period, by elapsed time t that occur within the 45 industries. The second, third and fourth panels repeat this exercise, beginning from the sixth, eleventh and sixteenth year of the 23-year run. (The ‘closest rival’ is defined as the second largest firm in the corresponding initial year). The expected number of crossings, and the confidence interval, for our benchmark process are also shown.

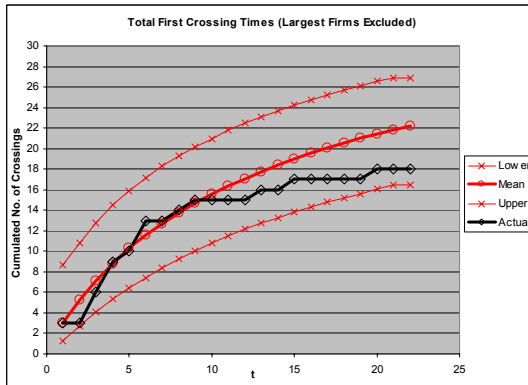
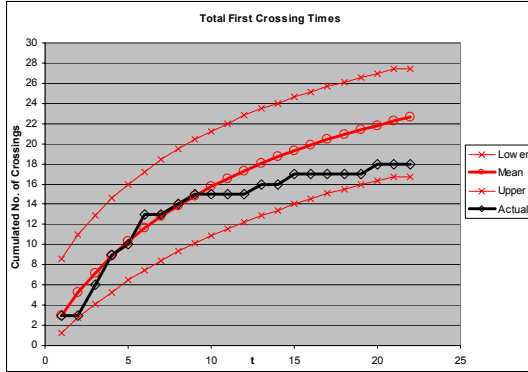


Figure 6 The top panel shows the number of crossings of the leading firm by any rival, by elapsed time t , that occur within the 45 industries. The lower panel shows the results for the 41-industry dataset in which 4 ‘highly concentrated’ industries are omitted.

7. Passage Times II

The analysis can be extended from the study of crossing times of the form t_{12} to the study of general crossing times (i.e. the initial leader is overtaken by any firm). Here, the reduction of the problem to one in which the distribution of crossing times is determined by a single number (‘normalized gap’) as no longer possible. We need instead to specify the full vector of initial market shares, and the volatility parameter for the industry. Monte Carlo estimates were constructed in this way, which specify the probability of any crossing during the interval $[1, 22]$, and the probabilities were summed, as before, over

industries to obtain an expected number of crossings. The results are shown in the top panel of Figure 6. Most crossings are, in practice, made by the second largest firm in the industry, and the results shown in Figure 6 are closely similar to those shown in the corresponding (top) panel of Figure 5.

It was noted above that the assumption of independence underlying the null hypothesis is invalid for the four most highly concentrated industries in the dataset. With this in mind, the exercise was repeated using the remaining 41 industries only. The results are shown in the lower panel of Figure 6, and are closely similar to those shown in the upper panel. (No crossings occur in these 4 industries, and they all feature a large initial gap in shares, and low volatility, so that the expected number of crossings under the null hypothesis is close to zero).

The overall conclusion is that there appears to be a tendency for fewer crossings (longer persistence of leadership) than predicted by the benchmark model. This observation raises an obvious question: could this tendency to be driven by some systematic (strategic) mechanism that operates across the general run of industries? Is there some ‘Chandlerian’ mechanism at work, for example, which could be interpreted by saying that current annual market shares are not a ‘sufficient statistic’ for the (superior) level of capability employed by (leading) firms? Or, on the other hand, does this tendency merely represent the overall average behaviour of a series of industries, each driven by its own idiosyncratic features? To investigate these questions, we turn to some case studies of industries in the sample.

8. Digging Deeper

A central argument of the present paper is that any simple stochastic model can easily be bettered as a representation of any one industry, by incorporating industry-specific features which will include a strategic representation of firms' competitive responses to market share changes. Once we aim at constructing a 'richer' model of this kind, however, we meet the problem that 'strategic effects' will turn on various features, some of which are intrinsically 'unobservable' as far as the outside economist is concerned.

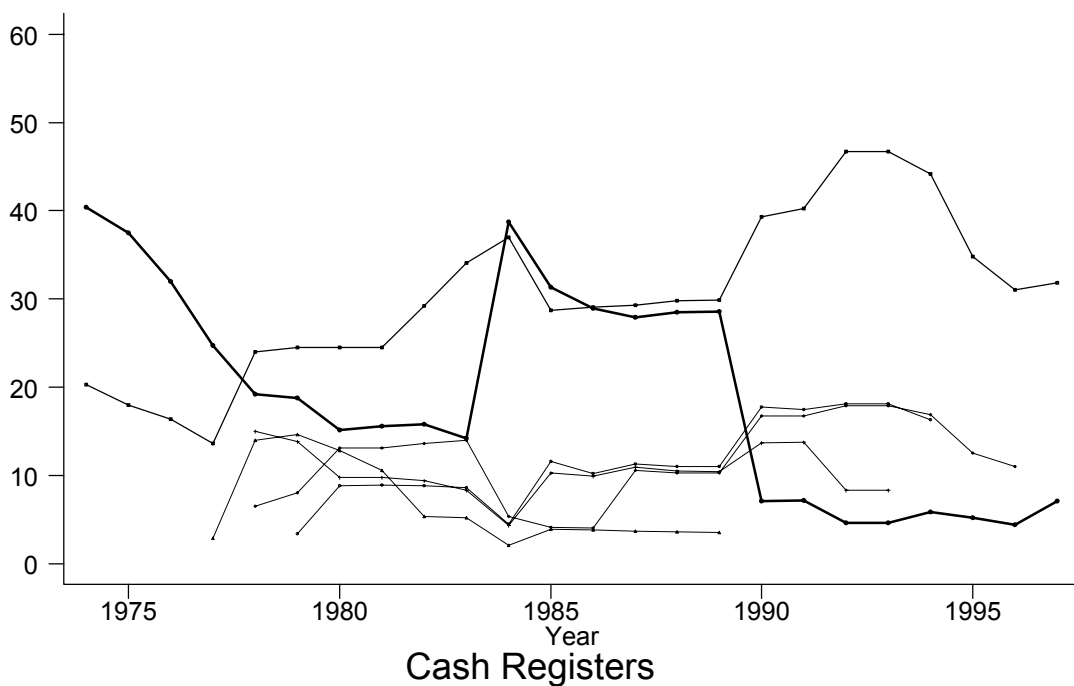
A more sophisticated model would retain exogenous shocks to underlying 'technology and tastes' parameters, but would extend firm's reactions beyond the price-quantity adjustments allowed for above, to deal with changes in marketing and/or R&D outlays aimed at raising (perceived) quality, and with the entry and exit of products. It is in respect of these latter adjustments that subtle differences appear across different industries, which seem to be driven by various factors, some of which are very difficult to measure, proxy or control for in empirical studies. Most importantly in the present context, they include the beliefs of agents as to their rivals' private information, and strategic responses.

In considering a model that allows for strategic responses on the marketing or R&D side, the key question of interest is: how do these strategic responses impinge on the degree of volatility of market shares, and on the evolution of shares over time?

One obvious factor that might impinge on market share volatility is the degree to which existing products are displaced by new products. We begin with two industries in which

the rate of product displacement was very high, but in which market share patterns were very different.

In the cash register industry, the technology changed continuously and dramatically over the 25-year period, as free-standing electromagnetic registers were first replaced by electronic types, and as these electronic types were in turn displaced by store-wide or company-wide computer




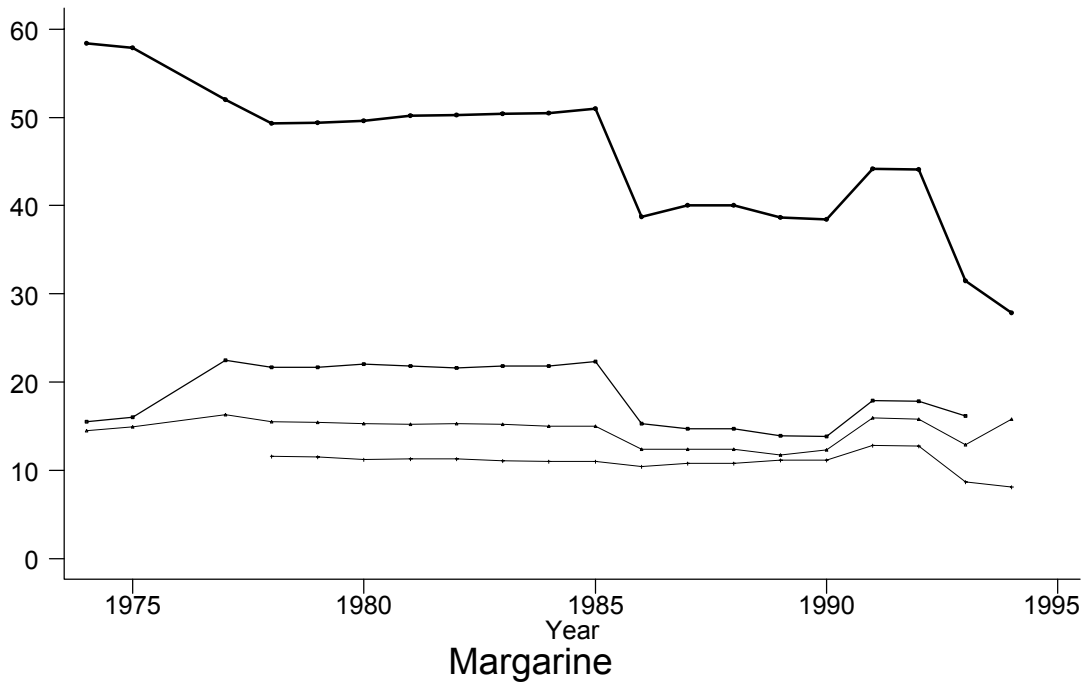


Figure 4(a). Market Shares

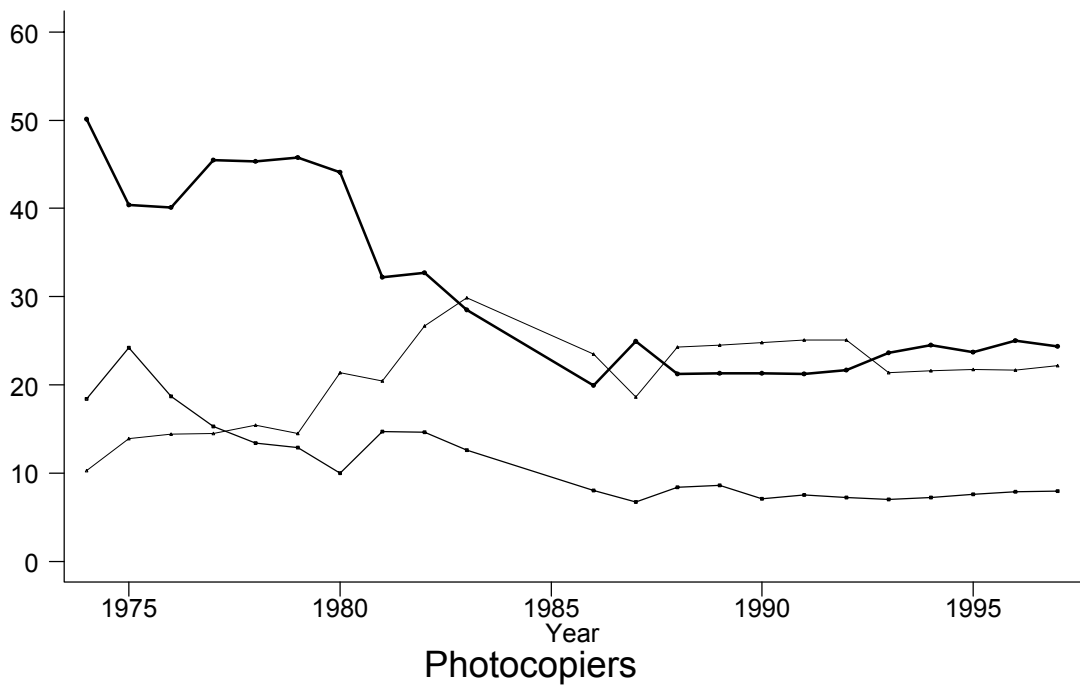


Figure 4(b). Market Shares




Figure 4(c). Market Shares

linked networks. The market share pattern was extremely volatile, as successive firms gained a relative technical advantage. (Figure 4(a))

A contrasting pattern arises in the margarine industry. Here, the industry was characterized (perhaps surprisingly) by a very rapid rate of introduction of new varieties (one manufacturer's 1990 brochure contained scores of varieties, which differed in form of packaging, choice of flavouring, hardness and texture, etc.). In spite of the high degree of new product introductions, market shares remained remarkably stable, as each successful innovation by any firm was immediately countered with a response by rivals, who quickly imitated successful products.

One interpretation of the different outcomes in these two industries lies in the different strategic responses of firms to rivals' successes. This difference might, however, simply reflect differences in the ease with which innovations can be imitated by rivals. To explore this latter idea, it is interesting to look at the evolution of market shares in the beer industry. Here, there are two periods of interest. The late 1970s was marked by what came to be known in the industry as the 'packaging wars'. Firms vied with each other in introducing new forms of packaging (bottles and cans of new sizes; plastic containers in odd and unusual shapes, and so on). Throughout this period, market shares remained quite stable. The second period is the 1980s, when a beer was marketed by the Asahi company, then the industry's fourth largest firm, under the name 'Asahi Dry'. Despite its initial success, rivals were slow to respond, and 'Asahi Dry' propelled the Asahi company to second place in the industry. (The market leader Kirin eventually imitated this strategy by marketing its 'Kirin Dry' product, whose sales remained below those of 'Asahi Dry' over the next decade). The question raised by this is: if we constructed a 'fully specified, strategic model' of the Japanese beer market, what variables accessible to the researcher could have predicted the non-impact of the packaging wars, as against the substantial impact of the 'Dry beers' marketing campaign? It would seem that the speed and effectiveness of rivals' responses differed in the two cases because of different beliefs on the part of rivals' as to the probable effectiveness of the innovator's strategy. What this suggests is that, just as the literature on dynamic oligopoly suggests, the size of the market share response, and so the level of market share volatility in the industry, will depend inter alia on the beliefs of agents – a factor that we must perforce treat as an unobservable in most settings.

What I want to suggest, then, is that while differences in firms' strategic responses to rivals' actions may be one of the factors that account for the different patterns of evolution of shares in different industries, these differences in strategic responses will depend delicately on factors that are difficult to account for on the basis of stable and observable 'industry characteristics'. To illustrate this point, it is of interest to consider two industries which display quite different patterns of market share dynamics at different periods. The photocopier industry, for example, offers an example in which a technically innovative follower (the third firm in 1974) gradually overtakes the leader; but the industry then moves to a stable setting for 15 years, with the innovator sharing first place with the original market leader. (Figure 4(b)). In the colour film market, there are again two phases; over the decade, the market leader's share rises at the expense of its nearest rival, but then shares stabilize and show little volatility over the next 15 years. (Figure 4(c)).

In contrast to these two cases, consider the market for pocket calculators, an industry that was marked, like the photocopiers and colour film industries, by a fast rate of technical innovation and new product introductions. (Figure 4(c)). Here, the top two firms escalated their innovatory activities and both their shares increased steadily, and in step, as the shares of all the smaller firms declined. It is easy to specify a suitable game theoretic model which has these features (see for example, Sutton (1991), Chapter 5); but it is not clear what observable 'industry characteristics' for 1974 could have predicted that this industry's market share pattern would have differed in this way from those seen in photocopiers or colour film – beyond attributing it to different 'stochastic realizations' of outcomes to the different firms' early innovatory efforts, and arguing that the initial 'accidental' successes induced different strategic choices thereafter.

These examples, taken together, suggest a serious caveat regarding the traditional ‘persistence of leadership’ debate. That debate has been conducted on the premise that there might be some general mechanism(s), either of a ‘Schumpeterian’ or ‘Chandlerian’ kind, that operate(s) across the general run of industries. What these examples suggest is that there are many mechanisms, some operating in one direction, others in another direction, so that when we test for some ‘bias’ in either direction we are (at best) assessing some kind of average outcome that will be highly sensitive to our selection of industries. It is in this (rather cautious) spirit that any conclusions as to a possible ‘Chandlerian’ bias in the present set of Japanese industries should be drawn.

9. Conclusions

This paper makes three points. The first relates to the limitations of game theoretic (strategic) models. The second relates to the use of scaling relationships for the variance of firm growth rates, and market shares. The third relates to the ‘persistence of leadership’ debate.

a. Game Theoretic Models

I have argued elsewhere (Sutton (1991, 1998)) that game theoretic models can lead us to a small number of robust predictions, which allow us to place limited restrictions (bounds) on market structure. Beyond these few robust results, however, outcomes will depend delicately on factors that are difficult to measure, proxy or control for in cross-

industry studies. In this setting, it can be of interest to examine ‘low-level’ representations of the data, of the kind attempted here.

b. Scaling Relationships

The recent literature regarding scaling relationships on firm growth rates, has focussed on the description of relationships, and on differences in views as to candidate explanations for such relationships (Stanley et al. (1996), Sutton (2002)). Little attention has been paid to the question of whether the characterization of such relationships is empirically useful. In this paper, I have argued that the characterization of a simple scaling relationship between a firm’s market share and the variance of changes in market share, permits a useful simplification in the description of market share dynamics, allowing the (limited) crossing-time problem addressed here to be reduced to the study of a simple random walk.

This scaling relationship also provides a useful criteria for model selection in the area of market share dynamics, as it is a feature of only one of the several standard models in the current economics literature.¹⁶

c. The persistence of ‘leadership’ problem

The claim of this paper in regard to the ‘persistence of leadership’ debate is a modest one.

We explore the properties of a benchmark model in which there is no systematic bias,

¹⁶ It is worth noting, however, that it is also consistent with certain models, popular in the marketing literature, in which individual consumers are attached to firms over successive periods, but where each consumer has a (small) probability of shifting allegiance to a new firm in any period. These models are purely statistical in nature, and do not rest on optimising by agents. For a model of this type based on maximizing behaviour by firms and consumers, see Sutton (1980).

whether of a ‘Chandlerian’ or ‘Schumpeterian’ kind, present. Such a model is consistent with a degree of persistence of leadership that might seem to be rather ‘long’ on the basis of intuition (Feller (1950)). The empirical evidence for the Japanese industries examined here is such as to suggest a degree of persistence that is somewhat in excess of that predicted under this benchmark model. This raises the question of whether there is any single systematic influence at work in driving this persistence. An examination of the drivers of market share changes in different industries, suggests that there is no single, systematic effect at work here. Rather, what we seem to be observing is an average over this sample of industries, of various industry specific mechanisms. Thus, we might or might not find that this pattern persists over a different or broader sample of Japanese manufacturing industries. This paper makes no claims in this regard. It would nonetheless be of some interest to see whether a similar pattern holds good for the general run of manufacturing industries in other countries.¹⁷

¹⁷ It might strike the reader that extending the data set to more industries and/or waiting for more data to accumulate for the present Japanese dataset might allow a sharper conclusion to be drawn. This was the author’s original reaction, five years ago; but now having expanded the range of industries to all those for which satisfactory data can be assembled and having added 5 years of data, I have decided to report the results at this point.

References

- Chandler, Alfred D. Jr. (1990), Scale and Scope: The Dynamics of Industrial Capitalism, Cambridge, Mass.: Harvard University Press.
- Ericson, Richard and Ariel Pakes, 1995, "Markov-Perfect Industry Dynamics: A framework for Empirical Work; Review of Economic Studies, vol. 62, pp. 53-82.
- Feller, William (1950), An Introduction to Probability Theory and its Applications, Vol. 1, (3rd ed.), New York, Wiley
- Fisher, F. M., J. J. McGowan and J. E. Greenwood, (1983), Folded, Spindled and Mutilated: Economic Analysis and U.S. v. IBM, Cambridge, Mass.: MIT Press.
- Harris, Christopher (1994), 'Dynamic Models of Competition', Review of Economic Studies Lecture, 1995, unpublished working paper, Nuffield College, Oxford.
- Little, Ian, (1962) 'Higgledy Piggledy Growth', Bulletin of the Oxford Institute of Statistics, vol. 24.
- Little, I. M. D. and Anthony C. Rayner, (1966) Higgledy Piggledy Growth Again: An Investigation of the Predictability of Company earnings and Dividends in the U.K. 1951-1961, Cambridge, M.A.: Harvard University Press.
- Mueller, D. (1986), Profits in the Long Run, Cambridge, Cambridge University Press.
- Stanley, R. Michael, L. A. N. Amaral, S. V. Buldyrev, S. Harlin, H. Leschorn, P. Maas, M. A. Salinger, H. E. Stanley (1996), 'Scaling Relationships in the Growth of Companies', Nature, vol. 319, pp. 804-806.
- Sutton, J. (1980), 'A Model of Stochastic Equilibrium in a Quasicompetitive Industry', Review of Economic Studies, vol. 47, pp. 705-722.
- Sutton, J. (1991), Sunk Costs and Market Structure, Cambridge, MIT Press.
- Sutton, J. (1998), Technology and Market Structure, Cambridge, MIT Press.
- Sutton, J. (2002), 'The Variance of Firm Growth Rates: The Scaling Puzzle', Physica A, vol. 313, pp. 577-590..

Appendix 1

Calculating the impact of a quality shock

We may take advantage of the fact that the system of equations (4) in the main text is linear in the p_j and u_j to deduce that a unit change in u_i will affect equilibrium prices $p_i, p_{i+1}, p_{i-1}, p_{i+2}, p_{i-2}, \dots$ by a constant amount, independently of the initial vector of qualities. Hence we may ease the notional burden in what follows by taking as a point of reference the case where all the u_j are initially zero, and all prices are equal to t/N . We now consider the impact on equilibrium prices of a unit rise in the quality of some one good holding all other qualities constant.

We will confine analysis in what follows to the case where the total number of products is odd (the even case can be treated similarly). Label the good whose quality has risen as good 0, its k -th neighbour to the right as good k , and k -th neighbour to the left as good $-k$. Denote the total number of products by $2n+1$; we then have that the index k runs from 1 to n , and good n has right hand neighbour $-n$. We denote the deviation of the quality-adjusted price $(p_j - u_j)$ from its initial level t/N by x_j , viz $x_j = \Delta(p_j - u_j)$. We have from the symmetry of equation (4) that the equilibrium price deviations satisfy $x_j = x_{-j}$ for all $j = 1, 2, \dots, n$. It therefore follows from (4) on writing $\Delta u_0 = 1, \Delta u_j = 0$ for all $j = 1, 2, \dots, n$ that the deviations in quality-adjusted prices x_j must satisfy the equations:

$$x_0 = \frac{1}{2} + \frac{1}{2}x_1 \tag{A1}$$

$$x_i = \frac{1}{4}x_{i-1} + \frac{1}{4}x_{i+1} \quad i = \pm(n-1) \tag{A2}$$

$$x_n = \frac{1}{4}x_{n-1} + \frac{1}{4}x_{-(n-1)} \quad (\text{A3})$$

with the convention that $x_{-n} \equiv x_n$ (Figure A1). Note that the x_j correspond to price changes for goods $\pm 1, \pm 2, \dots, n$; but for good 0, whose quality has risen by 1 unit, the change in price equals $1 - x_0$.




Figure A1: Labelling the Products

Now from symmetry, $x_{n-1} = x_{-(n-1)}$ so (A3) implies that

$$x_n = \frac{1}{2}x_{n-1} \quad (\text{A4})$$

while (A2) implies

$$4x_{n-1} = x_n + x_{n-2} \quad (\text{A5})$$

Using (A4) to substitute for x_n in (A5) and solving we have

$$x_{n-1} = \frac{1}{4 - \frac{1}{2}} x_{n-2} \quad (\text{A6})$$

We may now proceed iteratively

$$4x_{n-2} = x_{n-1} + x_{n-3}$$

Using (A6) to substitute for x_{n-1} we obtain

$$x_{n-2} = \frac{1}{4 - \frac{1}{4 - \frac{1}{2}}} x_{n-3} \quad (\text{A7})$$

For any given n , we may solve for the x_i by combining the relation between x_0 and x_1 derived in this manner with (A1). Take for example the case $n = 3$.

There are 6 products labelled $0, \pm 1, \pm 2, 3$. Equation (A7) implies that

$$x_1 = \frac{1}{4 - \frac{1}{4 - \frac{1}{2}}} x_0 = \frac{7}{26} x_0$$

while from (A1) we have

$$x_0 = \frac{1}{2} + \frac{1}{2} x_1$$

whence using (A7), (A6) and (A4) we obtain

$$x_0 = -\frac{26}{45} \left(\text{whence } \Delta p_i = \frac{19}{45} \right)$$

$$x_2 = \frac{2}{7}x_1 = -\frac{2}{45}$$

$$x_3 = \frac{1}{2}x_2 = -\frac{1}{45}$$

In the limit $n \rightarrow \infty$, the recursion relation illustrated by (A6), (A7) above becomes

$$x_{n-i} = \frac{1}{4 - \frac{1}{4 - \dots}} x_{n-i-1}$$

Solving for the repeated fraction we obtain

$$x_{n-i} = (2 - \sqrt{3})x_{n-i-1} \quad (\text{A8})$$

so that the changes decline geometrically as we move away from x_0 , each change being about one-third as big as its ('upper') neighbour.

Setting $i = n - 1$ in we have

$$x_1 = (2 - \sqrt{3})x_0 \quad (\text{A8})$$

Combining this with (A1) we obtain:

$$x_0 = -\frac{1}{\sqrt{3}}, \quad x_1 = -\frac{2-\sqrt{3}}{\sqrt{3}}, \quad \dots, \quad x_k = -\frac{(2-\sqrt{3})^k}{\sqrt{3}}, \dots$$

We may interpret this intuitively as follows. Recall that $x_0 = \Delta(p_0 - u_0) = \Delta p_0 - 1$. As the quality of product zero rises by 1 unit, its price rises by $1 - \frac{1}{\sqrt{3}}$ units, so that its quality-adjusted price falls by $\frac{1}{\sqrt{3}}$ units. There is a fall in the prices of all other products, the size of this change falling off geometrically as we move away from good zero.

To find the changes in quantities, we note that it follows from inspection of the demand function (equation (4) of the main text) that

$$\Delta q_0 = x_1 - x_0 = \Delta p_1 - \Delta(p_0 - u_0)$$

$$\Delta q_{-1} = \Delta q_1 = \frac{1}{2}[x_0 + x_2 - 2x_1] = \frac{1}{2}[\Delta(p_0 - u_0) + \Delta p_2 - 2\Delta p_1]$$

$$\Delta q_{-i} = \Delta q_i = \frac{1}{2}[x_{i+1} + x_{i-1} - 2x_i] = \frac{1}{2}[\Delta p_{i-1} + \Delta p_{i+1} - 2\Delta p_i] \quad i \geq 2$$

whence we obtain, on substituting for the x_j s, that

$$\Delta q_0 = -\frac{2-\sqrt{3}}{\sqrt{3}} + \frac{1}{\sqrt{3}} = 1 - \frac{1}{\sqrt{3}}$$

$$\Delta q_{-i} = \Delta q_i = -\frac{1}{\sqrt{3}}(2 - \sqrt{3})^i \text{ for } i \geq 1$$

It follows that the sales, and so the (volume) market share of good 0 rises, while that of all other goods fall, the size of the fall decreasing geometrically as we move away from good 0.

Appendix 2

Industries in the Dataset

- 1 Sugar
- 2 Frozen Food
- 3 Regular Coffee
- 4 Instant Coffee
- 5 Chocolate
- 6 Chewing Gum
- 7 Cola
- 8 Beer
- 9 Womens Clothing
- 10 Adhesives
- 11 Bath Soap
- 12 Toothpaste
- 13 Car Tyres/Tubes
- 14 Elevators
- 15 Escalators
- 16 Tin Cans
- 17 Gas Stoves
- 18 Oil Stoves
- 19 Airconditioner (Window)
- 20 Airconditioner (Package)
- 21 Cash Registers
- 22 English Typewriters
- 23 Pocket Calculators
- 24 Photocopiers
- 25 Refrigerators
- 26 Washing Machines
- 27 Vacuum Cleaners
- 28 Colour TVs
- 29 Cars
- 30 Buses
- 31 Trucks
- 32 Motorcycles
- 33 Optical Measuring Equipment
- 34 Analytical Equipment
- 35 Length and Precision Measuring Equipment
- 36 Electric Meters
- 37 Gas Meters
- 38 Water Meters
- 39 35mm Cameras
- 40 Spare Lenses for Cameras
- 41 Black and White Film
- 42 Colour Film
- 43 Pencils
- 44 Fountain Pens
- 45 Ball Point Pens